

Sunflowers and Testing Triangle-Freeness of Functions*

Ishay Haviv[†]

Ning Xie[‡]

Abstract

A function $f : \mathbb{F}_2^n \rightarrow \{0, 1\}$ is *triangle-free* if there are no $x_1, x_2, x_3 \in \mathbb{F}_2^n$ satisfying $x_1 + x_2 + x_3 = 0$ and $f(x_1) = f(x_2) = f(x_3) = 1$. In testing triangle-freeness, the goal is to distinguish with high probability triangle-free functions from those that are ε -far from being triangle-free. It was shown by Green that the query complexity of the canonical tester for the problem is upper bounded by a function that depends only on ε (GAFA, 2005), however the best known upper bound is a tower type function of $1/\varepsilon$. The best known lower bound on the query complexity of the canonical tester is $1/\varepsilon^{13.239}$ (Fu and Kleinberg, RANDOM, 2014).

In this work we introduce a new approach to proving lower bounds on the query complexity of triangle-freeness. We relate the problem to combinatorial questions on collections of vectors in \mathbb{Z}_D^n and to *sunflower conjectures* studied by Alon, Shpilka, and Umans (Comput. Complex., 2013). The relations yield that a refutation of the Weak Sunflower Conjecture over \mathbb{Z}_4 implies a super-polynomial lower bound on the query complexity of the canonical tester for triangle-freeness. Our results are extended to testing k -cycle-freeness of functions with domain \mathbb{F}_p^n for every $k \geq 3$ and a prime p . In addition, we generalize the lower bound of Fu and Kleinberg to k -cycle-freeness for $k \geq 4$ by generalizing the construction of uniquely solvable puzzles due to Coppersmith and Winograd (J. Symbolic Comput., 1990).

1 Introduction

The research on *property testing*, initiated by Rubinfeld and Sudan [31] and by Goldreich, Goldwasser, and Ron [21], is concerned with very efficient algorithms that distinguish with high probability objects which satisfy a given property from those that are far from satisfying it. Typically, one can think of an input object as a function from a domain D to a range R , and of a property \mathcal{P} as a subset of the function set $D \rightarrow R$. For a distance parameter ε , the goal of the randomized algorithm, called a *tester*, is to accept the functions of \mathcal{P} and to reject functions which are ε -far from \mathcal{P} , that is, disagree with any function in \mathcal{P} on at least ε -fraction of the inputs. In case that the functions of \mathcal{P} are always accepted, we say that the tester has *one-sided* error. The main objective in property testing is to minimize the number of queries that the tester makes to the input object. If the number of queries depends solely on the distance parameter ε , we say that the property is *strongly testable*.

*A preliminary version appeared in Proceedings of the 6th Innovations in Theoretical Computer Science conference (ITCS), 2015, pages 357-366.

[†]School of Computer Science, The Academic College of Tel Aviv-Yaffo, Tel Aviv 61083, Israel.

[‡]SCIS, Florida International University, Miami, FL 33199, USA. Research supported in part by NSF grant 1423034.
Email: nxie@cis.fiu.edu

Since the invention of the property testing model, many natural properties were shown to be strongly testable. A considerable amount of attention was given to testing *graph* properties, and the strongly testable dense graph properties were fully characterized [3, 11]. An important graph property testing problem is that of deciding if a given undirected graph is H -free, i.e., contains no subgraph isomorphic to H , or is ε -far from H -freeness, where H is a fixed graph. Whereas H -freeness is known to be strongly testable for every graph H , it turns out that the graph H significantly affects the dependence of the query complexity on ε . Alon proved in [1] that for every bipartite graph H , the one-sided error query complexity of testing H -freeness is polynomial in $1/\varepsilon$, and that for every non-bipartite graph H , it is super-polynomial in $1/\varepsilon$, namely, at least $(1/\varepsilon)^{\Omega(\log(1/\varepsilon))}$. Interestingly, the lower bound for the non-bipartite case relies on a construction of Behrend [7] of dense sets of integers with no 3-term arithmetic progressions (see also [32, 16]) and on an extension of this construction given in [1].

Kaufman and Sudan initiated in [24] a systematic study of testing *algebraic* properties of functions with domain \mathbb{F}^n for a field \mathbb{F} . They considered the class of *linear invariant* properties, those which are closed under all linear transformations of the domain, and asked for necessary and sufficient conditions for their strong testability (see [35, 8] for relevant surveys). This class includes the properties that can be described as freeness of solutions to (possibly infinite) systems of linear equations, which were shown to be strongly testable in [33] and in [27] (see also [9]). As opposed to the H -freeness property of graphs, it is not known which of these properties have query complexity polynomial in $1/\varepsilon$. The k -cycle-freeness properties, whose query complexity is the focus of the current work, fall into this category and are described next.

1.1 Testing k -Cycle-Freeness of Boolean Functions

Let n and $k \geq 3$ be integers, and let \mathbb{F}_p be the finite field of prime order p . A k -cycle of a function $f : \mathbb{F}_p^n \rightarrow \{0, 1\}$ is defined as k vectors $x_1, \dots, x_k \in \mathbb{F}_p^n$ satisfying $x_1 + \dots + x_k = 0$ and $f(x_i) = 1$ for every $1 \leq i \leq k$. In case that f has no k -cycles, we say that it is *k -cycle-free*. In the property testing problem of k -cycle-freeness over \mathbb{F}_p , the input is a function $f : \mathbb{F}_p^n \rightarrow \{0, 1\}$, and the goal is to distinguish with high probability k -cycle-free functions from those that are ε -far from every k -cycle-free function.

In the *multiple-function variant* of the k -cycle-freeness problem, the input is a k -tuple of functions $f_1, \dots, f_k : \mathbb{F}_p^n \rightarrow \{0, 1\}$, and a k -cycle is defined as k vectors $x_1, \dots, x_k \in \mathbb{F}_p^n$ satisfying $x_1 + \dots + x_k = 0$ and $f_i(x_i) = 1$ for every $1 \leq i \leq k$. The goal here is to distinguish k -cycle-free k -tuples of functions from those that are ε -far from k -cycle-freeness, that is, at least $\varepsilon \cdot p^n$ values returned by the functions f_1, \dots, f_k should be changed in order to make the k -tuple free of k -cycles. Clearly, the query complexity of the multiple-function variant of k -cycle-freeness is at least as large as that of the one-function variant. We observe, though, that whenever p does not divide k , the two variants of testing k -cycle-freeness are essentially equivalent. On the other hand, in case that p does divide k , the one-function variant of the problem seems to be easier, and has query complexity $O(1/\varepsilon)$ (see Section 2.1 for details). Therefore, in order to understand the query complexity of testing k -cycle-freeness in the one-function case, it suffices to understand it for the multiple-function case. The latter is more convenient to deal with while studying lower bounds, so from now on, unless otherwise specified, we refer to the multiple-function variant as

the k -cycle-freeness problem.

A natural tester for k -cycle-freeness over \mathbb{F}_p , known as the *canonical tester* of the problem, repeatedly picks independently and uniformly at random $k - 1$ vectors $x_1, \dots, x_{k-1} \in \mathbb{F}_p^n$ and checks if they form, together with $-x_1 - \dots - x_{k-1}$, a k -cycle of the functions f_1, \dots, f_k . If no k -cycle is found the tester accepts and otherwise it rejects. Despite the simplicity of this one-sided error tester, Green proved in [22] that for every k it has a constant probability of success for query complexity that depends only on ϵ , hence the k -cycle-freeness property is strongly testable. However, the query complexity achieved by Green has a huge dependence on ϵ , namely, it is a tower of twos whose height is polynomial in $1/\epsilon$. An improved upper bound on the tower's height follows from results of [19] and [26] (see [19, Section 5] and [23]).

The study of lower bounds on the query complexity of testing k -cycle-freeness was initiated by Bhattacharyya and Xie in [10], where the case of triangles over \mathbb{F}_2 was considered. They provided the first non-trivial lower bound of $1/\epsilon^{4.847}$ on the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_2 . They also showed that the query complexity of the canonical tester for k -cycle-freeness over \mathbb{F}_2 is polynomially related to that of more general one-sided error testers for the problem (see [10, Section 4] for details).

The proof technique of the above lower bound involved the notion, introduced in [10], of *perfect-matching-free* families of vectors (PMFs). Roughly speaking, a PMF (for triangles over \mathbb{F}_2) is a collection $\{(a_i, b_i, c_i)\}_{i \in [m]}$ of m triples of vectors in \mathbb{F}_2^n satisfying that $a_{i_1} + b_{i_2} + c_{i_3} = 0$ if and only if $i_1 = i_2 = i_3$. This means that the functions $f_1, f_2, f_3 : \mathbb{F}_2^n \rightarrow \{0, 1\}$ defined as the characteristic functions of the a_i 's, b_i 's, and c_i 's respectively, have m triangles which are pairwise disjoint. The distance of these functions from triangle-freeness is relatively large compared to the number of their triangles. Hence, they can be used to obtain lower bounds on the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_2 . The authors of [10] used a computer search to construct a PMF of vectors in \mathbb{F}_2^n of size (roughly) 1.67^n , and this allowed them to get their $1/\epsilon^{4.847}$ lower bound. Further, they showed that a PMF of size $(2 - o(1))^n$ implies a super-polynomial lower bound on the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_2 , and conjectured that such a PMF exists.

