Electronic Colloquium on Computational Complexity, Report No. 203 (2023)

Refuting approaches to the log-rank conjecture for XOR functions

Hamed Hatami* Kaave Hosseini† Shachar Lovett‡ Anthony Ostuni§

Abstract

The log-rank conjecture, a longstanding problem in communication complexity, has persistently eluded resolution for decades. Consequently, some recent efforts have focused on potential approaches for establishing the conjecture in the special case of XOR functions, where the communication matrix is lifted from a boolean function, and the rank of the matrix equals the Fourier sparsity of the function, which is the number of its nonzero Fourier coefficients.

In this note, we refute two conjectures. The first has origins in Montanaro and Osborne (arXiv'09) and is considered in Tsang et al. (FOCS'13), and the second one is due to Mande and Sanyal (FSTTCS'20). These conjectures were proposed in order to improve the best-known bound of Lovett (STOC'14) regarding the log-rank conjecture in the special case of XOR functions. Both conjectures speculate that the set of nonzero Fourier coefficients of the boolean function has some strong additive structure. We refute these conjectures by constructing two specific boolean functions tailored to each.

1 Introduction

The study of communication complexity seeks to determine the inherent amount of communication between multiple parties required to complete a computational task. Arguably, the most outstanding conjecture in the field is the *log-rank conjecture* of Lovász and Saks [LS93]. They suggest that the (deterministic) communication complexity of a two-party boolean function is upper bounded by the matrix rank over $\mathbb{R}$. More precisely,

Conjecture 1.1 (Log-rank conjecture [LS93])**.** *Let $f : X \times Y \to \{-1, 1\}$ be an arbitrary two-party boolean function. Then,*

$$\mathrm{CC}(f) \leq \mathrm{polylog}(\mathrm{rank}(f)),$$

where $\mathrm{CC}(f)$ is the communication complexity of f and $\mathrm{rank}(f)$ is the rank over $\mathbb{R}$ of the corresponding boolean matrix.

It is well-known that $\log(\mathrm{rank}(f)) \leq \mathrm{CC}(f)$ [MS82], so a positive resolution to Conjecture 1.1 would imply that the communication complexity of two-party boolean functions is determined by rank, up to polynomial factors.

*School of Computer Science, McGill University. Email: hatami@cs.mcgill.ca

†Department of Computer Science, University of Rochester. Email: kaave.hosseini@rochester.edu

‡Department of Computer Science and Engineering, UC San Diego. Research supported by NSF awards 1953928 and 2006443, and a Simons investigator award. Email: slovett@ucsd.edu

§Department of Computer Science and Engineering, UC San Diego. Research supported by NSF award 2006443. Email: aostuni@ucsd.edu

ISSN 1433-8092

To date, the best known bound is still exponentially far from that in Conjecture 1.1. Concretely, Lovett [Lov16] showed that $\mathrm{CC}(f) \leq O(\sqrt{\mathrm{rank}(f)} \log \mathrm{rank}(f))$. Very recently, Sudakov and Tomon posted a preprint improving the bound to $O(\sqrt{\mathrm{rank}(f)})$ [ST23]. In hopes of improving this result, many researchers have considered the special case of *XOR functions*, where $f_{\oplus}(x, y) = f(x + y)$ for a boolean function $f : \mathbb{F}_2^n \to \{-1, 1\}$ [MO09, ZS10, TWXZ13, STV14, HHL18, MS20]. The XOR setting has several convenient properties. For example, the eigenvalues of $f_{\oplus}$ correspond to the Fourier coefficients of f. Thus, $\mathrm{rank}(f_{\oplus}) = |\mathrm{supp}(\widehat{f})|$, the number of nonzero coefficients in f's Fourier expansion (also known as the *Fourier sparsity*). Additionally, Hatami, Hosseini, and Lovett [HHL18] proved a polynomial equivalence between $\mathrm{CC}(f_{\oplus})$ and the *parity decision tree* complexity of f, denoted $\mathrm{PDT}(f)$. Parity decision trees are defined similarly to standard decision trees, with the extra power that each node can query an arbitrary parity of input bits. These facts together imply that the log-rank conjecture for XOR functions can be restated as follows:

Conjecture 1.2 (XOR log-rank conjecture)**.** *Let* $f : \mathbb{F}_2^n \to \{-1, 1\}$*. Then,*

$$\mathrm{PDT}(f) \leq \mathrm{polylog}(|\mathrm{supp}(\widehat{f})|).$$

The best known bound, due to [TWXZ13, STV14], is $\mathrm{PDT}(f) \leq O(\sqrt{|\mathrm{supp}(\widehat{f})|})$, a mere log-factor improvement on the general case bound by Lovett [Lov16], and matched by the recent bound of Sudakov and Tomon [ST23].

