Communication is Bounded by Root of Rank

SHACHAR LOVETT, University of California, San Diego

We prove that any total boolean function of rank r can be computed by a deterministic communication protocol of complexity $O(\sqrt{r} \cdot \log(r))$. Equivalently, any graph whose adjacency matrix has rank r has chromatic number at most $2^{O(\sqrt{r} \cdot \log(r))}$. This gives a nearly quadratic improvement in the dependence on the rank over previous results.

Categories and Subject Descriptors: F.1.2 [**Computation by abstract devices**]: Modes of Computation—*Interactive and reactive computation*

General Terms: Theory

Additional Key Words and Phrases: Log-rank conjecture

ACM Reference Format:
Shachar Lovett. 2016. Communication is bounded by root of rank. J. ACM 63, 1, Article 1 (February 2016), 9 pages.
DOI: http://dx.doi.org/10.1145/2724704

1. INTRODUCTION

The *log-rank conjecture* proposed by Lovász and Saks [1988] suggests that, for any Boolean function $f : X \times Y \to \{-1, 1\}$, its deterministic communication complexity $\mathrm{CC}^{\mathrm{det}}(f)$ is polynomially related to the logarithm of the rank of the associated matrix, in which the rank is computed over the reals. Validity of this conjecture is one of the fundamental open problems in communication complexity. Very little progress has been made toward resolving it. The best upper bound, until recently, was

$$\mathrm{CC}^{\mathrm{det}}(f) \leq \log(4/3) \cdot \mathrm{rank}(f),$$

due to Kotlov [1997] (in fact, it only applies to a special case, corresponding to the chromatic number of graphs). In terms of lower bounds, Kushilevitz (unpublished; see Nisan and Wigderson [1994]) gave an example of a family of functions with $\mathrm{CC}^{\mathrm{det}}(f) \geq (\log \mathrm{rank}(f))^{\log_3 6}$. Recently, a conditional improvement was made by Ben-Sasson et al. [2012], who showed that assuming a number-theoretic conjecture (the polynomial Freiman-Ruzsa conjecture), $\mathrm{CC}^{\mathrm{det}}(f) \leq O(\mathrm{rank}(f)/\log \mathrm{rank}(f))$. In this article, we establish the following (unconditional) improved upper bound on the deterministic communication complexity.

THEOREM 1.1. *Let $f : X \times Y \to \{-1, 1\}$ be a Boolean function with rank r. Then there exists a deterministic protocol computing f that uses $O(\sqrt{r} \cdot \log r)$ bits of communication.*

This work is supported by the National Science Foundation CAREER award no. 135048.
Author's address: S. Lovett, Department of Computer Science and Engineering, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093, USA; email: slovett@cs.ucsd.edu.

DOI: http://dx.doi.org/10.1145/2724704

The log-rank conjecture can be equivalently formulated as the relation between the rank of the adjacency matrix of a graph and its chromatic number [Lovász and Saks 1988]. With this formulation, we derive a new bound on the chromatic number of a graph given its rank.

COROLLARY 1.2. *Let G be a graph whose adjacency matrix has rank r. Then, the chromatic number of G is at most* $2^{O(\sqrt{r}\log r)}$.

PROOF. Let $G = (V, E)$. Define a function $f : V \times V \to \{0, 1\}$ as $f(u, v) = 1_{(u,v)\in E}$. In particular, $f(v, v) = 0$ for all $v \in V$. Applying Theorem 1.1 (formally, to the function $(-1)^f$, which has a rank at most $r + 1$), there is a deterministic protocol computing f that uses $C = O(\sqrt{r}\log r)$ bits. For any $v \in V$, let $\pi(v) \in \{0, 1\}^C$ be the transcript of the protocol applied to inputs (v, v). Observe that, if $\pi(u) = \pi(v)$, then we must have $(u, v) \notin E$, as the protocol does not distinguish the inputs (u, v) from (v, v) and, in particular, $f(u, v) = f(v, v) = 0$. Hence, π defines a coloring of G with 2^C colors. □

1.1. Proof Overview

The proof is based on analyzing the discrepancy of Boolean functions. The discrepancy of a Boolean function f is given by

$$\mathrm{disc}(f) = \min_{\mu} \max_{R} \left| \sum_{(x,y)\in R} f(x, y)\mu(x, y) \right|,$$

where μ ranges over all distributions over $X \times Y$ and R ranges over all rectangles, $R = A \times B$ for $A \subset X, B \subset Y$. Discrepancy is a well-studied property in the context of communication complexity lower bounds; see, for example, Lokam [2009] for an excellent survey. It is known that low-rank matrices have a noticeable discrepancy [Linial et al. 2007; Linial and Shraibman 2009]: if f has rank r then

$$\mathrm{disc}(f) \geq \frac{1}{8\sqrt{r}}.$$

Discrepancy can be used to prove upper bounds as well. Linial et al. [2007] showed that functions of discrepancy δ have randomized (or quantum) protocols of complexity $O(1/\delta^2)$. Unfortunately, this does not give any improved bounds, in general, as there is always a trivial protocol using r bits. We show that the combination of high discrepancy and low rank implies an improved bound. The following lemma shows that, if f is a Boolean function with discrepancy δ, then there exists a large rectangle on which f is nearly monochromatic. In the following, we denote by $\mathbb{E}[f|R]$ the average value of f on a rectangle R.

