Forum of Mathematics, Sigma (2021), **Vol. 9**:e19 1–10
doi:10.1017/fms.2021.1

CAMBRIDGE UNIVERSITY PRESS

RESEARCH ARTICLE

High-entropy dual functions over finite fields and locally decodable codes

Jop Briët and Farrokh Labib

CWI, Science Park 123, 1098 XG Amsterdam, The Netherlands; E-mails: j.briet@cwi.nl and labib@cwi.nl.

Received: 5 November 2020; **Revised:** 11 December 2020; **Accepted:** 20 December 2020

2020 Mathematics subject classification: *Primary* – 11B30; *Secondary* – 11B25

Abstract

We show that for infinitely many primes p there exist dual functions of order k over $\mathbb{F}_p^n$ that cannot be approximated in L_∞-distance by polynomial phase functions of degree $k-1$. This answers in the negative a natural finite-field analogue of a problem of Frantzikinakis on L_∞-approximations of dual functions over $\mathbb{N}$ (a.k.a. multiple correlation sequences) by nilsequences.

1. Introduction

For $k \geq 2$, integer vector $\mathbf{i} = (i_1, \dots, i_k) \in \mathbb{Z}_{\geq 0}^k$ and finite abelian group G, the associated set of *order-k dual functions* is given by

$$\Delta_{\mathbf{i}} = \big\{\phi : y \mapsto \mathbb{E}_{x \in G} f_1(x + i_1 y) \cdots f_k(x + i_k y) \;\mid\; f_i : G \to \mathbb{D}\big\},$$

where $\mathbb{D}$ denotes the complex unit disc. For example, if $A \subseteq G$ is a subset, $\mathbf{i} = (0, 1, 2)$ and $f_i = 1_A$ for each $i \in [3]$, then $\phi(y)$ is the fraction of three-term arithmetic progressions in A with common difference y.

For applications in additive combinatorics and higher-order Fourier analysis, it is desirable to understand to what extent dual functions can be approximated by simpler functions. If $k = 2$, it follows from the Fourier inversion formula that one has the simple decomposition in terms of the characters:

$$\phi(y) = \sum_{\chi \in \widehat{G}} \alpha_\chi \chi\big((i_2 - i_1)y\big), \tag{1}$$

where $\|\alpha\|_{\ell_1} \leq 1$. Similar decompositions exist for higher-order dual functions thanks to deep 'inverse theorems' for the Gowers uniformity norms. Inverse theorems roughly show that if f has large U^k-norm, then f correlates with a function $\psi : G \to \mathbb{D}$ akin to a polynomial of degree at most $k-1$. Here the 'linear' ψ are precisely the characters. What exactly the 'higher-order characters' are depends on the group G. For finite vector spaces $\mathbb{F}_p^n$ with $p \geq k$, they are the *polynomial phase functions*

$$\psi(x) = e^{2\pi i P(x)/p},$$

where $P \in \mathbb{F}_p[x_1, \dots, x_n]$ is a polynomial of degree at most $k-1$ [27]. When $p < k$, one has to consider the larger class of nonclassical polynomials [28]. For the cyclic group $\mathbb{Z}_N$, they are the $(k-1)$-step nilsequences (of bounded complexity) [20]. Combined with the Hahn-Banach theorem, these inverse theorems imply that the decomposition (1) generalises for larger k in terms of higher-order characters of degree at most $k-1$ up to small L_1-error [19]. More precisely, in the finite-field setting, this amounts to the following.

Proposition 1.1. *Let $p \geq k+1$ be a prime and let $G = \mathbb{F}_p^n$. Then, for any $\varepsilon > 0$ and $\mathbf{i} \in \mathbb{Z}_{\geq 0}^k$, there is an $M = M(\varepsilon, k, p) > 0$ such that any dual function $\phi \in \Delta_{\mathbf{i}}$ can be decomposed as*

$$\phi = \sum_{i=1}^{r} \alpha_i \psi_i + \tau, \tag{2}$$

where $\alpha_1, \dots, \alpha_r \in \mathbb{C}$ satisfy $|\alpha_1| + \cdots + |\alpha_r| \leq M$, $\psi_1, \dots, \psi_r$ are polynomial phases of degree at most $k-1$ and $\|\tau\|_{L_1} \leq \varepsilon$.

Though facts like this (in particular over $\mathbb{Z}_N$) can be useful in higher-order Fourier analysis [19], for other applications in additive combinatorics it is preferable to have more precise control over the error function τ in (2). A natural finite-field analogue of a question raised by Frantzikinakis in [12, Problem 1] (see also [1]) asks whether this error function can be bounded *everywhere*; that is, whether Proposition 1.1 still holds with $\|\tau\|_{L_\infty} \leq \varepsilon$. The apparent expectation of a positive answer to Frantzikinakis's question motivated conjectures on a poorly understood probabilistic variant of Szemerédi's theorem on arithmetic progressions (cf. Section 1.1). Our main result, however, shows that in the finite-field setting, the answer is negative.

