$\mathbb{F}_p$ IS LOCALLY LIKE $\mathbb{C}$

CODRUŢ GROSU

ABSTRACT. Vu, Wood and Wood showed that any finite set S in a characteristic zero integral domain can be mapped to $\mathbb{F}_p$, for infinitely many primes p, while preserving finitely many algebraic incidences of S. In this note we show that the converse essentially holds, namely any small subset of $\mathbb{F}_p$ can be mapped to some finite algebraic extension of $\mathbb{Q}$, while preserving bounded algebraic relations. This answers a question of Vu, Wood and Wood. We give several applications, in particular we show that for small subsets of $\mathbb{F}_p$, the Szemerédi-Trotter theorem holds with optimal exponent 4/3, and we improve the previously best-known sum-product estimate in $\mathbb{F}_p$. We also give an application to an old question of Rényi. The proof of the main result is an application of elimination theory and is similar in spirit with the proof of the quantitative Hilbert Nullstellensatz.

1. INTRODUCTION

Suppose p is a prime and N a positive integer. In what follows $\mathbb{Z}_N$ denotes the additive group of integers modulo N, $\mathbb{F}_p$ the field with p elements and $\mathbb{Z}_{(p)}$ the localization of $\mathbb{Z}$ at (p), which is the same as the ring of fractions with denominator not divisible by p.

Let $k \geq 1$ be an integer, Z and W two abelian groups and $A \subseteq Z, B \subseteq W$ finite subsets. A bijection $\phi : A \to B$ is a *Freiman isomorphism of order k*, or simply F_k-isomorphism, if for any $a_1, \ldots, a_{2k} \in A$ we have

$$a_1 + \ldots + a_k = a_{k+1} + \ldots + a_{2k}$$

if and only if

$$\phi(a_1) + \ldots + \phi(a_k) = \phi(a_{k+1}) + \ldots + \phi(a_{2k}).$$

From the definition it follows that any F_{k+1}-isomorphism is also an F_k-isomorphism, and furthermore translation does not affect the isomorphism. An important property in additive combinatorics is that any finite subset of a torsion-free group is F_k-isomorphic to a subset of $\mathbb{Z}_N$, for any large enough N. This helps reducing general additive problems to $\mathbb{Z}_p$, where more powerful techniques, such as Fourier analysis, are available.

In the other direction it is well-known that small subsets of $\mathbb{Z}_p$, with p prime, are Freiman isomorphic to subsets of $\mathbb{Z}$.

Theorem 1 ([2]). *Let $A \subseteq \mathbb{Z}_p$, where p is a prime. If $|A| \leq \log_{2k} p$, then there exists a set of integers $A' \subset \mathbb{Z}$ such that the canonical homomorphism $\mathbb{Z} \to \mathbb{Z}_p$ induces an F_k-isomorphism of A' onto A.*

The result holds for $|A| \leq \log_{2k} p + \log_{2k} \log_{2k} p$ as well. In [2] it is also shown the existence of a set $A \subset \mathbb{Z}_p$ of cardinality at most $2 \log_k p + 1$ which is not F_k-isomorphic to any set of

integers. Assuming A has small doubling constant allows the theorem to hold for $|A| \leq cp$, for some $c > 0$. This is the *Freiman rectification principle* (see [2], [16]).

It is now a natural question if it is possible to preserve both the additive and multiplicative structure. In this direction we have the following result of Vu, Wood and Wood.

Theorem 2 ([37]). *Let S be a finite subset of a characteristic zero integral domain D, and let L be a finite set of non-zero elements in the subring $\mathbb{Z}[S]$ of D. There exists an infinite sequence of primes with positive relative density such that for each prime p in the sequence, there is a ring homomorphism $\phi_p : \mathbb{Z}[S] \to \mathbb{F}_p$ satisfying $0 \notin \phi_p(L)$.*

Here $\mathbb{Z}[S]$ is the smallest subring of D containing S.

It was asked by Vu, Wood and Wood [37] whether given a small enough set $A \subseteq \mathbb{F}_p$, it is possible to map A to some characteristic zero integral domain, while preserving algebraic incidences.

Let us first make a few observations. One could not always map A to $\mathbb{Z}$, as we may need, for example, to preserve identities of the form $y^2 + z^2 = 0$ with $y, z \neq 0$ for some $y, z \in A$. Also, we should allow only "bounded" algebraic incidences, as any identity of the form $py = 0$ with non-zero $y \in A$ can not be mapped in any characteristic zero integral domain. Therefore the following definitions make sense.

Let $k, t > 0$. A polynomial $f \in \mathbb{Z}[x_1, \ldots, x_n]$ is called (k, t)-bounded if $\|f\|_1 \leq k$, and its degree is at most t. Here $\|f\|_1$ represents the sum of the absolute values of the coefficients of f, and similarly we define $\|f\|_\infty$ to be the maximum of the absolute values of the coefficients of f. When we evaluate f at a point $(a_1, \ldots, a_n)$ with $a_i \in R$ for some ring R, all operations are carried in R, in the natural way. If $k = t$, we simply call f k-bounded. If $t = 1$, we say f is a k-bounded linear polynomial.

Now let R_1, R_2 be two rings and $A \subseteq R_1, B \subseteq R_2$ finite subsets. We call a bijection $\phi : A \to B$ a *Freiman ring-isomorphism of order k*, or simply F_k-ring-isomorphism, if $A = \{a_1, \ldots, a_n\}$ and for any k-bounded $f \in \mathbb{Z}[x_1, \ldots, x_n]$ we have

$$f(a_1, \ldots, a_n) = 0$$

if and only if

$$f(\phi(a_1), \ldots, \phi(a_n)) = 0.$$

Our main result is the following.

Theorem 3. *Let $k \geq 2$ be an integer, p be a prime and $A \subseteq \mathbb{F}_p$. If $|A| < \log_2 \log_{2k} \log_{2k^2} p - 1$ then there exists a finite algebraic extension K of $\mathbb{Q}$ of degree at most $(2k)^{2^{|A|}}$, a subset $A' \subset K$ and a homomorphism $\phi_p : \mathbb{Z}[A'] \to \mathbb{F}_p$ such that ϕ_p is an F_k-ring-isomorphism between A' and A.*

One can use the construction from [2] to see that for any $k \geq 2$ and any prime number p there exists a subset $A \subseteq \mathbb{F}_p$ of size $O(\log p)$, which is not F_k-ring-isomorphic to any subset of a characteristic zero integral domain. For $k \geq 3$ we can improve this bound to the following.

Theorem 4. *For any $k \geq 3$ and any prime number $p \geq 2^{32(k-1)^2 \log_2^2(16(k-1))}$ there exists a subset $A \subseteq \mathbb{F}_p$ of size $|A| \leq \frac{10}{k-1} \frac{\log_2 p}{\log_2 \log_2 p}$ which is not F_k-ring-isomorphic to any subset of a characteristic zero integral domain.*

It is an open problem if a better bound is possible. In this direction I would like to make the following conjecture.

Conjecture 5. *For any $k \geq 3$ there is an infinite sequence of prime numbers, such that for each prime p in the sequence, there exists a subset $A \subseteq \mathbb{F}_p$ of size $O(\log \log p)$ which is not F_k-ring-isomorphic to any subset of a characteristic zero integral domain.*

As explained in Section 5, this conjecture would have a positive answer if, for example, there are infinitely many Mersenne primes (primes of the form $2^n - 1$; this would follow from the Lenstra–Pomerance–Wagstaff conjecture), or infinitely many Fermat primes (primes of the form $2^{2^n} + 1$; this is a question of Eisenstein).

The proof of Theorem 3 uses elimination theory. This is not the first time when elimination theory is applied to additive combinatorics: similar techniques were used by Chang in the proof of Lemma 2.14 from [6]. We state this lemma below in an equivalent form.

Lemma 6 (Lemma 2.14, [6]). *Let $f_1, \ldots, f_s \in \mathbb{Z}[x_1, \ldots, x_n]$ be polynomials of degree at most t and $\| \cdot \|_\infty$-norm at most k. If the system*

$$f_1(x) = \ldots = f_s(x) = 0$$

has a solution $(a_1, \ldots, a_n) \in \mathbb{C}^n$, then it also has a solution $(b_1, \ldots, b_n)$, where each b_i is the root of an integer polynomial of degree at most C and $\| \cdot \|_\infty$-norm at most Ck^C, with $C := C(t, n, s)$ depending only on t, n and s.

This lemma is discussed by Tao on his blog [32], in particular he gives a proof of it using nonstandard analysis. Neither this proof nor the proof in [6] provides a bound on the constant C.

The proof of Lemma 6 from [6] shows in fact a bit more; namely that if we are further given a polynomial $g \in \mathbb{Z}[x_1, \ldots, x_n]$ which does not vanish at $(a_1, \ldots, a_n)$, and has degree at most t and $\| \cdot \|_\infty$-norm at most k, then it is possible to choose $(b_1, \ldots, b_n)$ such that $g(b_1, \ldots, b_n) \neq 0$. On close examination of the proof it turns out that translated into the correct setting it implies the following weak version of Theorem 3.

Theorem 7. *For any $k \geq 2$ there exists a function $\nu_k : \mathbb{N} \to \mathbb{N}$ with $\lim_{n \to \infty} \nu_k(n) = \infty$, such that the following holds. If p is a prime and $A \subseteq \mathbb{F}_p$ with $|A| \leq \nu_k(p)$ then there exists a finite algebraic extension K of $\mathbb{Q}$ and a subset $A' \subset K$ such that A' is F_k-ring-isomorphic with A.*

An upper bound for the constant C implies a lower bound for $\nu_k(n)$; however from the proof of Lemma 6 one can only extract a rather poor bound for C.

It is also important to note that Theorem 7 does not provide any bound on the degree of the field extension K, nor does it guarantee that the F_k-ring-isomorphism is the restriction of a genuine ring homomorphism, as in Theorem 3. In fact it is easy to construct an example of a Freiman ring-isomorphism ϕ between a subset $A' \subset \mathbb{C}$ and a subset $A \subset \mathbb{F}_p$ such that ϕ is not the restriction of any ring homomorphism between $\mathbb{Z}[A']$ and $\mathbb{F}_p$. Indeed, consider $A' := \{-\frac{1}{2}, 2\} \subset \mathbb{C}$ and $A := \{3, 7\} \subset \mathbb{F}_{11}$. The map ϕ sending $-\frac{1}{2}$ to 3 and 2 to 7 is an F_2-ring-isomorphism, but it is obviously not the restriction of a ring homomorphism between $\mathbb{Z}[A']$ and $\mathbb{F}_{11}$ (as any such homomorphism would send 2 to 2). Examples for arbitrarily large k and p can be constructed as well.

The rest of the paper is organized as follows.

We start by giving several applications of the main result to subsets of $\mathbb{F}_p$ of size $O(\log \log \log p)$. In Section 2, we use Theorem 3 to prove a Szemerédi-Trotter type theorem with optimal exponent 4/3. In Section 3, we apply Theorem 3 to improve the currently best-known sum-product estimate in $\mathbb{F}_p$. Finally, in Section 4 we give several estimates for sets with small doubling

constant. All these results are proved by transferring the corresponding theorem from $\mathbb{C}$ to $\mathbb{F}_p$ via Theorem 3. In all these applications only the existence of a Freiman ring-isomorphism between A and a subset of $\mathbb{C}$ is needed, and not the stronger conclusion of Theorem 3.