LEMMA 1.3. *Let* $f : X \times Y \to \{-1, 1\}$ *be a function with* $\mathrm{disc}(f) = \delta$. *Then, there exists a rectangle R of size*

$$|R| \geq 2^{-O(\delta^{-1}\cdot\log(1/\varepsilon))}|X \times Y|$$

such that $|\mathbb{E}[f|R]| \geq 1 - \varepsilon$.

In fact, we prove a more general lemma (Lemma 3.1), which holds under general distributions. Now, if f has a low rank, we apply Lemma 1.3 with $\varepsilon = 1/2r$ to deduce the existence of a large rectangle R with $|\mathbb{E}[f|R]| \geq 1 - 1/2r$. Next, we apply the following claim from Gavinsky and Lovett [2013], which shows that low-rank matrices that are nearly monochromatic contain large monochromatic rectangles.

CLAIM 1.4 ([GAVINSKY AND LOVETT 2013]). *Let $f : X \times Y \to \{-1, 1\}$ be a function with* $\mathrm{rank}(f) = r$ *and* $\mathbb{E}[f|R] \geq 1 - 1/2r$. *Then, there exists a subrectangle $R' \subset R$ of size $|R'| \geq |R|/4$ such that f is monochromatic on R'.*

Finally, we apply a theorem of Nisan and Wigderson [1994], who showed that, in order to establish that low-rank matrices have efficient deterministic protocols, it suffices to show that they have large monochromatic rectangles (which is what we just showed).

THEOREM 1.5 ([NISAN AND WIGDERSON 1994]). *Assume that, for any function $f : X \times Y \to \{-1, 1\}$ of $\mathrm{rank}(f) = r$, there exists a monochromatic rectangle of size $|R| \geq 2^{-c(r)}|X \times Y|$. Then, any Boolean function of rank r is computable by a deterministic protocol of complexity $O(\log^2 r + \sum_{i=0}^{\log r} c(r/2^i))$.*

As the proof in Nisan and Wigderson [1994] is shown only for the special case related to the log-rank conjecture, we include a proof sketch of Theorem 1.5 for general function $c(r)$ in Section 4.1. Theorem 1.1 now follows by setting $c(r) = O(\sqrt{r} \cdot \log(r))$.

1.2. Related Works

Lemma 3.1, which allows amplification of discrepancy bounds to obtain near-monochromatic rectangles, can also be derived from the rectangle bound, defined by Klauck [2003] and Jain and Klauck [2010]. Specifically, one first relates the discrepancy bound to a rectangle bound for error close to 1/2, then applies error reduction for the rectangle bound [Klauck 2003]. The proofs are very similar; we refer the interested reader to the original papers for details.

There are two recent works that also made progress on the log-rank conjecture. Tsang et al. [2013] studied the special case of XOR functions, which are functions of the form $f(x, y) = F(x \oplus y)$. For this case, they established results similar to Theorem 1.1. Although the results are similar, the techniques seem to be different. In particular, the main tool used in Tsang et al. [2013] is Fourier analysis, while our results are based on discrepancy. It would be interesting to understand if there are deeper connections between these techniques. In another recent work, Gavinsky and Lovett [2013], we show that in order to prove the log-rank conjecture, it suffices to show that any low-rank matrix has an efficient randomized protocol, a low information-cost protocol, or an efficient zero-communication protocol.

Article Organization. We give preliminary definitions in Section 2. We prove Lemma 1.3 in Section 3. We prove Theorem 1.1 in Section 4. We give a proof sketch of Theorem 1.5 in Section 4.1. We discuss a conjecture related to matrix rigidity in Section 5, and further open problems in Section 6.

2. PRELIMINARIES

For standard definitions in communication complexity, we refer the reader to Kushilevitz and Nisan [1997]. We give here only the basic definitions that we would require.