Very recently, Fu and Kleinberg [20] discovered an interesting connection between PMFs and the combinatorial objects known as *uniquely solvable puzzles* (USPs). The latter were introduced in the context of matrix multiplication algorithms and were explicitly defined by Cohn et al. in [12]. Coppersmith and Winograd [13] implicitly gave a probabilistic construction of n -dimensional USPs of size $(3/2^{2/3} - o(1))^n \approx 1.89^n$ that played a central role in their famous $O(n^{2.376})$ -time algorithm for multiplication of n by n matrices, whose running time was improved only two decades later [34, 37]. It was shown in [20] that every USP implies a PMF of the same cardinality, and this led to an improved lower bound of $1/\epsilon^{13.239}$ on the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_2 . However, it was observed in [12] that the USP construction of [13] is essentially optimal, hence it seems that the USP-based approach to proving lower bounds on testing triangle-freeness has been pushed to its limit, and, in particular, cannot yield super-polynomial lower bounds. Yet, a strengthened notion of USP, known as *strong USP*, was studied by Cohn et al. [12], who proved that if strong USPs of optimal size exist then the exponent of matrix multiplication is 2. A fascinating challenge, which was left open in [20], is to show that strong USPs might imply super-polynomial lower bounds on the query complexity of related testing problems.

In the current work we show that lower bounds on testing k -cycle-freeness might follow from the existence of certain collections of vectors in \mathbb{Z}_D^n . These collections are related to famous sunflower conjectures, which we turn to describe in the next section.

1.2 Sunflower Conjectures

A k -sunflower is a collection of k sets that have the same pairwise intersections. This notion was introduced in 1960 by Erdős and Rado [17], and besides being of great interest in combinatorics, it found applications in several areas of computer science, e.g., circuit complexity [30, 2], hardness of approximation [14], and property testing [4]. The main question regarding k -sunflowers is how large a collection of sets containing no k -sunflowers can be. It was shown in [17] that the size of any collection of subsets of size s of some universe U with no k -sunflowers is at most $s! \cdot (k - 1)^s$. The classical sunflower conjecture of Erdős and Rado asserts the following.

Conjecture 1.1 (Classical Sunflower Conjecture [17]). *For every $k > 0$ there exists a constant c_k , such that every collection of at least c_k^s subsets of size s of some universe U contains a k -sunflower.*

The above conjecture is still open even for the special case of $k = 3$. For this case, Kostochka showed an improved upper bound of $c \cdot s! \cdot \left(\frac{30 \ln \ln \ln s}{\ln \ln s}\right)^s$ for some constant $c > 0$ [25]. Erdős and Szemerédi [18] presented in 1978 the following conjecture on 3-sunflowers inside $[n]$, and proved that Conjecture 1.1, even restricted to $k = 3$, implies it.

Conjecture 1.2 (Sunflower Conjecture in $\{0, 1\}^n$ [18]). *There exists an $\varepsilon > 0$, such that every collection \mathcal{F} of subsets of $[n]$ ($n \geq 2$) of size $|\mathcal{F}| \geq 2^{(1-\varepsilon)n}$ contains a 3-sunflower.*

In a recent paper, Alon, Shpilka, and Umans [5] have studied a new notion of sunflowers over $\mathbb{Z}_D = \{1, \dots, D\}$ and several related sunflower conjectures. Following their definition, we say that k vectors v_1, \dots, v_k in \mathbb{Z}_D^n form a k -sunflower if for every $i \in [n]$ it holds that the i th entries $(v_1)_i, \dots, (v_k)_i$ of the vectors are either all equal or all distinct. It was shown in [5] that Conjecture 1.2 can be *equivalently* formulated in terms of sunflowers of vectors as follows.

Conjecture 1.3 (Sunflower Conjecture in \mathbb{Z}_D^n [5]). *There exist $\varepsilon > 0$, D_0 and n_0 , such that for every $D \geq D_0$ and $n \geq n_0$, every collection \mathcal{F} of vectors in \mathbb{Z}_D^n of size $|\mathcal{F}| \geq D^{(1-\varepsilon)n}$ contains a 3-sunflower.*

The above conjecture, just like Conjecture 1.2, is widely believed to be true. Still, one might wonder if its assertion holds for small values of D . It is stated below for a specific integer D .

Conjecture 1.4 (Weak Sunflower Conjecture over \mathbb{Z}_D [5]). *There exist $\varepsilon > 0$ and n_0 , such that for every $n \geq n_0$, every collection \mathcal{F} of vectors in \mathbb{Z}_D^n of size $|\mathcal{F}| \geq D^{(1-\varepsilon)n}$ contains a 3-sunflower.*

Of special importance is the Weak Sunflower Conjecture over \mathbb{Z}_3 , which refers to the maximum possible size of a collection of vectors in the group \mathbb{Z}_3^n with no 3-term arithmetic progressions (or, equivalently, non-trivial triples of vectors with zero sum modulo 3). The largest known construction of such collections has cardinality c^n for $c \approx 2.217$ [15]. An upper bound of $2 \cdot 3^n / n$ was shown by Meshulam in [29] (see [28] for a generalization of his result), and this was recently improved by Bateman and Katz to $O(3^n / n^{1+\varepsilon})$ for some constant $\varepsilon > 0$ [6].

Given the similarity between Conjecture 1.3 and the Weak Sunflower Conjecture over \mathbb{Z}_D , one might guess that the latter is true for small values of D . In fact, it was observed in [5] that for every $D \geq 3$, the assertion of the Weak Sunflower Conjecture over \mathbb{Z}_D implies Conjecture 1.3. Nevertheless, Conjecture 1.3 seems to be much more likely to hold than Conjecture 1.4 for small values of D . For example, as explained in [5], the case of $D = 3$ can be viewed as a variant of the assertion that collections of $D^{(1-\varepsilon)n}$ vectors in \mathbb{Z}_D^n must contain a 3-term arithmetic progression modulo D . However, a result of Salem and Spencer [32] implies that the latter is false for large values of D , namely, for $D > 2^{2/\varepsilon}$.

1.3 Our Contribution

In this work we introduce a new approach to proving lower bounds on the query complexity of testing k -cycle-freeness over \mathbb{F}_p for general $k \geq 3$ and primes p . To do so, we show that certain collections of vectors in \mathbb{Z}_D^n , which are related to some of the sunflower conjectures described above, can be used to obtain perfect-matching-free vector families. For example, for the special case of triangle-freeness over \mathbb{F}_2 , it is shown that a large collection of vectors in \mathbb{Z}_4^n containing no 3-sunflowers implies a large perfect-matching-free family over \mathbb{F}_2 , thus implying a lower bound on testing triangle-freeness. In case that the size of the collection is $(4 - o(1))^n$, it yields a super-polynomial lower bound, as stated below.

Theorem 1.5. *If the Weak Sunflower Conjecture (Conjecture 1.4) over \mathbb{Z}_4 is false, then the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_2 for distance ε is super-polynomial in $1/\varepsilon$.*

As alluded to before, a refutation of the Weak Sunflower Conjecture over \mathbb{Z}_D for small values of D would not be overly surprising, thus a super-polynomial lower bound on testing triangle-freeness over \mathbb{F}_2 might stem from the above theorem. Yet, even if the Weak Sunflower Conjecture over \mathbb{Z}_4 is true, large collections of vectors in \mathbb{Z}_4^n containing no 3-sunflowers might provide improvements on the known lower bounds. Specifically, our results imply that any such collection of size $(c - o(1))^n$ for $c > 9/2^{4/3} \approx 3.57$ beats the best known lower bound of [20], but for $c < 4$ the obtained lower bound is only polynomial in $1/\varepsilon$ (see Theorem 3.4).

Using the connection shown in [10] between the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_2 and that of general (possibly adaptive) one-sided error testers for the problem, it follows that:

Corollary 1.6. *If the Weak Sunflower Conjecture (Conjecture 1.4) over \mathbb{Z}_4 is false, then the query complexity of every one-sided error tester for triangle-freeness over \mathbb{F}_2 for distance ε is super-polynomial in $1/\varepsilon$.*

We then generalize Theorem 1.5 in a couple of ways. First, we obtain the following extension to triangle-freeness over \mathbb{F}_p , where p is an arbitrary prime.

Theorem 1.7. *For every prime p , if the Weak Sunflower Conjecture (Conjecture 1.4) over \mathbb{Z}_{p^2} is false, then the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_p for distance ε is super-polynomial in $1/\varepsilon$.*

Theorem 1.7 implies that for every prime p , a refutation of a certain Weak Sunflower Conjecture over \mathbb{Z}_D (for $D = p^2$) implies a super-polynomial lower bound on the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_p . Therefore, the unlikely event that for *some* prime p the query complexity is polynomial, implies Conjecture 1.3. On the other hand, it was shown in [5] that a conjecture of Coppersmith and Winograd, which was shown in [13] to imply that the matrix multiplication exponent is 2, implies that Conjecture 1.3 is false. Hence, the conjecture of [13] implies, if true, a super-polynomial lower bound on the number of queries made by the canonical tester for testing triangle-freeness over \mathbb{F}_p for every prime p .

We note that for $p = 3$ one can show a stronger statement than that of Theorem 1.7. Indeed, in this case a super-polynomial lower bound follows quite easily from a refutation of the Weak Sunflower Conjecture over \mathbb{Z}_3 (which can be only weaker than its refutation over \mathbb{Z}_9 ; see Section 3). Interestingly, Alon et al. [5] studied a variant of this conjecture, called the *Multicolored* Sunflower Conjecture over \mathbb{Z}_3 , and related it to the notion of *strong* uniquely solvable puzzles. It turns out that this multicolored conjecture coincides with our question on perfect-matching-free families over \mathbb{F}_3 , and that their results imply an (unconditional) lower bound of $1/\varepsilon^{7.298}$ on the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_3 . Moreover, we use a result of [5] and the connection observed here to obtain that if the conjecture of [12] that strong USPs of optimal size exist is true, then the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_3 is super-polynomial. This gives, in a sense, an affirmative answer to a question posed in [20].