Theorem 1.2. *For infinitely many primes p, there is a $k = k(p) \in \mathbb{N}$ and an integer vector $\mathbf{i} \in \mathbb{Z}_{\geq 0}^k$ such that (2) cannot hold with $\|\tau\|_{L_\infty} \leq \varepsilon$.*

Special cases of Theorem 1.2 show that for $k = 3$ and $p = 2^t - 1$ a Mersenne prime, the decomposition (2) requires polynomial phases of degree at least t for fixed ε, M and $\|\tau\|_{L_\infty} \leq \varepsilon$. The largest known Mersenne prime as of January 2018 has $t = 77,232,917$ [16].

1.1. Locally decodable codes and random Szemerédi

The examples behind Theorem 1.2 originate from constructions of special types of error-correcting codes called *locally decodable codes* (LDCs). These codes have the property that any single encoded message symbol can be retrieved from a code word with good probability by reading only a tiny number of code word symbols, even if the code word is partially corrupted. LDCs originated in complexity theory [2–4] and cryptography [10] and were defined in the context of channel coding in [21]. They have since found many other applications in computer science and mathematics; for instance, in fault-tolerant distributed storage systems [17] and Banach space geometry [9]. We refer to [18, 31] for extensive surveys.

Despite their ubiquity, LDCs are poorly understood. Of particular interest is the trade-off between the code word length N as a function of message length k when the *query complexity* – the number of probed code word symbols – and alphabet size are constant. The Hadamard code is a 2-query LDC of length $N = 2^{O(k)}$ and this length is optimal in the 2-query regime [22]. For $q \geq 3$, the best-known lower bounds show that any q-query LDC has at least polynomial length $k^{1+1/(\lceil q/2 \rceil - 1) - o(1)}$ [22, 29]. The family of Reed-Muller codes, which generalise the Hadamard code, were for a long time the best-known examples, giving q-query LDCs of length $\exp(O(k^{1/(q-1)}))$.

In a breakthrough result, Yekhanin [30] constructed an entirely new family of vastly shorter LDCs. For each Mersenne prime $p = 2^t - 1$, he gave a 3-query LDC of length $N \leq \exp(O(k^{1/t}))$. The construction uses a family of k homomorphisms from $\mathbb{F}_p^n$ to the multiplicative subgroup of $\mathbb{F}_{2^t}$. The homomorphisms are constructed using a family of *matching vectors* $(u_i, v_i)_{i \in [k]}$, which are pairs of

https://doi.org/10.1017/fms.2021.1 Published online by Cambridge University Press

orthogonal vectors in $\mathbb{F}_p^n$ such that the inner products $\langle u_i, v_j \rangle$ with $i \neq j$ belong to a special subset of $\mathbb{F}_p^*$. It is this construction that forms the basis for Theorem 1.2.

Subsequently, Efremenko [11] constructed much larger matching vector families over $\mathbb{Z}_m^n$ for composite moduli m and used Yekhanin's framework to give the first 3-query LDCs of subexponential length $N \leq \exp(\exp(O\sqrt{\log k / \log\log k}))$. But huge gaps persist between the best-known upper and lower bounds for constant-query LDCs.

In contrast with other combinatorial objects such as expander graphs, the probabilistic method has so far not been successfully used to beat the best explicit LDC constructions. In [6], a probabilistic framework was given that could in principle yield best possible LDCs, albeit nonconstructively. A special instance of this framework connects LDCs with a probabilistic version of Szemerédi's theorem alluded to above. The setup for this is as follows.

For a finite abelian group G of size $N = |G|$, let $D \subseteq G$ be a random subset where each element is present with probability ρ independent of all others. For $k \geq 3$ and $\varepsilon \in (0, 1)$, let E be the event that every subset $A \subseteq G$ of size $|A| \geq \varepsilon |G|$ contains a proper k-term arithmetic progression with common difference in D. If $\rho = 1$, then it follows from the Density Hales-Jewett Theorem [15] that E holds with probability 1 provided N is large enough in terms of k and ε. It is an open problem to determine the smallest value of ρ – which we will denote by ρ_k – such that $\Pr[E] \geq \frac{1}{2}$. This value will depend on ε too, but we will suppress this in the notation and assume that ε is a fixed constant. It is also assumed that N is large enough so that ρ_k exists.

