Hyperplane covers of finite spaces and applications

János Nagy[*] Péter Pál Pach[†] István Tomon[‡]

Abstract

Given a prime p and positive integer n, let $f_p(n)$ denote the minimal number of hyperplanes in an irredundant covering of $\mathbb{F}_p^n$ such that the normal vectors of the hyperlanes span the whole space. The function $f_p(n)$ appears in connection to several longstanding conjectures in linear algebra and group theory, such as the Alon-Jaeger-Tarsi conjecture, the Additive basis conjecture, and a conjecture of Pyber on irredundant coset covers of abelian groups. We prove that

$$f_p(n)=\Omega\left(\frac{\log p}{\log\log p}\cdot n\right),$$

and use this result to make substantial progress on each of the aforementioned conjectures.

1 Introduction

Let p be a prime (or prime power) and n be a positive integer. At least how many hyperplanes are needed to cover the finite space $\mathbb{F}_p^n$? Without further restrictions, the answer is trivially p, as one can take all p translates of any given hyperplane, and this is optimal. However, as is shown by the celebrated result of Alon and Füredi [1], the answer changes drastically if we remove a single point of $\mathbb{F}_p^n$. In this case, at least $(p-1)n$ hyperplanes are needed, and this bound is the best possible. Since, hyperplane covers with various restrictions are extensively studied, see e.g. [12, 22, 23].

In this paper, we study hyperplane covers originating from a work of Szegedy [26]. Given a hyperplane H, let $H^\perp$ denote a normal vector of H, that is, some $v\in\mathbb{F}_p^n\setminus\{0\}$ such that H is given by the equation $\langle v,x\rangle+t=0$ for some $t\in\mathbb{F}_p$ (note that v is not unique, but this will cause no issue later). A covering of a set with a collection of its subsets is *irredundant*, if no proper subcollection is a covering. We are interested in the minimum number of hyperplanes in an irredundant covering $\mathcal{H}$ of $\mathbb{F}_p^n$ such that $\operatorname{span}\langle H^\perp : H\in\mathcal{H}\rangle=\mathbb{F}_p^n$. Let us denote this minimum by $f_p(n)$. As observed by Szegedy [26], the problem of estimating $f_p(n)$ is closely related to several longstanding conjectures in linear algebra and group theory. These applications serve as the main motivation for studying this problem. We discuss these connections in later subsections. First, let us explore some properties of $f_p(n)$.

Note that $f_p(n)\geq n$ trivially holds by the condition that the normal vectors of the hyperplanes span an n-dimensional space. This can be easily improved to $f_p(n)\geq n+1$, which is then tight for $p=2$. For $p\geq 3$, it seems already highly difficult to prove that $f_p(n)\geq(1+\varepsilon_p)n$ with some $\epsilon_p>0$, and such bounds would already have some unexpected consequences. Our main theorem is the following lower bound.

[*] Alfréd Rényi Institute of Mathematics and MTA-BME Lendület Arithmetic Combinatorics Research Group, ELKH, *email*: **janomo4@gmail.com**

[†] MTA-BME Lendület Arithmetic Combinatorics Research Group, ELKH, Department of Computer Science and Information Theory, Budapest University of Technology and Economics, *email*: **ppp@cs.bme.hu**

[‡] Umeå University, *e-mail*: **istvan.tomon@umu.se**

Theorem 1.1. *For every prime p and positive integer n,*

$$f_p(n) \geq (1 - o(1)) \cdot \frac{\log p}{\log \log p} \cdot n,$$

where the $o(1)$ error term depends only on p. Moreover, if $p \geq 5$, then there exists $\varepsilon_p > 0$ such that $f_p(n) \geq (1 + \varepsilon_p)n$.

Moreover, for proper prime powers, we have the following slightly stronger result.

Theorem 1.2. *For every prime power $q = p^\alpha$ and positive integer n,*

$$f_q(n) \geq (1 - o(1)) \cdot \frac{\log q}{\log \log p} \cdot n,$$

where the $o(1)$ error term depends only on p.

In the planar case, we have $f_q(2) = q + 1$. The upper bound follows by considering the $q + 1$ lines going through the origin, while the lower bound is a simple exercise, which we omit here. In general, we establish the following upper bound.

Proposition 1.3. *Let q be a prime power. Then $f_q(n) \leq \lceil \frac{n}{2} \rceil \cdot q + 1$.*

In particular, we prove that $f_q(n) - 1$ is subadditive, which then implies by Fekete's lemma that $\lim_{n\to\infty} \frac{f_q(n)}{n}$ exists for every q. There is still a large gap between our lower and upper bound on this limit, and we believe the upper bound should be closer to the truth. Therefore, we propose the following conjecture.

Conjecture 1.4. *There exists $c > 0$ such that for every prime (or prime power) p and integer n,*

$$f_p(n) \geq cpn.$$

In the upcoming subsections, we discuss the applications of our lower bound.

1.1 The Alon-Jaeger-Tarsi conjecture

The Alon-Jaeger-Tarsi conjecture [3, 8] states that if $p \geq 5$ is a prime and $M \in \mathbb{F}_p^{n\times n}$ is an invertible matrix, then there exists some vector $x \in \mathbb{F}_p^n$ such that neither x, nor Mx has a zero coordinate. Alon and Tarsi [3] proved that the conjecture holds if p is a proper prime power. However, the conjecture remained open for every prime p until recently, when the first two authors [14] proved that it holds if p is sufficiently large.

Despite the lack of early progress, DeVos [5] proposed a substantial strengthening of the Alon-Jaeger-Tarsi conjecture, which he coined as the Choosability conjecture. A matrix $M \in \mathbb{F}_p^{n\times n}$ is (a, b)-*choosable* if for all subsets $X_1, \dots, X_n, Y_1, \dots, Y_n \subset \mathbb{F}_p$ such that $|X_i| = a, |Y_i| = b$ for $i \in [n]$, there exists a vector $x \in X_1 \times \dots \times X_n$ such that $Mx \in Y_1 \times \dots \times Y_n$.

Conjecture 1.5. *If $M \in \mathbb{F}_p^{n\times n}$ is invertible, then M is $(k + 2, p - k)$-choosable for every $k \in [p - 2]$.*

Observe that if M is $(p - 1, p - 1)$-choosable, then it satisfies the desired condition of the Alon-Jaeger-Tarsi conjecture. The first two authors [14] proved that if $p \geq 61$, $p \neq 79$, then M is $(p - 1, p - 1)$-choosable. But how is this problem related to hyperplane covers? As we will show later, the inequality $f_p(n) > 2kn$ implies that every M is $(p-k, p-k)$-choosable (see Lemma 6.1). Therefore, we get the following corollary of Theorem 1.1 and Theorem 1.2.

Theorem 1.6. *For every prime power $q = p^\alpha$ and positive integer n, every invertible $M \in \mathbb{F}_p^{n\times n}$ is $(q-k, q-k)$-choosable if $k \leq (\frac{1}{2} - o(1)) \cdot \frac{\log q}{\log\log p}$, where the $o(1)$ error term depends only on p.*

Another far reaching generalization of the Alon-Jaeger-Tarsi conjecture was proposed by the first two authors [14]. The following theorem resolves exactly this.

Theorem 1.7. *Let $k \geq 2$ be a positive integer, then there exists $q_0 = q_0(k)$ such that the following holds for every positive integer n. Let $q > q_0$ be a prime a power, and let $M_1, \dots, M_k \in \mathbb{F}_q^{n\times n}$ be invertible matrices. Then there exists $x \in \mathbb{F}_q^n$ such that the vectors $M_1x, \dots, M_kx$ have no zero coordinates.*

Note that $q_0(k)$ indeed needs to grow with k, in particular $q_0(k) \geq k+1$. Otherwise, if $n > k \geq q$, then let $M_i \in \mathbb{F}_q^{n\times n}$ for $i \in [k]$ be defined as $(M_ix)(i) = x(i)$ and $(M_ix)(j) = x(i) - x(j)$ for $j \in [n]\setminus\{i\}$. If $x \in \mathbb{F}_q^n$, then x has two equal coordinates among the first $q+1$ coordinates, say $x(j) = x(j')$ with $j < j' \leq q+1$, in which case $(M_jx)(j') = 0$. Later, we provide a common extension of Theorem 1.6 and Theorem 1.7, which the interested reader can find as Theorem 6.2.

1.2 The Additive Basis conjecture

If p is a prime, n is a positive integer, and $A \subset \mathbb{F}_p$, a multiset $B \subset \mathbb{F}_p^n$ is called an *A-basis*, if every vector $w \in \mathbb{F}_p^n$ can be written as $w = \sum_{v\in B} \alpha_v v$, where $\alpha_v \in A$ for every $v \in B$. Also, an *additive basis* is a $\{0,1\}$-basis. Clearly, if B_0 is a linear basis, and B is the union of $p-1$ copies of B_0, then B is an additive basis (here and later, union is taken as a multiset). The Additive Basis conjecture of Jaeger, Linial, Payan, and Tarsi [11] asks whether there exists a constant $c_1(p)$ (possibly $c_1(p) = p$) such that the union of $c_1(p)$ linear bases is always an additive basis. More precisely, this conjecture first appeared in a work of Alon, Linial, and Meshulam [2], who proved that the union of $\Omega(p\log n)$ bases is an additive basis in $\mathbb{F}_p^n$, but in [2], the conjecture is attributed to [11].