The above results are also extended to testing k -cycle-freeness for every $k \geq 3$. We show how lower bounds on the query complexity of the canonical tester for k -cycle-freeness over \mathbb{F}_p follow from the existence of certain collections of vectors in \mathbb{Z}_D^n for an appropriate choice of D . Namely, we are interested in collections of vectors in \mathbb{Z}_D^n for $D = p^{k-1}$, satisfying that for every k vectors in the collection (not all equal) there is some $i \in [n]$ for which the k vectors have exactly *two* distinct symbols in their i th entries. Notice that for $k = 3$ this simply means that the collection contains no 3-sunflowers. As before, for vector collections of optimal size $(D - o(1))^n$, the obtained lower bound on the query complexity turns out to be super-polynomial (see Section 3).

Finally, we show that the lower bound of Fu and Kleinberg [20] on testing triangle-freeness over \mathbb{F}_2 can be generalized to testing k -cycle-freeness over \mathbb{F}_p .

Theorem 1.8. *For every $k \geq 3$ and a prime p , the query complexity of the canonical tester for k -cycle-freeness over \mathbb{F}_p for distance ε is $\Omega(1/\varepsilon^{g(k)-o(1)})$ for*

$$g(k) = \frac{k - 1 - H(1/k) / \log_2 p}{1 - H(1/k) / \log_2 p},$$

where H stands for the binary entropy function.

The proof of Theorem 1.8 relies on a delicate extension of a construction of Coppersmith and Winograd [13] of uniquely solvable puzzles. This construction is based on a construction of Behrend [7] which found great interest in additive combinatorics. Interestingly, our construction requires an extension of Behrend's result, given in [1], that was used there for proving lower bounds on testing H -freeness of graphs.

1.4 Outline

The rest of the paper is organized as follows. In Section 2 we provide a background on the k -cycle-freeness problem, relate its one-function and multiple-function variants, present the notion of perfect-matching-free families of vectors (PMFs), and show how they imply lower bounds on the query complexity of the problem. In Section 3 we prove that PMFs can be constructed using certain collections of vectors in \mathbb{Z}_D^n and derive relations to sunflower conjectures, including Theorem 1.7. Finally, in Section 4, we prove Theorem 1.8.

2 Testing k -Cycle-Freeness of Boolean Functions

Let n and $k \geq 3$ be integers, and let \mathbb{F}_p be the finite field of prime order p . A k -cycle of k functions $f_1, \dots, f_k : \mathbb{F}_p^n \rightarrow \{0, 1\}$ is defined as k vectors $x_1, \dots, x_k \in \mathbb{F}_p^n$ satisfying $x_1 + \dots + x_k = 0$ and $f_i(x_i) = 1$ for every $1 \leq i \leq k$. If a k -tuple of functions (f_1, \dots, f_k) has no k -cycles, we say that it is k -cycle-free. Its *distance* from k -cycle-freeness is defined as

$$\min_{(g_1, \dots, g_k)} \sum_{i=1}^k \text{dist}(f_i, g_i),$$

where the minimum is over all the k -cycle-free k -tuples of functions (g_1, \dots, g_k) , and $\text{dist}(f, g)$ denotes the fraction of points at which the functions f and g disagree. We say that a k -tuple of functions is ε -far from k -cycle-freeness if its distance from k -cycle-freeness is at least ε .

In the property testing problem of k -cycle-freeness over \mathbb{F}_p , the input is a k -tuple of functions $f_1, \dots, f_k : \mathbb{F}_p^n \rightarrow \{0, 1\}$, and the goal is to accept k -cycle-free k -tuples of functions with probability at least $2/3$ and to reject k -tuple of functions which are ε -far from k -cycle-freeness with probability at least $2/3$. The *canonical tester* for k -cycle-freeness over \mathbb{F}_p repeatedly picks uniformly and independently k vectors with zero sum and checks if they form a k -cycle of the input functions. If no k -cycle is found the tester accepts and otherwise it rejects.

2.1 Multiple-function vs. One-function

As mentioned before, one might consider the one-function variant of the k -cycle-freeness testing problem. A k -cycle of a function $f : \mathbb{F}_p^n \rightarrow \{0, 1\}$ is defined as k vectors that sum to the zero vector and are all mapped by f to 1. The input of the one-function variant is a single function $f : \mathbb{F}_p^n \rightarrow \{0, 1\}$, and the goal is to decide if f is k -cycle-free or ε -far from every k -cycle-free function. The canonical tester for k -cycle-freeness is naturally extended to the one-function case.

We observe that whenever p does not divide k , every k -tuple of functions can be transformed to a single function with the same number of k -cycles, a similar domain size, and a similar distance from k -cycle freeness. This implies that, in this case, the canonical testers for the multiple-function and the one-function variants of the problem have essentially the same query complexity.

Lemma 2.1. *Let n be a positive integer, let $k \geq 3$ and p be fixed integers such that p is a prime that does not divide k , and let $\alpha > 0$ be a real number. Suppose that the k -tuple of functions $f_1, \dots, f_k : \mathbb{F}_p^n \rightarrow \{0, 1\}$ is ε_1 -far from k -cycle-freeness and that the canonical tester for k -cycle-freeness needs to make $q = \Omega(1/\varepsilon_1^\alpha)$*

queries to (f_1, \dots, f_k) . Then, there exists a function $f : \mathbb{F}_p^{n+k-1} \rightarrow \{0, 1\}$, such that f is ε_2 -far from k -cycle-freeness for $\varepsilon_2 = \varepsilon_1 / p^{k-1}$, and the canonical tester needs to make $\Omega(1/\varepsilon_2^\alpha)$ queries to f .

Proof: Given the k -tuple of functions (f_1, \dots, f_k) , define $f : \mathbb{F}_p^{n+k-1} \rightarrow \{0, 1\}$ as follows. For all $y \in \mathbb{F}_p^n$ and $z \in \mathbb{F}_p^{k-1}$, let

$$f(y, z) = \begin{cases} f_i(y), & \text{if } z = e_i \text{ for } 1 \leq i \leq k-1, \\ f_k(y), & \text{if } z = -e_1 - \dots - e_{k-1}, \\ 0, & \text{otherwise,} \end{cases}$$

where e_i denotes the vector whose entries are all 0 except the i th which is 1.

First, observe that the only way to choose k vectors (repetitions are allowed) from the set

$$\{e_1, \dots, e_{k-1}, -e_1 - \dots - e_{k-1}\},$$

so that their sum is the zero vector over \mathbb{F}_p , is to choose each of the vectors exactly once (because p does not divide k). This implies that all the k -cycles of f have exactly one point in each of the subfunctions f_1, \dots, f_k . Hence there exists a bijection between the k -cycles of f_1, \dots, f_k and those of f . Since (f_1, \dots, f_k) is ε_1 -far from k -cycle-freeness, it follows that f is ε_2 -far from k -cycle-freeness for $\varepsilon_2 = \varepsilon_1 / p^{k-1}$.

Let N_{cycles} be the number of k -cycles of (f_1, \dots, f_k) and of f . Since the query complexity of the canonical tester on a k -tuple of functions (resp. function) is proportional to the inverse of the number of k -cycles of the input k -tuple (resp. function), the query complexity on f is $\Omega(q')$ for

$$q' = p^{(n+k-1)(k-1)} / N_{\text{cycles}} = \Theta(p^{n(k-1)} / N_{\text{cycles}}) = \Theta(q) = \Omega(1/\varepsilon_1^\alpha) = \Omega(1/\varepsilon_2^\alpha).$$

■

In case that the prime p divides k , the one-function variant of k -cycle-freeness over \mathbb{F}_p is quite easy. The reason is that in this case every vector in the support of a function $f : \mathbb{F}_p^n \rightarrow \{0, 1\}$, taken with multiplicity k , forms a k -cycle of f . Thus, the problem reduces to deciding if the input function is the zero constant function or is ε -far from it. The tester that given a function f picks uniformly and independently $O(1/\varepsilon)$ random vectors in \mathbb{F}_p^n and accepts if and only if they are all mapped by f to 0 implies the following.

Claim 2.2. *Let $k \geq 3$ be an integer divisible by a prime p . Then, for every $\varepsilon > 0$, there is a one-sided error tester for the one-function variant of k -cycle-freeness over \mathbb{F}_p for distance ε with query complexity $O(1/\varepsilon)$.*

One may ask if a similar result can be shown once we consider only *non-trivial* cycles of f , that is, k vectors, not all equal, that sum to zero and are all mapped by f to 1. It turns out that if p divides k , $O(1/\varepsilon)$ queries are still sufficient to decide if a given function $f : \mathbb{F}_p^n \rightarrow \{0, 1\}$ is free of non-trivial k -cycles or ε -far from this property. The reason is that the density of such functions turns out to be very small, as follows from the following (special case of a) theorem of Liu and Spencer [28].

Theorem 2.3 ([28]). *For every n and a prime $p \geq 3$, if $A \subseteq \mathbb{F}_p^n$ contains no p vectors, not all equal, whose sum is the zero vector, then $|A| = o(p^n)$.*

Using the above theorem, it can be easily observed that if p divides k and $f : \mathbb{F}_p^n \rightarrow \{0, 1\}$ is free of non-trivial k -cycles, then it is $o(1)$ -close to the zero constant function. Thus, by the same tester that was used for Claim 2.2, we get query complexity $O(1/\varepsilon)$ and an “almost” one-sided error, that is, functions that are free of non-trivial k -cycles are accepted with probability that tends to 1 where n tends to infinity.

2.2 Perfect-Matching-Free Families

We now define the notion of *local perfect-matching-free* vector families, which can be used to obtain lower bounds on the query complexity of the canonical tester for k -cycle-freeness over \mathbb{F}_p .