In [6] it is shown that there exist k-query LDCs of message length $\Omega(\rho_k N)$ and code word length $O(N)$. As such, Szemerédi's theorem with random differences, in particular lower bounds on ρ_k, can be used to show the existence of LDCs. Conversely, this connection indirectly implies the best-known upper bounds on ρ_k for all $k \geq 3$, given by $N^{-(1-o(1))/\lceil k/2 \rceil}$ [7, 13]. However, a conjecture of Frantzikinakis et al. [14] states that over $\mathbb{Z}_N$ we have $\rho_k \ll_k N^{-1} \log N$ for all k, which would be best possible. Truth of this conjecture would imply that over this group, Szemerédi's theorem with random differences cannot give LDCs better than the Hadamard code. For finite fields, Altman [1] showed that this conjecture is false. In particular, over $\mathbb{F}_p^n$ for p odd, he proved that $\rho_3 \gg p^{-n} n^2$; generally, $\rho_k \gg p^{-n} n^{k-1}$ holds when $p \geq k + 1$ [5]. In turn, these bounds are conjectured to be optimal for the finite-field setting, which would imply that over finite fields, Szemerédi's theorem with random differences cannot give LDCs better than Reed-Muller codes.

These conjectures appear to be motivated mainly by the possibility of an L_∞-version of Proposition 1.1 (and analogous variants over $\mathbb{Z}_N$) with dual functions based on 3-term progressions. Theorem 1.2 falls short of obstructing this route to obtaining optimal bounds in the finite-field setting for two reasons. First, our examples do not include 'arithmetic-progression dual functions,' those with $\mathbf{i} = (0, 1, \ldots, (k-1))$; in fact, in the Appendix we show that our current framework cannot give such examples. Second, even if we had such examples, they do not appear to imply any new lower bounds on ρ_k. Nevertheless, we do not expect arithmetic progressions to be exceptional patterns for which there are no such examples.

Remark 1.3. Ideas behind Theorem 1.2 recently inspired similar examples in the integer setting for 3-term progressions in joint work of the first author and Green [8].

2. Preliminaries

We will identify the set of maps $G \to \mathbb{C}$ with $\mathbb{C}^G$. For a polynomial $P(x) = \sum_{\iota=0}^{t} c_\iota x^\iota$, define its *support* $\mathbf{i}(P)$ to be the sequence of degrees $\iota \in \mathbb{Z}_{\geq 0}$ such that $c_\iota \neq 0$, arranged in increasing order. The *support size* is the length of $\mathbf{i}(P)$. We will use some basic facts from the theory of finite fields, for which we refer to [23]. The Minkowski sum of two sets $A, B \subseteq \mathbb{C}^n$ is the set given by

$$A + B = \{a + b \mid a \in A,\, b \in B\}.$$

https://doi.org/10.1017/fms.2021.1 Published online by Cambridge University Press

We will use the following slight generalisation of the notion of the convex hull, where we allow for complex coefficients. For a compact set $A \subseteq \mathbb{C}^n$, define

$$\mathrm{Conv}_{\mathbb{C}}(A) = \Big\{ \sum_{a\in A} \alpha_a a \;\mid\; \alpha_a \in \mathbb{C} \quad \forall a \in A, \quad \sum_{a\in A} |\alpha_a| \le 1 \Big\}.$$

For a finite set $A \subseteq \mathbb{D}^n$ and $\varepsilon, M \in (0,\infty)$, define $\mathcal{N}(A,\varepsilon,M)$ to be the smallest size of a finite set $B \subseteq M\mathbb{D}^n$ such that

$$A \subseteq \mathrm{Conv}_{\mathbb{C}}(B) + \varepsilon\mathbb{D}^n.$$

Then, for any $a \in A$, there is a $b \in \mathrm{Conv}_{\mathbb{C}}(B)$ such that $\|a-b\|_{\ell_\infty} \le \varepsilon$ and so $\mathcal{N}(A,\varepsilon,M)$ is a restricted form of the covering number of A relative to the ℓ_∞ distance. Note that for $I \subseteq [n]$, the projection of A to the set of coordinates I, given by $A_I = \{(a_i)_{i\in I} \mid a \in A\}$, is contained in $\mathrm{Conv}_{\mathbb{C}}(B_I) + \varepsilon\mathbb{D}^I$. Because $|B| \ge |B_I|$, it follows that

$$\mathcal{N}(A,\varepsilon,M) \ge \mathcal{N}(A_I,\varepsilon,M). \tag{3}$$

3. Covering numbers from hypercubes

We will use the following lemma, which shows that containment of a high-dimensional hypercube implies a large restricted covering number.

Lemma 3.1. *Let $c > 0$, $z \in \mathbb{C}$ be a complex number such that $\Re(z) \le 0$ and let $S \subseteq \mathbb{C}^k$ be a finite set such that $\{c,z\}^k \subseteq S$. Then, for any $\varepsilon \in (0, \frac{c}{2})$ and $M > 0$, we have that*

$$\log_2\big(\mathcal{N}(S,\varepsilon,M)\big) \gg_{c,\varepsilon,M} k.$$

Proof. Let θ be a uniformly distributed $\{-1,1\}^k$-valued random vector. For a compact set $A \subseteq \mathbb{C}^k$, define

$$w(A) = \mathbb{E}\max_{a\in A} |\langle a,\theta\rangle|.$$

We use the following basic properties:

1. If $A \subseteq B$, then $w(A) \le w(B)$.
2. For a finite set $A \subseteq \mathbb{C}^k$, it holds that $w(\mathrm{Conv}_{\mathbb{C}}(A)) = w(A)$.
3. For $A, B \subseteq \mathbb{C}^k$ finite, it holds that $w(A+B) \le w(A) + w(B)$.