Szegedy [26] proposed a weakening of this conjecture, which is referred to as the Weak Additive Basis conjecture. This states that for every prime $p \geq 3$ there exists a constant $c_2(p)$ such that if $B \subset \mathbb{F}_p^n$ is the union of $c_2(p)$ bases, then B is a $\{1, \dots, p-1\}$-basis. Szegedy observed that if the inequality $f_p(n) \geq (1+\varepsilon_p)n$ holds with some $\varepsilon_p > 0$, then the Weak additive basis conjecture also holds for p. Hence, our Theorem 1.1 resolves this conjecture for every $p \geq 5$. In the case $p = 3$, the Weak additive basis conjecture is equivalent to the Additive basis conjecture, which then reamins open. However, we can establish the following substantial strengthening of the Weak additive basis conjecture for large primes.

Theorem 1.8. *Let $p \geq 5$ be a prime, and n be a positive integer. There exists $A \subset \mathbb{F}_p$ of size $(1+o(1))\log_2 p$ such that the union of p bases in $\mathbb{F}_p^n$ is an A-basis.*

This theorem does not directly follow from our bounds on $f_p(n)$, it is rather a byproduct of the proof Theorem 1.1. One can also ask whether the previous theorem can be extended for proper prime powers $q = p^\alpha$. Clearly, the Additive basis conjecture does not hold in this case: if B is the union of any number of copies of a basis B_0, then the only elements that can be expressed as a 0-1 linear combinations of elements of B are in $\mathbb{F}_p \cdot B_0$. On the other hand, a reasonable conjecture to make is the following.

Conjecture 1.9. *Let $q = p^\alpha$, then there exists $c > 0$ such that the following holds for every integer n. Let $A_0 \subset \mathbb{F}_q$ be a basis of $\mathbb{F}_q$ over $\mathbb{F}_p$, and let $A = A_0 \cup \{0\}$. Then the union of c bases in $\mathbb{F}_q^n$ is an A-basis.*

We can establish the following weaker variant of this conjecture, following Theorem 1.8.

Theorem 1.10. *Let $q = p^\alpha$ be a prime power, and n be a positive integer. There exists $A \subset \mathbb{F}_q$ of size $(1+o(1))\log_2 q$ such that the union of $\alpha \cdot p$ bases in $\mathbb{F}_q^n$ is an A-basis.*

The Additive Basis conjecture is closely related to a celebrated conjecture of Jaeger [9] about the existence of modulo k-orientations in graphs, which in particular is motivated by an old problem of Tutte (see e.g. [24]) on nowhere zero 3-flows. A *modulo k-orientation* in a graph G is an orientation of the edges such that for every vertex, the in- and outdegree are equal modulo k. Let $c_3(k)$ denote the smallest constant such that any $c_3(k)$-edge-connected graph has a modulo k-orientation. Jaeger's Circular Flow conjecture [9] states that $c_3(k) = 2k-2$ for every odd k. On the other hand, a weak version of this conjecture [10] asks whether $c_3(k)$ exists at all. The weak version of the conjecture was proved by Thomassen [27], and in a subsequent paper, Lovász, Thomassen, Wu, and Zhang [13] established the upper bound $c_3(k) \leq 3k-3$. However, the strong version of the conjecture was disproved by Han, Li, Zhu, and Wang [7] for $k \geq 7$. In [11], it is demonstrated that if the Additive Basis conjecture is true, then $c_3(p) \leq 2c_1(p)$ (more precisely, this was shown for $p = 3$, but the same argument works in general).

1.3 Coset covers of abelian groups

An old result of Neumann [16] states that if G is a group and $\{H_i x_i : i \in [k]\}$ is an irredundant covering of G with cosets (where H_i is a subgroup of G, and $x_i \in G$), then the index $|G : \bigcap_{i\in[k]} H_i|$ is finite, and in [17], Neumann proved that $|G : \bigcap_{i\in[k]} H_i|$ is bounded by a function of k. Therefore, it makes sense to define $f(k)$ denoting the maximum of $|G : \bigcap_{i\in[k]} H_i|$, where G is a group and $\{H_i x_i : i \in [k]\}$ is an irredundant covering of G with cosets. Similarly, define $g(k)$ to be the maximum of $|G : \bigcap_{i\in[k]} H_i|$ if $\{H_i : i \in [k]\}$ is an irredundant covering of G with subgroups. Tomkinson [28] proved that $f(k) = k!$ for every k, the lower bound achieved by the symmetric group, while $\Omega(3^{2k/3}) = g(k) < (k-2)^3(k-3)!$ for $k \geq 5$. It is a longstanding open problem whether g grows at most exponentially.

Pyber [20] observed that in case G is an elementary p-group, the latter question is closely related to the problem of covering a group by abelian subgroups. He conjectured that an exponential upper bound should hold at least in this case. See e.g. [19] for further details about the problem of covering by abelian subgroups. To this end, define $f_A(k)$ and $g_A(k)$ the same way as $f(k)$ and $g(k)$, respectively, with the additional restriction that G is abelian. Szegedy [26] conjectured that $f_A(k) = 2^{O(k)}$ might also be true, immediately implying $g_A(k) = 2^{O(k)}$ and the conjecture of Pyber as well. Note that if true, then this bound is the best possible up to the constant hidden in the $O(.)$ notation. For example, one can take $G = \mathbb{Z}_2^n$, which has a covering with $n+1$ subgroups, whose intersection is trivial.

As the reader may observe, this problem is of very similar flavour to our question about hyperplane covers. Indeed, Szegedy proved that if the inequality $f_p(n) \geq \Omega(n \log p)$ holds for every prime p and integer n, then $g_A(k) = f_A(k) = 2^{O(k)}$. Unfortunately, our Theorem 1.1 is not strong enough to prove this, but we can still establish the following substantial improvement over the non-abelian case.

Theorem 1.11. *There exists $c > 0$ such that the following holds. Let G be an abelian group, and let $\{H_i x_i : i \in [k]\}$ be an irredundan covering of G with cosets. Then*

$$|G : \bigcap_{i\in[k]} H_i| \leq e^{ck \log\log k}.$$

Organization. Our paper is organized as follows. In the next section, we introduce some notation, and discuss basic notions of linear algebra and the discrete Fourier transform. Then, in section 3, we define group rings and prove several results about their properties. Then, Section 4 contains most of the combinatorial ideas required for the proof. In Section 5, we prove Theorem 1.1 and Proposition 1.3; the proof of Theorem 1.2 is delayed until Section 8. In Section 6, we prove our results related the the Alon-Jaeger-Tarsi conjecture, that is, Theorem 1.7. In Section 7, we prove our main results about additive bases, that is, Theorems 1.8 and 1.10. In Section 8, we consider coset coverings of groups, and prove Theorem 1.11, and also Theorem 1.2. We finish our paper with some concluding remarks.

2 Preliminaries

2.1 Linear algebra

Let us introduce some notation. Given a prime power q, an integer n, a vector $v \in \mathbb{F}_q^n$ and $t \in \mathbb{F}_p$, let $H_{q,n}(v,t)$ denote the hyperplane $\{x \in \mathbb{F}_q^n : \langle v, x\rangle + t = 0\}$. If q and n are clear from the context, we simply write $H(v,t)$.

Given a multiset of vectors $V \subset \mathbb{F}_q^n$, let

$$\ker(V) = \left\{ w \in \mathbb{F}_q^V : \sum_{v \in V} w(v)v = 0 \right\},$$

and let $\dim(V)$ be the dimension of the vector space spanned by the elements of V. Note that

$$\dim(V) + \dim(\ker(V)) = |V|.$$

2.2 Discrete Fourier transform

Given a function $h : \mathbb{F}_p^n \to \mathbb{C}$, the *(discrete) Fourier transform of* h is the function $F(h) : \mathbb{F}_p^n \to \mathbb{C}$ defined as

$$F(h)(y) = \sum_{x \in \mathbb{F}_p^n} e^{\frac{2\pi i}{p}\langle x,y\rangle} h(x).$$

We will use the following basic properties of the Fourier transform. For a function $h : \mathbb{F}_p^n \to \mathbb{C}$, we have $h \equiv 0$ if and only if $F(h) \equiv 0$.

The *convolution* of two functions $h, h' : \mathbb{F}_p^n \to \mathbb{C}$ is the function $h * h' : \mathbb{F}_p^n \to \mathbb{C}$ defined as

$$(h * h')(y) = \sum_{x \in \mathbb{F}_p^n} h(x)h'(y - x).$$

We have the following identity: $F(h * h') = F(h) \cdot F(h')$.