Definition 2.4. An (n, m) local perfect-matching-free family (PMF) for k -cycles over \mathbb{F}_p is a collection

$$\{(x_i^{(1)}, x_i^{(2)}, \dots, x_i^{(k)})\}_{i \in [m]},$$

such that for every $i \in [m]$, $x_i^{(1)}, x_i^{(2)}, \dots, x_i^{(k)}$ are k vectors in \mathbb{F}_p^n whose sum is zero, and for every $i_1, i_2, \dots, i_k \in [m]$, if the sum of the vectors $x_{i_1}^{(1)}, x_{i_2}^{(2)}, \dots, x_{i_k}^{(k)}$ is zero then $i_1 = i_2 = \dots = i_k$. The local PMF capacity for k -cycles over \mathbb{F}_p is the largest constant c for which there exist $(n, (c - o(1))^n)$ local PMFs for k -cycles over \mathbb{F}_p for infinitely many values of n .

Two remarks are in order.

Remark 2.5. If the local PMF capacity for k -cycles over \mathbb{F}_p is c , then there exist $(n, (c - o(1))^n)$ local PMFs for k -cycles over \mathbb{F}_p for every sufficiently large value of n (and not only for infinitely many of them). To see this, for every n , denote by m_n the largest integer for which there exists an (n, m_n) local PMF for k -cycles over \mathbb{F}_p . Since $m_{n+n'} \geq m_n \cdot m_{n'}$, we may apply Fekete’s lemma (see, e.g., [36, Lemma 11.6]) to show that the limit of $m_n^{1/n}$, as n tends to infinity, exists and equals the capacity c .

Remark 2.6. Our definition of local PMFs is slightly different from the definition of PMFs given in [10]. Namely, the requirement in the definition of PMFs in [10] is that for every k permutations π_1, \dots, π_k of $[m]$, either $\pi_1 = \dots = \pi_k$, or there exists an $i \in [m]$ for which the sum $x_{\pi_1(i)}^{(1)} + \dots + x_{\pi_k(i)}^{(k)}$ is nonzero. Clearly, every local PMF is a PMF. Whereas the other direction does not hold, it is easy to see that the local PMF capacity for k -cycles over \mathbb{F}_p equals the PMF capacity for k -cycles over \mathbb{F}_p . For completeness, we include a short proof (which resembles that of [12, Proposition 6.3]), and throughout the paper we prefer to consider the notion of local PMFs, mainly for simplicity of presentation.

Claim 2.7. The local PMF capacity for k -cycles over \mathbb{F}_p equals the PMF capacity for k -cycles over \mathbb{F}_p .

Proof: Clearly, the PMF capacity for k -cycles over \mathbb{F}_p is at least as large as the local PMF capacity for k -cycles over \mathbb{F}_p . For the other direction, let \mathcal{F} be an (n, m) PMF for k -cycles over \mathbb{F}_p for $m = (c - o(1))^n$. For every permutation π of $[m]^k$ consider the k -tuple of vectors of length nm , obtained by concatenating the m k -tuples of vectors in \mathcal{F} ordered according to π . Let \mathcal{G} be the collection of all the k -tuples obtained this way. Observe that \mathcal{G} is an $(nm, m!)$ local PMF for k -cycles over \mathbb{F}_p and that

$$m! = m^{(1-o(1))m} = (c - o(1))^{(1-o(1))nm} = (c - o(1))^{nm}.$$

Thus, the local PMF capacity for k -cycles over \mathbb{F}_p is at least c , and we are done. ■

The following lemma and corollary show how local PMFs imply lower bounds for testing k -cycle-freeness. Similar statements were shown in [10], and we include here the proofs for completeness.

Lemma 2.8. *Let $k \geq 3$ be an integer, and let p be a prime. Suppose that there exists an (n, m) local PMF for k -cycles over \mathbb{F}_p . Then, the query complexity of the canonical tester for k -cycle-freeness over \mathbb{F}_p for distance ε on n variable functions is $\Omega(1/\varepsilon^\alpha)$ for $\varepsilon = m/p^n$ and $\alpha = \frac{k-1-(\log_p m)/n}{1-(\log_p m)/n}$.*

Proof: Let $\mathcal{F} = \{(x_i^{(1)}, x_i^{(2)}, \dots, x_i^{(k)})\}_{i \in [m]}$ be an (n, m) local PMF for k -cycles over \mathbb{F}_p . For every $1 \leq j \leq k$, let $f_j : \mathbb{F}_p^n \rightarrow \{0, 1\}$ be the characteristic function of the set $\{x_i^{(j)}\}_{i \in [m]}$. By definition of local PMFs, the number of k -cycles of the k -tuple of functions (f_1, \dots, f_k) is m , and these cycles are pairwise disjoint. Hence, in order to remove all the m cycles, one has to change at least m values of the functions, so this k -tuple is ε -far from k -cycle-freeness for $\varepsilon = \frac{m}{p^n}$. On the other hand, the probability that one iteration of the canonical tester, applied to (f_1, \dots, f_k) , finds a k -cycle is $\frac{m}{p^{(k-1)n}}$, so its query complexity is $\Omega(q)$, for

$$q = \frac{p^{(k-1)n}}{m} = p^{(k-1)n - \log_p m} = (1/\varepsilon)^{\frac{k-1-(\log_p m)/n}{1-(\log_p m)/n}}.$$

■

Corollary 2.9. *Let $k \geq 3$ and p be fixed integers, such that p is prime. If the local PMF capacity for k -cycles over \mathbb{F}_p is c , then for every $d < c$, the query complexity of the canonical tester for k -cycle-freeness over \mathbb{F}_p for distance ε is $\Omega(1/\varepsilon^\alpha)$ where $\alpha = \frac{k-1-\log_p d}{1-\log_p d}$. Furthermore, for every sufficiently small ε there exists an $n_0 = n_0(\varepsilon)$ such that for every $n \geq n_0$ the lower bound holds for k -tuples of n variable functions that depend on all of their input variables. In particular, if the local PMF capacity for k -cycles over \mathbb{F}_p is p , then the query complexity of the canonical tester for k -cycle-freeness over \mathbb{F}_p for distance ε is super-polynomial in $1/\varepsilon$.*

Proof: Let c denote the local PMF capacity for k -cycles over \mathbb{F}_p , and take an arbitrary $d < c$. Using Remark 2.5, for every sufficiently large n there exists an $(n, \lceil d^n \rceil)$ local PMF for k -cycles over \mathbb{F}_p . Now, for a given sufficiently small ε , let $n_0 = n_0(\varepsilon)$ be the largest integer satisfying $\varepsilon \leq \frac{\lceil d^{n_0} \rceil}{p^{n_0}}$. For this n_0 there exists an $(n_0, \lceil d^{n_0} \rceil)$ local PMF for k -cycle-freeness over \mathbb{F}_p . By Lemma 2.8, the corresponding k -tuple of functions $f_1, \dots, f_k : \mathbb{F}_p^{n_0} \rightarrow \{0, 1\}$ is ε -far from k -cycle-freeness and requires $\Omega(1/\varepsilon^\alpha)$ queries of the canonical tester for α as in the statement of the corollary.

It remains to show that the above lower bound can be extended to k -tuples of functions with domain \mathbb{F}_p^n for every $n \geq n_0$. For every $1 \leq j \leq k$ define the function $g_j : \mathbb{F}_p^n \rightarrow \{0, 1\}$ such that $g_j(y) = 1$ if and only if $y = (x, z)$ for $x \in \mathbb{F}_p^{n_0}$ and $z \in \mathbb{F}_p^{n-n_0}$ satisfying $f_j(x) = 1$ and $\sum_{i=1}^{n-n_0} z_i = 0$.¹ The k -tuple of functions (g_1, \dots, g_k) has at least $\varepsilon \cdot p^{n_0} \cdot (p^{n-n_0-1})^{k-1}$ k -cycles, and every vector of these cycles belongs to $(p^{n-n_0-1})^{k-2}$ of the cycles. Therefore, (g_1, \dots, g_k) is ε' -far from k -cycle-freeness for $\varepsilon' = \varepsilon/p = \Theta(\varepsilon)$ and requires query complexity $\Omega(1/\varepsilon^\alpha)$, thus the required lower bound holds for every sufficiently small distance parameter. In addition, it is easy to verify that the k -tuple (g_1, \dots, g_k) depends on all of its input variables, as required.

¹This is a slight generalization of a construction due to Jakob Nordström (Private communication, 2010).

Finally, observe that if the local PMF capacity for k -cycles over \mathbb{F}_p is p , then for every $\alpha > 0$, the query complexity of the canonical tester for k -cycle-freeness over \mathbb{F}_p for some distance ε is $\Omega(1/\varepsilon^\alpha)$, thus it is super-polynomial in $1/\varepsilon$. \blacksquare

We turn to define (strong) uniquely solvable puzzles (USPs). Then we state a theorem of Alon et al. [5] that says that strong USPs imply local PMFs for triangles over \mathbb{F}_3 (in their language, collections of ordered 3-sunflowers in $\mathbb{Z}_3^n \times \mathbb{Z}_3^n \times \mathbb{Z}_3^n$ containing no *multicolored* sunflowers).

Definition 2.10. *An n -dimensional uniquely solvable puzzle (USP) is a collection of vectors $\{x_i\}_{i \in [m]}$ in \mathbb{Z}_3^n satisfying that for every three permutations π_1, π_2, π_3 of $[m]$, either $\pi_1 = \pi_2 = \pi_3$, or there exist $i \in [m]$ and $j \in [n]$ for which at least two of $(x_{\pi_1(i)})_j = 1$, $(x_{\pi_2(i)})_j = 2$, and $(x_{\pi_3(i)})_j = 3$ hold. A strong USP is defined similarly replacing the “at least two” by “exactly two”. The (strong) USP capacity is the largest constant c for which there exist n -dimensional (strong) USPs of size $(c - o(1))^n$ for infinitely many values of n .*

Theorem 2.11 ([5]). *If the strong USP capacity is at least c , then the local PMF capacity for triangles over \mathbb{F}_3 is at least $2^{2/3} \cdot c$.*

It is known that the strong USP capacity is at least $2^{2/3}$ [12, Proposition 3.8]. Hence, by Theorem 2.11, the local PMF capacity for triangles over \mathbb{F}_3 is at least $2^{4/3}$. By Corollary 2.9, it follows that the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_3 for distance ε is at least $1/\varepsilon^{7.298}$. Cohn, Kleinberg, Szegedy, and Umans conjectured that the strong USP capacity is $3/2^{2/3}$ and proved in [12] that their conjecture implies that the exponent of matrix multiplication is 2. By Theorem 2.11, if their conjecture is true then the local PMF capacity for triangles over \mathbb{F}_3 is 3, and the latter yields, by Corollary 2.9, a super-polynomial lower bound on the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_3 .