1.1 Folding

Folding is a fundamental concept in the analysis of the additive structure of a function's Fourier support. Let

$$\mathcal{S} = \mathrm{supp}(\widehat{f}) = \{\gamma \in \mathbb{F}_2^n : \widehat{f}(\gamma) \neq 0\} \quad \text{and} \quad \mathcal{S} + \gamma = \{s + \gamma : s \in \mathcal{S}\}.$$

If $(s_1, s_2), (s_3, s_4) \in \binom{\mathcal{S}}{2}$ satisfy $s_1 + s_2 = s_3 + s_4 = \gamma$, we say the pairs (s_1, s_2) and (s_3, s_4) *fold* in the direction γ.

Analyzing folding directions is useful in constructing efficient PDTs in the context of Conjecture 1.2. In particular, when a function f is restricted according to the result of some parity query γ, all pairs of elements in $\mathcal{S}$ that fold in the direction γ collapse to a single term in the restricted function $f|_{\gamma}$'s Fourier support. Iterating this process until the restricted function is constant yields a PDT whose depth depends on the number of iterations performed and, thus, on the size of the folding directions queried. Indeed, this is the general strategy used to prove the aforementioned closest result to Conjecture 1.2 [TWXZ13, STV14].

1.2 Refuting a greedy approach

An approach dating back to [MO09] seeks to prove Conjecture 1.2 through the existence of a single large folding direction. They conjectured that there always exists γ_1, γ_2 such that $|(\mathcal{S} + \gamma_1) \cap (\mathcal{S} + \gamma_2)| \geq |\mathcal{S}|/K$ for some constant $K > 1$. This yields the following $O(\log |\mathcal{S}|)$-rounds greedy approach: query $\gamma_1 + \gamma_2$ and consider the function restricted to the query response. This restriction decreases the Fourier sparsity by a constant factor, so the function must become constant in $O(\log |\mathcal{S}|)$ rounds. This implies the strong upper bound of

$$\mathrm{PDT}(f) \leq O\left(\log |\mathcal{S}|\right).$$

However, O'Donnell, Wright, Zhao, Sun, and Tan [OWZ+14] constructed a function with communication complexity $\Omega(\log(|\mathcal{S}|)^{\log_3(6)})$; hence one can not take K to be a constant. However, to prove the log-rank conjecture, taking $K = O(\operatorname{poly}\log(|\mathcal{S}|))$ suffices, and this choice of K remained plausible up to date. Such an approach is mentioned in both [TWXZ13] and [MS20], and a similar approach was used to verify the log-rank conjecture for many cases of functions lifted with AND (rather than XOR) gadgets [KLMY21]. We strongly refute this conjecture.

Theorem 1.3. *For infinitely many n, there is a function $f : \mathbb{F}_2^n \to \{-1, 1\}$ such that for $\mathcal{S} = \operatorname{supp}(\widehat{f})$, it holds*

$$|(\mathcal{S} + \gamma_1) \cap (\mathcal{S} + \gamma_2)| \leq O\left(|\mathcal{S}|^{5/6}\right)$$

for all distinct $\gamma_1, \gamma_2 \in \mathbb{F}_2^n$.

Remark 1.4. Observe that this theorem implies the greedy method cannot obtain a bound better than $\mathrm{PDT}(f) = \widetilde{O}(|\mathcal{S}|^{1/6})$. In fact, a more careful analysis can rule out bounds better than $\mathrm{PDT}(f) = \widetilde{O}(|\mathcal{S}|^{1/5})$ (see Remark 2.7).

The functions used in Theorem 1.3 are a variant of the addressing function using disjoint (affine) subspaces. While we believe the specific construction is novel, the concept of using functions defined with disjoint subspaces has previously appeared in the literature in this context. Most notably, Chattopadhyay, Garg, and Sherif used XOR functions based on this idea in the pursuit of stronger counterexamples to a more general version of the log-rank conjecture [CGS21].

1.3 Refuting a randomized approach

Rather than simply looking for a large folding direction, a recent work of Mande and Sanyal [MS20] attempts to address Conjecture 1.2 through a deeper understanding of the additive structure of the spectrum of boolean functions. They showed the following conjecture on the number of nontrivial folding directions yields a polynomial improvement to the state-of-the-art upper bound for the XOR log-rank conjecture via a randomized approach.