3 Group rings

In this section, we study the group rings $\mathbb{C}[\mathbb{F}_p^n]$ and $\mathbb{F}_p[\mathbb{F}_p^n]$, and establish connections between hyperplane covers and group ring products. This will give the main theoretical background needed for our proofs. Most of the material covered in this section is a reformulation or clarification of results presented in [14]. First, let us formally introduce the notion of group ring.

Given an additive group G and a ring R, the *group ring* $R[G]$ is the ring of formal expressions $\sum_{v \in G} r_v g^v$, where $r_v \in R$, and g is a formal variable. (Here, the formal exponentiation g^v is used to turn the additive structure of G into a multiplicative one.) Addition and multiplication are defined in the natural way, that is,

$$\left(\sum_{v\in G} r_v g^v\right) + \left(\sum_{v\in G} r'_v g^v\right) = \sum_{v\in G}(r_v + r'_v)g^v,$$

and

$$\left(\sum_{v\in G} r_v g^v\right) \cdot \left(\sum_{v\in G} r'_v g^v\right) = \sum_{v\in G}\left(\sum_{w\in G} r_w r'_{v-w}\right) g^v.$$

Note that an element $h = \sum_{v\in G} r_v g^v \in R[G]$ corresponds to the function $h^* : G \to R$ defined as $h^*(v) = r_v$. Then, the product $h_1 \cdot h_2$ corresponds to the convolution $h_1^* * h_2^*$.

We will study minimally vanishing products in the group rings $\mathbb{C}[\mathbb{F}_p^n]$ and $\mathbb{F}_p[\mathbb{F}_p^n]$. To this end, we introduce the following definition.

Definition 1. A multiset $V \subset \mathbb{F}_p^n$ is $\mathbb{F}_p$-*vanishing* if

$$\prod_{v\in V}(1-g^v) = 0$$

in $\mathbb{F}_p[\mathbb{F}_p^n]$. Also, say that V is $\mathbb{F}_p$-*irredundant* if V is $\mathbb{F}_p$-vanishing, but no proper subset of V is $\mathbb{F}_p$-vanishing.

Furthermore, V is $\mathbb{C}$-*vanishing* if there exists $(t_v)_{v\in V} \in (\mathbb{F}_p)^V$ such that

$$\prod_{v\in V}(1-e^{\frac{2\pi i t_v}{p}}g^v) = 0$$

in $\mathbb{C}[\mathbb{F}_p^n]$. Also, say that V is $\mathbb{C}$-*irredundant* if there exists $(t_v)_{v\in V} \in (\mathbb{F}_p)^V$ such that $\prod_{v\in V}(1 - e^{\frac{2\pi i t_v}{p}}g^v) = 0$, but no proper subset $V' \subset V$ satisfies $\prod_{v\in V'}(1 - e^{\frac{2\pi i t_v}{p}}g^v) = 0$. (Note that this is not equivalent to saying that no proper subset of V is $\mathbb{C}$-vanishing.)

In [14], it was shown that if V is $\mathbb{C}$-vanishing, then it is also $\mathbb{F}_p$-vanishing, however, we will not use this fact. The next lemma connects irredundant sets of vectors with irredundant hyperplane covers.

Lemma 3.1. *Let $V \subset \mathbb{F}_p^n$ be a multiset. Then some translates of the hyperplanes $H(v,0)$ for $v \in V$ form an irredundant cover of $\mathbb{F}_p^n$ if and only if V is $\mathbb{C}$-irredundant.*

Proof. For $v \in \mathbb{F}_p^n$ and $t \in \mathbb{F}_p$, let $h_{v,t} = 1 - e^{\frac{2\pi i t}{p}}g^v \in \mathbb{C}[\mathbb{F}_p^n]$. Then $h_{v,t}^* : \mathbb{F}_p^n \to \mathbb{C}$ is the function defined as $h_{v,t}^*(0) = 1$, $h_{v,t}^*(v) = -e^{\frac{2\pi i t}{p}}$, and $h_{v,t}^*(x) = 0$ for $x \in \mathbb{F}_p^n \setminus \{0, v\}$. Consider the discrete Fourier transform of $h_{v,t}^*$, that is $F(h_{v,t}^*)$. It is easy to calculate that $F(h_{v,t}^*)(y) = 1 - e^{\frac{2\pi i}{p}(t+\langle y,v\rangle)}$ for every $y \in \mathbb{F}_p^n$. But then the set of vectors on which $F(h_{v,t}^*)$ vanishes is exactly the hyperplane $H(v,t)$.

Therefore, if for each $v \in V$ we assign some $t_v \in \mathbb{F}_p^n$, then the hyperplanes $H(v, t_v)$ for $v \in V$ form a covering of $\mathbb{F}_p^n$ if and only if

$$\prod_{v\in V} F(h_{v,t_v}^*) \equiv 0.$$

This is equivalent to the convolution of the functions h_{v,t_v}^* for $v \in V$ being 0, which is further equivalent to $\prod_{v\in V} h_{v,t} = 0$. Therefore, V is $\mathbb{C}$-irredundant if and only if some translates of the hyperplanes $H(v,0)$ for $v \in V$ form an irredundant cover of $\mathbb{F}_p^n$. □

One of the first (and very few) applications of group ring identities in combinatorics is the celebrated theorem of Olson [18] about vanishing sums in abelian groups whose order is a prime power. Olson's proof relied on an identity, which we state (and prove, for completeness) in a somewhat weaker form.

Lemma 3.2. *Let $V \subset \mathbb{F}_p^n$ be a multiset of size at least $(p-1)n+1$. Then V is $\mathbb{F}_p$-vanishing.*

Proof. Let $e_1, \dots, e_n$ be a basis of $\mathbb{F}_p^n$. Note that if $v \in \mathbb{F}_p^n$, then we can write

$$1 - g^v = \sum_{i=1}^n f_{v,i} \cdot (1 - g^{e_i})$$

with suitable $f_{v,1}, \dots, f_{v,n} \in \mathbb{F}_p[\mathbb{F}_p^n]$. Indeed, if $v = \sum_{i=1}^n b_i e_i$, then

$$1 - g^v = \sum_{i=1}^n g^{b_1 e_1 + \dots + b_{i-1} e_{i-1}} \cdot (1 - g^{b_i e_i}),$$

so we can take $f_{v,i} = g^{b_1 e_1 + \dots + b_{i-1} e_{i-1}} \cdot (1 + g^{e_i} + \dots + g^{(b_i - 1) e_i})$. Consider the product

$$\prod_{v \in V} (1 - g^v) = \prod_{v \in V} \left(\sum_{i=1}^n f_{v,i} (1 - g^{e_i}) \right).$$

After expanding the outer brackets on the right hand side, we get a sum, whose every term has the form $f(1-g^{e_1})^{\alpha_1} \dots (1 - g^{e_n})^{\alpha_n}$, where $\alpha_1 + \dots + \alpha_n = |V| > (p-1)n$ and $f \in \mathbb{F}_p[\mathbb{F}_p^n]$. Therefore, in each such term, at least one of the α_i's is at least p. But $(1 - g^{e_i})^p = 0$, so every term evaluates to 0. □

Interestingly, we can establish a $\mathbb{C}$ analogue of this lemma with slightly weaker bounds, with a geometric proof. We present this in order to uncover further interesting connections between the group rings $\mathbb{C}[\mathbb{F}_p^n]$ and $\mathbb{F}_p[\mathbb{F}_p^n]$.

Lemma 3.3. *Let $V \subset \mathbb{F}_p^n$ be a multiset of size at least $np \log p$. Then V is $\mathbb{C}$-vanishing.*

Proof. By Lemma 3.1, it is enough to prove that for every $v \in V$ we can choose $t_v \in \mathbb{F}_p$ such that the hyperplanes $H(v, t_v)$ form a covering of $\mathbb{F}_p^n$. Choose each t_v randomly and independently from the uniform distribution on $\mathbb{F}_p$. Given $x \in \mathbb{F}_p^n$, we have $\mathbb{P}(x \in H(v, t_v)) = \frac{1}{p}$, so

$$\mathbb{P}\left(x \notin \bigcup_{v \in V} H(v, t_v) \right) \leq \left(1 - \frac{1}{p}\right)^{|V|} < e^{-|V|/p} \leq p^{-n}.$$

Therefore, with positive probability, there is a choice $\{t_v\}_{v \in V}$ such that $\bigcup_{v \in V} H(v, t_v) = \mathbb{F}_p^n$. □

The proof of Lemma 3.2 builds on the observation that if in the product $\prod_{v \in V}(1 - g^v) = \sum_y c_y g^y$ the constant term vanishes (that is, $c_0 = 0$), then V must contain a nonempty subset whose elements sum to 0. However, we will show that the whole product $\prod_{v \in V}(1 - g^v)$ being zero carries much more information about V. In particular, all of our main results rely on the following key lemma, which we state after providing the following key definition.