In Section 5 we give an application of Theorem 3 to an old question of Rényi. In this case we will make essential use of the upper bound on the degree of the algebraic extension K in Theorem 3.

In Section 6 we show, as an example for the general strategy, how to preserve bounded linear polynomials. In Section 7 we gather all the necessary results from elimination theory. Finally, Section 8 is devoted to the proof of Theorem 3 and in Section 9 we prove Theorem 4.

We conclude with some further remarks concerning Freiman isomorphisms and Lemma 6.

Remark. After completion of this work I was informed by Pierre Simon that one can use the arithmetic Nullstellensatz stated in [22] to prove a good lower bound for the function ν_k in Theorem 7. With his idea, my own computations show that one can take $\nu_k(p) = \Omega(\frac{\log\log p}{\log\log\log p})$. This would improve the upper bound for n in Theorems 11, 13, 15 and 17 below to $O(\frac{\log\log p}{\log\log\log p})$.

Moreover, in his blog post *Rectification and the Lefschetz principle* [33], Tao presented a short proof of the following version of Theorem 3.

Theorem 8. *Let $k, n \geq 1$. If $\mathbb{F}$ is a field of characteristic at least $C_{k,n}$ for some $C_{k,n}$ depending only on k and n, and A is a subset of $\mathbb{F}$ of cardinality n, then there exists a map $\phi: A \to A'$ into a subset A' of the complex numbers which is a Freiman ring-isomorphism of order k.*

The proof uses non-standard analysis, and hence does not offer any bound on $C_{k,n}$. However, unlike Theorem 3, it also applies to fields of prime power order.

2. The Szemerédi-Trotter theorem

The well-known Szemerédi-Trotter theorem gives a tight upper bound on the number of incidences between a finite set of lines and a finite set of points in $\mathbb{R} \times \mathbb{R}$. This was extended to the complex plane $\mathbb{C}^2$ by Tóth.

Theorem 9 ([35]). *Let $\mathcal{P}$ and $\mathcal{L}$ be sets of points and lines in $\mathbb{C}^2$, with cardinalities $|\mathcal{P}|, |\mathcal{L}| \leq n$. Then there is a positive absolute constant c such that*

$$|\{(p, l) \in \mathcal{P} \times \mathcal{L} : p \in l\}| \leq cn^{4/3}.$$

Tóth's paper is still unpublished; but very recently Zahl gave a different proof of Theorem 9 in [38]. Unfortunately, Zahl's paper is also still unpublished. However, if we allow an $\varepsilon > 0$ error in the exponent, and the constant c to depend on ε, then in this form Theorem 9 follows from a generalization of the Szemerédi-Trotter theorem to algebraic varieties due to Solymosi and Tao [31].

The problem of establishing a similar bound in $\mathbb{F}_p$ has been considered before ([4], [17]). We have the following result, due to Helfgott and Rudnev.

Theorem 10 ([17]). *Let p be a prime number, and $\mathcal{P}$ and $\mathcal{L}$ sets of points and lines in $\mathbb{F}_p^2$, with $|\mathcal{P}|, |\mathcal{L}| \leq n$ and $n < p$. Then there is a positive absolute constant c such that*

$$|\{(p, l) \in \mathcal{P} \times \mathcal{L} : p \in l\}| \leq cn^{\frac{3}{2}-\delta},$$

with $\delta = \frac{1}{10678}$.

The best (still unpublished) bound to date for $n < p$ is due to Jones [19], who proved that one can take $\delta = \frac{1}{662} - o(1)$ in the above.

We show that one can achieve optimal exponent $4/3$ in Theorem 10 provided n is sufficiently small compared to p.

Theorem 11. *Let p be a prime number, and $\mathcal{P}$ and $\mathcal{L}$ sets of points and lines in $\mathbb{F}_p^2$, with $|\mathcal{P}|, |\mathcal{L}| \leq n$ and $5n < \log_2 \log_6 \log_{18} p - 1$. Then there is a positive absolute constant c such that*

$$|\{(p, l) \in \mathcal{P} \times \mathcal{L} : p \in l\}| \leq cn^{4/3}.$$

Moreover, this inequality is sharp up to the constant c.

Proof. We may assume w.l.o.g. that $|\mathcal{P}| = |\mathcal{L}| = n$, by adding some points and lines if necessary. Let $\mathcal{P} = \{(x_i, y_i) : 1 \leq i \leq n\}$. By uniquely parametrizing each line $l \in \mathcal{L}$ defined by $a_i y + b_i x + c_i = 0$, by the ordered triple (a_i, b_i, c_i), let $\mathcal{L} = \{(a_i, b_i, c_i) : 1 \leq i \leq n\}$. Now form the set $A := \cup_{i=1}^n \{x_i, y_i, a_i, b_i, c_i\}$. As $|A| \leq 5n$, we may apply Theorem 3 to find a subset $A' \subset \mathbb{C}$ and an F_3-ring-isomorphism ϕ between A and A'. By definition we have

$$a_j y_i + b_j x_i + c_j = 0 \Leftrightarrow \phi(a_j)\phi(y_i) + \phi(b_j)\phi(x_i) + \phi(c_j) = 0, \forall 1 \leq i, j \leq n,$$

hence the number of incidences between $\mathcal{P}$ and $\mathcal{L}$ in $\mathbb{F}_p^2$ is the same as the number of incidences between $\phi(\mathcal{P})$ and $\phi(\mathcal{L})$ in $\mathbb{C}$. Note that $\phi(\mathcal{P})$ and $\phi(\mathcal{L})$ have cardinality exactly n as ϕ is bijective. Hence by Theorem 9, the number of incidences is $O(n^{4/3})$, as desired.

To show that the bound is sharp, we use a standard construction that proves sharpness of the Szemerédi-Trotter theorem in $\mathbb{R}^2$. Let $r := \lfloor \frac{1}{2} n^{1/3} \rfloor$. We set $\mathcal{P}$ to be the points of the lattice $[r] \times [2r^2]$ in $\mathbb{F}_p^2$, and $\mathcal{L}$ to be all lines $y = mx + b$, with $1 \leq m \leq r, 1 \leq b \leq r^2$. Then every line from $\mathcal{L}$ is incident with exactly r points from $\mathcal{P}$, for a total of $r^4 = \Theta(n^{4/3})$ incidences. $\square$

One can now combine Theorem 11 with Theorem 2 to generalize Theorem 9 to any characteristic zero integral domain. As this statement can be proved directly with no recurse to Theorem 11, we do not discuss it here (see Theorem 2.3 and Lemma 7.1 from [37] for more details).

3. Sum-product estimates in $\mathbb{F}_p$

Suppose R is a commutative ring and $A \subset R$ a finite subset. We can define the sumset $A + A := \{a + b : a, b \in A\}$ and the product $A \cdot A := \{ab : a, b \in A\}$. Intuitively, the quantities $|A + A|$ and $|A \cdot A|$ can not both be small. The prototype theorem is a lower bound of the form $\max\{|A + A|, |A \cdot A|\} \geq c|A|^{1+\varepsilon_R}$, where $c > 0$ is an absolute constant and ε_R depends on the ring R. The first sum-product estimate is due to Erdős and Szemerédi [13] for the case $R = \mathbb{Z}$ and it was followed by numerous improvements and generalizations ([11], [24], [14], [7], [30]). For $R = \mathbb{C}$, the best-known value $\varepsilon_{\mathbb{C}} = \frac{3}{11} - o(1)$ was for many years given by a result of Solymosi [29]. Using a beautiful geometric argument, Konyagin and Rudnev [21] have very recently improved this to $\varepsilon_{\mathbb{C}} = \frac{1}{3} - o(1)$, thus matching the lower bound for the reals.

Theorem 12 ([21])**.** *Suppose $A \subset \mathbb{C}$. Then there is a positive absolute constant c such that*

$$|A + A| + |A \cdot A| \geq c|A|^{1 + \frac{1}{3} - o(1)}. \tag{1}$$

Bourgain, Katz and Tao [4] showed that a sum-product theorem holds in $\mathbb{F}_p$. Substantial work has gone into finding the best value for $\varepsilon_{\mathbb{F}_p}$. Garaev [15] showed that for $|A| < \sqrt{p}$ one can take $\varepsilon_{\mathbb{F}_p} = \frac{1}{14} - o(1)$. Katz and Shen [20] improved this to $\frac{1}{13} - o(1)$, and then Bourgain and Garaev [3] showed that $\frac{1}{12} - o(1)$ is in fact possible. Li [23] later removed the $o(1)$ term. The best result to date is due to Rudnev [26], who showed that

$$|A + A| + |A \cdot A| \geq c|A|^{1 + \frac{1}{11} - o(1)}, \tag{2}$$

whenever $|A| < \sqrt{p}$.

We now improve (2) for small A.

Theorem 13. *Let p be a prime number and $A \subseteq \mathbb{F}_p$ with $|A| < \log_2 \log_8 \log_{32} p - 1$. Then*

$$|A + A| + |A \cdot A| \geq c|A|^{1 + \frac{1}{3} - o(1)},$$

for some positive absolute constant c.

Proof. We apply Theorem 3 to find a subset $A' \subset \mathbb{C}$ and an F_4-ring-isomorphism ϕ between A and A'. Then $|\phi(A) + \phi(A)| = |A + A|$ and $|\phi(A) \cdot \phi(A)| = |A \cdot A|$. By (1) applied to $A' = \phi(A)$, the theorem follows. $\square$

4. ESTIMATES FOR SETS WITH SMALL DOUBLING CONSTANT

We gather in this section several miscellaneous results for the case when A has small doubling constant. We first have the following result, due to Solymosi.

Theorem 14 ([29]). *If $A \subset \mathbb{C}$ and $|A| = n$ with $|A + A| \leq Cn$, then $|A \cdot A| \geq cn^2 / \log n$.*

This transfers immediately to $\mathbb{F}_p$ as follows.

Theorem 15. *If $A \subseteq \mathbb{F}_p$ and $|A| = n < \log_2 \log_8 \log_{32} p - 1$ with $|A + A| \leq Cn$, then $|A \cdot A| \geq cn^2 / \log n$.*

The proof is similar to that of Theorem 13 and we omit it. We also have the following result due to Chang [6].

Theorem 16. *Let $A \subset \mathbb{C}$ with $|A| = n$ and $|A + A| \leq Cn$, for some $C > 0$. Then the following holds.*

  (i) *If $0 \notin A$ then $|A^{-1} + A^{-1}| > \exp^{-C' \frac{\log n}{\log \log n}} n^2$, for some C' depending only on C.*

  (ii) *If $f(x) \in \mathbb{C}[x]$ is a polynomial of degree $t \geq 2$ then $|f(A) + f(A)| > \exp^{-C' \frac{\log n}{\log \log n}} n^2$, for some $C' := C'(C, t)$.*

Here $A^{-1} = \{a^{-1} : a \in A\}$ and $f(A) = \{f(a) : a \in A\}$. The proof of Theorem 16 uses algebraic methods, in particular Lemma 6, but also relies crucially on facts specific to $\mathbb{C}$. We now transfer this theorem to small subsets of $\mathbb{F}_p$.