Conjecture 1.5 ([MS20])**.** *There are constants $\alpha, \beta \in (0, 1)$ such that for every boolean function $f : \mathbb{F}_2^n \to \{-1, 1\}$, for $\mathcal{S} = \operatorname{supp}(\widehat{f})$, it holds*

$$\Pr_{\gamma_1,\gamma_2 \in \mathcal{S}}\left[|(\mathcal{S} + \gamma_1) \cap (\mathcal{S} + \gamma_2)| > |\mathcal{S}|^{\beta}\right] \geq \alpha.$$

In fact, Mande and Sanyal conjectured that one can take $\beta = \frac{1}{2} - o(1)$. The conjecture might seem plausible given the numerous results on the additive structure of the spectrum of boolean functions. However, we strongly refute it, as well:

Theorem 1.6. *For any constant $c > 0$ and infinitely many n, there is a function $f : \mathbb{F}_2^n \to \{-1, 1\}$ such that for $\mathcal{S} = \operatorname{supp}(\widehat{f})$, it holds*

$$\Pr_{\gamma_1,\gamma_2 \in \mathcal{S}}\left[|(\mathcal{S} + \gamma_1) \cap (\mathcal{S} + \gamma_2)| > c\right] = o(1).$$

Overview We prove more precise versions of Theorem 1.3 in Section 2 and Theorem 1.6 in Section 3.

2 One excellent folding direction

A large folding direction implies the existence of a parity query whose answer substantially simplifies the resulting restricted function. This suggests the following greedy approach to resolve the XOR log-rank conjecture: if we can always find distinct γ_1, γ_2 such that $|(\mathcal{S} + \gamma_1) \cap (S + \gamma_2)| \geq \Omega(|\mathcal{S}| \,/\, \mathrm{polylog}(|\mathcal{S}|))$, then querying $\gamma_1 + \gamma_2$ and recursing on the appropriate restriction of f will force f to be constant in $\mathrm{polylog}(|\mathcal{S}|)$ rounds.

We refute this strategy by proving a precise version of Theorem 1.3.

Theorem 2.1. *For $n = 2^k + 7k$ with $k \in \mathbb{N}^{\geq 3}$, there is a function $f : \mathbb{F}_2^n \to \{-1, 1\}$ such that for $\mathcal{S} = \mathrm{supp}(\widehat{f})$, it holds $|\mathcal{S}| \geq 2^{6k}$, and yet $|(\mathcal{S} + \gamma_1) \cap (\mathcal{S} + \gamma_2)| \leq 2^{5k+4}$ for all distinct $\gamma_1, \gamma_2 \in \mathbb{F}_2^n$.*

Our construction has the following form.

Example 2.2 (Subspace addressing)**.** *Let $A_1, \ldots, A_{2^k} \subset \mathbb{F}_2^{7k}$ be disjoint affine subspaces of dimension $2k$. Define $f : \mathbb{F}_2^{7k+2^k} \to \{-1, 1\}$ by*

$$f(x, y) = \begin{cases} (-1)^{y_i} & x \in A_i \\ 1 & x \notin A_1 \cup \cdots \cup A_{2^k} \end{cases},$$

where $x \in \mathbb{F}_2^{7k}$ and $y \in \mathbb{F}_2^{2^k}$.

We choose A_i's randomly and show that the resulting function f has the suitable properties we need with high probability. Let $V^{\perp} = \{w : \langle w, v \rangle = 0 \text{ for all } v \in V\}$ be the orthogonal complement of a subspace V.

Lemma 2.3. *Suppose the random function f is constructed by picking random affine subspaces $A_1, \cdots, A_{2^k} \subset \mathbb{F}_2^{7k}$ as follows: for each $i \in [2^k]$, choose vectors $a_i, v_i^1, \cdots, v_i^{2k} \in \mathbb{F}_2^{7k}$ uniformly and independently, and let $V_i = \langle v_i^1, \cdots, v_i^{2k} \rangle$ and $A_i = V_i + a_i$. Then with probability $1 - 2^{-k+2}$, all of the following hold:*

- *(a)* $\forall i$, $\dim(V_i) = 2k$.
- *(b)* $\forall i \neq j$, $A_i \cap A_j = \emptyset$.
- *(c)* $\forall i \neq j$, $V_i \cap V_j = \{0\}$.
- *(d)* *For all nonzero $v \in \mathbb{F}_2^{7k}$,* $|\{i : v \in V_i^{\perp}\}| \leq 7$.

Proof. For brevity, let $m = 7k$.