Definition 2. A subspace $W < \mathbb{F}_p^N$ is *versatile* if for every $x \in \mathbb{F}_p^N$ with no zero coordinates and every index $j \in [N]$ there exists $w \in W$ such that $w(i) \in \{-x(i), 0, x(i)\}$ for every $i \in [N]$, and $w(j) = x(j)$.

Lemma 3.4. *Let $R \in \{\mathbb{C}, \mathbb{F}_p\}$, and let $V \subset \mathbb{F}_p^n$ be an R-irredundant multiset. Then* $\ker(V)$ *is versatile.*

In order to prove this, we use the following simple but crucial observation: being $\mathbb{C}$-irredundant or $\mathbb{F}_p$-irredundant is a projective property.

Lemma 3.5. *Let $R \in \{\mathbb{C}, \mathbb{F}_p\}$, let $V \subset \mathbb{F}_p^n$ be a multiset and $(a_v)_{v \in V} \in (\mathbb{F}_p^*)^V$. Set $V' = \{a_v v : v \in V\}$. Then V is R-irredundant if and only if V' is R-irredundant.*

Proof. In case $R = \mathbb{F}_p$, this follows easily from the identity

$$(1 - g^{av}) = (1 - g^v) \cdot \left(\sum_{i=0}^{a-1} g^{iv} \right).$$

Indeed, writing $f = \prod_{v \in V} (\sum_{i=0}^{a_v - 1} g^{iv})$ and $f' = \prod_{v \in V} (\sum_{i=0}^{a_v^{-1} - 1} g^{i a_v v})$, we have $f \cdot \prod_{v \in V} (1 - g^v) = \prod_{v \in V} (1 - g^{a_v v})$ and $\prod_{v \in V} (1 - g^v) = f' \cdot \prod_{v \in V} (1 - g^{a_v v})$. Hence, $\prod_{v \in V} (1 - g^v) = 0$ if and only if $\prod_{v \in V} (1 - g^{a_v v}) = 0$.

Similarly, in case $R = \mathbb{C}$, this follows from the identity

$$(1 - \lambda^{at} g^{av}) = (1 - \lambda^t g^v) \cdot \left(\sum_{i=0}^{a-1} \lambda^{it} g^{iv} \right).$$

Also, Lemma 3.1 gives a more geometric reason why this is true. We have that V is $\mathbb{C}$-irredundant if and only if there is an irredundant covering of $\mathbb{F}_p^n$ with hyperplanes such that V is a set of normal vectors of the hyperplanes. But multiplying the elements of V by any non-zero scalar does not change whether they are normal vectors. □

Proof of Lemma 3.4. We prove this only for $R = \mathbb{C}$, as almost the same proof applies for $R = \mathbb{F}_p$. We need to show that for every $x \in \mathbb{F}_p^V$ with no zero coordinates and $j \in V$, there exists $w \in \ker(V)$ such that $w(j) = x(j)$ and $w(v) \in \{-x(v), 0, x(v)\}$ for every $v \in V \setminus \{j\}$. Therefore, let us fix some $x \in \mathbb{F}_p^V$ with no zero coordinates and $j \in V$.

Define the multiset $V' = \{x(v)v : v \in V\} \subset \mathbb{F}_p^n$, then by Lemma 3.5, V' is also $\mathbb{C}$-irredundant. Therefore, there exists $(t_v)_{v \in V} \in \mathbb{C}^V$ such that

$$\prod_{v \in V} (1 - e^{\frac{2\pi i t_v}{p}} g^{x(v)v}) = 0,$$

but

$$\prod_{v \in V \setminus \{j\}} (1 - e^{\frac{2\pi i t_v}{p}} g^{x(v)v}) = \sum_{z \in \mathbb{F}_p^n} c_z g^z \neq 0.$$

Observe that if $c_z \neq 0$, then z is the sum of some elements of $V' \setminus \{x(j)j\}$. Fix $z_0 \in \mathbb{F}_p^n$ such that $c_{z_0} \neq 0$. Note that

$$0 = \prod_{v \in V} (1 - e^{\frac{2\pi i t_v}{p}} g^{x(v)v}) = (1 - e^{\frac{2\pi i t_j}{p}} g^{x(j)j}) \sum_{z \in \mathbb{F}_p^n} c_z g^z = \sum_{z \in \mathbb{F}_p^n} (c_z - e^{\frac{2\pi i t_j}{p}} c_{z - x(j)j}) g^z.$$

Hence, $c_{z_0 - x(j)j} \neq 0$. This implies that there exist $(s_v)_{v \in V \setminus \{j\}}, (s'_v)_{v \in V \setminus \{j\}} \in \{0,1\}^{V \setminus \{j\}}$ such that

$$z_0 = \sum_{v \in V \setminus \{j\}} s_v x(v) v$$

and

$$z_0 - x(j)j = \sum_{v \in V \setminus \{j\}} s'_v x(v)v.$$

Subtracting the first equality from the second, we get

$$x(j)j = \sum_{v \in V \setminus \{j\}} (s_v - s'_v)x(v)v.$$

This shows that the vector $w \in \mathbb{F}_p^V$ defined as $w(j) = x(j)$ and $w(v) = (s'_v - s_v)x(v)$ for $v \in V \setminus \{j\}$ is in $\ker(V)$. But as $s'_v - s_v \in \{-1, 0, 1\}$, w has the property that $w(j) = x(j)$ and $w(v) \in \{-x(v), 0, x(v)\}$ for $v \in V \setminus \{j\}$, finishing the proof. □

4 Versatile subspaces

In this section, we study properties of versatile subspaces. In particular, we show that they must be close to full-dimensional. The bounds presented here will serve as the main tools in the proofs our theorems. Let us start with a definition.

Definition 3. A set $A \subset \mathbb{F}_p$ is an *arithmetic set* if A is nonempty, and for every $a \in A$ there exists $b \in \mathbb{F}_p \setminus \{0\}$ such that $a - b, a + b \in A$.

In other words, A is an arithmetic set if A is nonempty and every element of A is the middle element of a 3-term arithmetic progression contained in A. The minimum size of arithmetic sets in $\mathbb{F}_p$ has been extensively studied, and it is proved in [4, 15] that this minimum is $(1 + o(1)) \log_2 p$. In the case of small primes, we point out that if $p \geq 5$, then any $p - 1$ element subset is an arithmetic set. The main connection between arithmetic sets and versatile subspaces is summarized in the following lemma.

Lemma 4.1. *Let $A \subset \mathbb{F}_p$ be an arithmetic set, and let $W < \mathbb{F}_p^N$ be a versatile subspace. Then*

$$W + A^N = \mathbb{F}_p^N.$$

Proof. It is enough to prove that for every $y \in \mathbb{F}_p^N$, we have $(y + W) \cap A^N \neq \emptyset$. Let $z \in y + W$ be a vector with the most coordinates contained in A, and let $I \subset [n]$ be the set of indices i such that $z(i) \in A$. If $I = [n]$, we are done, so we can assume that there exists some $j \in [n] \setminus I$ such that $z(j) \notin A$.

Define the vector $x \in \mathbb{F}_p^N$ as follows. Let $x(j)$ be any number such that $z(j) + x(j) \in A$. For $i \in I$, as $z(i) \in A$, there exists $b \neq 0$ such that $z(i) - b, z(i) + b \in A$; set $x(i) = b$. For $i \in [n] \setminus (I \cup \{j\})$, choose $x(i) \neq 0$ arbitrarily. Then x has no zero coordinates, so there exists $w \in W$ such that $w(j) = x(j)$ and $w(i) \in \{-x(i), 0, x(i)\}$ for every $i \in [N]$. Let $z' = z + w$, then $z' \in z + W = y + W$, $z'(j) \in A$ and $z'(i) \in A$ for every $i \in I$. Therefore, z' has more coordinates in A than z, contradicting the maximality of z. This finishes the proof. □

From this, one can get the following lower bound on the dimension of versatile subspaces.

Corollary 4.2. *Let s be the size of the smallest arithmetic set in $\mathbb{F}_p$. If $W < \mathbb{F}_p^N$ is versatile, then*

$$\dim(W) \geq \left(1 - \frac{\log s}{\log p}\right) N.$$

Proof. Let A be an arithmetic set of size s. By Lemma 4.1, we have $W + A^N = \mathbb{F}_p^N$. This implies that $|W| \cdot |A|^N \geq p^N$, so $|W| \geq (p/s)^N$. Hence, $\dim(W) \geq \log_p (p/s)^N$, which is equivalent to the desired inequality. □

It would be interesting to find tight bounds on the dimension of versatile subspaces. We believe that $\dim(W) \geq \left(1 - \frac{c}{\log p}\right) N$ should also hold with some constant $c > 0$.

5 Irredundant hyperplane covers

In this section, we prove Theorem 1.1, and Proposition 1.3. The proof of Theorem 1.2 will be completed in a later section. Theorem 1.1 follows immediately from the next result.