Theorem 17. *Let $A \subseteq \mathbb{F}_p$ with $|A| = n$ and $|A + A| \leq Cn$, for some $C > 0$. Then the following holds.*

  (i) *Suppose $2n < \log_2 \log_8 \log_{32} p - 1$ and $0 \notin A$. Then $|A^{-1} + A^{-1}| > \exp^{-C' \frac{\log n}{\log \log n}} n^2$, for some C' depending only on C.*

  (ii) *Let $f(x) \in \mathbb{Z}[x]$ be a k-bounded polynomial of degree at least 2. If $n < \log_2 \log_{8k} \log_{32k^2} p - 1$ then $|f(A) + f(A)| > \exp^{-C' \frac{\log n}{\log \log n}} n^2$, for some $C' := C'(C, k)$.*

Proof. We first prove (i).

We apply Theorem 3 to find a subset $A' \subset \mathbb{C}$ and an F_4-ring-isomorphism ϕ between $A \cup A^{-1}$ and A'. Then $|\phi(A)| = n$, $|\phi(A) + \phi(A)| = |A + A|$ and $|\phi(A^{-1}) + \phi(A^{-1})| = |A^{-1} + A^{-1}|$. Moreover, all identities of the form $a^{-1}a = 1, a \in A$, must be preserved by the ring-isomorphism, and hence $\phi(a^{-1}) = \phi(a)^{-1}, \forall a \in A$. Then by applying Theorem 16, (i), the result follows.

We now prove (ii).

We apply Theorem 3 to find a subset $A' \subset \mathbb{C}$ and an F_{4k}-ring-isomorphism ϕ between A and A'. Then $|\phi(A)| = n$ and $|\phi(A) + \phi(A)| = |A + A|$. We further have

$$f(\phi(a)) + f(\phi(b)) - f(\phi(c)) - f(\phi(d)) = 0 \Leftrightarrow f(a) + f(b) - f(c) - f(d) = 0,$$

for any $a, b, c, d \in A$, as ϕ is an F_{4k}-ring-isomorphism. Hence $|f(\phi(A)) + f(\phi(A))| = |f(A) + f(A)|$. Then by applying Theorem 16, (ii), the result follows. $\square$

5. A QUESTION OF RÉNYI

Let K be a field of characteristic zero. For a polynomial $f \in K[x]$ we define $N(f)$ to be the number of non-zero terms of f. For $k \geq 1$, let

$$Q_K(k) = \min_{f \in K[x]: N(f) = k} N(f^2). \tag{3}$$

As reported by Erdős [12], it was first asked by Rédei if $Q_{\mathbb{R}}(k) < k$ is possible, and Rényi [25] later constructed an example showing $Q_{\mathbb{Q}}(29) \leq 28$. Rényi made several conjectures about the behaviour of $Q_{\mathbb{R}}(k)$.

He conjectured that $\lim_{k \to \infty} \frac{Q_{\mathbb{R}}(k)}{k} = 0$, and this was proved by Erdős [12], who in fact showed that $Q_{\mathbb{Q}}(k) < ck^{1-\varepsilon}$, for some positive absolute constants c and ε.

Rényi further conjectured that $\lim_{k \to \infty} Q_{\mathbb{R}}(k) = \infty$, and this was proved many years later by Schinzel [27], using a very ingenious argument. Schinzel showed that $Q_K(k) \geq c \log \log k$, for some positive absolute constant c and any field K of characteristic zero. This lower bound was not improved for another 20 years, until recently Schinzel and Zannier [28], by an adaptation of the original method of Schinzel, proved that $Q_K(k) \geq c \log k$, for some positive absolute constant c.

Erdős [12] asked for the determination of the order of $Q_{\mathbb{R}}(k)$, and the general belief seems to be that $Q_{\mathbb{R}}(k)$ should be closer to the upper bound than the lower bound. Despite some work in this direction ([36], [10]), a solution to this problem seems at present out of reach.

From the definition we see that for any $k \geq 1$,

$$Q_{\mathbb{C}}(k) \leq Q_{\mathbb{R}}(k) \leq Q_{\mathbb{Q}}(k). \tag{4}$$

It is less known that Rényi [25] (see also [12]) asked whether equality holds in (4) everywhere for any k, and this problem seems to have received little attention.

For any $k \geq 1$ it also holds that

$$Q_{\mathbb{C}}(k) \leq Q_K(k) \leq Q_{\mathbb{Q}}(k), \tag{5}$$

for any finite algebraic extension K of $\mathbb{Q}$, and thus if we have equality in (4), then we also have equality in (5). In view of this we have the following result.

Theorem 18. *For any $k \geq 3$ there exists a finite algebraic extension K of $\mathbb{Q}$ such that $Q_{\mathbb{C}}(k) = Q_K(k)$, with degree at most k^{2^k}, if k is even, and at most $(k+1)^{2^k}$, if k is odd.*

Proof. Set $s := \lfloor \frac{k+1}{2} \rfloor$. Note that $s \geq 2$.

Let $f \in \mathbb{C}[x]$ be a polynomial with k non-zero terms minimizing $N(f^2)$. Suppose $f = a_0 + a_1 x^{n_1} + \ldots + a_{k-1} x^{n_{k-1}}$ and set $A := \{a_0, \ldots, a_{k-1}\} \subset \mathbb{C}$.

We now apply Theorem 2 in order to find a sufficiently large prime p (compared to k) and a homomorphism $\phi : \mathbb{Z}[A] \to \mathbb{F}_p$ which is an F_s-ring-isomorphism between A and $\phi(A)$. We then apply Theorem 3 to the set $\phi(A)$ in order to find a finite algebraic extension K of $\mathbb{Q}$ of degree at most $(2s)^{2^k}$, a subset $B \subset K$ and a map ψ between $\phi(A)$ and B, which is an F_s-ring-isomorphism. Then $\psi \circ \phi$ is an F_s-ring-isomorphism between A and B by construction.

Let $g = (\psi \circ \phi)(a_0) + (\psi \circ \phi)(a_1)x^{n_1} + \ldots + (\psi \circ \phi)(a_{k-1})x^{n_{k-1}}$. Then $g \in K[x]$ and $N(g) = k$. As any coefficient of g^2 is given by a polynomial with integer coefficients of degree at most 2 and $\| \cdot \|_1$-norm at most s, evaluated at $((\psi \circ \phi)(a_0), (\psi \circ \phi)(a_1), \ldots, (\psi \circ \phi)(a_{k-1}))$, we see that $N(g^2) = N(f^2)$. Consequently $Q_K(k) \leq N(f^2) = Q_{\mathbb{C}}(k)$, thus proving the theorem. $\square$

Remark. Lemma 29 below shows that K can in fact be chosen of degree at most 4^{2^k}.

6. PRESERVING THE ADDITIVE STRUCTURE

For comparison reasons we start by sketching a proof of Theorem 1, following [2].

Proof of Theorem 1. We first choose $0 < t < p$ such that multiplying every element of A by t (modulo p) results in a set $A^* \subseteq \{-\lfloor \frac{p}{2k} \rfloor, \ldots, \lfloor \frac{p}{2k} \rfloor\}$. The existence of t follows from the Kronecker approximation theorem (Corollary 3.2.5, [34]). Let $m \in \mathbb{Z}$ be such that $mt \equiv 1 \pmod{p}$. We multiply every element of A^* by m to obtain A'. Then the canonical homomorphism maps A' onto A, and one easily sees that this is also an F_k-isomorphism. $\square$

We will now consider the problem of preserving bounded linear polynomials. As we allow non-zero constant terms, we will have to find a proof different from that of Theorem 1.

We first prove an inequality.

Lemma 19. *Suppose $M = (m_{ij})$ is an $n \times n$ matrix with entries $m_{ij} \in \mathbb{Z}[x_1, \ldots, x_r]$. If for any i, $\sum_j \|m_{ij}\|_1 \leq k$, then $\|\det(M)\|_1 \leq k^n$. Furthermore, for any matrix M with integer entries, $|\det(M)|$ is at most the product of the $\| \cdot \|_1$-norms of the rows.*

Proof. We use the easily verified inequality $\|fg\|_1 \leq \|f\|_1 \|g\|_1$, which holds for any $f, g \in \mathbb{Z}[x_1, \ldots, x_r]$, to see that

$$\|\det(M)\|_1 \leq \sum_{\pi \in S_n} \|m_{1\pi(1)}\|_1 \ldots \|m_{n\pi(n)}\|_1 \leq \sum_{1 \leq i_1, \ldots, i_n \leq n} \|m_{1i_1}\|_1 \ldots \|m_{ni_n}\|_1$$

$$= (\sum_j \|m_{1j}\|_1) \ldots (\sum_j \|m_{nj}\|_1) \leq k^n.$$

$\square$

The last statement of Lemma 19 is also a consequence of Hadamard's inequality.

We now have the following technical result.

Lemma 20. *Let $k > 1$ be an integer and p be a prime. Suppose $A = \{a_1, \ldots, a_n\} \subseteq \mathbb{Z}_p$ and let $\mathcal{L}_1, \mathcal{L}_2 \subset \mathbb{Z}[x_1, \ldots, x_n]$ be collections of k-bounded linear polynomials, such that any $f \in \mathcal{L}_1$ is zero when evaluated at $(a_1, \ldots, a_n)$, and any $f \in \mathcal{L}_2$ is non-zero when evaluated at $(a_1, \ldots, a_n)$. If $|A| < \log_k p - 1$, then there exists $A' = \{b_1, \ldots, b_n\} \subset \mathbb{Z}_{(p)}$ such that the canonical homomorphism $\mathbb{Z}_{(p)} \to \mathbb{Z}_p$ maps b_i to a_i, and $f(b_1, \ldots, b_n) = 0$ for $f \in \mathcal{L}_1$, $f(b_1, \ldots, b_n) \neq 0$ for $f \in \mathcal{L}_2$.*

This directly implies Theorem 1, with almost the same bound.

Corollary 21. *Let $k \geq 1$ be an integer and p be a prime. Then for any $A \subseteq \mathbb{Z}_p$ with $|A| < \log_{2k} p - 1$ there exists $A' \subset \mathbb{Z}$ F_k-isomorphic with A via the canonical homomorphism.*

Proof. We consider all linear polynomials in $n := |A|$ variables having $\|\cdot\|_1$-norm at most $2k$, and split them into $\mathcal{L}_1$ and $\mathcal{L}_2$ according to the result of evaluation with elements from A. This includes all polynomials used in the definition of the usual Freiman isomorphism. Applying Lemma 20, we get a subset $A' \subset \mathbb{Z}_{(p)}$, which by definition must be F_k-isomorphic with A via the canonical homomorphism. Multiplying all values of A' by a large enough integer, which is 1 modulo p and cleares all denominators, will ensure that A' lies in $\mathbb{Z}$, while still being F_k-isomorphic with A via the canonical homomorphism. $\qquad\square$

Proof of Lemma 20. We can express $\mathcal{L}_1$ as the system $M\mathbf{x} = \mathbf{b}$, for some $m \times n$ matrix M and vector $\mathbf{b}$. We then form the augmented matrix $M' = (M|\mathbf{b})$. By assumption, the $\|\cdot\|_1$-norm of any row of M' is at most k.