It follows from the first property that

$$w(S) \ge w(\{c,z\}^k) \ge \frac{ck}{2}. \tag{4}$$

For the second inequality, observe that for fixed $\theta \in \{-1,1\}^k$, we have

$$\begin{aligned}\max_{a\in\{c,z\}^k} |\langle a,\theta\rangle| &\ge \Big| \sum_{i:\theta_i=1} c - \sum_{i:\theta_i=-1} z \Big| \\ &\ge \Big| \Re\Big(\sum_{i:\theta_i=1} c - \sum_{i:\theta_i=-1} z \Big) \Big| \\ &\ge c|\{i\in[k] \mid \theta_i = 1\}|.\end{aligned}$$

Averaging over θ then gives the result.

https://doi.org/10.1017/fms.2021.1 Published online by Cambridge University Press

Let $B \subseteq M\mathbb{D}^k$ be a finite set such that $S \subseteq \mathrm{Conv}_{\mathbb{C}}(B) + \varepsilon\mathbb{D}^k$. Let $l = |B|$ and $p = \log_2 l$. By the second property of w, Jensen's inequality and the Khintchine inequality [24, Chapter 5],

$$\begin{aligned} w\big(\mathrm{Conv}_{\mathbb{C}}(B)\big) &= \mathbb{E}\max_{b\in B}|\langle b,\theta\rangle| \\ &\le \mathbb{E}\Big(\sum_{b\in B}|\langle b,\theta\rangle|^p\Big)^{\frac{1}{p}} \\ &\le \Big(\sum_{b\in B}\mathbb{E}|\langle b,\theta\rangle|^p\Big)^{\frac{1}{p}} \\ &\ll \sqrt{p}\Big(|B|\max\{\|b\|_{\ell_2}^p \mid b\in B\}\Big)^{\frac{1}{p}} \\ &\ll M\sqrt{k\log l}. \end{aligned}$$

We also have $w(\varepsilon\mathbb{D}^k) = \varepsilon k$. Because $S \subseteq \mathrm{Conv}_{\mathbb{C}}(B) + \varepsilon\mathbb{D}^k$, the second and third properties of w and (4) then give

$$\frac{ck}{2} \le w(S) \le w(\mathrm{Conv}_{\mathbb{C}}(B)) + \varepsilon\mathbb{D}^k) \ll M\sqrt{k\log_2 l} + \varepsilon k.$$

Rearranging the left- and right-hand sides now gives the claim. □

4. Locating high-dimensional hypercubes

Here we show that for certain primes p and some integer vectors $\mathbf{i}$, the dual functions in $\Delta_{\mathbf{i}}$ over $\mathbb{F}_p^n$ contain high-dimensional hypercubes.

Proposition 4.1. *Let p, r be distinct primes, let $t = \mathrm{ord}_p(r)$ and let $G = \mathbb{F}_p^n$. Suppose there exists a polynomial $P(x) \in \mathbb{F}_r[x]$ that has a root in $\mathbb{F}_{r^t}^*$ of order p and such that $P(1) \neq 0$. Then, there exists a $z \in \mathbb{C}$ with $\Re(z) \le 0$ and a set $D \subseteq G$ of size $|D| \gg_p n^t$ such that*

$$\{z, 1\}^D \subseteq \Delta^D_{\mathbf{i}(P)}.$$

The proof of this proposition relies on the following result due to Yekhanin, which is implicit in [30] (and shown explicitly in [25]). We include a proof for completeness.

Theorem 4.2 (Yekhanin). *Let p, r be distinct primes and $t := \mathrm{ord}_p(r)$. For integer $m > p - 1$, let*

$$k = \binom{m}{p-1} \qquad \text{and} \qquad n = \binom{m + \frac{p-1}{t} - 1}{\frac{p-1}{t}}.$$

Let

$$P(x) = \sum_{\iota=0}^{s} c_\iota x^\iota \in \mathbb{F}_r[x]$$

be a polynomial with a root $\gamma \in \mathbb{F}_{r^t}^*$ of order p. Then, for each $i \in [k]$ there exists a function $f_i : \mathbb{F}_p^n \to \mathbb{F}_{r^t}$ and vectors $d_i, w_i \in \mathbb{F}_p^n$ such that for every $x \in \mathbb{F}_p^n$, we have

$$\sum_{\iota=0}^{s} c_\iota f_i(x + \iota d_j) = \begin{cases} \gamma^{\langle x, w_i\rangle} P(1) & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

https://doi.org/10.1017/fms.2021.1 Published online by Cambridge University Press

Proof. For a $(p-1)$-element subset $S \subseteq [m]$, define the vectors $u_S = 1_S$ and $v_S = 1_{[m]} - u_S$ in $\mathbb{F}_p^m$. Then, $\langle u_S, v_T \rangle = 0$ if and only if $S = T$. Let $l = \frac{p-1}{t}$. Then, for $a \in \mathbb{F}_p^*$, we have $a^l \in \{r^q \mid q = 0, 1, \ldots, p-1\}$.