Theorem 5.1. *Let p be a prime, and let s be the size of the smallest arithmetic set in $\mathbb{F}_p$. Then $f_p(n) \geq \frac{\log p}{\log s} n$.*

Proof. Let $\mathcal{H}$ be an irredundant covering of $\mathbb{F}_p^n$ with N hyperplanes such that $\mathrm{span}\langle H^\perp : H \in \mathcal{H}\rangle = \mathbb{F}_p^n$. Let V be the multiset of normal vectors of the elements of $\mathcal{H}$, then $\dim(V) = n$, and V is $\mathcal{C}$-irredundant by Lemma 3.1. Let $W = \ker(V) < \mathbb{F}_p^V$, then W is versatile by Lemma 3.4. Hence, $\dim(W) \geq \left(1 - \frac{\log s}{\log p}\right) N$ by Corollary 4.2, where $s = (1 + o(1)) \log_2 p$ is the size of the smallest arithmetic set. But

$$N = |V| = \dim(V) + \dim(\ker(V)) \geq n + \left(1 - \frac{\log s}{\log p}\right) N,$$

from which we get $N \geq \frac{\log p}{\log s} n$. □

Proof of Theorem 1.1. Recalling that $s \leq p - 1$ if $p \geq 5$, and also $s = (1 + o(1)) \log_2 p$ in general, Theorem 5.1 implies the desired result. □

In applications, the condition that the normal vectors of the hyperplanes span the whole space is usually too much to ask. Instead, we use the following simple proposition.

Proposition 5.2. *Let $\mathcal{H}$ be an irredundant covering of $\mathbb{F}_q^n$ with hyperplanes, and let $k = \dim\{H^\perp : H \in \mathcal{H}\}$. Then $|\mathcal{H}| \geq f_q(k)$.*

Proof. Let $U = \mathrm{span}\langle H^\perp : H \in \mathcal{H}\rangle$ and $W = U^\perp = \{x \in \mathbb{F}_p^n : \forall u \in U, \langle x, u\rangle = 0\}$. Then $\dim(W) = n - k$. We can write $\mathbb{F}_p^n = W \oplus V$ with some $V < \mathbb{F}_p^n$, where $\dim(V) = k$. For each $H \in \mathcal{H}$, let H' be the projection of H onto V, and let $\mathcal{H}' = \{H' : H \in \mathcal{H}\}$. Then H' is a hyperplane in V, and $\mathcal{H}'$ is an irredundant covering of V such that $\mathrm{span}\langle H'^\perp : H' \in \mathcal{H}'\rangle = V$. Hence, $|\mathcal{H}| = |\mathcal{H}'| \geq f_p(k)$. □

Later, we will establish the following variant of Theorem 5.1 for prime powers, which then immediately implies Theorem 1.2.

Theorem 5.3. *Let $q = p^\alpha$ be a prime power, and let s be the size of the smallest arithmetic set in $\mathbb{F}_p$. Then $f_q(n) \geq \frac{\log q}{\log s} n$.*

Now let us turn to the proof of the upper bound on $f_q(n)$. First, we show that the function $f_q(n) - 1$ is subadditive. By Fekete's lemma, this also implies that the limit $\lim_{n\to\infty} \frac{f_q(n)}{n}$ exists for every q.

Lemma 5.4. *Let n, m be positive integers. Then $f_q(m + n) \leq f_q(m) + f_q(n) - 1$.*

Proof. Let $V \cong \mathbb{F}_q^m$, $W \cong \mathbb{F}_q^n$. Let $M = f_q(m)$, $N = f_q(n)$, let $H_i = H_{m,q}(v_i, t_i)$ for $i \in [M]$ be an irredundant covering of V with $\dim(\{v_i : i \in [M]\}) = m$, and let $H'_i = H_{n,q}(w_i, u_i)$ for $i \in [N]$ be an irredundant covering of W with $\dim(\{w_i : i \in [N]\}) = n$. We have $M \geq m + 1$, so we may assume that $v_1, \dots, v_{M-1}$ span V, and similarly $w_1, \dots, w_{N-1}$ span W.

Write the elements of $V \oplus W \cong \mathbb{F}_q^{m+n}$ as (v, w), where $v \in V$ and $w \in W$. Consider the following $M + N - 1$ hyperplanes in $V \oplus W$.

- $H''_i = H_{m+n,q}((v_i, 0), t_i)$ for $i \in [M-1]$,
- $H''_{i+M-1} = H_{m+n,q}((0, w_i), u_i)$ for $i \in [N-1]$,
- $H''_{M+N-1} = H_{m+n,q}((v_M, w_N), t_M + u_N)$.

We show that $\mathcal{H}'' = \{H''_1, \dots, H''_{M+N-1}\}$ form an irredundant covering of $V \oplus W$, and their normal vectors span $V \oplus W$. The latter follows as $(v_1, 0), \dots, (v_{M-1}, 0)$ and $(0, w_1), \dots, (0, w_{N-1})$ already span $V \oplus W$.

Now let us show that $\mathcal{H}''$ is a covering. Let $(x, y) \in V \oplus W$. If x is covered by H_i for some $i \in [M-1]$, then (x, y) is covered by H''_i. Also, if y is covered by H'_i for some $i \in [N-1]$, then (x, y) is covered by H_{i+M-1}. Otherwise, x is covered by H_M and y is covered by H'_N, in which case (x, y) is covered by H''_{N+M-1}.

Finally, we show that $\mathcal{H}''$ is irredundant. For $i \in [M]$, there exists $x_i \in V$ that is only covered by H_i. Also, for $i \in [N]$, there exists $y_i \in W$ that is only covered by H'_i. But then for $i = 1, \dots, M-1$, the vector (x_i, y_N) is only covered by H''_i. For $i = 1, \dot{,} N-1$, the vector (x_M, y_i) is only covered by H''_{i+M-1}. Finally, the vector (x_N, y_N) is only covered by H''_{M+N-1}. This shows that $\mathcal{H}''$ is irredundant. □

Proof of Proposition 1.3. Trivially, we have $f_q(1) = q$. Also, as we discussed in the Introduction, $f_q(2) = q + 1$, as taking all lines through the origin gives an irredundant covering of $\mathbb{F}_q^2$ with $q + 1$ lines, and this is optimal. But then by Lemma 5.4, if n is even, we have,

$$f_q(n) - 1 \leq \frac{n}{2} \cdot (f_q(2) - 1) = \frac{n}{2} q$$

and if n is odd, then

$$f_q(n) - 1 \leq (f_1(q) - 1) + \frac{n-1}{2} \cdot (f_q(2) - 1) = \frac{n+1}{2} q - 1.$$

□

6 The Alon-Jaeger-Tarsi conjecture

In this section, we prove Theorem 1.7, which turns out to be a simple consequence of Theorem 1.2. In particular, we establish a connection between the function $f_q(n)$ and the stronger choosability version of Theorem 1.7.

Lemma 6.1. *Let q be a prime power, k, r be positive integers. If the inequality $f_q(N) > krN$ holds for every positive integer N, then the following holds for every positive integer n. Let $M_1, \dots, M_k \in \mathbb{F}_q^{n \times n}$ be invertible matrices. Then given $k \cdot n$ sets $X_{i,j} \subset \mathbb{F}_q$ of size $q - r$ for $(i, j) \in [k] \times [n]$, there exists $x \in \mathbb{F}_q^n$ such that $(M_i x)(j) \in X_{i,j}$ for $(i, j) \in [k] \times [n]$.*

Proof. Let $v_{i,j}$ denote the j'th row of M_i. Note that $(M_i x)(j) = \langle v_{i,j}, x\rangle$. Suppose the theorem does not hold, that is, for every x, there exists $(i,j) \in [k] \times [n]$ such that $(M_i x)(j) \notin X_{i,j}$. Let $J_0 = [k] \times [n] \times X_{i,j}$, and for every $(i,j,t) \in J_0$, define the hyperplane $H_{i,j,t} = H(v_{i,j}, -t)$. Then the collection of hyperplanes $\{H_{i,j,t} : (i,j,t) \in J\}$ forms a covering of $\mathbb{F}_q^n$. Therefore, one can select a subcollection which forms an irredundant covering of $\mathbb{F}_q^n$, let $J \subset J_0$ be the corresponding set of indices.

Let $d = \dim\{v_{i,j,t} : (i,j,t) \in J\}$. We claim that $d \geq \frac{|J|}{rk}$. Indeed, by the pigeonhole principle, at least $\frac{|J|}{rk}$ of the elements of J have the same first coordinate $i = i_0$, and the same last coordinate $t = t_0$. But then the vectors $v_{i_0,j}$, where $(i_0, j, t_0) \in J$, are distinct rows of the invertible matrix M_{i_0}, so they are linearly independent. This indeed implies $d \geq \frac{|J|}{rk}$.

Hence, by Proposition 5.2, we have $|J| \geq f_q(d) > krd \geq |J|$, contradiction. □

From this, we immediately deduce the following stronger choosability version of Theorem 1.7, as promised in the Introduction.