(a) Fix $i \in [2^k]$. The probability that vectors $v_i^1, \cdots, v_i^{2k}$ are linearly independent is at least

$$\frac{2^m - 1}{2^m} \cdot \frac{2^m - 2}{2^m} \cdot \frac{2^m - 2^2}{2^m} \cdot \cdots \cdot \frac{2^m - 2^{2k-1}}{2^m} \geq (1 - 2^{2k-m})^m \geq 1 - m2^{2k-m}.$$

Hence the probability that there is $i \in [2^k]$ for which $v_i^1, \cdots, v_i^{2k}$ are not linearly independent is at most $m2^{3k-m} = 7k2^{-4k} \leq 2^{-k}$.

(b) Fix $i \neq j$. The probability that $A_i \cap A_j \neq \emptyset$ is at most $2^{2k}2^{2k-m} = 2^{4k-m}$. Hence, the probability that there are $i \neq j$ with $A_i \cap A_j \neq \emptyset$ is at most $2^{2k}2^{4k-m} = 2^{-k}$.

(c) Fix $i \neq j$. The probability that $V_i \cap V_j \neq \{0\}$ is at most $(2^{2k}-1)2^{2k-m} \leq 2^{4k-m}$. Hence, the probability that there are $i \neq j$ with $V_i \cap V_j \neq \emptyset$ is at most $2^{2k}2^{4k-m} = 2^{-k}$.

(d) The probability that a fixed nonzero vector $v \in \mathbb{F}_2^{7k}$ is orthogonal to at least t subspaces among $V_1, \cdots, V_{2^k}$ is at most $\binom{2^k}{t}2^{-2tk} \leq 2^{-tk}$. Taking $t = 8$ and union bounding over all $2^{7k}-1$ options for v shows that the probability that there is v for which $|\{i : v \in V_i^\perp\}| \geq 8$ is at most 2^{-k}.

By the union bound, the probability that any of items (a) to (d) are not satisfied is at most $4 \cdot 2^{-k} \leq 2^{-k+2}$. □

We will assume from now on that f is chosen randomly so that Lemma 2.3 holds, and set $\mathcal{S} = \mathrm{supp}(\widehat{f})$. It remains to prove there is no large folding direction. First, we give a lower bound on the size of Fourier support of f.

Claim 2.4. $|\mathcal{S}| \geq 2^{6k}$.

Proof. We can express f as

$$\begin{aligned} f(x,y) &= \mathbb{1}_{(A_1\cup\cdots\cup A_{2^k})^c}(x) + \sum_{i=1}^{2^k} \mathbb{1}_{A_i}(x)\cdot(-1)^{y_i} \\ &= 1 - \sum_{i=1}^{2^k} \mathbb{1}_{A_i}(x) + \sum_{i=1}^{2^k} \mathbb{1}_{A_i}(x)\cdot(-1)^{y_i} \\ &= 1 + \sum_{i=1}^{2^k} \mathbb{1}_{A_i}(x)\cdot((-1)^{y_i}-1). \end{aligned}$$

Note that the Fourier support of the function $\mathbb{1}_{A_i}(x)$ is $V_i^\perp \subset \mathbb{F}_2^{7k}$, and of $\mathbb{1}_{A_i}(x)\cdot(-1)^{y_i}$ is $V_i^\perp + e_i$, where e_i is the i-th basis vector in the standard basis for $\mathbb{F}_2^{2^k}$ embedded in the space $\mathbb{F}_2^{7k} \times \mathbb{F}_2^{2^k}$. Since the affine subspaces $V_i^\perp + e_i$ are disjoint and also $\left(V_i^\perp + e_i\right) \cap \left(V_i^\perp\right) = \emptyset$ the coefficients of characters in $V_i^\perp + e_i$ will not be canceled. Hence, we get that

$$\bigcup_{i=1}^{2^k} \left(V_i^\perp + e_i\right) \subset \mathcal{S}$$

and hence $|\mathcal{S}| \geq 2^k \cdot 2^{7k-2k} = 2^{6k}$. □

We also need the following claim.