Consider the expansion of the polynomial $Q(x) \in \mathbb{F}_p[x_1, \ldots, x_m]$ given by

$$Q(x) = (x_1 + \cdots + x_m)^l = \sum_{\beta \in \mathcal{M}_l} c_\beta x^\beta,$$

where $\mathcal{M}_l := \{\beta \in \mathbb{Z}_{\geq 0}^m \mid \sum_{i=1}^m \beta_i = l\}$ and $x^\beta := \prod_{i=1}^m x_i^{\beta_i}$. For each subset $S \subseteq [m]$ of size $p-1$, define the vectors $w_S = (u_S^\beta)_{\beta \in \mathcal{M}_l}$ and $d_S = (c_\beta v_S^\beta)_{\beta \in \mathcal{M}_l}$. Because $x^\beta y^\beta = (x \circ y)^\beta$, where $\circ$ denotes the coordinate-wise product, we have that

$$\langle w_S, d_T \rangle = Q(u_S \circ v_T) = \langle u_S, v_T \rangle^l.$$

By the above, this equals zero if $S = T$ and a power of r otherwise. Moreover, the vectors w_S and d_S have dimension $|\mathcal{M}_l| = \binom{m+l-1}{l}$.

Define $f_S : \mathbb{F}_p^n \to \mathbb{F}_{r^t}^*$ by

$$f_S(x) = \gamma^{\langle x, w_S \rangle}.$$

Note that this a homomorphism, because γ has order p. Then,

$$\begin{aligned}\sum_{\iota=0}^{s} c_\iota f_S(x + \iota d_S) &= \gamma^{\langle x, w_S \rangle} \sum_{\iota=0}^{s} c_\iota \gamma^{\iota \langle d_S, w_S \rangle} \\ &= \gamma^{\langle x, w_S \rangle} \sum_{\iota=0}^{s} c_\iota \\ &= \gamma^{\langle x, w_S \rangle} P(1).\end{aligned}$$

If $S \neq T$, then $\langle d_T, w_S \rangle = r^q \bmod p$ for some integer q and, therefore,

$$\begin{aligned}\sum_{\iota=0}^{s} c_\iota f_S(x + \iota d_T) &= \gamma^{\langle x, w_S \rangle} \sum_{\iota=0}^{s} c_\iota \gamma^{\iota \langle d_T, w_S \rangle} \\ &= \gamma^{\langle x, w_S \rangle} \sum_{\iota=0}^{s} c_\iota \gamma^{\iota r^q} \\ &= \gamma^{\langle x, w_S \rangle} P(\gamma)^{r^q} \\ &= 0.\end{aligned}$$

This completes the proof. □

Proof of Proposition 4.1. Let $P(x) \in \mathbb{F}_r[x]$ be as in Proposition 4.1 and let $\gamma \in \mathbb{F}_{r^t}^*$ be a pth root of unity such that $P(\gamma) = 0$. Let $f_i : \mathbb{F}_p^n \to \mathbb{F}_{r^t}^*$ and $d_i, w_i \in \mathbb{F}_p^n$ be as in Theorem 4.2. Let $\chi : \mathbb{F}_{r^t} \to \mathbb{C}$ be a nontrivial additive character such that the complex number

$$z := \mathbb{E}_{c \in \mathbb{F}_p} \chi\big(\gamma^c P(1)\big)$$

satisfies $\Re(z) \leq 0$. To see that such a character exists, observe that by orthogonality of the characters,

$$\mathbb{E}_{\chi \in \widehat{\mathbb{F}_{r^t}}} \mathbb{E}_{c \in \mathbb{F}_p} \chi\big(\gamma^c P(1)\big) = \mathbb{E}_{c \in \mathbb{F}_p} \Big(\mathbb{E}_{\chi \in \widehat{\mathbb{F}_{r^t}}} \chi\big(\gamma^c P(1)\big) \Big) = 0.$$

https://doi.org/10.1017/fms.2021.1 Published online by Cambridge University Press