Theorem 6.2. *Let k, r be positive integers, let $q = p^\alpha$ be a prime power, and let s be the size of the smallest arithmetic set in $\mathbb{F}_p$. Suppose that $s^{kr} < q$, and let $M_1, \dots, M_k \in \mathbb{F}_p^{n\times n}$ be invertible matrices. Then given $k \cdot n$ sets $X_{i,j} \subset \mathbb{F}_p$ of size $p - r$ for $(i,j) \in [k] \times [n]$, there exists $x \in \mathbb{F}_p^n$ such that $(M_i x)(j) \in X_{i,j}$ for $(i,j) \in [k] \times [n]$.*

Proof. By Lemma 6.1, it is enough to prove that $f_q(N) > krN$ holds for every positive integer N. By Theorem 5.3, we have $f_q(N) \geq \frac{\log q}{\log s} N$. Hence, the condition $s^{kr} < q$ guarantees that $kr < \frac{\log q}{\log s}$, finishing the proof. □

Proof of Theorem 1.7. Apply Theorem 6.2 with parameters $r = 1$, and $X_{i,j} = \mathbb{F}_q \setminus \{0\}$ for $(i,j) \in [k] \times [n]$. Note that $s \leq 2\log_2 q$ by Lemma 4.1, so setting $q_0(k)$ sufficiently large, we have $s^{kr} < q$ satisfied for $q > q_0$. This finishes the proof. □

7 Additive bases

In this section, we prove Theorems 1.8 and 1.10. First, we prove the following result, which then almost immediately implies Theorems 1.8.

Theorem 7.1. *Let p be a prime, and let $A \subset \mathbb{F}_p$ be an arithmetic set. If $B \subset \mathbb{F}_p^n$ is the union of at least p bases, then B is an A-basis.*

Proof. We proceed by induction on n. In the base case $n = 0$ there is nothing to prove, so suppose that $n \geq 1$. Let $B \subset \mathbb{F}_p^n$ be such that B is the union of at least p bases, then $|B| \geq pn \geq (p-1)n + 1$. Hence, by Lemma 3.2, B is $\mathbb{F}_p$-vanishing. But then B contains an $\mathbb{F}_p$-irredundant subset V.

Let $T = \langle V\rangle$, and $S = \mathbb{F}_p^n/T \cong \mathbb{F}_p^{n-\dim(T)}$, so that $\mathbb{F}_p^n = S \oplus T$. Then every $x \in \mathbb{F}_p^n$ can be written as $x = (x_S, x_T)$ with $x_S \in S$ and $x_T \in T$. Let $B' = \{b_S : b \in B \setminus V\} \subset S$. Note that if C is a linear basis contained in B, then $C' = \{b_S : b \in C \setminus V\}$ contains a linear basis of S, so B' contains the union of p linear bases of S. Therefore, by our induction hypothesis applied to S, B' is an A-basis. Equivalently, for every $x \in \mathbb{F}_p^n$, there exists $\alpha_0 \in A^{B\setminus V}$ such that

$$x_S = \sum_{v\in B\setminus V} \alpha_0(v) v_S.$$

Now let $W = \ker(V) < \mathbb{F}_p^V$. Then W is versatile by Lemma 3.4, so $W + A^V = \mathbb{F}_p^V$ by Lemma 4.1. Let

$$x' = x_T - \sum_{v\in B\setminus V} \alpha_0(v)v_T \in T.$$

As $x' \in T = \langle V\rangle$, there exists $\beta \in \mathbb{F}_p^V$ such that $x' = \sum_{v\in V} \beta(v)v$. But then there exists $\gamma \in W = \ker(V)$ such that $\beta + \gamma \in A^V$. Writing $\alpha_1 = \beta + \gamma$, we have $x' = \sum_{v\in V} \alpha_1(v)v$. Let $\alpha \in \mathbb{F}_p^B$ be the vector defined as $\alpha(v) = \alpha_0(v)$ if $v \in B \setminus V$, and $\alpha(v) = \alpha_1(v)$ if $v \in V$. Then $\alpha \in A^V$ and $x = \sum_{v\in B} \alpha(v)v$, finishing the proof. □

Proof of Theorem 1.8. As there exists an arithmetic set $A \subset \mathbb{F}_p$ of size $(1 + o(1)) \log_2 p$, Theorem 7.1 implies the desired result. □

In the rest of this section, we prove Theorem 1.10. Let $q = p^\alpha$, and let $\lambda_1, \dots, \lambda_\alpha$ be a basis of $\mathbb{F}_q$ over $\mathbb{F}_p$. Define the linear map

$$x \in \mathbb{F}_q \mapsto \widetilde{x} \in \mathbb{F}_p^s$$

by mapping λ_i to the i-th unit vector. Also, extend this map as $v \in \mathbb{F}_q^n \mapsto \widetilde{v} \in \mathbb{F}_p^{sn}$ in the natural coordinate-wise manner.

We prepare the proof with two lemmas, the first of which is a well known result of Rado [21]

Lemma 7.2. *Let V be a vectorspace and let $A_1, \dots, A_k \subset V$. If for every $I \subset [k]$, we have*

$$\dim\left(\bigcup_{i\in I} A_i\right) \geq |I|,$$

then there exists $v_i \in A_i$ for $i \in [k]$ such that $v_1, \dots, v_k$ are linearly independent.

Lemma 7.3. *Let $q = p^\alpha$, $B_1, \dots, B_\alpha$ be bases of $\mathbb{F}_q^n$, and let $B = \bigcup_{i=1}^\alpha B_i$ (as a multiset). Then for every $x \in B$ there exists $i_x \in [\alpha]$ such that $\{\widetilde{\lambda_{i_x} x} : x \in B\}$ is a basis of $\mathbb{F}_p^{\alpha n}$.*

Proof. For $x \in \mathbb{F}_q^n$, let $A(x) = \{\widetilde{\lambda_i x} : i \in [\alpha]\} \subset \mathbb{F}_p^{\alpha n}$. Note that for any $X \subset B$, we have $\dim(\bigcup_{x\in X} A(x)) \geq |X|$. This is true as there exists $i \in [\alpha]$ such that $|B_i \cap X| \geq |X|/\alpha$, and then

$$\dim\left(\bigcup_{x\in X\cap B_i} A(x)\right) \geq \alpha|X \cap B_i| \geq |X|.$$

But then we can apply Lemma 7.2 to conclude that there exists $v_x = \widetilde{\lambda_{i_x} x} \in A(x)$ for $x \in B$ such that $C = \{v_x : x \in B\}$ is linearly independent. As $|C| = \alpha n$, C is also a basis of $\mathbb{F}_p^{\alpha n}$. □

Now we are ready to prove Theorem 1.10.

Proof of Theorem 1.10. Let $B_1 \dots, B_{\alpha p}$ be bases of $\mathbb{F}_q^n$. For $j = 1, \dots, p$, apply Lemma 7.3 to the bases $B_{(j-1)\alpha+1}, \dots, B_{j\alpha}$, and write $\mathcal{B}_j = B_{(j-1)\alpha+1} \cup \dots \cup B_{j\alpha}$ (as multiset). Then for $x \in \mathcal{B}_j$, there exists $i_x \in [\alpha]$ such that $C_j = \{\widetilde{\lambda_{i_x} x} : x \in \mathcal{B}_j\}$ is a basis in $\mathbb{F}_p^{\alpha n}$.

Therefore, by Theorem 1.8, there exists $0 \in A' \subset \mathbb{F}_p$ of size at most $(1 + o(1)) \log p$ such that $C = \bigcup_{j\in[p]} C_j$ is an A'-basis. We show that if $A = A' \cdot \{\lambda_i : i \in [\alpha]\}$, then $B = \bigcup_{i=1}^{\alpha p} B_i$ is an A-basis. Let $w \in \mathbb{F}_q^n$, then for $v \in C$ there exists $c_v \in A'$ such that

$$\widetilde{w} = \sum_{v\in C} c_v v.$$

For every $v \in C$, there is a unique $x \in B$ such that $\widetilde{\lambda_{i_x} x} = v$, so writing $c_x = c_v$, we have

$$\widetilde{w} = \sum_{x \in B} c_x \cdot \widetilde{\lambda_{i_x} x}.$$

Equivalently,

$$w = \sum_{x \in B} c_x \lambda_{i_x} x.$$

Here, $c_x \lambda_{i_x} \in A$, finishing the proof. □

8 Coset covers

In this section, we prove Theorem 1.11. For a group G, let $\phi(G)$ denote the smallest k for which there exists an irredundant coset cover $\{H_i x_i : i \in [k]\}$ such that $\bigcap_{i \in [k]} H_i$ is trivial. Note that Theorem 1.11 is equivalent with the statement that for every finite abelian group G, we have $\phi(G) = \Omega(\log |G| / \log \log \log |G|)$. In particular, we prove the following.

Theorem 8.1. *Let G be a finite abelian group and let $p_1^{n_1} \dots p_m^{n_m}$ be the prime factorization of $|G|$. Then*

$$\phi(G) \geq 1 + \sum_{i \in [m]} f_{p_i}(n_i) - 1.$$

We prepare the proof of this theorem with several statements.