3 Sunflower Conjectures vs. Local PMFs

In this section we prove that local PMFs for k -cycles over \mathbb{F}_p can be constructed from certain collections of vectors in \mathbb{Z}_D^n , some of which are related to sunflower conjectures of Alon et al. [5]. The idea behind the construction is quite simple: every vector of these collections is mapped to a k -tuple of vectors, in a way that every symbol of \mathbb{Z}_D is replaced by certain k vectors over \mathbb{F}_p . An important property of the transformation is that it preserves optimal capacity, namely, vector collections of optimal size $(D - o(1))^n$ are transformed to local PMFs of optimal capacity p .

Theorem 3.1. *Let $k \geq 3$ be an integer, and let p be a prime. Assume that for infinitely many values of n there exists a collection \mathcal{F} of $(c - o(1))^n$ vectors in $\mathbb{Z}_{p^{k-1}}^n$ such that for every k vectors $v_1, \dots, v_k \in \mathcal{F}$, not all equal, there exists an $i \in [n]$ for which*

$$|\{(v_1)_i, \dots, (v_k)_i\}| = 2.$$

Then, the local PMF capacity for k -cycles over \mathbb{F}_p is at least $c^{1/(k-1)}$.

We need the following lemma and the corollary that follows it. We use here the notation $A^{(\ell)}$ to denote the ℓ th column of a matrix A .

Lemma 3.2. For every prime p and a positive integer k , there exists a collection of p^k matrices A_1, A_2, \dots, A_{p^k} in $\mathbb{F}_p^{k \times k}$ such that for every $1 \leq i \neq j \leq p^k$ and every non-empty set $I \subseteq [k]$ it holds that

$$\sum_{\ell \in I} A_i^{(\ell)} \neq \sum_{\ell \in I} A_j^{(\ell)}.$$

Proof: Denote $q = p^k$, and let $\alpha_1, \dots, \alpha_q$ be the q elements of the field \mathbb{F}_q . Let $enc : \mathbb{F}_q \rightarrow \mathbb{F}_p^k$ be the natural encoding of the elements of \mathbb{F}_q as distinct vectors in \mathbb{F}_p^k . This encoding is linear, that is, $enc(0) = (0, \dots, 0)$ and $enc(x + y) = enc(x) + enc(y)$ for every $x, y \in \mathbb{F}_q$. Let β be a generator of the multiplicative group \mathbb{F}_q^* , and notice that the set $\{1, \beta, \beta^2, \dots, \beta^{k-1}\}$ is linearly independent over \mathbb{F}_p .

Now, for every $1 \leq i \leq q$, define the k by k matrix A_i as the matrix whose columns are

$$enc(\alpha_i), enc(\alpha_i \cdot \beta), enc(\alpha_i \cdot \beta^2), \dots, enc(\alpha_i \cdot \beta^{k-1}).$$

To prove that the collection A_1, A_2, \dots, A_q satisfies the requirement, take $1 \leq i \neq j \leq q$ and a non-empty set $I \subseteq [k]$. Assume, for the sake of contradiction, that the matrices A_i and A_j have the same sum of columns that correspond to indices in I . Viewing these sums as elements of \mathbb{F}_q , it follows that

$$\alpha_i \cdot \sum_{\ell \in I} \beta^{\ell-1} = \alpha_j \cdot \sum_{\ell \in I} \beta^{\ell-1}.$$

By linear independence, it follows that $\sum_{\ell \in I} \beta^{\ell-1}$ is nonzero, thus $\alpha_i = \alpha_j$, a contradiction. \blacksquare

Corollary 3.3. For every prime p and a positive integer k , there exists a collection of p^k matrices B_1, B_2, \dots, B_{p^k} in $\mathbb{F}_p^{k \times (k+1)}$ such that

1. for every $1 \leq i \leq p^k$, the sum of the columns of B_i is the zero vector, and
2. for every $1 \leq i \neq j \leq p^k$ and every non-empty set $I \subset [k+1]$, the sum

$$\sum_{\ell \in I} B_i^{(\ell)} + \sum_{\ell \in [k+1] \setminus I} B_j^{(\ell)}$$

is nonzero.

Proof: By Lemma 3.2, there exists a collection of p^k matrices A_1, A_2, \dots, A_{p^k} in $\mathbb{F}_p^{k \times k}$ such that for every $1 \leq i \neq j \leq p^k$ and every non-empty set $I \subseteq [k]$ it holds that

$$\sum_{\ell \in I} A_i^{(\ell)} \neq \sum_{\ell \in I} A_j^{(\ell)}.$$

For every $1 \leq i \leq p^k$ define the k by $k+1$ matrix B_i as the matrix whose columns are

$$A_i^{(1)}, A_i^{(2)}, \dots, A_i^{(k)}, - \sum_{\ell=1}^k A_i^{(\ell)}.$$

The collection B_1, B_2, \dots, B_{p^k} trivially satisfies Item 1. To prove that Item 2 is also satisfied, take $1 \leq i \neq j \leq p^k$ and a non-empty set $I \subset [k+1]$, and assume by contradiction that the sum of

the columns of B_i corresponding to indices in I and the columns of B_j corresponding to indices in $[k+1] \setminus I$ is the zero vector. Assume without loss of generality that $k+1 \notin I$, and use Item 1 to observe that

$$\sum_{\ell \in I} A_i^{(\ell)} = \sum_{\ell \in I} B_i^{(\ell)} = - \sum_{\ell \in [k+1] \setminus I} B_j^{(\ell)} = \sum_{\ell \in I} B_j^{(\ell)} = \sum_{\ell \in I} A_j^{(\ell)},$$

in contradiction to the property of the collection A_1, A_2, \dots, A_{p^k} . \blacksquare

Now, equipped with Corollary 3.3, we are ready to show how certain collections of vectors in \mathbb{Z}_D^n can be transformed to local PMFs over \mathbb{F}_p and to prove Theorem 3.1.

Proof of Theorem 3.1: Let $\mathcal{F} \subseteq \mathbb{Z}_{p^{k-1}}^n$ be a collection of $(c - o(1))^n$ vectors satisfying the condition given in the theorem. By Corollary 3.3, applied with $k-1$, there exists a collection of p^{k-1} matrices $B_1, B_2, \dots, B_{p^{k-1}}$ in $\mathbb{F}_p^{(k-1) \times k}$ satisfying that (1) for every $1 \leq i \leq p^{k-1}$, the sum of the columns of B_i is the zero vector, and (2) for every $1 \leq i \neq j \leq p^{k-1}$ and every non-empty set $I \subset [k]$, the sum $\sum_{\ell \in I} B_i^{(\ell)} + \sum_{\ell \in [k] \setminus I} B_j^{(\ell)}$ is nonzero.

Consider the function $f : \mathbb{Z}_{p^{k-1}}^n \rightarrow \mathbb{F}_p^{n(k-1) \times k}$ that maps every vector $v \in \mathbb{Z}_{p^{k-1}}^n$ to the concatenation of the n matrices $B_{v_1}, B_{v_2}, \dots, B_{v_n}$. We claim that the set $\mathcal{G} = \{f(v) \mid v \in \mathcal{F}\}$ is an $(n(k-1), (c - o(1))^n)$ local PMF for k -cycles over \mathbb{F}_p , thus the PMF capacity for k -cycles over \mathbb{F}_p is at least $c^{1/(k-1)}$.

To see this, first observe that property (1) of the matrices $B_1, B_2, \dots, B_{p^{k-1}}$ implies that the sum of the k columns of every $f(v)$ is the zero vector. Second, let $f(v_1), \dots, f(v_k)$ be k elements, not all equal, of \mathcal{G} . Our goal is to prove that the sum of the vectors $f(v_1)^{(1)}, \dots, f(v_k)^{(k)}$ is nonzero. Since f is injective, the vectors v_1, \dots, v_k are not all equal, so there is an $i \in [n]$ for which

$$|\{(v_1)_i, \dots, (v_k)_i\}| = 2.$$

Hence, the i th blocks (of length $k-1$) of the matrices $f(v_1), \dots, f(v_k)$ contain exactly two distinct matrices B_j and $B_{j'}$. By property (2), the sum of the i th blocks of the vectors $f(v_1)^{(1)}, \dots, f(v_k)^{(k)}$ is nonzero, hence the sum of these vectors is also nonzero, and we are done. \blacksquare

Theorem 3.1 gives us a method to derive lower bounds on the query complexity of the canonical tester for k -cycle-freeness over \mathbb{F}_p from certain collections of vectors in $\mathbb{Z}_{p^{k-1}}^n$. In the special case of $k=3$, the vectors are in $\mathbb{Z}_{p^2}^n$, and for every three vectors, not all equal, there is a coordinate in which the three symbols are not all equal and are not all distinct. This exactly means that the three vectors do not form a 3-sunflower (see Section 1.2), yielding the following result.