The existence of the desired character then follows by averaging. For each $a \in \{0,1\}^k$ and $\iota \in \mathbf{i}(P)$, define $F_a^\iota : \mathbb{F}_p^n \to \mathbb{C}$ by

$$F_a^\iota(x) = \chi\Big(c_\iota \sum_{j=1}^{k} a_j f_j(x)\Big). \tag{5}$$

Based on these functions, we define the dual function $\phi_a : \mathbb{F}_p^n \to \mathbb{D}$ by

$$\phi_a(y) = \mathbb{E}_{x\in\mathbb{F}_p^n} \prod_{\iota\in\mathbf{i}(P)} F_a^\iota(x+\iota y). \tag{6}$$

Then,

$$\begin{aligned}\phi_a(d_i) &= \mathbb{E}_{x\in\mathbb{F}_p^n}\chi\Big(\sum_{j=1}^{k} a_j \sum_{\iota\in\mathbf{i}(P)} c_\iota f_j(x+\iota d_i)\Big)\\ &= \mathbb{E}_{x\in\mathbb{F}_p^n}\chi(a_i\gamma^{\langle x,w_i\rangle}P(1))\\ &= \mathbb{E}_{c\in\mathbb{F}_p}\chi(a_i\gamma^c P(1)).\end{aligned}$$

The last expectation equals 1 if $a_i = 0$ and z if $a_i = 1$ and, therefore,

$$\{1,z\}^k \subseteq \{(\phi(d))_{d\in D} \mid \phi \in \Delta_{\mathbf{i}(P)}\}.$$

Because $k \geq (\frac{m}{p})^{p-1}$, $n \leq (\frac{2etm}{p})^{\frac{p-1}{t}}$ and $t \leq p-1$, we have $k \gg_p n^t$. □

5. Sparse polynomials over $\mathbb{F}_2$

The following lemma supplies infinitely many primes and polynomials that can be used in Proposition 4.1.

Lemma 5.1. *For infinitely many primes p, there is an irreducible polynomial $P(x) \in \mathbb{F}_2[x]$ with support size at most $t = \mathrm{ord}_p(2)$ and a root in $\mathbb{F}_{2^t}^*$ of order p.*

To prove Lemma 5.1, we use some basic theory of cyclotomic polynomials (see, for example, [23, Chapter 2]). Let r be a prime and $n \in \mathbb{N}$ not divisible by r. Recall that a primitive nth root of unity over $\mathbb{F}_r$ is a generator of the nonzero elements of the splitting field of the polynomial $x^n - 1$ over $\mathbb{F}_r$. Then, for any such root of unity ζ, the nth cyclotomic polynomial is given by

$$\Phi_n(x) = \prod_{\gcd(s,n)=1} (x - \zeta^s),$$

where the product is over $s \in \{1,\dots,n\}$ such that $\gcd(s,n) = 1$. The following lemma gives the properties of cyclotomic polynomials we need.

Lemma 5.2. *Let r be a prime, $n \in \mathbb{N}$ not divisible by r. Then, the coefficients of $\Phi_n(x)$ lie in $\mathbb{F}_r$. Moreover, if n is a prime, then $\Phi_n(x)$ factors into $(n-1)/\mathrm{ord}_n(r)$ distinct monic irreducible polynomials, all of which have degree exactly $\mathrm{ord}_n(r)$.*

For an integer $k \geq 2$, denote by $p(k)$ the largest prime number that divides k. We will use the following result of Stewart [26].

Lemma 5.3 (Stewart). *For all n large enough, we have*

$$p(2^n - 1) > n \exp\left(\frac{\log n}{104 \log\log n}\right).$$

https://doi.org/10.1017/fms.2021.1 Published online by Cambridge University Press

Proof of Lemma 5.1. By Lemma 5.3, for $p = p(2^n - 1)$ and n sufficiently large, we have $\mathrm{ord}_p(2) \le n < (p-1)/2$. Hence, there are infinitely many primes p such that $t := \mathrm{ord}_p(2) \le (p-1)/2$. For such a p, consider the pth cyclotomic polynomial $\Phi_p(x)$ over $\mathbb{F}_2$. By Lemma 5.2, $\Phi_p(x)$ factors into $(p-1)/t$ distinct monic irreducible polynomials over $\mathbb{F}_2$ of degree exactly t. Because over $\mathbb{F}_2$ there is only one polynomial of degree t with support size $t+1$, there must be an irreducible factor with support of size at most t. Let $P(x)$ be such a factor. Then, because $P(x)|\Phi_p(x)$, its roots lie in the set of pth roots of unity in $\mathbb{F}_{2^t}$. □

Remark 5.4. For Mersenne primes $p = 2^t - 1$, there are polynomials over $\mathbb{F}_2$ with support size 3 that meet the conditions of Proposition 4.1. Indeed, because in this case any pth root of unity ζ in $\mathbb{F}_{2^t}$ is a generator of $\mathbb{F}_{2^t}^*$ and because $1 + \zeta \neq 0$, there exists an s such that $P(x) = 1 + x + x^s$ satisfies $P(1) = 1$ and $P(\zeta) = 0$.