Definition 4. A coset cover $\{H_i x_i : i \in [k]\}$ of an abelian group G is *efficient*, if it is irredundant, $\bigcap_{i \in [k]} H_i$ is trivial, and H_i is a maximal subgroup of G for $i \in [k]$.

Observe that if $G = \mathbb{F}_p^n$, then an efficient coset cover is exactly an irredundant covering with hyperplanes such that the normal vectors of the hyperplanes span the whole space. Therefore, every efficent coset cover of $\mathbb{F}_p^n$ contains at least $f_p(n)$ elements.

Lemma 8.2. *If G has an efficient coset cover, then $G \cong \mathbb{F}_p^n$ for some prime p and $n \in \mathbb{Z}^+$.*

Proof. Let $\{H_i x_i : i \in [k]\}$ be an efficient coset cover of G. By the fundamental theorem of finite abelian groups, we can write $G = G_1 \oplus \dots \oplus G_m$, where $|G_1|, \dots, |G_m|$ are powers of distinct primes. First, we show that $m = 1$. Suppose that $m > 1$, and for $i \in [m]$, $a \in G$, let $\pi_i(a)$ denote the projection of a into G_i. As H_i is a maximal subgroup of G, there is a unique $\tau(i) \in [m]$ such that $\pi_{\tau(i)}(H_i)$ is a maximal subgroup of $G_{\tau(i)}$, and $\pi_j(H_i) = G_j$ for $j \in [m] \setminus \{\tau(i)\}$. For $j \in [m]$, let $J_j \subset [k]$ be the set of indices i such that $\tau(i) = j$. Note that J_j is nonempty for every $j \in [m]$, otherwise $G_j < \pi_j(\bigcap_{i \in [k]} H_i)$. Furthermore, $\{H_i x_i : i \in J_j\}$ does not cover at least one element $a \in G$, so it does not cover any element $b \in G$ with $\pi_j(b) = \pi_j(a) =: \alpha_j$. But then $\{H_i x_i : i \in [k]\}$ does not cover $(\alpha_1, \dots, \alpha_m)$, contradiction.

Now we can assume that $|G|$ is a power of some prime p. The intersection of all maximal subgroups of G, denoted by $\mathrm{Fr}(G)$, is called the Frattini subgroup [6]. It is known that if G is a p-group, then $\mathrm{Fr}(G)$ is the smallest normal subgroup N such that $G/N \cong \mathbb{F}_p^n$ for some n. Note that if G has an efficient coset cover, then $\mathrm{Fr}(G)$ is trivial, therefore, $G \cong \mathbb{F}_p^n$ for some $n \in \mathbb{Z}^+$. □

The proof of Theorem 8.1 follows closely an argument of Szegedy [26]. We will use the following simple claim repeatedly.

Claim 8.3. *Let $\{H_i x_i : i \in [k]\}$ be an irredundant coset cover of the group G. Then for every $j \in [k]$, we have $\bigcap_{i\in[k]} H_i = \bigcap_{i\in[k]\setminus\{j\}} H_i$.*

Proof. Let $X = G\setminus(\bigcup_{i\in[k]\setminus\{j\}} H_i x_i)$. Then X is nonempty, as $\{H_i x_i : i \in [k]\}$ is irredundant. But then X is the union of cosets of $\bigcap_{i\in[k]\setminus\{j\}} H_i$. As $X \subset H_j x_j$, we must have $\bigcap_{i\in[k]\setminus\{j\}} H_i \subset H_j$, finishing the proof. □

For every $N \in \mathbb{Z}^+$ with prime factorization $N = p_1^{n_1} \dots p_m^{n_m}$, define $\lambda(N) = \sum_{i=1}^m f_{p_i}(n_i) - 1$. Then Theorem 8.1 is equivalent with the statement that $\phi(G) \geq \lambda(|G|) + 1$.

Claim 8.4. *If a, b are positive integers, then $\lambda(ab) \leq \lambda(a) + \lambda(b)$.*

Proof. This follows from the subadditivity of $f_p(n) - 1$, that is, Lemma 5.4. □

Proof of Theorem 8.1. We prove by induction on $|G|$ that $\phi(G) \geq \lambda(|G|) + 1$. In case $|G| = 1$, the statement is trivial, so suppose that $|G| \geq 2$. Let $k = \phi(G)$, and let $\mathcal{C} = \{H_i x_i : i \in [k]\}$ be an irredundant coset cover of G such that $\bigcap_{i\in[k]} H_i$ is trivial.

Let M be the number of non-maximal subgroups among $H_1, \dots, H_k$. We will also proceed by induction on M. In case $M = 0$, the coset covering $\{H_i x_i : i \in [k]\}$ is also efficient, so $G \cong \mathbb{F}_p^n$ for some prime p and $n \in \mathbb{Z}^+$ by Lemma 8.2. Hence, $k \geq f_p(n) = \lambda(|G|) + 1$, and we are done.

Therefore, we can assume that $M \geq 1$, and without loss of generality, H_k is not a maximal subgroup of G. Replace H_k with some maximal subgroup $H_k' < G$ containing H_k. Let $\mathcal{C}' = \{H_i x_i : i \in [k-1]\} \cup \{H_k' x_k\}$, then $\mathcal{C}'$ is a coset covering, and $H_k' \cap \bigcap_{i\in[k-1]} H_i$ is trivial by Claim 8.3. Note that there are $M - 1$ non-maximal subgroups among $H_1, \dots, H_{k-1}, H_k'$, so if $\mathcal{C}'$ is irredundant, we are done by our induction hypothesis.

Therefore, we can assume that $\mathcal{C}'$ is not irredundant, so, without loss of generality, there exists $\ell \leq k-2$ such that $\mathcal{C}'' = \{H_i x_i : i \in [\ell]\} \cup \{H_k' x_k\}$ is an irredundant cover of G. As this cover contains less than $k = \phi(G)$ cosets, we must have that $B = H_k' \cap \bigcap_{i\in[\ell]} H_i$ is nontrivial. Therefore, using our first induction hypothesis, we get

$$\ell + 1 \geq \phi(G/B) \geq \lambda(|G/B|) + 1 \geq \lambda(|G|) - \lambda(|B|) + 1.$$

Here, the last inequality follows by Claim 8.4. For $t = 0, 1, \dots,$ we define the sequence of 4-tuples (B_t, X_t, Y_t, I_t), where $B_t < G$ is a subgroup, $X_t, Y_t \subset G$ are subsets and $I_t \subset [k] \setminus [\ell]$ is an index set, in such a way that the following properties hold.

(i) $X_t = \bigcup_{y\in Y_t} B_t y$,

(ii) $\{X_t \cap (H_i x_i) : i \in I_t\}$ is an irredundant cover of X_t,

(iii) $B_t \cap \bigcap_{i\in I_t} H_i$ is trivial,

(iv) $k - |I_t| \geq \lambda(|G|) - \lambda(|B_t|) + t$.

Set $B_0 := B$, $X_0 := G \setminus \bigcup_{i\in[\ell]} H_i x_i$, and $I_0 := [k] \setminus [\ell]$. Note that $B_0 = \bigcap_{i\in[\ell]} H_i$ holds by Claim 8.3, so X_0 is a union of cosets of B_0. Therefore, there exists $Y_0 \subset G$ such that $X_0 = \bigcup_{y\in Y_0} B_0 y$. Also, we have $k - |I_0| = \ell \geq \lambda(|G|) - \lambda(|B_0|)$, so the choice (B_0, X_0, Y_0, I_0) satisfies (i)-(iv). If B_t, X_t, Y_t, I_t are already defined satisfying the above properties, we proceed as follows. Suppose that B_t is non-trivial, then I_t, X_t, Y_t are nonempty. (Note that for the initial step $t = 0$ these indeed hold.) By (iii), there exists $j \in I_t$ such that H_j does not contain B_t, and by (ii), there exists some $x \in X_t$ which is only covered by $H_j x_j$. Let $y \in Y_t$ be such that $x \in B_t y$, and let $J \subset I_t$ be a set of indices such that

$\{H_i x_i : i \in J\}$ is an irredundant cover of $B_t y$. Note that $j \in J$. Set $B_{t+1} := B_t \cap \bigcap_{i\in J} H_i$, then $|J| \geq \phi(B_t/B_{t+1}) \geq \lambda(|B_t|) - \lambda(|B_{t+1}|) + 1$ by our induction hypothesis. Set $I_{t+1} := I_t \setminus J$. Observe that we have $k - |I_{t+1}| = k - |I_t| + |J| \geq \lambda(|G|) - \lambda(|B_{t+1}|) + t + 1$.

If B_{t+1} is a trivial subgroup, then we stop. Note that in this case (iv) implies $k \geq k - |I_{t+1}| \geq \lambda(|G|) + 1$, finishing our proof.