The system $\mathcal{L}_1$ is solvable in a field $\mathbb{F}$ if and only if $\mathrm{rk}_{\mathbb{F}} M = \mathrm{rk}_{\mathbb{F}} M'$. We will show that this is the case in $\mathbb{Q}$.

As the rank of M' is the maximum size of one of its square submatrices with non-zero determinant, we see that $\mathrm{rk}_{\mathbb{Q}} M' \geq \mathrm{rk}_{\mathbb{F}_p} M'$. On the other hand, let M'_1 be any square submatrix of M' of full rank in $\mathbb{Q}$. By Lemma 19, $|\det(M'_1)| \leq k^{n+1} < p$. Hence $\det(M'_1)$ is also non-zero in $\mathbb{F}_p$, and consequently $\mathrm{rk}_{\mathbb{Q}} M' \leq \mathrm{rk}_{\mathbb{F}_p} M'$. But then M' has the same rank t in $\mathbb{Q}$ and in $\mathbb{F}_p$. Similarly we obtain that M has the same rank in both $\mathbb{Q}$ and $\mathbb{F}_p$. However, the system $\mathcal{L}_1$ is solvable in $\mathbb{F}_p$, and so we must have $t = \mathrm{rk}\, M \leq n$. Consequently $\mathcal{L}_1$ is solvable in $\mathbb{Q}$. This is nevertheless not enough for our purposes; we must further show that a solution A' with the desired properties exists.

We may assume w.l.o.g. that

$$M = \begin{pmatrix} M_1 & M_2 \\ M_3 & M_4 \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{pmatrix}$$

where M_1 is a square matrix of full rank $t = \mathrm{rk}\, M$ in both $\mathbb{Q}$ and $\mathbb{F}_p$, and $\mathbf{b}$ is partitioned accordingly. Let M_1^* be the adjoint of M_1.

We get

$$\begin{pmatrix} M_1^* & 0 \\ 0 & I \end{pmatrix} \begin{pmatrix} M_1 & M_2 \\ M_3 & M_4 \end{pmatrix} x = \begin{pmatrix} \det(M_1)I & M_1^*M_2 \\ M_3 & M_4 \end{pmatrix} x = \begin{pmatrix} M_1^*b_1 \\ b_2 \end{pmatrix}.$$

By Lemma 19, $|\det(M_1)| \leq k^n$.

Consequently we can express the first t variables in terms of the last $n-t$ variables, involving fractions with denominator bounded by $k^n < p$. By letting $b_i := a_i$ and replacing x_i with b_i in these equations for $t + 1 \leq i \leq n$, we obtain values $b_1, \ldots, b_t$ in $\mathbb{Z}_{(p)}$ for $x_1, \ldots, x_t$ such that b_i is mapped to a_i by the canonical homomorphism, for any $1 \leq i \leq n$. Furthermore, as $\mathrm{rk}_{\mathbb{Q}} M' = t$, by replacing x_i with b_i in the last $m - t$ equations we obtain the identity $0 = 0$ in $\mathbb{Q}$ everywhere.

We conclude that $A' := \{b_1, \ldots, b_n\}$ is a solution for $\mathcal{L}_1$ in $\mathbb{Z}_{(p)}$. Furthermore, no polynomial $f \in \mathcal{L}_2$ can be zero when evaluated at A', for otherwise it would also be zero modulo p, hence 0 when evaluated at A, a contradiction. Then we are done. $\qquad\square$

7. Resultants, subresultants and the gcd

As in the case of linear polynomials, we must bound the complexity of solving a system of multivariate polynomials. We gather in this section all the tools required for the proof.

In what follows we shall introduce and make substantial use of subresultants, an alternative to Euclid's algorithm for computing the greatest common divisor of two polynomials. This approach will be essential in obtaining any reasonable quantitative bound in Theorem 3, as Euclid's algorithm leads to an explosive growth of the coefficients involved in the polynomial division.

Suppose A is an integral domain. If $A \subseteq B$, B is a commutative ring and $b \in B$, we shall denote by ev_b the evaluation homomorphism $\mathrm{ev}_b : A[x] \to B$ mapping $f(x)$ to $f(b)$. If $0 \neq a \in A$, we shall denote by $A[\frac{1}{a}]$ the ring of polynomials $A[x]$ evaluated at $\frac{1}{a}$. This is the same as the ring of fractions of A with respect to $\{a^n : n \geq 0\}$, and is sometimes denoted by A_a. If B is another integral domain and $\phi : A \to B$ is a homomorphism, ϕ extends to a homomorphism from $A[x]$ to $B[x]$, which we shall also denote by ϕ.

Let $f, g \in A[x]$. We say $g|f$ if there exists $h \in A[x]$ such that $f = hg$. Hence $h|0$ for any $h \in A[x]$, but 0 divides only 0. Moreover if A is a unique factorization domain (UFD), then $\gcd_A(f, g)$ is well-defined. Here we use the conventions $\gcd_A(h, 0) = \gcd_A(0, h) = h$, for any polynomial h. Note that $\gcd_A(f, g)$ is unique only up to a unit of A. If no confusion may occur, we shall drop the subscript A. Furthermore if $f_1, \ldots, f_m \in A[x]$ we let $\gcd(f_1, \ldots, f_m)$ denote their greatest common divisor, where for $m = 1$ this is by convention f_1.

We also make the convention $\deg(0) = -\infty$.

We shall need the following easy fact.

Lemma 22. *Suppose $A \subseteq B$ are integral domains, $f, g \in A[x]$ non-zero and $g|f$ in $B[x]$. Then $g|f$ in $A[\frac{1}{\gamma}]$, where γ is the leading coefficient of g.*

Proof. By replacing A with $A[\frac{1}{\gamma}]$ and B with $B[\frac{1}{\gamma}]$, we may suppose $\frac{1}{\gamma} \in A$.

Assume $p := \deg(f), q := \deg(g)$ and $a \neq 0$ is the leading coefficient of f. By assumption, $f = hg$, for some $h \in B[x]$.

We prove by induction on $\deg(h) \geq 0$ that $h \in A[x]$.

Let $c \neq 0$ be the leading coefficient of h. Note that $\deg(h) = p - q$. Then $c\gamma = a$, and so $c = \frac{a}{\gamma} \in A$. If $\deg(h) = 0$, we are done, otherwise $f - cx^{p-q}g = (h - cx^{p-q})g$, and so by induction $h - cx^{p-q} \in A[x]$. Thus the claim is proved. $\square$

Now let A be an integral domain, $f, g \in A[x]$ be non-zero polynomials and suppose $f = a_p x^p + \ldots + a_0, g = b_q x^q + \ldots + b_0$ with $a_p, b_q \neq 0$. The *Sylvester matrix* of f and g is the $(p+q) \times (p+q)$ matrix

$$
S_{f,g} := \begin{pmatrix}
a_p & \ldots & a_0 & & & \\
 & \ddots & & & \ddots & \\
 & & a_p & \ldots & & a_0 \\
b_q & \ldots & b_0 & & & \\
 & \ddots & & & \ddots & \\
 & & b_q & \ldots & & b_0
\end{pmatrix},
$$

where the first q lines are formed by shifting the first row to the right, and the last p lines are formed by shifting the $(q+1)$th row to the right. If $p = q = 0$, we define $S_{f,g} = (1)$. The *resultant* of f and g, denoted by $\mathrm{res}(f, g)$, is the determinant of $S_{f,g}$. We also define

$\operatorname{res}(0,h) = \operatorname{res}(h,0) = 0$, for any polynomial h, so that the resultant is now properly defined for any two polynomials in $A[x]$.

The main application of resultants is to determine when two polynomials have a common root.

Theorem 23 (Proposition 4.16, [1]). *Suppose A is a UFD and $f, g \in A[x]$ are non-zero. Then $\gcd(f,g)$ is non-constant if and only if $\operatorname{res}(f,g) = 0$.*

Unfortunately we will have to deal with more than two polynomials and more than one variable. We therefore make the following definition, following [18].

Let $f_1, \ldots, f_m \in A[x_1, \ldots, x_n], m \geq 1$. Let $y_3, \ldots, y_m$ be new indeterminates and define $A' := A[x_2, \ldots, x_n], A'' := A'[y_3, \ldots, y_m]$. Let F_1, F_2 be polynomials in $A''[x_1]$ defined as follows:

$$F_1 := f_1 \tag{6}$$
$$F_2 := f_2 + y_3 f_3 + \ldots + y_m f_m.$$

If $m = 1$, we take $F_2 := 0$. We define the resultant of the polynomials $f_1, \ldots, f_m$ in terms of x_1, denoted by $\operatorname{res}_{x_1}(f_1, \ldots, f_m)$, as the resultant of F_1 and F_2. Note that this is a polynomial in $x_2, \ldots, x_n$ and $y_3, \ldots, y_m$.

We first have a lemma.

Lemma 24. *Suppose A is a UFD. Then $\gcd_{A'}(f_1, \ldots, f_m) = \gcd_{A''}(F_1, F_2)$.*

Proof. If $m = 1$, this is true by definition. So assume $m \geq 2$.

By hypothesis A' and A'' are both UFD. Now if $g := \gcd_{A'}(f_1, \ldots, f_m)$ and $g' := \gcd_{A''}(F_1, F_2)$ then $g | g'$, as $g | F_1$ and $g | F_2$. Furthermore $g' \in A'$, because $g' | f_1$. Giving values $y_i = 0$ we see that $g' | f_2$. Also if we let $y_j = 1$ and $y_i = 0, i \neq j$, we see that $g' | f_2 + f_j$. Hence $g' | f_j, 2 \leq j \leq m$. Then $g' | g$ and the claim follows. $\square$

Theorem 25. *Assume A is a field and let K be its algebraic closure. Let $(a_2, \ldots, a_n) \in K^{n-1}$ and suppose that the leading coefficient of x_1 in $f_1 \in A[x_1, \ldots, x_n]$, a polynomial in $x_2, x_3, \ldots, x_n$, does not vanish when replacing x_2 with a_2, x_3 with $a_3, \ldots, x_n$ with a_n. Then there exists an $a_1 \in K$ such that $(a_1, \ldots, a_n)$ is a common zero for $f_1, \ldots, f_m$ if and only if $\operatorname{res}_{x_1}(f_1, \ldots, f_m)(a_2, \ldots, a_n) = 0$.*

Proof. We replace x_i by $a_i, 2 \leq i \leq n$, in F_1 and F_2. Then the degree of F_1 stays the same, but the degree of F_2 may decrease with some amount $r \geq 0$.

If $F_2 = 0$ then by definition $\operatorname{res}_{x_1}(f_1, \ldots, f_m)(a_2, \ldots, a_n) = 0$. As $\deg_{x_1}(F_1) \geq 1$ and a_1 can be taken to be any root of F_1, the claim trivially holds.

So assume $F_2 \neq 0$. By definition of the Sylvester matrix we know that

$$\operatorname{res}_{x_1}(f_1, \ldots, f_m)(a_2, \ldots, a_n) = c^r \operatorname{res}_{x_1}(f_1(a_2, \ldots, a_n), \ldots, f_m(a_2, \ldots, a_n)).$$

where $0 \neq c \in K$ is the leading coefficient of x_1 in $f_1(a_2, \ldots, a_n)$. Thus by replacing f_i with $f_i(a_2, \ldots, a_n), 1 \leq i \leq m$, and A with K, we may suppose w.l.o.g. that $n = 1$.