Claim 2.5. *Suppose $W_1, W_2 \subset \mathbb{F}_2^n$ are two linear subspaces such that $W_1 \cap W_2 = \{0\}$. Then for all $x \in \mathbb{F}_2^n$,*

$$|W_1^\perp \cap (W_2^\perp + x)| = 2^{n-\dim W_1 - \dim W_2}.$$

Proof. Suppose $\dim(W_1) = d_1$ and $\dim(W_2) = d_2$. Without loss of generality, assume that $W_1 = \mathbb{F}_2^{d_1} \times 0^{d_2} \times 0^{n-d_1-d_2}$ and $W_2 = 0^{d_1} \times \mathbb{F}_2^{d_2} \times 0^{n-d_1-d_2}$. Note that $W_1^\perp = 0^{d_1} \times \mathbb{F}_2^{d_2} \times \mathbb{F}_2^{n-d_1-d_2}$ and $W_2^\perp = \mathbb{F}_2^{d_1} \times 0^{d_2} \times \mathbb{F}_2^{n-d_1-d_2}$. Pick an arbitrary $x = (x_1, x_2, x_3) \in \mathbb{F}_2^{d_1} \times \mathbb{F}_2^{d_2} \times \mathbb{F}_2^{n-d_1-d_2}$. Then $W_2^\perp + (x_1, x_2, x_3) = \mathbb{F}_2^{d_1} \times \{x_2\} \times \mathbb{F}_2^{n-d_1-d_2}$ and $W_1^\perp \cap (W_2^\perp + x) = 0^{d_1} \times \{x_2\} \times \mathbb{F}_2^{n-d_1-d_2}$ has the claimed size. □

Finally, Theorem 2.1 follows from claim below.

Claim 2.6. *For all distinct* $\gamma_1, \gamma_2 \in \mathbb{F}_2^{7k+2^k}$*, we have*

$$|(\mathcal{S}+\gamma_1) \cap (\mathcal{S}+\gamma_2)| \le 2^{5k+4}.$$

Proof. First, note that it suffices to prove the claim for all distinct $\gamma_1, \gamma_2 \in \mathcal{S}$, since if $s_1+\gamma_1 = s_2+\gamma_2$ for $s_1, s_2 \in \mathcal{S}$, it must be that $\gamma_1+\gamma_2 = s_1+s_2 \in \mathcal{S}+\mathcal{S}$. Pick an arbitrary non-zero $\gamma = \gamma_1+\gamma_2$ for $\gamma_1, \gamma_2 \in \mathcal{S}$. Remember that

$$|(\mathcal{S}+\gamma_1) \cap (\mathcal{S}+\gamma_2)| = |\mathcal{S} \cap (\mathcal{S}+\gamma)| \quad \text{and} \quad \mathcal{S} \subseteq \left(\bigcup_{i=1}^{2^k} V_i^\perp\right) \cup \left(\bigcup_{i=1}^{2^k} (V_i^\perp + e_i)\right).$$

Hence $\mathcal{S} \cap (\mathcal{S}+\gamma) \subseteq A \cup B \cup C$, where

$$\begin{aligned} A &= \bigcup_{i,j} \left(V_i^\perp \cap (V_j^\perp + \gamma)\right) \\ B &= \bigcup_{i,j} \left(V_i^\perp \cap (V_j^\perp + e_j + \gamma)\right) \\ C &= \bigcup_{i,j} \left((V_i^\perp + e_i) \cap (V_j^\perp + e_j + \gamma)\right). \end{aligned}$$

Let $|\cdot|$ denote the Hamming weight of a vector. Decompose $\gamma = (\gamma_x, \gamma_y)$ where $\gamma_x \in \mathbb{F}_2^{7k}$ and $\gamma_y \in \mathbb{F}_2^{2^k}$. Observe that $|\gamma_y| \le 2$, since (as noted above) we may assume $\gamma \in \mathcal{S}+\mathcal{S}$ without loss of generality.

Case 1: $|\gamma_y| = 0$.

Note that in this case $B = \emptyset$ and $C = \bigcup_i \left((V_i^\perp + e_i) \cap (V_i^\perp + e_i + \gamma_x)\right)$. Overall, we get

$$\begin{aligned} |\mathcal{S} \cap (\mathcal{S}+\gamma_x)| &\le \left|\bigcup_{i,j} \left(V_i^\perp \cap (V_j^\perp + \gamma_x)\right)\right| + \left|\bigcup_i \left((V_i^\perp + e_i) \cap (V_i^\perp + e_i + \gamma_x)\right)\right| \\ &= \left|\bigcup_{i,j} \left(V_i^\perp \cap (V_j^\perp + \gamma_x)\right)\right| + \left|\bigcup_i \left(V_i^\perp \cap (V_i^\perp + \gamma_x)\right)\right| \\ &\le \sum_{i \ne j} |V_i^\perp \cap (V_j^\perp + \gamma_x)| + 2\sum_i |V_i^\perp \cap (V_i^\perp + \gamma_x)| \\ &\le \sum_{i \ne j} |V_i^\perp \cap (V_j^\perp + \gamma_x)| + 2 \cdot 2^{5k} \cdot |\{i : \gamma_x \in V_i^\perp\}| \end{aligned}$$

To bound the first term, note that $V_i \cap V_j = \{0\}$ for all $i \ne j$ (by item (c) of Lemma 2.3). Using Claim 2.5 we have that $|V_i^\perp \cap (V_j^\perp + \gamma_x)| = 2^{7k-\dim(V_i)-\dim(V_j)} = 2^{7k-2k-2k} = 2^{3k}$.