If B_{t+1} is non-trivial, then I_{t+1} and $X_{t+1} := X_t \setminus (\bigcup_{i\in J} H_i x_i)$ are nonempty (by using (iii) and (ii), respectively). Also, X_{t+1} is the union of cosets of B_{t+1}, so there exists $Y_{t+1} \subset G$ such that $X_{t+1} = \bigcup_{y \in Y_{t+1}} B_{t+1} y$. Hence, (i)-(iv) are satisfied for $(B_{t+1}, X_{t+1}, Y_{t+1}, I_{t+1})$ as well. Note that B_{t+1} is a proper subgroup of B_t, so the sequence stops after a finite number of steps, giving the desired result. □

This finishes the proof of Theorem 8.1. By using our lower bounds on $f_p(n)$, this theorem immediately implies that if p is the largest prime divisor of $|G|$, then $\phi(G) \geq \Omega(\frac{\log |G|}{\log\log p})$. In case $|G| = e^{\Omega(p)}$, this gives our desired bound $\phi(G) \geq \Omega(\frac{\log |G|}{\log\log\log |G|})$. However, in case $|G| = e^{O(p)}$, there is an even simpler argument.

Lemma 8.5. *Let G be a finite abelian group, and let p be a prime divisor of $|G|$. Then $\phi(G) \geq p$.*

Proof. Let $k = \phi(G)$, and let $\{H_i x_i : i \in [k]\}$ be an irredundant coset covering of G such that $\bigcap_{i \in [k]} H_i$ is trivial. Furthermore, let $B < G$ be the unique maximal p-subgroup of G. Without loss of generality $B \not< H_1$. There exists some $a \in G$ which is only covered by $H_1 x_1$, let By be the coset of B containing a. Then for any $i \in [k]$, we have $|H_i x_i \cap By| \leq |B|/p$, so $\{H_i x_i : i \in [k]\}$ must contain at least p cosets in order to cover By. □

Proof of Theorem 1.11. Let $G' = G / \bigcap_{i\in[k]} H_i$, let $N = |G'|$, and write $N = p_1^{n_1} \dots p_m^{n_m}$, where $p := p_1 > \dots > p_m$. If $N < e^p$, then using Lemma 8.5, we get $k \geq \phi(G') \geq p$, which gives $N < e^k$. On the other hand, if $N \geq e^p$, then we can use Theorem 8.1 and Theorem 1.1 to conclude that

$$k \geq \phi(G') > \lambda(N) = \sum_{i=1}^{m} f_{p_i}(n_i) - 1 \geq \sum_{i=1}^{m} c \frac{\log p_i}{\log\log p_i} n_i \geq \frac{c}{\log\log p} \sum_{i=1}^{m} n_i \log p_i = \frac{c \log N}{\log\log p},$$

where $c > 0$ is some small absolute constant. From this, we get the desired result $N \leq e^{c' k \log\log k}$, where $c' > 0$ is some further constant. □

We finish this section with the proof of Theorem 1.2.

Proof of Theorem 1.2. Let $N = f_q(n)$, and let $\mathcal{H} = \{H(v_i, t_i) : i \in [N]\}$ be an irredundant covering of $\mathbb{F}_q^n$ with N hyperplanes such that $\dim\{v_i : i \in [N]\} = n$. The latter is equivalent to the statement that $\bigcap_{i\in[N]} H(v_i, 0)$ is trivial. But then we can apply Theorem 8.1 to the additive group of $\mathbb{F}_q^n$, and Theorem 1.1 to conclude that

$$N \geq \phi(\mathbb{F}_q^n) \geq f_p(\alpha n) \geq (1 - o(1)) \frac{\log p}{\log\log p} \cdot (\alpha n) = (1 - o(1)) \frac{\log q}{\log\log p} n,$$

finishing the proof. □

9 Concluding remarks

In this paper, we proved that $\lim_{n\to\infty} \frac{f_p(n)}{n}$ exists for every prime (or prime power) p, and is between $(1-o(1))\frac{\log p}{\log\log p}$ and $\frac{p}{2}$. While we believe the upper bound should be closer to the truth, improving the lower bound to $\Omega(\log p)$ would be already interesting. Indeed, such a mild improvement already implies the conjectures of Pyber and Szegedy about irredundant coset covers, discussed in Section 1.3. One approach to achieve this improvement is to show that every versatile subspace (see Definition 2) in $\mathbb{F}_p^N$ has dimension at least $(1-O(\frac{1}{\log p}))N$.

Acknowledgments

J. N. and P. P. P. was supported by the Lendület program of the Hungarian Academy of Sciences (MTA). P. P. P. was also partially supported by the National Research, Development and Innovation Office NKFIH (Grant Nr. K124171 and K129335).

References

[1] N. Alon, Z. Füredi, Covering the cube by affine hyperplanes, European J. Combin. 14 (2) (1993), 79–83.

[2] N. Alon, N. Linial, and R. Meshulam, Additive bases of vector spaces over prime fields, J. Combin. Theory Ser. A 57 (1991), 203–210.

[3] N. Alon and M. Tarsi, A nowhere-zero point in linear mappings, Combinatorica 9 (4) (1989), 393–395.

[4] J. Browkin, B. Diviš, and A. Schinzel, Addition of sequences in general fields, Monatsh. Math. 82 (4) (1976), 261–268.

[5] M. DeVos, Matrix choosability, J. Combin. Theory Ser. A 90 (1) (2000), 197–209.

[6] G. Frattini, Intorno alla generazione dei gruppi di operazioni, Accademia dei Lincei, Rendiconti. (4) I (1885), 281–285, 455–457.

[7] M. Han, J. Li, Y. Wu, and C.-Q. Zhang, Counterexamples to Jaeger's Circular Flow Conjecture, J. Combin. Theory Ser. B 131 (2018), 1–11.

[8] F. Jaeger, Problem presented in the 6th Hungar. Comb. Coll., Eger, Hungary 1981, and: Finite and Infinite Sets (eds.: Hajnal, A., Lovász, L., Sós, V. T.). North Holland, Amsterdam, 1982 II, 879.

[9] F. Jaeger, On circular flows in graphs, in: Finite and Infinite Sets, Eger, 1981, in: Colloquia Mathematica Societatis János Bolyai, vol. 37, North-Holland, 1984, 391–402.

[10] F. Jaeger, Nowhere-zero flow problems, in "Selected Topics in Graph Theory" (L. Beineke and R. Wilson, Eds.), Vol. 3, pp. 91–95, Academic Press, London/New York, 1988.

[11] F. Jaeger, N. Linial, C. Payan, and M. Tarsi, Group Connectivity of Graphs—A Nonhomogeneous Analogue of Nowhere-Zero Flow Properties, J. Combin. Theory Ser. B 56 (1992), 165–182.

[12] N. Linial, J. Radhakrishnan, Essential covers of the cube by hyperplanes, J. of Combin. Theorys Ser. A 109 (2) (2005), 331–338.

[13] L. M. Lovász, C. Thomassen, Y. Wu, and C-Q. Zhang, Nowhere-zero 3-flows and modulo k-orientations, J. of Combin. Theory Ser. B 103 (2013), 587–598.

[14] J. Nagy, and P. P. Pach, The Alon-Jaeger-Tarsi conjecture via group ring identities, arXiv preprint, arXiv:2107.03956.

[15] Z. Nedev, An algorithm for finding a nearly minimal balanced set in $\mathbb{F}_p$, Mathematics of Computation 78 (268) (2009), 2259–2267.

[16] B. H. Neumann, Groups covered by permutable subsets, J. London Math. Soc. 29 (1954), 236–243.

[17] B. H. Neumann, Groups covered by finitely many cosets, Publ. Math. Debrecen 3 (1954), 227–242.

[18] J. E. Olson, An addition theorem modulo p, J. Combin. Theory 5 (1968), 45–52.

[19] L. Pyber, The number of pairwise non-commuting elements and the index of the center in a finite group, J. London Math. Soc. 35(2) (1987), 287–295.

[20] L. Pyber, How abelian is a finite group? in: The Mathematics of Paul Erdős, vol. I, Springer-Verlag, Heidelberg, 1996, 372–384.

[21] R. Rado, A theorem on independence relations, The Quarterly Journal of Mathematics 1 (1942), 83–89.

[22] M. E. Saks, Slicing the hypercube, London Mathematical Society Lecture Note Series, 211–256, Cambridge University Press, 1993.

[23] L. Sauermann, and Y. Wigderson, Polynomials that vanish to high order on most of the hypercube, Journal of the London Mathematical Society 106.3 (2022), 2379–2402.

[24] R. Steinberg, Grötzsch's theorem dualized, M. Math. thesis, University of Waterloo, Ontario, Canada, 1976.

[25] E. G. Straus, Differences of residues (mod p), J. Number Theory 8 (1) (1976), 40–42.

[26] B. Szegedy, Coverings of abelian groups and vector spaces, J. Combin. Theory Ser. A 114 (1) (2007), 20–34.

[27] C. Thomassen, The weak 3-flow conjecture and weak circular flow conjecture, J. Combin. Theory Ser. B 102 (2012), 521–529.

[28] M.J. Tomkinson, Groups covered by finitely many cosets or subgroups, Comm. Algebra 15 (4) (1987), 845–859.