Theorem 3.4. *Let p be a prime, and assume that for infinitely many values of n there exists a collection of $(c - o(1))^n$ vectors in $\mathbb{Z}_{p^2}^n$ containing no 3-sunflowers. Then, for every $d < \sqrt{c}$, the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_p for distance ε is $\Omega(1/\varepsilon^\alpha)$ where $\alpha = \frac{2 - \log_p d}{1 - \log_p d}$.*

Proof: By Theorem 3.1, the assumption implies that the local PMF capacity for triangles over \mathbb{F}_p is at least \sqrt{c} . Corollary 2.9 completes the proof. \blacksquare

Observe that if the Weak Sunflower Conjecture over \mathbb{Z}_D (Conjecture 1.4) is false for $D = p^2$, it follows from the above theorem that the query complexity of the canonical tester for triangle-freeness over \mathbb{F}_p for distance ε is super-polynomial in $1/\varepsilon$, confirming Theorem 1.7.

For the special case of $p = 3$, it is easy to get a local PMF for triangles from a collection of vectors in \mathbb{Z}_3^n containing no 3-sunflowers. Indeed, for such a collection \mathcal{F} , the set $\{(x, x, x)\}_{x \in \mathcal{F}}$ forms a local PMF for triangles over \mathbb{F}_3 of the same size. Thus, in case that the Weak Sunflower Conjecture over \mathbb{Z}_3 is false, a super-polynomial lower bound on the query complexity of the canonical tester follows. Note that this assumption can be only weaker than the assumption that the Weak Sunflower Conjecture over \mathbb{Z}_9 is false. The reason is that given a collection of $(9 - o(1))^n$ vectors in \mathbb{Z}_9^n containing no 3-sunflowers one can replace every symbol of \mathbb{Z}_9 in the vectors by its base-3 representation to obtain a collection of $(3 - o(1))^{2n}$ vectors in \mathbb{Z}_3^{2n} containing no 3-sunflowers.

We note that for $k \geq 4$ the property required in Theorem 3.1 from the collection $\mathcal{F} \subseteq \mathbb{Z}_D^n$ does not coincide with freeness of k -sunflowers. Indeed, the collection should satisfy that for every k vectors $v_1, \dots, v_k \in \mathcal{F}$, not all equal, there exists a coordinate in which they contain exactly 2 distinct symbols. On the other hand, freeness of k -sunflowers means that for every such k vectors, there exists a coordinate in which the number of distinct symbols is in the range from 2 to $k - 1$.

It is natural to ask if one can relate local PMFs for k -cycles to collections of vectors with no k -sunflowers for $k \geq 4$. It seems, though, that the proof technique of Theorem 3.1 cannot achieve this in a way that preserves optimal capacity. To see this, observe that such an extension requires a mapping from \mathbb{Z}_D to D matrices in $\mathbb{F}_p^{\ell \times k}$ for $D = p^\ell$ (because the transformation increases the length of the vectors by a factor of ℓ , and capacity D should be mapped to capacity p). The matrices returned by this mapping have to satisfy the following two properties: (1) the sum of the columns of each of the matrices should be zero, and (2) for every k of these matrices B_1, \dots, B_k , not all equal and not all distinct, the sum of the vectors $B_1^{(1)}, \dots, B_k^{(k)}$ should be nonzero. However, it is not difficult to show that such a collection of matrices does not exist for $k \geq 4$. First observe that the p^ℓ matrices should contain all the p^ℓ distinct vectors of \mathbb{F}_p^ℓ in each of their k columns, since otherwise two of the matrices contradict property (2). Now, take two arbitrary distinct matrices B_i and B_j , and consider the sum, say, of the first $k - 2$ columns of B_i and the $(k - 1)$ th column of B_j . The unique vector that completes this sum to zero is the k th column of one of the p^ℓ matrices, so we again contradict property (2).

4 A Lower Bound on Testing k -Cycle-Freeness

As mentioned before, the best known lower bound on the query complexity of the canonical tester for testing triangle-freeness over \mathbb{F}_2 for distance ε is $1/\varepsilon^{13.239}$, as was shown by Fu and Kleinberg [20]. Their proof is crucially based on a construction of uniquely solvable puzzles of Copper-Smith and Winograd [13], which employs a construction of Behrend [7] of dense sets of integers with no 3-term arithmetic progressions. In this section we generalize the lower bound of [20] to testing k -cycle-freeness over \mathbb{F}_p for every $k \geq 3$ and a prime p .

We need here a few notations. For a vector $v \in \mathbb{Z}_k^n$ we denote by $v|_j = \{i \in [n] \mid v_i = j\}$ the set of coordinates at which v has symbol $j \in \mathbb{Z}_k$. Note that for every vector $v \in \mathbb{Z}_k^n$, the sets $v|_1, \dots, v|_k$ form a partition of $[n]$. In case that the sets $v|_1, \dots, v|_k$ have the same size we

say that the vector v is *balanced*. Finally, let H stand for the binary entropy function, defined by $H(p) = -p \log_2 p - (1-p) \log_2 (1-p)$ for $0 \leq p \leq 1$.

Let us start with a generalization of the construction of [13], stated below. Note that the case of $k = 3$ gives the construction of [13], which implies a uniquely solvable puzzle as was defined for the purpose of fast matrix multiplication (see Definition 2.10).

Theorem 4.1. *For every fixed integer $k \geq 3$ and a sufficiently large n , there exists a collection \mathcal{F} of $(2^{H(1/k)} - o(1))^{nk}$ balanced vectors in \mathbb{Z}_k^{nk} such that for every k vectors $v_1, \dots, v_k \in \mathcal{F}$, the sets $v_1|_1, \dots, v_k|_k$ form a partition of $[n \cdot k]$ if and only if $v_1 = \dots = v_k$.*

Remark 4.2. *The cardinality of \mathcal{F} in Theorem 4.1 is optimal up to the $o(1)$ term. Indeed, the requirement on \mathcal{F} implies that the, say, $v|_1$'s for $v \in \mathcal{F}$ are distinct subsets of size n of $[n \cdot k]$, so $|\mathcal{F}| \leq \binom{nk}{n} \approx 2^{H(1/k)nk}$.*

In the proof of Theorem 4.1 we use the following extension of Behrend's result [7].

Lemma 4.3 (Lemma 3.1 in [1]). *For every fixed integer $r \geq 2$ and every positive integer m , there exists a set $B \subseteq [m]$ of size*

$$|B| \geq \frac{m}{e^{10\sqrt{\log m \log r}}}$$

with no non-trivial² solutions to the equation $x_1 + x_2 + \dots + x_r = r \cdot x_{r+1}$.

Proof of Theorem 4.1: For a sufficiently large n denote $N = n \cdot k$, and let M be the smallest prime which satisfies

$$M \geq c(k) \cdot \binom{n(k-1)}{n, \dots, n}^{1/(k-2)}, \quad (1)$$

where $c = c(k)$ is a constant that depends solely on k and will be determined later. As is well known, M is at most twice its lower bound in (1). By Lemma 4.3, applied with $m = \lfloor M/(k-1) \rfloor$, there exists a set $B = \{b_1, \dots, b_{|B|}\} \subseteq [m]$ of size $|B| = m^{1-o(1)} = M^{1-o(1)}$ with no non-trivial solutions to the equation

$$x_1 + x_2 + \dots + x_{k-1} = (k-1) \cdot x_k. \quad (2)$$

Since B is contained in $[m]$, it contains no non-trivial solutions to Equation (2) taken modulo M as well.

Consider the set \mathcal{I} of all the subsets of $[N]$ of size n , and identify the sets in \mathcal{I} with their characteristic vectors in $\{0, 1\}^N$. Let w_1, \dots, w_N and c_1, \dots, c_k be integers chosen at random uniformly and independently from \mathbb{F}_M , and denote $w = (w_1, \dots, w_N)$. For these numbers we define k mappings $\beta_1, \dots, \beta_k : \mathcal{I} \rightarrow \mathbb{F}_M$ as follows. For $1 \leq j \leq k-1$, β_j is defined by

$$\beta_j(I) = \langle w, I \rangle + c_j \pmod{M},$$

and β_k is defined by

$$\beta_k(I) = \left(\langle w, [N] \setminus I \rangle + \sum_{j=1}^{k-1} c_j \right) / (k-1) \pmod{M}.$$

²A *trivial* solution is one that satisfies $x_1 = \dots = x_{r+1}$.

The construction involves two steps. First, for every $1 \leq i \leq |B|$, let L_i denote the set of all k -tuples of sets $(I_1, \dots, I_k) \in \mathcal{I}^k$ satisfying $I_1 \cup \dots \cup I_k = [N]$ and $\beta_j(I_j) = b_i$ for every $1 \leq j \leq k$. Second, remove from every L_i all the k -tuples $(I_1, \dots, I_k) \in L_i$ that share some set I_j with other k -tuples in L_i , that is, satisfy $I_j = J_j$ for some j and $(J_1, \dots, J_k) \in L_i$. We denote by $L'_i \subseteq L_i$ the obtained set.

Every partition $(I_1, \dots, I_k) \in \mathcal{I}^k$ of $[N]$ can be naturally encoded by a balanced vector v in \mathbb{Z}_k^N defined by $v|_j = I_j$ for every $1 \leq j \leq k$. Define $\mathcal{F} \subset \mathbb{Z}_k^N$ to be the set of partitions in the union $\cup_{1 \leq i \leq |B|} L'_i$ encoded as vectors in \mathbb{Z}_k^N . We first show that \mathcal{F} satisfies the property required in Theorem 4.1, and then analyze its expected size.