To bound the second term, by item (d) of Lemma 2.3, we have that $|\{i : \gamma_x \in V_i^\perp\}| \le 7$. Overall, we get that

$$|\mathcal{S} \cap (\mathcal{S}+\gamma)| \le 2^{2k} \cdot 2^{3k} + 7 \cdot 2^{5k+1} \le 2^{5k+4}.$$

Case 2: $|\gamma_y| = 1$. Suppose $\gamma_y = e_i$ for some i.

In this case, $A = C = \emptyset$ and $B = V_i^{\perp} \cap (V_i^{\perp} + e_i + \gamma_y)$. Hence,

$$|\mathcal{S} \cap (\mathcal{S} + \gamma)| \leq |V_i^{\perp} \cap (V_i^{\perp} + e_i + \gamma_y)| \leq |V_i^{\perp}| = 2^{5k},$$

Case 3: $|\gamma_y| = 2$. This is similar to Case 2. □

Remark 2.7. We have chosen parameters for simplicity of exhibition; however, by choosing the original disjoint affine subspaces from $\mathbb{F}_2^{(6+\varepsilon)k}$ rather than $\mathbb{F}_2^{7k}$, a similar analysis rules out any bounds stronger than $\mathrm{PDT}(f) = \widetilde{O}(|\mathcal{S}|^{1/5})$ resulting from this greedy method.

3 Many good folding directions

Rather than hoping for one large folding direction, [MS20] sought many nontrivial ones. In this section, we refute their conjecture (Conjecture 1.5) with the following quantified version of Theorem 1.6.

Theorem 3.1. *For $n = 2^d - 1$ with $d \in \mathbb{N}$, there is a function $f : \mathbb{F}_2^n \to \{-1, 1\}$ such that for $\mathcal{S} = \mathrm{supp}(\widehat{f})$, it holds*

$$\Pr_{\gamma_1, \gamma_2 \in \mathcal{S}} \left[|(\mathcal{S} + \gamma_1) \cap (\mathcal{S} + \gamma_2)| \geq 2^{k+2} \right] \leq 2^{-k} + 2^{1-d} \qquad \forall k \geq 1.$$

Let T be a full binary decision tree of depth d. There are $n = 2^d - 1$ internal nodes indexed by $[2^d - 1]$, where we query (distinct) x_i at node i. Each of the largest depth internal nodes v is adjacent to two leaves: -1 and 1, corresponding to $v = 0$ and $v = 1$, respectively. Let $f : \mathbb{F}_2^n \to \{-1, 1\}$ be the resulting function. For example, the following decision tree corresponds to f for $n = 7$.

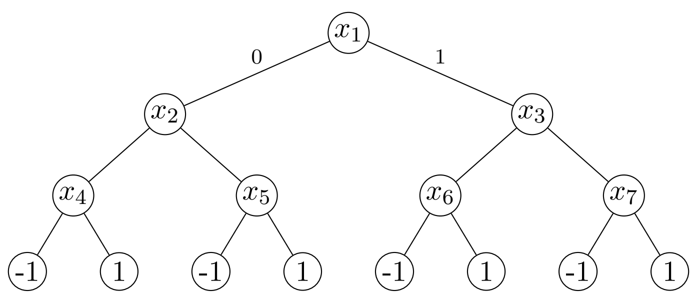

Suppose the leaves are indexed by $[2^d]$. Then f can be written as

$$\sum_{i \in [2^d]} \mathrm{sign}(L_i) \cdot \mathbb{1}_{L_i}, \tag{1}$$

where $\mathbb{1}_{L_i}$ denotes the indicator function of the inputs that result in leaf i, and $\mathrm{sign}(L_i) \in \{-1, 1\}$ is the output at leaf i. Let P_i be the ordered set of coordinates that are queried to reach the leaf i. Then for input $x = (x_1, \ldots, x_n) \in \mathbb{F}_2^n$, we can write

$$\mathbb{1}_{L_i}(x) = \prod_{t \in P_i} \left(\frac{1 + (-1)^{a_t + x_t}}{2} \right) = \frac{1}{2^d} \left(\sum_{P \subseteq P_i} (-1)^{\sum_{j \in P} a_j} \cdot (-1)^{\sum_{j \in P} x_j} \right),$$

where $a_t \in \mathbb{F}_2$ is the output of node t on the path P_i.