By Lemma 24, $\gcd(f_1, \ldots, f_m) = \gcd(F_1, F_2)$, and hence a_1 exists iff $\gcd(F_1, F_2)$ is non-constant. But by Theorem 23 this happens iff $\operatorname{res}(F_1, F_2) = \operatorname{res}_{x_1}(f_1, \ldots, f_m)$ is zero, hence the claim holds. $\square$

For a different proof of Theorem 25 see Theorem 6.1, [18].

We now turn to subresultants.

Let A be an integral domain, $f, g \in A[x]$ non-zero as before and again suppose $f = a_p x^p + \ldots + a_0, g = b_q x^q + \ldots + b_0$ with $a_p, b_q \neq 0$. The *subresultant sequence* for f and g is a list of polynomials $S_i(f, g) := \sum_{j=0}^{i} s_{ij}(f, g) x^j, 0 \leq i \leq \min\{p, q\}$, where $s_{ij}(f, g)$ is the determinant of the matrix built with rows $1, \ldots, q - i$ and $q + 1, \ldots, q + p - i$ of $S_{f,g}$, and columns $1, 2, \ldots, p + q - 2i - 1, p + q - i - j$ of $S_{f,g}$. This is well-defined except when $i = p = q$. Thus when $p = q \neq 0$ we set $S_q(f, g) = g$ and define s_{qj} in the obvious way. For $p = q = 0$ we set $S_0(f, g) = 1$.

Due to technical reasons we define the subresultant sequence also for the case when one of f or g (but not both) is 0. If $g = 0$, we let $S_i(f, g) := S_i(f, f), 0 \leq i \leq \deg(f)$, and we proceed similarly if $f = 0$.

We now have the following result.

Theorem 26. *Suppose A is a UFD and $f, g \in A[x]$ are not both zero. If $k \geq 0$ is minimal such that $s_{kk}(f, g) \neq 0$ then there exists non-zero $u, v \in A$ such that $u \gcd(f, g) = v S_k(f, g)$.*

In a similar form, Theorem 26 was already known in the 19th Century. Collins [9] introduced the terminology of subresultants, leading to the modern formulation of Theorem 26, in conjuction with the problem of efficiently computing the gcd of two polynomials. The theory was subsequently refined and simplified by Brown and Traub [5]. A good exposition of the theory of subresultants and a proof of Theorem 26 can be found in [1] (see also [8] and [5]).

In the proof of the main result we will encounter rings which are not UFD, and so we will not be able to apply Theorem 26 directly. We deal with this situation below.

Let A be an integral domain and $f_1, \ldots, f_m \in A[x], m \geq 1$. We define F_1 and F_2 as in (6). We first make a simple observation.

Lemma 27. *Assume $A \subseteq K \subseteq \overline{K}$, where $K, \overline{K}$ are fields, and $\overline{K}$ is algebraically closed. Suppose $G := \gcd_K(f_1, \ldots, f_m)$ has degree $\delta \geq 1$, and let $b_1, \ldots, b_d$ be the distinct roots of G in $\overline{K}$, each appearing with multiplicity $\mu_i, 1 \leq i \leq d$. Then*

$$S_\delta(F_1, F_2) = \ell \prod_{i=1}^{d} (x - b_i)^{\mu_i}, \tag{7}$$

where ℓ is the leading coefficient of $S_\delta(F_1, F_2)$ as a polynomial in x.

As $\delta \geq 1$ we have $\deg(F_1) = \deg(f_1) \geq 1$ and so $S_i(F_1, F_2)$ is well-defined (nevertheless it may happen that F_2 is 0 if $m = 1$). Further recall that $S_i(F_1, F_2)$ is a polynomial in $y_3, \ldots, y_m$ and x.

Proof of Lemma 27. By Lemma 24, $G = \gcd_{K[y_3, \ldots, y_m]}(F_1, F_2)$. Hence by Theorem 26, there are non-zero $u, v \in K[y_3, \ldots, y_m]$ such that $uG = vS_\delta(F_1, F_2)$. But for any $1 \leq i \leq d$, $(x - b_i)^{\mu_i} | uG$ in $\overline{K}[y_3, \ldots, y_m, x]$. Hence $(x - b_i)^{\mu_i} | S_\delta(F_1, F_2), 1 \leq i \leq d$. As $S_\delta(F_1, F_2)$ has degree exactly δ as a polynomial in x, (7) must hold, thus proving the lemma. $\square$

The main consequence of Theorem 26 is the following.

Lemma 28. *Suppose $A \subseteq \mathbb{C}, G := \gcd_\mathbb{C}(f_1, \ldots, f_m)$ has degree $\delta \geq 1$, $\ell := s_{\delta\delta}(F_1, F_2)$ and $\phi : A \to \mathbb{F}_p$ is a homomorphism such that*

$$\deg_x(\phi(F_1)) = \deg_x(F_1), \quad \deg_x(\phi(F_2)) = \deg_x(F_2) \quad and \quad \phi(\ell) \neq 0. \tag{8}$$

Then for any root $b' \in \mathbb{F}_p$ of $\gcd_{\mathbb{F}_p}(\phi(f_1), \ldots, \phi(f_m))$ there exists a root b of G and a homomorphism $\Phi : A[b] \to \mathbb{F}_p$ such that the following diagram commutes

$$
\begin{array}{ccc}
A[x] & \xrightarrow{\mathrm{ev}_b} & A[b] \\
\phi \downarrow & & \downarrow \Phi \\
\mathbb{F}_p[x] & \xrightarrow{\mathrm{ev}_{b'}} & \mathbb{F}_p
\end{array}
\tag{9}
$$

Proof. By definition the case $m = 1$ is equivalent to the case $m = 2$ where $f_2 = f_1$, and so we will assume w.l.o.g. that $m \geq 2$ and $F_2 \neq 0$. Let $G' := \gcd_{\mathbb{F}_p}(\phi(f_1), \ldots, \phi(f_m))$.

As $\deg_x(\phi(F_1)) = \deg_x(F_1)$ and $\deg_x(\phi(F_2)) = \deg_x(F_2)$, we have $\phi(S_i(F_1, F_2)) = S_i(\phi(F_1), \phi(F_2))$. Hence by Theorem 26 and the fact that $\phi(\ell) \neq 0$, we have $\deg(G) = \deg(G') = \delta \geq 1$.

Let b' be any root of G' in $\mathbb{F}_p$. By Lemma 27 we have $\phi(S_\delta(F_1, F_2))(b') = 0$. Define $\psi := ev_{b'} \circ \phi : A[x] \to \mathbb{F}_p$.

Let $b_1, \ldots, b_d$ be the distinct roots of G in $\mathbb{C}$, each appearing with multiplicity $\mu_i, 1 \leq i \leq d$. Assume for a contradiction that for any root b_i of G there is no homomorphism Φ making the diagram (9) commutative. This means $\ker ev_{b_i} \not\subseteq \ker \psi$, so there exists a polynomial $g_i \in A[x]$ such that $g_i(b_i) = 0$, but $(\phi \circ g_i)(b') \neq 0$.

Define

$$
H := \ell \prod_{i=1}^{d} g_i^{\mu_i}.
$$

Then $H \in A[x, y_3, \ldots, y_m]$. As $\phi(\ell) \neq 0$, we have $\phi(H)(b') \neq 0$ in $\mathbb{F}_p[y_3, \ldots, y_m]$. But by Lemma 27,

$$
S_\delta(F_1, F_2) = \ell \prod_{i=1}^{d} (x - b_i)^{\mu_i}
$$

in $\mathbb{C}[x, y_3, \ldots, y_m]$. Then $S_\delta(F_1, F_2) | H$ in $\mathbb{C}[x, y_3, \ldots, y_m]$. Hence by Lemma 22, $S_\delta(F_1, F_2) | H$ in $A[x, y_3, \ldots, y_m, \frac{1}{\ell}]$. But $\phi(\ell) \neq 0$, so ϕ extends to a homomorphism

$$
\phi : A[x, y_3, \ldots, y_m, \frac{1}{\ell}] \to \mathbb{F}_p[x, y_3, \ldots, y_m].
$$

This implies $\phi(S_\delta(F_1, F_2)) | \phi(H)$. As $\phi(S_\delta(F_1, F_2))(b') = 0$, we obtain $\phi(H)(b') = 0$, a contradiction. This finishes the proof of the lemma. $\qquad\square$

8. Preserving both the additive and multiplicative structure

We have the following technical result.

Lemma 29. *Let $k, t \geq 2$ be integers and p be a prime. Suppose $A = \{a_1, \ldots, a_n\} \subseteq \mathbb{F}_p$ and let $\mathcal{L}_1, \mathcal{L}_2 \subset \mathbb{Z}[x_1, \ldots, x_n]$ be collections of (k, t)-bounded polynomials, such that any $f \in \mathcal{L}_1$ is zero when evaluated at $(a_1, \ldots, a_n)$, and any $f \in \mathcal{L}_2$ is non-zero when evaluated at $(a_1, \ldots, a_n)$. If*

$$
|A| < \log_2 \log_{2t} \log_{2kt} p - 1
\tag{10}
$$

then there exists a finite algebraic extension K of $\mathbb{Q}$ of degree at most $(2t)^{2^n}$ and a subset $A' = \{b_1, \ldots, b_n\} \subset K$ such that $f(b_1, \ldots, b_n) = 0$ for $f \in \mathcal{L}_1$, and $f(b_1, \ldots, b_n) \neq 0$ for $f \in \mathcal{L}_2$. Furthermore, the map $\phi_p : \mathbb{Z}[A'] \to \mathbb{F}_p$ sending b_i to a_i is a ring homomorphism.

Proof. We first give a rough overview of the proof.

The proof has three steps.

In the first step we eliminate the variables one by one. We start with the collection of polynomials $\mathcal{L}^0 := \mathcal{L}_1$ and we compute the resultant R_1 in terms of x_1. We then form a new collection of polynomials $\mathcal{L}^1$ in $x_2, \ldots, x_n$ by taking the coefficients of the y-monomials in R_1. By Theorem 25, there is at least one choice for x_1 iff there exists a common solution to the polynomials in $\mathcal{L}^1$. We then eliminate x_2 and proceed further in the same manner to construct collections $\mathcal{L}^i$. After at most n steps we have eliminated all variables, and only constant polynomials remain. However, the same procedure could have been carried over in $\mathbb{F}_p$, with the same starting collection of polynomials, and there it is guaranteed that a solution exists. Hence if the final constants are less than p, they must in fact be 0, and so a solution exists in $\mathbb{C}$ as well.

In the second step we go back, trying to determine the b_i's. Suppose for example that we have only polynomials in one variable, say x_n, and we know that a common root exists. Then their gcd is non-constant, and we can use Lemma 28 to pick one of the roots of the gcd as b_n. The hypothesis of Lemma 28 will be satisfied by adding some more polynomials to $\mathcal{L}^i$ in the first step. We then adjoin b_n to $\mathbb{Q}$, replace x_n by b_n, and proceed similarly to determine b_{n-1}. Theorem 25 will ensure that once $b_{i+1}, \ldots, b_n$ are picked, there is still a choice for b_i.