Claim 4.4. *For every k vectors $v_1, \dots, v_k \in \mathcal{F}$, the sets $v_1|_1, \dots, v_k|_k$ form a partition of $[N]$ if and only if $v_1 = \dots = v_k$.*

Proof: It is clear that if $v_1 = \dots = v_k$ then $v_1|_1, \dots, v_k|_k$ form a partition of $[N]$. For the other direction, consider k vectors $v_1, \dots, v_k \in \mathcal{F}$, and denote $I_j = v_j|_j$ for every $1 \leq j \leq k$. Assume by contradiction that the vectors are not all equal and that $(I_1, \dots, I_k) \in \mathcal{I}^k$ is a partition of $[N]$. For every $1 \leq j \leq k$ denote $b_j = \beta_j(I_j)$, and observe that

$$\sum_{j=1}^{k-1} b_j = \sum_{j=1}^{k-1} \beta_j(I_j) = \sum_{j=1}^{k-1} \langle w, I_j \rangle + \sum_{j=1}^{k-1} c_j = \langle w, [N] \setminus I_k \rangle + \sum_{j=1}^{k-1} c_j = (k-1) \cdot \beta_k(I_k) = (k-1) \cdot b_k,$$

where all the equalities hold modulo M . This implies that the numbers $b_1, \dots, b_k \in B$ satisfy Equation (2) modulo M , hence by our choice of B , they must be all equal. Therefore, the vectors v_1, \dots, v_k correspond to k partitions in the same set L'_i . This implies that the partition (I_1, \dots, I_k) belongs to L_i and shares a subset of \mathcal{I} with each of the partitions that correspond to the vectors v_1, \dots, v_k . However, this implies that all these partitions were not added to L'_i in the second step of the construction. Hence all the v_i 's are equal, in contradiction. ■

We turn to analyze the expected size of the collection \mathcal{F} . We start with the size of the sets L_i (before performing the second step of the construction).

Claim 4.5. *For every $1 \leq i \leq |B|$, the expected size of the set L_i is $\binom{nk}{n, \dots, n} \cdot M^{-(k-1)}$.*

Proof: Fix $1 \leq i \leq |B|$, and let $(I_1, \dots, I_k) \in \mathcal{I}^k$ be a partition of $[N]$. Recall that (I_1, \dots, I_k) is added to L_i if $\beta_j(I_j) = b_i$ for every $1 \leq j \leq k$. We claim that this happens with probability $M^{-(k-1)}$. Indeed, the $k-1$ events $\beta_j(I_j) = b_i$, $1 \leq j \leq k-1$, are independent, each of them occurs with probability M^{-1} , and once they all occur, it follows that $\beta_k(I_k) = b_i$ as well. The number of partitions of $[N]$ in \mathcal{I}^k is $\binom{nk}{n, \dots, n}$, so by linearity of expectation the claim follows. ■

We now turn to estimate the expected number of partitions that are removed from every L_i in the second step of the construction. To do so, we have to consider the probability of two distinct partitions in \mathcal{I}^k that share some subset to belong to L_i . However, this probability depends on the specific pair of partitions. Indeed, sharing more than one subset of the partitions, or even a certain union of the subsets, might increase this probability. Hence, we need the following definition.

Definition 4.6. Let $t_1 \leq \dots \leq t_\ell$ be ℓ positive integers satisfying $\sum_{r=1}^\ell t_r = k$. We say that two partitions (I_1, \dots, I_k) and (J_1, \dots, J_k) of $[N]$ in \mathcal{I}^k are (t_1, \dots, t_ℓ) -similar if there exists a partition of $[k]$ into ℓ sets T_1, \dots, T_ℓ of sizes t_1, \dots, t_ℓ respectively, such that for some permutation $\pi : [k] \rightarrow [k]$,

$$\cup_{i \in T_r} I_i = \cup_{i \in T_r} J_{\pi(i)} \quad (3)$$

for every $1 \leq r \leq \ell$, and, in addition, no refinement of the partition T_1, \dots, T_ℓ satisfies (3) for any permutation π .

Claim 4.7. Let (I_1, \dots, I_k) and (J_1, \dots, J_k) be distinct (t_1, \dots, t_ℓ) -similar partitions of $[N]$ in \mathcal{I}^k for some ℓ positive integers $t_1 \leq \dots \leq t_\ell$ satisfying $\sum_{r=1}^\ell t_r = k$. Then, for every $1 \leq i \leq |B|$, the following holds.

1. For $1 \leq \ell \leq k - 1$, the probability that the two partitions are in L_i is at most $M^{-(k-1)} \cdot M^{-(k-\ell)}$.
2. For $\ell = k$, the probability that the two partitions are in L_i is at most $M^{-(k-1)} \cdot M^{-1}$.

Proof: Let (I_1, \dots, I_k) and (J_1, \dots, J_k) be distinct (t_1, \dots, t_ℓ) -similar partitions of $[N]$ in \mathcal{I}^k , and let T_1, \dots, T_ℓ and π be the corresponding partition and permutation of $[k]$ as in Definition 4.6.

For Item 1, we fix the values of c_1, \dots, c_k and analyze the probability that the two partitions (I_1, \dots, I_k) and (J_1, \dots, J_k) are in L_i over the random choice of w_1, \dots, w_N . First, notice that the $k - 1$ events $\beta_j(I_j) = b_i$ for $1 \leq j \leq k - 1$ are independent, occur with probability M^{-1} each, and imply the event $\beta_k(I_k) = b_i$. So the probability that $(I_1, \dots, I_k) \in L_i$ is $M^{-(k-1)}$. For $1 \leq r \leq \ell$, let A_r denote the event that the equalities $\beta_{\pi(j)}(J_{\pi(j)}) = b_i$ hold for every $j \in T_r$. It can be shown that for every $1 \leq r \leq \ell$, the probability that A_r occurs conditioned on the event $(I_1, \dots, I_k) \in L_i$ and on A_1, \dots, A_{r-1} is $M^{-(t_r-1)}$. Indeed, since no refinement of the partition T_1, \dots, T_ℓ satisfies the condition of Definition 4.6, it follows that $t_r - 1$ of the equalities of A_r are independent, occur with probability M^{-1} each, and might imply the last one. This can be verified by observing that every vector J_j that corresponds to such an equality is linearly independent of the vectors that correspond to the previously considered equalities. Hence, the probability that the two k -tuples are in L_i is at most

$$M^{-(k-1)} \cdot M^{-\sum_{r=1}^\ell (t_r-1)} = M^{-(k-1)} \cdot M^{-(k-\ell)}.$$

For Item 2, take $\ell = k$ and notice that in this case, (J_1, \dots, J_k) is a permutation of (I_1, \dots, I_k) , so we have $t_1 = \dots = t_\ell = 1$. The probability that $(I_1, \dots, I_k) \in L_i$ is again $M^{-(k-1)}$. Since the two partitions are distinct, there are distinct j, j' for which $I_j = J_{j'}$. Assume, without loss of generality, that $j' < k$. By the randomness of the choice of $c_{j'}$, the probability that $\beta_{j'}(J_{j'}) = b_i$, conditioned on $(I_1, \dots, I_k) \in L_i$, is M^{-1} . Therefore, the probability that both the partitions are in L_i is at most $M^{-(k-1)} \cdot M^{-1}$. ■

Claim 4.8. For every $1 \leq i \leq |B|$, the expected size of the set L'_i is at least $\frac{1}{2} \cdot \binom{nk}{n, \dots, n} \cdot M^{-(k-1)}$.

Proof: By Claim 4.5, the expected size of L_i is $\binom{nk}{n, \dots, n} \cdot M^{-(k-1)}$. We turn to bound from above the expected number of partitions removed from L_i in the second step. Notice that if two partitions of $[N]$ in \mathcal{I}^k share a subset then they are (t_1, \dots, t_ℓ) -similar for some $\ell \geq 2$ positive integers $t_1 \leq \dots \leq t_\ell$ satisfying $\sum_{r=1}^\ell t_r = k$ (in fact, we also know that at least one of the t_r 's equals 1). Therefore, we turn to bound the expected number of (ordered) pairs of (t_1, \dots, t_ℓ) -similar partitions of $[N]$ in \mathcal{I}^k for some t_1, \dots, t_ℓ as above.

We start with the case $2 \leq \ell \leq k-1$. Fix a partition (I_1, \dots, I_k) of $[N]$ in \mathcal{I}^k , and ℓ positive integers $t_1 \leq \dots \leq t_\ell$ satisfying $\sum_{r=1}^\ell t_r = k$. The number of partitions (J_1, \dots, J_k) of $[N]$ in \mathcal{I}^k which are (t_1, \dots, t_ℓ) -similar to (I_1, \dots, I_k) , associated with certain partition T_1, \dots, T_ℓ and permutation π of $[k]$, is at most

$$\binom{n \cdot t_1}{n, \dots, n} \cdots \binom{n \cdot t_\ell}{n, \dots, n} \leq \binom{n(k-\ell+1)}{n, \dots, n}.$$

By Item 1 of Claim 4.7, the probability that two such partitions are in L_i is at most $M^{-(k-1)} \cdot M^{-(k-\ell)}$. It follows that the expected total number of pairs of (t_1, \dots, t_ℓ) -similar partitions in L_i for some $t_1 \leq \dots \leq t_\ell$ as above ($2 \leq \ell \leq k-1$) is at most

$$k^{O(k)} \cdot \binom{nk}{n, \dots, n} \cdot \binom{n(k-\ell+1)}{n, \dots, n} \cdot M^{-(k-1)} \cdot M^{-(k-\ell)},$$

where the $k^{O(k)}$ term counts all the possible choices of numbers t_1, \dots, t_ℓ , partitions T_1, \dots, T_ℓ , and permutations π of $[k]$. Now, consider the case $k = \ell$, that is, $t_1 = \dots = t_\ell = 1$. Using Item 2 of Claim 4.7, the expected number of $(1, \dots, 1)$ -similar partitions in L_i is at most

$$k^{O(k)} \cdot \binom{nk}{n, \dots, n} \cdot M^{-(k-1)} \cdot M^{-1}.$$

Therefore, by linearity of expectation, the expected number of partitions removed from L_i is at most

$$\binom{nk}{n, \dots, n} \cdot M^{-(k-1)} \cdot \left(k^{O(k)} \cdot M^{-1} + \sum_{\ell=2}^{k-1} k^{O(k)} \cdot \binom{n(k-\ell+1)}{n, \dots, n} \cdot M^{-(k-\ell)} \right).$$