To find the Fourier support $\mathcal{S} = \mathrm{supp}(\widehat{f})$, it remains to determine which terms "survive" cancellation in Equation (1). Let $\mathcal{N}(i)$ be the index of the internal node adjacent to leaf i. Observe that when $\mathcal{N}(i) = \mathcal{N}(j)$ for $i \neq j$ (so $\mathrm{sign}(L_i) = -\mathrm{sign}(L_j)$),

$$\begin{aligned} 2^d(\mathrm{sign}(L_i) \cdot \mathbb{1}_{L_i}(x) + \mathrm{sign}(L_j) \cdot \mathbb{1}_{L_j}(x)) &= \mathrm{sign}(L_i) \sum_{P \subseteq P_i} (-1)^{\sum_{t\in P} a_t} \cdot (-1)^{\sum_{t \in P} x_t} \\ &\quad - \mathrm{sign}(L_i) \sum_{P \subseteq P_j} (-1)^{\sum_{t\in P} a_t} \cdot (-1)^{\sum_{t \in P} x_t} \\ &= 2 \cdot \mathrm{sign}(L_i) \cdot \sum_{P \subseteq P_i,\, \mathcal{N}(i) \in P} (-1)^{\sum_{t\in P} a_t} \cdot (-1)^{\sum_{t \in P} x_t}, \end{aligned}$$

since $x_{\mathcal{N}(i)}$ is the only x value that P_i and P_j disagree on. That is, each term in f's expansion must contain $\mathcal{N}(i)$ for some i. Moreover, once these cancellations are made, $\mathbb{1}_{L_i}$ does not interact with $\mathbb{1}_{L_j}$ for $\mathcal{N}(i) \neq \mathcal{N}(j)$, since no term can contain both $\mathcal{N}(i)$ and $\mathcal{N}(j)$. In summary,

$$\mathcal{S} = \bigcup_{i \in [2^d]} \{s : s \subseteq P_i \text{ and } \mathcal{N}(i) \in s\}.$$

Let $\gamma_1, \gamma_2 \in \mathcal{S}$. By our observation on the structure of $\mathcal{S}$, they have the form $\gamma_1 = \alpha_1 \dot\cup \{\mathcal{N}(i)\}, \gamma_2 = \alpha_2 \dot\cup \{\mathcal{N}(j)\}$ for some $i, j \in [2^d]$. We are interested in the number of pairs $(\beta_1, \beta_2) \in \mathcal{S} \times \mathcal{S}$ such that $\gamma_1 + \gamma_2 = \beta_1 + \beta_2$. It will suffice to focus on the setting $\mathcal{N}(i) \neq \mathcal{N}(j)$ since this occurs with overwhelming probability. In this case, the quantity $|(\mathcal{S} + \gamma_1) \cap (\mathcal{S} + \gamma_2)|$ depends only on the depth of the lowest common ancestor of P_i and P_j.

Claim 3.2. *If $|(\mathcal{S} + \gamma_1) \cap (\mathcal{S} + \gamma_2)| \geq 2^{k+2}$, then the lowest common ancestor of P_i and P_j is at depth at least k.*

Proof. We will show the contra-positive. Suppose the lowest common ancestor a of P_i and P_j is at depth $\ell < k$, and suppose $\beta_1, \beta_2 \in \mathcal{S}$ satisfy $\beta_1 + \gamma_1 = \beta_2 + \gamma_2$. Without loss of generality, assume $\mathcal{N}(i) \in \beta_1$ and $\mathcal{N}(j) \in \beta_2$. Then β_1 and β_2 must be a subset of the elements in the paths P_i and P_j, respectively.

First, consider each element $E \in P_i \cap P_j$, which is all those above (and including) a. If $E \in \gamma_1 + \gamma_2$, then $E \in \beta_1 + \beta_2$ only if E is in precisely one of β_1, β_2. Likewise, if $E \notin \gamma_1 + \gamma_2$, then $E \notin \beta_1 + \beta_2$ only if E is in neither or both of β_1, β_2. In either case, we have two options for each E.