Note that once the homomorphism ϕ_p is constructed, the conditions imposed by $\mathcal{L}_2$ are automatically satisfied. For if $f \in \mathcal{L}_2$ then $\phi_p(f(b_1, \ldots, b_n)) = f(a_1, \ldots, a_n) \not\equiv 0 \,(\mathrm{mod}\ p)$, hence $f(b_1, \ldots, b_n) \neq 0$ as well.

In the last step we will estimate the degree of the extension.

We now present the proof in detail.

Step 1. We let $u_0 := k, v_0 := t$ and for any $1 \leq i \leq n$ we define u_i and v_i inductively by

$$u_i := u_{i-1}^{2v_{i-1}} v_{i-1}^{v_{i-1}},$$
$$v_i := 2v_{i-1}^2.$$

We shall prove in Step 3 that for $0 \leq i \leq n$ we have

$$u_i < p. \tag{11}$$

Assume for the moment that this is indeed the case. For $0 \leq i \leq n$ let $\sigma_i : \mathbb{Z}[x_{i+1}, \ldots, x_n] \to \mathbb{F}_p[x_{i+1}]$ be the homomorphism mapping x_j to $a_j, i+1 < j \leq n$. We similarly define $\sigma : \mathbb{Z}[x_1, \ldots, x_n] \to \mathbb{F}_p$ as the homomorphism mapping x_j to a_j for all $1 \leq j \leq n$.

We will construct by induction on $i \geq 0$ sets $\mathcal{L}_1 = \mathcal{L}^0, \mathcal{L}^1, \ldots, \mathcal{L}^r, r \leq n$, such that $\mathcal{L}^i \subset \mathbb{Z}[x_{i+1}, \ldots, x_n]$ is a collection of (u_i, v_i)-bounded polynomials satisfying $\sigma(f) = 0$ for any $f \in \mathcal{L}^i, 0 \leq i \leq r$. Furthermore, it will be necessary at every step $i < r$ to slightly modify the set $\mathcal{L}^i$ into another one $\mathcal{A}_i$ by altering some of the polynomials. $\mathcal{A}_i$ will still contain only (u_i, v_i)-bounded polynomials f verifying $\sigma(f) = 0$.

The construction of the sets $\mathcal{L}^i$ will be done in three stages, indicated by the bold letters **(A), (B)** and **(C)**.

For $i = 0$, by assumption $\mathcal{L}^0$ is a collection of (u_0, v_0)-bounded polynomials mapped to 0 by σ.

Now suppose $n \geq i \geq 0$ and we have constructed $\mathcal{L}^i$. If $i = n$ or $\mathcal{L}^i$ is empty or $\{0\}$, we set $r = i$ and stop. Otherwise, let $\mathcal{L}^i = \{f_1, \ldots, f_m\}$ and $f_j = \sum_{\ell=0}^{d_j} c_{j\ell} x_{i+1}^\ell$. By assumption we have $i \leq n-1$.

(A) For any $1 \leq j \leq m$ and $\deg_{x_{i+1}}(\sigma_i(f_j)) < \ell \leq \deg_{x_{i+1}}(f_j)$ we put $c_{j\ell}$ into $\mathcal{L}^{i+1}$.

We then set $d'_j = \deg_{x_{i+1}}(\sigma_i(f_j))$ and define

$$f'_j := \sum_{\ell=0}^{d'_j} c_{j\ell} x_{i+1}^\ell.$$

Note that $d'_j \neq 0$, otherwise $\sigma(f_j) = \sigma_i(f_j) \neq 0$, a contradiction.

Let $\mathcal{A}_i := \{f'_1, \ldots, f'_m\}$. Clearly every polynomial in $\mathcal{A}_i$ is still (u_i, v_i)-bounded. Furthermore if x_{i+1} does not appear in any polynomial in $\mathcal{A}_i$, then $\mathcal{A}_i$ contains only 0 by the above. In this case there is nothing else to be done.

So assume w.l.o.g. that x_{i+1} appears in f'_1. Let

$$F_1 := f'_1$$
$$F_2 := f'_2 + y_3 f'_3 + \ldots + y_m f'_m$$

for unknowns $y_3, \ldots, y_m$, where $F_2 := 0$ if $m = 1$.

(B) We take $\mathrm{res}_{x_{i+1}}(f'_1, \ldots, f'_m)$ and put into $\mathcal{L}^{i+1}$ the coefficient of every monomial in y_j, which must be a polynomial in $x_{i+2}, \ldots, x_n$.

Set $R_1 := \mathbb{Z}[x_{i+2}, \ldots, x_n, y_3, \ldots, y_m]$ and $R_2 := \mathbb{F}_p[y_3, \ldots, y_m]$. Note that σ_i induces a homomorphism between $R_1[x_{i+1}]$ and $R_2[x_{i+1}]$. We have $F_1, F_2 \in R_1[x_{i+1}]$ and by (A), $\deg_{x_{i+1}}(\sigma_i(F_1)) = \deg_{x_{i+1}}(F_1)$ and $\deg_{x_{i+1}}(\sigma_i(F_2)) = \deg_{x_{i+1}}(F_2)$. So let $q_1 := \deg_{x_{i+1}}(F_1)$ and $q_2 := \deg_{x_{i+1}}(F_2)$. By assumption, $q_1 \geq 1$.

Let $\delta \geq 0$ be minimal such that $\sigma_i(s_{\delta\delta}(F_1, F_2)) \neq 0$, where $s_{k\ell}$ are the coefficients of the subresultant sequence.

(C) We put into $\mathcal{L}^{i+1}$ the coefficients of $s_{jj}(F_1, F_2)$ (polynomials in $x_{i+2}, \ldots, x_n$), for $1 \leq j < \delta$. For $j = 0$ this has already been done, as $s_{00}(F_1, F_2) = \mathrm{res}_{x_{i+1}}(f'_1, \ldots, f'_m)$ by definition.

The construction of $\mathcal{L}^{i+1}$ is now over. We must show that any polynomial in $\mathcal{L}^{i+1}$ is indeed (u_{i+1}, v_{i+1})-bounded.

This is certainly the case for the polynomials added in stage (A). So consider the stage (B) of the construction.

Fix an arbitrary monomial M in $y_3, \ldots, y_m$ of degree at most q_1. This has a coefficient g in $\mathrm{res}_{x_{i+1}}(f'_1, \ldots, f'_m)$ and we must estimate $\|g\|_1$ and $\deg(g)$. Since $q_1, q_2 \leq v_i$, the degree of g is at most $(q_1 + q_2)v_i \leq 2v_i^2 = v_{i+1}$, as desired.

Now let $2 \leq j_1, j_2, \ldots, j_{q_1} \leq m$ and define $S_{F_1, F_2}(j_1, \ldots, j_{q_1})$ by writing on line $q_2 + k'$ of S_{F_1, F_2}, instead of the coefficients of F_2, the corresponding coefficients of $f'_{j_{k'}}, 1 \leq k' \leq q_1$. Then g is a sum of $\det(S_{F_1, F_2}(j_1, \ldots, j_{q_1}))$, for certain q_1-tuples $j_1, j_2, \ldots, j_{q_1}$ depending on M. The number of such q_1-tuples is

$$\binom{q_1}{\deg(M)} \frac{\deg(M)!}{\deg_{y_3}(M)! \ldots \deg_{y_m}(M)!} \leq q_1^{\deg(M)} \leq v_i^{v_i}.$$

Recall that $\|f'_j\|_1 \leq u_i, 1 \leq j \leq m$. So by Lemma 19 applied to $S_{F_1, F_2}(j_1, \ldots, j_{q_1})$ (a square matrix of size $q_1 + q_2 \leq 2v_i$), we obtain $\|\det(S_{F_1, F_2}(j_1, \ldots, j_{q_1}))\|_1 \leq u_i^{2v_i}$. Hence $\|g\|_1 \leq u_i^{2v_i} v_i^{v_i} = u_{i+1}$, as desired.

Finally, as subresultants are defined using submatrices of S_{F_1, F_2}, all the above estimates apply to subresultants as well. Hence any polynomial added to $\mathcal{L}^{i+1}$ in stage (C) is also (u_{i+1}, v_{i+1})-bounded. Consequently any polynomial in $\mathcal{L}^{i+1}$ is (u_{i+1}, v_{i+1})-bounded, as claimed.

We must further check that σ maps all the polynomials in $\mathcal{L}^{i+1}$ to 0. This is certainly the case with the polynomials added in stages (A) and (C) of the construction. As $\deg_{x_{i+1}}(\sigma_i(f_j')) = \deg_{x_{i+1}}(f_j'), 1 \leq j \leq m$, we have that $\mathrm{res}_{x_{i+1}}(\sigma_i(f_1'), \ldots, \sigma_i(f_m')) = \sigma_i(\mathrm{res}_{x_{i+1}}(f_1', \ldots, f_m'))$. By Theorem 25 and the fact that the polynomials $\sigma_i(f_j')$ have the common root a_{i+1}, we obtain $\mathrm{res}_{x_{i+1}}(\sigma_i(f_1'), \ldots, \sigma_i(f_m')) = 0$. This shows that all the polynomials added in stage (B) of the construction are indeed mapped to 0 by σ. Thus the induction step is verified.

Step 2. If $\mathcal{L}^r$ is empty, all the sets $\mathcal{L}^i$ were empty, in particular $\mathcal{L}^0 = \mathcal{L}_1 = \emptyset$. Then we set $b_i = a_i, 1 \leq i \leq n$, take ϕ_p to be the canonical homomorphism, and we are done.

So we may assume that $\mathcal{L}^r$ is non-empty. Let $f \in \mathcal{L}^r$. By construction f is an integer constant at most u_r in absolute value, and $u_r < p$ by (11). However, $\sigma(f) = 0$, and as σ is a homomorphism, we must have $f = 0$. Hence $\mathcal{L}^r = \{0\}$.

By decreasing induction on $r \geq i \geq 0$ we shall find algebraic numbers $b_{i+1}, \ldots, b_n$ such that for any $f \in \mathcal{L}^i$, $f(b_{i+1}, \ldots, b_n) = 0$, and furthermore the map $\phi_p^i : \mathbb{Z}[b_{i+1}, \ldots, b_n] \to \mathbb{F}_p$, sending b_j to $a_j, i < j \leq n$, is a well-defined homomorphism.

For any $j > r$, we let b_j be the integer in $\{0, 1, \ldots, p-1\}$ satisfying $b_j \equiv a_j \pmod{p}$. Then $\phi_p^r = \sigma|_{\mathbb{Z}}$ is a homomorphism. As $\mathcal{L}^r = \{0\}$, the base case $i = r$ is verified.

Now assume $0 \leq i < r$ and we have found $b_{i+2}, \ldots, b_n$ satisfying the induction hypothesis.

Suppose $\mathcal{L}^i = \{f_1, \ldots, f_m\}$ and $\mathcal{A}_i = \{f_1', \ldots, f_m'\}$. We replace $x_{i+2}, \ldots, x_n$ with their values b_j in the polynomials $f_1, \ldots, f_m$ and $f_1', \ldots, f_m'$. By (A), $f_j = f_j'$ and furthermore $\deg_{x_{i+1}}(\phi_p^{i+1}(f_j)) = \deg_{x_{i+1}}(f_j), 1 \leq j \leq m$. If x_{i+1} does not appear in any of these polynomials, then all of them are in fact 0. In this case we let b_{i+1} be the integer in $\{0, 1, \ldots, p-1\}$ satisfying $b_{i+1} \equiv a_{i+1} \pmod{p}$. We have $\phi_p^{i+1}(b_{i+1}) = a_{i+1}$. Thus $\phi_p^i = \phi_p^{i+1}$ is a well-defined homomorphism, and the claim holds.