6. Proof of Theorem 1.2

Let $p, t, P(x)$ be as in Lemma 5.1, so that P has support size $k \le t$. Let $\mathbf{i} = \mathbf{i}(P)$. Because P is irreducible, $P(1) \neq 0$ and so it satisfies the conditions of Proposition 4.1. Fix $\varepsilon \in (0, \frac{1}{2})$ and $M \in (0, \infty)$. Suppose that Proposition 1.1 held with $\|\tau\|_{L_\infty} \le \varepsilon$, which is to say that

$$\Delta_{\mathbf{i}} \subseteq \mathrm{Conv}_{\mathbb{C}}\left(M \cdot \{\text{polynomial phases of degree} \le k-1\}\right) + \varepsilon \mathbb{D}^{\mathbb{F}_p^n}.$$

Then, because there are at most $p^{O(n^{k-1})}$ polynomial phase functions of degree at most $k-1$ (one for each n-variate polynomial of degree at most $k-1$), this implies that

$$\log_2 \mathcal{N}(\Delta_{\mathbf{i}}, \varepsilon, M) \ll_p n^{k-1} \ll_p n^{t-1}. \tag{7}$$

At the same time, Proposition 4.1, Lemma 3.1 and Property (3) give

$$\log_2 \mathcal{N}(\Delta_{\mathbf{i}}, \varepsilon, M) \gg_{p,\varepsilon,M} n^t.$$

This contradicts (7) for large n and finishes the proof of Theorem 1.2.

Appendix A. On the possible arithmetic patterns

Here we show that our construction cannot give examples for dual functions corresponding to arithmetic progressions. Let p, r be primes and $t = \mathrm{ord}_p(r)$. Suppose that for some $k, s \in \mathbb{N}$, there is a polynomial $P(x) \in \mathbb{F}_r[x]$ of the form

$$P(x) = \sum_{\iota=0}^{k-1} c_\iota x^{\iota s}$$

such that $P(1) \neq 0$ and $P(x)$ has a root in $\mathbb{F}_{r^t}^*$ of order p. Then, the functions defined as in (5) and (6) belong to the set of dual functions corresponding to the progression $\mathbf{i} = (0, s, 2s, \ldots, (k-1)s)$ and generate in a hypercube of dimension at least n^t. We show that $k \ge t+1$, which means that this does not contradict an L_∞-version of Proposition 1.1.

First note that s cannot be a multiple of p,because for any $\gamma \in \mathbb{F}_{r^t}^*$ of order p we would have $\gamma^s = 1$, which implies that $P(\gamma) = P(1) \neq 0$. It follows that for any such γ, the element γ^s also has order p and does not equal 1. Define the polynomial

$$Q(x) = \sum_{\iota=0}^{k-1} c_\iota x^\iota \in \mathbb{F}_r[x].$$

https://doi.org/10.1017/fms.2021.1 Published online by Cambridge University Press

Then, this polynomial has a root α in $\mathbb{F}_{r^t}^*$ of order p (where $\alpha = \gamma^s$), satisfies $Q(1) = P(1) \neq 0$ and has degree $k-1$. We claim that $k-1 \geq \mathrm{ord}_p(r)$. If Q is reducible, then it has a factor of degree strictly less than $k-1$ that has the same properties. So assume that Q is irreducible. Let $K = \mathbb{F}_r(\alpha)$ be the simple algebraic extension of $\mathbb{F}_r$ obtained by adjoining α. Then K is isomorphic to $\mathbb{F}_{r^{k-1}}$. Because α lies in $\mathbb{F}_{r^{k-1}}$ and has order p, it follows that $p \mid r^{k-1}-1$. But this implies that $k-1 \geq \mathrm{ord}_p(r) = t$.

Acknowledgements. We thank Xuancheng Shao for introducing us to the problem of estimating entropy numbers of dual functions and for showing us Proposition 1.1; Daniel Altman, Nikos Frantzikinakis and Ben Green for helpful discussions; Igor Shparlinski for pointing us to the reference [26] and an anonymous referee for pointing out an error in a previous claim on a strengthening of the main result under the assumption of the generalised Riemann hypothesis. The authors were supported by the Gravitation grant NETWORKS-024.002.003 from the Dutch Research Council (NWO).

Conflict of Interest: None.