Now consider each element $E \in P_i$ below a. By assumption, $E \notin P_j$. Thus, if $E \in \gamma_1 + \gamma_2$, it must be that $E \in \gamma_1$ and $E \notin \gamma_2$. For $\beta_1 + \beta_2$ to contain E, we must likewise have $E \in \beta_1$ and $E \notin \beta_2$. Similarly, if $E \notin \gamma_1 + \gamma_2$, it cannot be in γ_1 or γ_2. Thus, it is not in β_1 or β_2 either. An identical argument for $E \in P_j$ shows that we only have one way to account for elements in the paths P_i or P_j below a.

Doubling to compensate for the cases where $\mathcal{N}(j) \in \beta_1$ and $\mathcal{N}(i) \in \beta_2$, we find the number of options for $(\beta_1, \beta_2) \in \mathcal{S} \times \mathcal{S}$ such that $\beta_1 + \gamma_1 = \beta_2 + \gamma_2$ is at most $2^{\ell+2} < 2^{k+2}$. □

Theorem 3.1 follows quickly from the claim. The probability that P_i and P_j have a common ancestor at depth at least k is at most 2^{-k}, so

$$\begin{aligned}
&\Pr_{\gamma_1,\gamma_2\in\mathcal{S}}\left[|(\mathcal{S}+\gamma_1)\cap(\mathcal{S}+\gamma_2)|\geq 2^{k+2}\right]\\
&\leq \Pr_{\gamma_1,\gamma_2\in\mathcal{S}}\left[|(\mathcal{S}+\gamma_1)\cap(\mathcal{S}+\gamma_2)|\geq 2^{k+2}\;\middle|\;\mathcal{N}(\gamma_1)\neq\mathcal{N}(\gamma_2)\right]+2^{1-d}\\
&\leq 2^{-k}+2^{1-d},
\end{aligned}$$

where we overload notation by letting $\mathcal{N}(\gamma)=\mathcal{N}(i)\in\gamma$.

4 Acknowledgements

This work was done in part while the authors were visiting the Simons Institute for the Theory of Computing. A.O. thanks Daniel M. Kane for a number of helpful discussions. We also thank anonymous reviewers for useful comments on an earlier version of this manuscript.

References

[CGS21] Arkadev Chattopadhyay, Ankit Garg, and Suhail Sherif, *Towards stronger counterexamples to the log-approximate-rank conjecture*, Proceedings of the 41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS), 2021, pp. 175–190.

[HHL18] Hamed Hatami, Kaave Hosseini, and Shachar Lovett, *Structure of protocols for XOR functions*, SIAM Journal on Computing **47** (2018), no. 1, 208–217.

[KLMY21] Alexander Knop, Shachar Lovett, Sam McGuire, and Weiqiang Yuan, *Log-rank and lifting for AND-functions*, Proceedings of the 53rd Annual ACM Symposium on Theory of Computing (STOC), 2021, pp. 197–208.

[Lov16] Shachar Lovett, *Communication is bounded by root of rank*, Journal of the ACM (JACM) **63** (2016), no. 1, 1–9.

[LS93] László Lovász and Michael Saks, *Communication complexity and combinatorial lattice theory*, Journal of Computer and System Sciences **47** (1993), no. 2, 322–349.

[MO09] Ashley Montanaro and Tobias Osborne, *On the communication complexity of XOR functions*, arXiv preprint arXiv:0909.3392 (2009).

[MS82] Kurt Mehlhorn and Erik M Schmidt, *Las Vegas is better than determinism in VLSI and distributed computing*, Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing (STOC), 1982, pp. 330–337.

[MS20] Nikhil S Mande and Swagato Sanyal, *On parity decision trees for Fourier-sparse boolean functions*, 40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS), Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.

[OWZ+14] Ryan O'Donnell, John Wright, Yu Zhao, Xiaorui Sun, and Li-Yang Tan, *A composition theorem for parity kill number*, 2014 IEEE 29th Conference on Computational Complexity (CCC), IEEE, 2014, pp. 144–154.

[ST23] Benny Sudakov and István Tomon, *Matrix discrepancy and the log-rank conjecture*, 2023.

[STV14] Amir Shpilka, Avishay Tal, and Ben Lee Volk, *On the structure of boolean functions with small spectral norm*, Proceedings of the 5th Conference on Innovations in Theoretical Computer Science (ITCS), 2014, pp. 37–48.

[TWXZ13] Hing Yin Tsang, Chung Hoi Wong, Ning Xie, and Shengyu Zhang, *Fourier sparsity, spectral norm, and the log-rank conjecture*, 2013 IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2013, pp. 658–667.

[ZS10] Zhiqiang Zhang and Yaoyun Shi, *On the parity complexity measures of boolean functions*, Theoretical Computer Science **411** (2010), no. 26-28, 2612–2618.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il