Choosing the constant $c(k)$ in (1) to be sufficiently large, we have $k^{O(k)} \cdot M^{-1} \leq \frac{1}{2k}$. We turn to prove that for every $2 \leq \ell \leq k-1$,

$$k^{O(k)} \cdot \binom{n(k-\ell+1)}{n, \dots, n} \cdot M^{-(k-\ell)} \leq \frac{1}{2k'} \tag{4}$$

as this implies, combined with Claim 4.5, that the expected size of L'_i is at least

$$\binom{nk}{n, \dots, n} \cdot M^{-(k-1)} - \frac{1}{2} \cdot \binom{nk}{n, \dots, n} \cdot M^{-(k-1)} = \frac{1}{2} \cdot \binom{nk}{n, \dots, n} \cdot M^{-(k-1)}.$$

For (4), observe that our choice of M satisfies

$$k^{O(k)} \cdot \binom{n(k-\ell+1)}{n, \dots, n} \cdot M^{-(k-\ell)} \leq \frac{1}{2k} \cdot \binom{n(k-\ell+1)}{n, \dots, n} \cdot \binom{n(k-1)}{n, \dots, n}^{-\frac{k-\ell}{k-2}} \leq \frac{1}{2k'}$$

where the first inequality holds for a sufficiently large constant $c(k)$ in (1), and the second follows from the inequality $\binom{n(k-\ell+1)}{n, \dots, n}^{1/(k-\ell)} \leq \binom{n(k-1)}{n, \dots, n}^{1/(k-2)}$ that holds for every $\ell \geq 2$ by monotonicity of the geometric mean. \blacksquare

Finally, using Claim 4.8, we conclude that there exists a choice of w_1, \dots, w_N and c_1, \dots, c_k for which

$$|\mathcal{F}| = \sum_{i=1}^{|B|} |L'_i| \geq \frac{1}{2} \cdot \binom{nk}{n, \dots, n} \cdot M^{-(k-1)} \cdot |B| \geq \binom{nk}{n, \dots, n} \cdot M^{-(k-2)-o(1)} \geq \binom{nk}{n}^{1-o(1)}.$$

By standard estimations of Binomial coefficients, this completes the proof of Theorem 4.1. \blacksquare

Corollary 4.9. For every $k \geq 3$ and a prime p , the local PMF capacity for k -cycles over \mathbb{F}_p is at least $2^{H(1/k)}$.

Proof: By Theorem 4.1, for every sufficiently large n , there exists a collection \mathcal{F} of $(2^{H(1/k)} - o(1))^{nk}$ balanced vectors in \mathbb{Z}_k^{nk} such that for every k vectors $v_1, \dots, v_k \in \mathcal{F}$, the sets $v_1|_1, \dots, v_k|_k$ form a partition of $[n \cdot k]$ if and only if $v_1 = \dots = v_k$. For every vector $v \in \mathcal{F}$ consider the k -tuple of vectors, whose first $k - 1$ vectors are the characteristic vectors of $v|_1, v|_2, \dots, v|_{k-1}$, and the last one is the characteristic vector of $[n] \setminus v|_k$ multiplied by -1 (modulo p). Observe that the collection of all k -tuples obtained in this way from the vectors of \mathcal{F} is an $(nk, |\mathcal{F}|)$ local PMF for k -cycles over \mathbb{F}_p . Hence, the local PMF capacity for k -cycles over \mathbb{F}_p is at least $2^{H(1/k)}$, as required. ■

The above corollary, combined with Lemma 2.8, completes the proof of Theorem 1.8.

References

- [1] N. Alon. Testing subgraphs in large graphs. *Random Struct. Algorithms*, 21(3-4):359–370, 2002. Preliminary version in FOCS’01.
- [2] N. Alon and R. B. Boppana. The monotone circuit complexity of boolean functions. *Combinatorica*, 7(1):1–22, 1987.
- [3] N. Alon, E. Fischer, I. Newman, and A. Shapira. A combinatorial characterization of the testable graph properties: It’s all about regularity. *SIAM J. Comput.*, 39(1):143–167, 2009. Preliminary version in STOC’06.
- [4] N. Alon, R. Hod, and A. Weinstein. On active and passive testing. *CoRR*, abs/1307.7364, 2013.
- [5] N. Alon, A. Shpilka, and C. Umans. On sunflowers and matrix multiplication. *Computational Complexity*, 22(2):219–243, 2013. Preliminary version in CCC’12.
- [6] M. Bateman and N. H. Katz. New bounds on cap sets. *J. Amer. Math. Soc.*, 25(2):585–613, 2012.
- [7] F. A. Behrend. On sets of integers which contain no three terms in arithmetical progression. *Proc. National Academy of Sciences USA*, 32(12):331–332, 1946.
- [8] A. Bhattacharyya. Guest column: On testing affine-invariant properties over finite fields. *SIGACT News*, 44(4):53–72, 2013.
- [9] A. Bhattacharyya, E. Grigorescu, and A. Shapira. A unified framework for testing linear-invariant properties. *Random Struct. Algorithms*, 46(2):232–260, 2015. Preliminary version in FOCS’10.
- [10] A. Bhattacharyya and N. Xie. Lower bounds for testing triangle-freeness in Boolean functions. *Computational Complexity*, 24(1):65–101, 2015. Preliminary version in SODA’10.
- [11] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós, B. Szegedy, and K. Vesztegombi. Graph limits and parameter testing. In *STOC*, pages 261–270, 2006.

- [12] H. Cohn, R. D. Kleinberg, B. Szegedy, and C. Umans. Group-theoretic algorithms for matrix multiplication. In *FOCS*, pages 379–388, 2005.
- [13] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. Symb. Comput.*, 9(3):251–280, 1990. Preliminary version in STOC’87.
- [14] I. Dinur and S. Safra. On the hardness of approximating minimum vertex cover. *Annals of Mathematics*, 162(1):439–485, 2005. Preliminary version in STOC’02.
- [15] Y. Edel. Extensions of generalized product caps. *Des. Codes Cryptography*, 31(1):5–14, 2004.
- [16] M. Elkin. An improved construction of progression-free sets. *Israel J. of Math.*, 184(1):93–128, 2011. Preliminary version in SODA’10.
- [17] P. Erdős and R. Rado. Intersection theorems for systems of sets. *J. London Math. Soc.*, 35:85–90, 1960.
- [18] P. Erdős and E. Szemerédi. Combinatorial properties of systems of sets. *J. Comb. Theory, Ser. A*, 24(3):308–313, 1978.
- [19] J. Fox. A new proof of the graph removal lemma. *Annals of Mathematics*, 174(1):561–579, 2011.
- [20] H. Fu and R. Kleinberg. Improved lower bounds for testing triangle-freeness in boolean functions via fast matrix multiplication. In *RANDOM*, pages 669–676, 2014.
- [21] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *J. ACM*, 45(4):653–750, 1998. Preliminary version in FOCS’96.
- [22] B. Green. A Szemerédi-type regularity lemma in Abelian groups. *Geom. and Funct. Anal.*, 15(2):340–376, 2005.
- [23] P. Hatami, S. Sachdeva, and M. Tulsiani. An arithmetic analogue of Fox’s triangle removal argument. *CoRR*, abs/1304.4921, 2013.
- [24] T. Kaufman and M. Sudan. Algebraic property testing: the role of invariance. In *STOC*, pages 403–412, 2008.
- [25] A. V. Kostochka. A bound of the cardinality of families not containing Δ -systems. In *The Mathematics of Paul Erdős II, Algorithms and Combinatorics*, volume 14, pages 229–235. Springer, Berlin, 1997.
- [26] D. Král’, O. Serra, and L. Vena. A combinatorial proof of the removal lemma for groups. *J. Comb. Theory, Ser. A*, 116(4):971–978, 2009.
- [27] D. Král’, O. Serra, and L. Vena. A removal lemma for systems of linear equations over finite fields. *Israel J. of Math.*, 187(1):193–207, 2012.
- [28] Y.-R. Liu and C. V. Spencer. A generalization of Meshulam’s theorem on subsets of finite Abelian groups with no 3-term arithmetic progression. *Des. Codes Cryptography*, 52(1):83–91, 2009.

- [29] R. Meshulam. On subsets of finite Abelian groups with no 3-term arithmetic progressions. *J. Comb. Theory, Ser. A*, 71(1):168–172, 1995.
- [30] A. A. Razborov. Lower bounds for the monotone complexity of some boolean functions. *Soviet Mathematics Doklady*, 31(2):354–357, 1985.
- [31] R. Rubinfeld and M. Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM J. Comput.*, 25(2):252–271, 1996. Preliminary version in SODA’92.
- [32] R. Salem and D. C. Spencer. On sets of integers which contain no three terms in arithmetical progression. *Proc. National Academy of Sciences USA*, 28(12):561–563, 1942.
- [33] A. Shapira. A proof of Green’s conjecture regarding the removal properties of sets of linear equations. *J. London Math. Soc.*, 81(2):355–373, 2010. Preliminary version in STOC’09.
- [34] A. J. Stothers. *On the complexity of matrix multiplication*. PhD thesis, The University of Edinburgh, 2010.
- [35] M. Sudan. Guest column: Testing linear properties: some general theme. *SIGACT News*, 42(1):59–80, 2011.
- [36] J. H. van Lint and R. M. Wilson. *A course in combinatorics*. Cambridge University Press, second edition, 2001.
- [37] V. V. Williams. Multiplying matrices faster than Coppersmith-Winograd. In *STOC*, pages 887–898, 2012.