References

[1] D. Altman, 'On Szemerédi's theorem with differences from a random set', *Acta Arith.* **195** (2020), 97–108.
[2] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy, 'Proof verification and the hardness of approximation problems', *J. ACM.* **45**(3) (1998), 501–555.
[3] S. Arora and S. Safra, 'Probabilistic checking of proofs: a new characterization of NP', *J. ACM* **45**(1) (1998), 70–122.
[4] M. Blum and S. Kannan, 'Designing programs that check their work', *J. ACM* **42**(1) (1995), 269–291.
[5] J. Briët, 'Subspaces of tensors with high analytic rank', *Online J. Anal. Comb.* 2020, arXiv:1908.04169.
[6] J. Briët, Z. Dvir, and S. Gopi, 'Outlaw distributions and locally decodable codes', *Theory Comput.* **15**(12) (2019), 1–24.
[7] J. Briët and S. Gopi, 'Gaussian width bounds with applications to arithmetic progressions in random settings', *Int. Math. Res. Not.* **2020**(22) November 2020, 8673–8696.
[8] J. Briët and B. Green, 'Multiple correlation sequences not approximable by nilsequences, 2020, preprint arXiv:2010.14960.
[9] J. Briët, A. Naor, and O. Regev, 'Locally decodable codes and the failure of cotype for projective tensor products', *Electron. Res. Announc. Math. Sci.* **19** (2012), 120–130.
[10] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan, 'Private information retrieval', *J. ACM* **45**(6) (1998), 965–982.
[11] K. Efremenko, '3-Query locally decodable codes of subexponential length', *SIAM J. Comput.* **41**(6) (2012), 1694–1703.
[12] N. Frantzikinakis 'Some open problems on multiple ergodic averages', *Bull. Hellenic Math. Soc.* **60** (2016), 41–90.
[13] N. Frantzikinakis, E. Lesigne, and M. Wierdl, 'Random sequences and pointwise convergence of multiple ergodic averages', *Indiana Univ. Math. J.*, **61**(2) (2012), 585–617.
[14] N. Frantzikinakis, E. Lesigne, and M. Wierdl, 'Random differences in Szemerédi's theorem and related results', *J. Anal. Math.* **130** (2016), 91–133.
[15] H. Furstenberg and Y. Katznelson, 'A density version of the Hales-Jewett theorem', *J. Anal. Math.* **57**(1) (1991), 64–119.
[16] GIMPS. Great Internet Mersenne prime search, December 21, 2018. URL: https://www.mersenne.org/.
[17] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin, 'On the locality of codeword symbols', *IEEE Trans. Inform. Theory* **58**(11) (2012), 6925–6934.
[18] S. Gopi, Locality in Coding Theory, PhD Thesis (Princeton University, 2018).
[19] W. T. Gowers, 'Decompositions, approximate structure, transference, and the Hahn–Banach theorem', *Bull. Lon. Math. Soc.* **42**(4) (2010), 573–606.
[20] B. Green, T. Tao, and T. Ziegler, 'An inverse theorem for the Gowers U^{s+1} $[N]$-norm', *Ann. Math.* Second Series, **176**(2) (2012), 1231–1372.
[21] J. Katz and L. Trevisan, 'On the efficiency of local decoding procedures for error-correcting codes', in *Proc. 32nd STOC* (ACM Press, New York, NY, United States 2000), 80–86.
[22] I. Kerenidis and R. de Wolf, 'Exponential lower bound for 2-query locally decodable codes via a quantum argument', *J. Comput. System Sci.* **69**(3) (2004), 395–420.
[23] R. Lidl and H. Niederreiter, *Finite Fields*, Vol. 20 (Cambridge University Press, The Edinburgh Building, Cambridge CB2 8RU, UK 1997, Published in New York, USA).
[24] V. D. Milman and G. Schechtman, *Asymptotic Theory of Finite-Dimensional Normed Spaces*, Vol. 1200 of *Lecture Notes in Mathematics* (Springer-Verlag, Berlin, 1986). With an appendix by M. Gromov.
[25] P. Raghavendra, 'A note on Yekhanin's locally decodable codes', *Electron. Coll. Comput. Complex.* **14**(16) (2007), 1–8.
[26] C. L. Stewart, 'On divisors of Lucas and Lehmer numbers', *Acta Math.* **211**(2) (2013), 291–314.
[27] T. Tao and T. Ziegler, 'The inverse conjecture for the Gowers norm over finite fields via the correspondence principle', *Anal. PDE* **3**(1) (2010), 1–20.
[28] T. Tao and T. Ziegler, 'The inverse conjecture for the Gowers norm over finite fields in low characteristic', *Ann. Comb.* **16**(1) (2012), 121–188.

https://doi.org/10.1017/fms.2021.1 Published online by Cambridge University Press

[29] D. Woodruff, 'New lower bounds for general locally decodable codes', *Electron. Coll. Comput. Complex.* **14**(6) (2007), 1–19.
[30] S. Yekhanin, 'Towards 3-query locally decodable codes of subexponential length', *J. ACM* **55**(1) (2008), 1:1–1:16.
[31] S. Yekhanin, 'Locally decodable codes', *Found. Trends Theor. Comput. Sci.* **6**(3) (2012), 139–255.

https://doi.org/10.1017/fms.2021.1 Published online by Cambridge University Press