So assume x_{i+1} appears in f_1. Here we use the same indexing scheme as in Step 1; in particular, f_1 corresponds to the polynomial f_1' selected in Step 1.

By (B) and Theorem 25, at least one choice b_{i+1} for x_{i+1} exists, such that replacing x_{i+1} with this value vanishes all polynomials in $\mathcal{L}^i$. In other words, $G := \gcd_{\mathbb{C}}(f_1, \ldots, f_m)$ has degree $\delta \geq 1$.

Now recall our construction of F_1 and F_2. By (A), $\deg_{x_{i+1}}(\phi_p^{i+1}(F_2)) = \deg_{x_{i+1}}(F_2)$. Let $\ell := s_{\delta\delta}(F_1, F_2)$. By (C), Lemma 24 and Theorem 26 applied to F_1 and F_2 in $\mathbb{C}[x_{i+1}, y_3, \ldots, y_m]$, we see that $\phi_p^{i+1}(\ell) \neq 0$.

Hence the hypothesis of Lemma 28 is satisfied for the ring $A := \mathbb{Z}[b_{i+2}, \ldots, b_n]$, the polynomials $f_1, \ldots, f_m$ and the homomorphism $\phi := \phi_p^{i+1}$. This implies that for the root a_{i+1} of $\gcd_{\mathbb{F}_p}(\phi_p^{i+1}(f_1), \ldots, \phi_p^{i+1}(f_m))$ there exists a root b_{i+1} of G and a homomorphism $\phi_p^i : \mathbb{Z}[b_{i+1}, \ldots, b_n] \to \mathbb{F}_p$ making the diagram (9) commutative. Then ϕ_p^i still maps b_j to a_j for $i + 1 < j \leq n$. Furthermore by construction, replacing x_{i+1} with b_{i+1} in the polynomials in $\mathcal{L}^i$ vanishes all of them. This proves the induction step.

Continuing in this way we obtain all algebraic numbers $b_1, \ldots, b_n$ and in the last step $\phi_p := \phi_p^0$ maps b_j to a_j as desired.

Step 3. We now compute the degree of the extension and verify (11).

First note that $r \leq n$ and $v_i = 2^{2^i-1}t^{2^i}, 0 \leq i \leq n$. Then the degree of the extension is at most

$$\prod_{i=0}^{r-1} v_i \leq \prod_{i=0}^{n-1} 2^{2^i-1}t^{2^i} \leq 2^{2^n-(n+1)}t^{2^n} \leq (2t)^{2^n}.$$

Further note that

$$\prod_{i=0}^{n-1} 2v_i \leq 2^{2^n-1} t^{2^n}.$$

We also have $u_0 = k$ and

$$u_{i+1} = u_i^{2v_i} v_i^{v_i}, \tag{12}$$

and so by iterating (12), and using the above estimates, we obtain

$$
\begin{aligned}
u_n &= u_{n-1}^{2v_{n-1}} v_{n-1}^{v_{n-1}} \\
&= u_{n-2}^{(2v_{n-2})(2v_{n-1})} v_{n-2}^{v_{n-2}(2v_{n-1})} v_{n-1}^{v_{n-1}} \\
&= \ldots \\
&= \exp\left\{ \left(\prod_{i=0}^{n-1} 2v_i\right) \log k + \sum_{i=0}^{n-1} v_i(2v_{i+1})\ldots(2v_{n-1}) \log v_i \right\} \\
&\leq \exp\left\{ \left(\prod_{i=0}^{n-1} 2v_i\right) (\log k + n \log v_{n-1}) \right\} \\
&\leq \exp\left\{ 2^{2^n-1} t^{2^n} (\log k + n \log(2t)^{2^{n-1}}) \right\} \\
&\leq k^{2^{2^n} t^{2^n}} (2t)^{n 2^{n-1} 2^{2^n-1} t^{2^n}} \\
&\leq k^{(2t)^{2^n}} (2t)^{2^{2^n+2n-2} t^{2^n}} \\
&\leq k^{(2t)^{2^n}} (2t)^{(2t)^{2^{n+1}}} \\
&\leq (2kt)^{(2t)^{2^{n+1}}}.
\end{aligned}
$$

Thus the condition $u_n < p$ is satisfied if $n < \log_2 \log_{2t} \log_{2kt} p - 1$. This shows that (11) holds, and hence the proof is finished. $\qquad\square$

Proof of Theorem 3. We consider all k-bounded polynomials in $n := |A|$ variables, and we split them into $\mathcal{L}_1$ and $\mathcal{L}_2$ according to the result of evaluation with elements from A. Applying Lemma 29, we get a finite algebraic extension K of $\mathbb{Q}$ of degree at most $(2k)^{2^n}$, a subset $A' \subset K$ and a homomorphism $\phi_p : \mathbb{Z}[A'] \to \mathbb{F}_p$ which by definition is an F_k-ring-isomorphism between A' and A. This proves the theorem. $\qquad\square$

9. SHARPNESS OF THE MAIN RESULT

In this section we prove Theorem 4. For $k \geq 2, t \geq 1$ we say that a positive integer r is (k,t)-*constructible in at most n steps* if there exists a sequence of non-negative integers $0 = a_0, a_1, \ldots, a_m = r, m \leq n$, such that for any $i \geq 1, a_i = f_i(a_0, \ldots, a_{i-1})$, with $f_i \in \mathbb{Z}[x_0, \ldots, x_{i-1}]$ a (k,t)-bounded polynomial.

The main step is to prove the following lemma.

Lemma 30. *Let $k \geq 2$. Any $p \geq 2^{32(k \log_2(16k))^2}$ is (k,k)-constructible in at most $\frac{10}{k} \frac{\log_2 p}{\log_2 \log_2 p}$ steps, and moreover this is sharp up to a constant not depending on k.*

Proof. Let $p \geq 2^{32(k \log_2(16k))^2}$ arbitrary. We first note the following inequality:

$$\log_2 \log_2 p \geq 2 \log_2(k \log_2 \log_2 p). \tag{13}$$

Indeed, this is true if $\log_2 p \geq k^2(\log_2 \log_2 p)^2$, which in turn is true if $\log p \geq \frac{2k^2}{\log 2}(\log \log p)^2$. By derivation this holds whenever $p \geq 2^{32(k \log_2(16k))^2} \geq e^{8(k \log_2(16k))^2}$.

Now set

$$s := \left\lceil \log_2 \left(\frac{\log_2 p}{k \log_2 \log_2 p} \right) \right\rceil \quad \text{and} \quad N := \lfloor \log_2 p \rfloor .$$

Note that $s \geq 1$, as $\log_2 p > k \log_2 \log_2 p$ by (13).

Consider the base-2 representation $(b_0 b_1 \ldots b_N)$ of p, with b_0 being the least significant bit. We break it into $\ell := \left\lceil \frac{N+1}{sk} \right\rceil \geq 1$ contiguous subsequences $(b_0 b_1 \ldots b_{sk-1}), \ldots, (b_{(\ell-1)sk} b_{(\ell-1)sk+1} \ldots b_N)$, all of them except possibly the last one of length sk, defining in base-2 numbers $p_0, p_1, \ldots, p_{\ell-1}$. Note that

$$p = \sum_{i=0}^{\ell-1} 2^{ski} p_i$$

and $p_i < 2^{sk}, 0 \leq i < \ell$. We further write

$$p_i = \sum_{j=0}^{k-1} 2^{sj} p_{ij},$$

with $0 \leq p_{ij} < 2^s$.

We now define the sequence $a_0, \ldots, a_{2^s+\ell+2(\ell-1)}$ as follows.

We start by setting $a_0 := 0$ and $a_i := a_{i-1} + 1, 1 \leq i \leq 2^s$. Note that $a_i = i, 1 \leq i \leq 2^s$. For any $0 \leq i \leq \ell - 1$ we define

$$a_{2^s+1+i} := \sum_{j=0}^{k-1} a_{2^s}^j a_{p_{ij}}.$$

Hence $a_{2^s+1+i} = p_i$. For any $1 \leq i \leq \ell - 1$ we further define $a_{2^s+\ell+2(i-1)+1}$ and $a_{2^s+\ell+2(i-1)+2}$ as follows:

$$a_{2^s+\ell+2(i-1)+1} := \begin{cases} a_{2^s}^k, & \text{if } i = 1, \\ a_{2^s+\ell+2(i-2)+1} a_{2^s+\ell+1}, & \text{otherwise.} \end{cases}$$

$$a_{2^s+\ell+2(i-1)+2} := \begin{cases} a_{2^s+\ell+1} a_{2^s+2} + a_{2^s+1}, & \text{if } i = 1, \\ a_{2^s+\ell+2(i-1)+1} a_{2^s+i+1} + a_{2^s+\ell+2(i-2)+2}, & \text{otherwise.} \end{cases}$$

Hence

$$a_{2^s+\ell+2(i-1)+1} = 2^{ski},$$

$$a_{2^s+\ell+2(i-1)+2} = \sum_{j=0}^{i} 2^{skj} p_j.$$

In particular, $a_{2^s+\ell+2(\ell-1)} = p$. Hence p is (k,k)-constructible in at most $2^s + 3\ell - 2$ steps. But

$$2^s + 3\ell - 2 \le 2^s + 1 + 3\frac{N+1}{sk}$$

$$\le 2^s + 4\frac{N}{sk}, \quad \text{as } sk + 3N + 3 \le 4N,$$

$$\le \frac{2}{k}\frac{\log_2 p}{\log_2 \log_2 p} + \frac{4}{k}\frac{\log_2 p}{\log_2 \log_2 p - \log_2(k \log_2 \log_2 p)}$$

$$\le \frac{10}{k}\frac{\log_2 p}{\log_2 \log_2 p}, \quad \text{by (13).}$$

This proves the first part of the lemma. To show that this bound is essentially best possible, we fix n and count the number of positive integers (k,k)-constructible in at most n steps.

First note that for given $\ell \ge 1$, the number of monomials in ℓ variables $x_1, \ldots, x_l$ of degree at most k is $\binom{\ell+k}{k} \le (kl)^k$. Hence the number of (k,k)-bounded polynomials in ℓ variables is at most $3^k\binom{\ell+k}{k} \le (3k\ell)^k$, as any such polynomial is a sum of k monomials in ℓ variables of degree at most k, with coefficients $1, -1$ or 0.

Now to any number which is (k,k)-constructible in at most n steps corresponds a sequence of (k,k)-bounded polynomials $f_1, \ldots, f_m, m \le n$, such that f_i is a polynomial in i variables. Thus the number of integers (k,k)-constructible in at most n steps is upper bounded by the number of such sequences, which for $n \ge 3k$ is at most

$$\prod_{i=1}^{n}(3ki)^k \le (3k)^{kn}n^{kn} \le n^{2kn}.$$

However if p is given, then for $n \le \frac{\log p}{2k \log \log p}$ we have

$$n^{2kn} \le \left(\frac{\log p}{2k \log \log p}\right)^{\frac{\log p}{\log \log p}} < p.$$

Hence not all numbers between 1 and p are (k,k)-constructible in at most $\frac{\log p}{2k \log \log p}$ steps. This finishes the proof of the lemma. $\qquad\square$

Proof of Theorem 4. Given $p \ge 2^{32(k-1)^2 \log_2^2(16(k-1))}$ a prime number, we apply Lemma 30 to find a sequence of non-negative integers $0 = a_0, \ldots, a_n = p, n \le \frac{10}{k-1}\frac{\log_2 p}{\log_2 \log_2 p}$, which shows that it is $(k-1, k-1)$-constructible. Let $A' := \{a_0, a_1, a_2, \ldots, a_n\}$. Taking the residues modulo p of the numbers in A' we obtain a set $A \subseteq \mathbb{F}_p$ of size at most n.

Now suppose for a contradiction that there exists an F_k-ring-isomorphism ϕ of A into an integral domain R of characteristic 0.

There is a natural embedding of $\mathbb{Z}$ into R, and we can identify $\mathbb{Z}$ with the image of this embedding. Let $x_i \in A$ be the image of a_i in $\mathbb{F}_p, 0 \le i \le n$. By induction on $i \ge 0$ we see that $\phi(x_i)$ must equal a_i.

This is certainly the case for $x_0 = 0$. For $i \ge 1$ there exists a $(k-1)$-bounded polynomial f_i such that $a_i = f_i(a_0, \ldots, a_{i-1})$. Hence $f_i(x_0, \ldots, x_{i-1}) - x_i = 0$ in $\mathbb{F}_p$. As this is a k-bounded polynomial, it must be preserved by ϕ. Therefore the induction hypothesis implies $\phi(x_i) = a_i$, as claimed.

However, A has size at most n, while A' has size $n+1$. Therefore the image of ϕ can not contain the whole of A', a contradiction. This proves the theorem. $\qquad\square$

The proof of Lemma 30 tells us that for given $M \geq 1$ there are only $(\log M)^{O(\log \log \log M)}$ positive integers less than M which are $(2,2)$-constructible in $O(\log \log M)$ steps. Nevertheless any Mersenne prime (a prime p of the form $2^n - 1$) is $(2,2)$-constructible in $O(\log n) = O(\log \log p)$ steps, by using the base-2 representation of n and an approach similar to that of Lemma 30. Furthermore any Fermat prime (a prime p of the form $2^{2^n} + 1$) is $(2,2)$-constructible in $O(n) = O(\log \log p)$ steps. Thus the existence of infinitely many such primes would imply Conjecture 5. Unfortunately proving or disproving such a statement seems at present to be an unreachable goal.

10. Concluding remarks

Remark 1. Theorem 4 does not cover the case $k = 2$, and in fact here I believe, but can not prove, that the correct bound is $\Theta(\log p)$; that is, any subset $A \subseteq \mathbb{F}_p$ of size $O(\log p)$ is F_2-ring-isomorphic to a subset of $\mathbb{C}$. Neither the proof of Theorem 1 nor that of Lemma 20 properly adapt to this situation, as one would have to work over the multiplicative group $\mathbb{F}_p^*$ of order $p - 1$.

Remark 2. Lemma 29 implies the following weaker version of Lemma 6: under the hypothesis of Lemma 6, there exists a solution $(b_1, \ldots, b_n) \in K^n$ to the polynomials $f_1, \ldots, f_s$, where K is a finite algebraic extension of $\mathbb{Q}$ of degree at most $(2t)^{2^n}$. Indeed, suppose each f_i has degree at most t and $\|\cdot\|_\infty$-norm at most k. Then each f_i is $(k(nt)^t, t)$-bounded. Fix $A := \{a_1, \ldots, a_n\}$, the coordinates of a complex solution of the system of polynomials $\{f_i : 1 \leq i \leq s\}$. We first apply Theorem 2 in order to find a sufficiently large prime p (compared to n, k and t) and a homomorphism $\phi : \mathbb{Z}[A] \to \mathbb{F}_p$. We then apply Lemma 29 to the collections $\mathcal{L}_1 := \{f_1, \ldots, f_s\}$ and $\mathcal{L}_2 := \emptyset$, in order to find a finite algebraic extension K of degree at most $(2t)^{2^n}$, a subset $A' \subset K$ and a map ψ between $\phi(A)$ and A'. Then $((\psi \circ \phi)(a_i))_{i=1}^n$ are the coordinates of a solution $(b_1, \ldots, b_n) \in K^n$ of the system of polynomials $\{f_i : 1 \leq i \leq s\}$.

Remark 3. In view of Theorem 11 one may ask what is the largest number $n(p)$ of points and lines in $\mathbb{F}_p^2$ for which the upper bound $cn(p)^{4/3}$ on the number of incidences holds. I have only proved $n(p) = \Omega(\log \log \log p)$, and I am not aware of any non-trivial upper bound for this function.

References

[1] S. Basu, R. Pollack and M. Roy, *Algorithms in Real Algebraic Geometry (Algorithms and Computation in Mathematics)* (Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006).

[2] Y. Bilu, V. F. Lev and I. Z. Rusza, 'Rectification Principles in Additive Number Theory', *Discrete Comput. Geom.* 19 (1998) 343–353.

[3] J. Bourgain and M. Z. Garaev, 'On a variant of sum-product estimates and explicit exponential sum bounds in prime fields', *Math. Proc. Cambridge Philos. Soc.* 146 (2009) 1–21.

[4] J. Bourgain, N. H. Katz and T. Tao, 'A sum-product estimate in finite fields and applications', *Geom. Funct. Anal.* 14 (2004) 27–57.

[5] W. S. Brown and J. F. Traub, 'On Euclid's Algorithm and the Theory of Subresultants', *J. ACM* 18 (1971) 505–514.

[6] M. Chang, 'Factorization in generalized arithmetic progressions and applications to the Erdős-Szemerédi sum-product problems', *Geom. Funct. Anal.* 13 (2003) 720–736.

[7] M. Chang, 'A sum-product estimate in algebraic division algebras', *Israel J. Math.* 150 (2005) 369–380.

[8] G. E. Collins, 'Polynomial Remainder Sequences and Determinants', *Amer. Math. Monthly* 73 (1966) 708–712.

[9] G. E. Collins, 'Subresultants and reduced polynomial remainder sequences', *J. ACM* 14 (1967) 128–142.

[10] D. Coppersmith and J. Davenport, 'Polynomials whose powers are sparse', *Acta Arith.* LVIII (1991) 79–87.

[11] G. Elekes, 'On the number of sums and products', *Acta Arith.* LXXXI (1997) 365–367.

[12] P. Erdős, 'On the number of terms of the square of a polynomial', *Nieuw Arch. Wiskunde* 23 (1949) 63–65.

[13] P. Erdős and E. Szemerédi, 'On sums and products of integers', *Studies in Pure Math.* (Birkhäuser, Basel, 1983) 213–218.

[14] K. Ford, 'Sums and Products from a Finite Set of Real Numbers', *Ramanujan J.* 2 (1998) 59–66.

[15] M. Z. Garaev, 'An explicit sum-product estimate in $\mathbb{F}_p$', *Intern. Math. Res. Notices* 2007 (2007) 1–11.

[16] B. Green and I. Z. Rusza, 'Sets with small subsets and rectification', *Bull. London Math. Soc.* 38 (2006) 43–52.

[17] H. A. Helfgott and M. Rudnev, 'An explicit incidence theorem in $\mathbb{F}_p$', *Mathematika* 57 (2011) 135–145.

[18] R. Hermann, *Linear Systems Theory and Introductory Algebraic Geometry* (Math Sci Press, Brookline, Mass., 1974).

[19] T. G. F. Jones, 'Further improvements to incidence and Beck-type bounds over prime finite fields', Preprint, Available online at `http://arxiv.org/abs/1206.4517`, 2012.

[20] N. H. Katz and C. Shen, 'A slight improvement to Garaev's sum product estimate', *Proc. Amer. Math. Soc.* 136 (2008) 2499–2504.

[21] S. V. Konyagin and M. Rudnev, 'On New Sum-Product Type Estimates', *SIAM J. Discrete Math.* 27 (2013) 973–990.

[22] T. Krick, L. M. Pardo and M. Sombra, 'Sharp estimates for the arithmetic Nullstellensatz', *Duke Math. J.* 109 (2001) 521–598.

[23] L. Li, 'Slightly improved sum-product estimates in fields of prime order', *Acta Arith.* 147 (2011) 153–160.

[24] M. B. Nathanson, 'On sums and products of integers', *Proc. Amer. Math. Soc.* 125 (1997) 9-16.

[25] A. Rényi, 'On the minimal number of terms of the square of a polynomial', *Hungarica Acta Math.* 1 (1947), 30–34, reprinted in *Selected papers of Alfréd Rényi*, vol. 1, 42–47, edited by Paul Turán (Akadémiai Kiadó, Budapest, 1976).

[26] M. Rudnev, 'An improved sum-product inequality in fields of prime order', *Intern. Math. Res. Notices* 2012 (2012) 3693–3705.

[27] A. Schinzel, 'On the number of terms of a power of a polynomial', *Acta Arith.* XLIX (1987) 55–70.

[28] A. Schinzel and U. Zannier, 'On the number of terms of a power of a polynomial', *Atti Accad. Naz. Lincei Cl. Sci. Fis. Mat. Natur. Rend. Lincei Mat. Appl.* 20 (2009) 95–98.

[29] J. Solymosi, 'On the number of sums and products', *Bull. London Math. Soc.* 37 (2005) 491–494.

[30] J. Solymosi, 'Bounding multiplicative energy by the sumset', *Adv. Math.* 222 (2009) 402–408.

[31] J. Solymosi and T. Tao, 'An Incidence Theorem in Higher Dimensions', *Discrete Comput. Geom.* 48 (2012) 255–280.

[32] T. Tao, 'Polynomial bounds via nonstandard analysis', 5 July 2011, `http://terrytao.wordpress.com/2011/07/05/polynomial-bounds-via-nonstandard-analysis/`.

[33] T. Tao, 'Rectification and the Lefschetz principle', 14 March 2013, `http://terrytao.wordpress.com/2013/03/14/rectification-and-the-lefschetz-principle/`.

[34] T. Tao and V. H. Vu, *Additive Combinatorics* (Cambridge University Press, New York, USA, 2009).

[35] C. D. Tóth, 'The Szemerédi-Trotter Theorem in the Complex Plane', Preprint, Available online at `http://arxiv.org/abs/math/0305283`, 2003.

[36] W. Verdenius, 'On the number of terms of the square and the cube of polynomials', *Indag. Math.* 11 (1949) 459–465.

[37] V. H. Vu, M. M. Wood and P. M. Wood, 'Mapping incidences', *J. London Math. Soc.* 84 (2011) 433–445.

[38] J. Zahl, 'A Szemerédi-Trotter type theorem in $\mathbb{R}^4$', Preprint, Available online at `http://arxiv.org/abs/1203.4600`, 2012.

Institut für Mathematik, Freie Universität Berlin, Arnimallee 3-5, D-14195 Berlin, Germany
  E-mail address: grosu.codrut@gmail.com