Let $f : X \times Y \to \{-1, 1\}$ be a total Boolean function, where X and Y are finite sets. If μ is a distribution over $X \times Y$, then we denote by $\mathbb{E}_\mu[f] = \sum_{x,y} \mu(x, y) f(x, y)$ the average of f under μ. A *rectangle* is a set $R = A \times B$ for $A \subset X$, $B \subset Y$. We denote by $\mathbb{E}[f|R]$ the average of f under the uniform distribution over R, and more generally by $\mathbb{E}_\mu[f|R]$ the average of f under the conditional distribution of μ conditioned to be in R. A rectangle is *monochromatic* if $f(x, y) = 1$ for all $x, y \in R$ or $f(x, y) = -1$ for all $x, y \in R$.

The *rank* of f is the rank (over the reals) of its associated $X \times Y$ matrix. The *discrepancy* of f with respect to a distribution μ on $X \times Y$ is the maximal bias achieved by a rectangle,

$$\mathrm{disc}_\mu(f) \stackrel{\mathrm{def}}{=} \max_{\text{rectangle } R} \left| \sum_{(x,y)\in R} \mu(x,y) f(x,y) \right|.$$

The discrepancy of f is the minimal discrepancy possible over all possible distributions μ,

$$\mathrm{disc}(f) \stackrel{\mathrm{def}}{=} \min_\mu \mathrm{disc}_\mu(f).$$

Note that discrepancy is a hereditary property. That is, if R is a rectangle, then the discrepancy of f restricted to R is at least the original discrepancy of f. Similarly, low rank is a hereditary property, as ranks of submatrices cannot exceed the rank of the original matrix. We will rely on the following theorem, which lower bounds the discrepancy of functions with low rank. The following theorem follows from Corollary 3.1 and Lemma 4.2 in Linial et al. [2007]; see also Theorem 3.1 in Linial and Shraibman [2009].

THEOREM 2.1. *Let $f : X \times Y \to \{-1, 1\}$ be a function with rank r. Then,* $\mathrm{disc}(f) \geq 1/8\sqrt{r}$.

For completeness, we sketch the proof of Theorem 2.1; the reader is referred to the original papers for the details. Let A be the $X \times Y$ Boolean matrix associated with f, with $A_{x,y} = f(x, y)$. First, one shows the assumption that $\mathrm{rank}(A) = r$ implies that the γ_2 norm of A is $O(\sqrt{r})$. That is, there exists a factorization of A as $A_{x,y} = \langle u_x, v_y \rangle$, where u_x, v_y are vectors in Euclidean space that satisfy $1 \leq \|u_x\|_2 \|v_y\|_2 \leq O(\sqrt{r})$ for all $x \in X, y \in Y$ (see Lemma 4.2 in Linial et al. [2007]). Then, the Grothendick inequality implies that the discrepancy of f is at least $\Omega(1/\sqrt{r})$ (see Theorem 3.1 in Linial and Shraibman [2009]).

3. AN AMPLIFICATION LEMMA

Our main technical lemma is the following lemma, which shows that any Boolean function with high discrepancy contains a large rectangle that is nearly monochromatic.

LEMMA 3.1. *Let $f : X \times Y \to \{-1, 1\}$ be a function with* $\mathrm{disc}(f) = \delta$. *Then, for any $\varepsilon > 0$ and any distribution μ over $X \times Y$, there exists a rectangle R with*

$$\mu(R) \geq 2^{-O(\delta^{-1}\cdot\log(1/\varepsilon))}$$

such that $|\mathbb{E}_\mu[f|R]| \geq 1 - \varepsilon$.

We note that Lemma 1.3 from the introduction is a special case of Lemma 3.1, for which μ is chosen to be the uniform distribution. Our original proof of Lemma 3.1 used an iterative amplification step. After giving a talk on this result in the Banff complexity workshop, Salil Vadhan suggested to us a simplified proof, which avoids the iterative step by applying Yao's mini-max principle. We present his proof below.

PROOF. Let us assume without loss of generality that $\mathbb{E}_\mu[f] \geq 0$; otherwise, apply the lemma to $-f$. Let σ be any distribution over $X \times Y$ such that $\mathbb{E}_\sigma[f] = 0$. By assumption, there exists a rectangle R_1 such that

$$\left| \sum_{(x,y)\in R_1} \sigma(x,y) f(x,y) \right| \geq \delta.$$

Let $R_1 = A \times B$ and define $A' = X \setminus A$, $B' = Y \setminus B$. Consider the four rectangles

$$R_1 = A \times B, R_2 = A' \times B, R_3 = A \times B', R_4 = A' \times B'.$$

As $\sum_{(x,y)\in X\times Y} \sigma(x, y) f(x, y) = \mathbb{E}_\sigma[f] = 0$, there must exist a rectangle $R \in \{R_1, R_2, R_3, R_4\}$ such that

$$\sum_{(x,y)\in R} \sigma(x, y) f(x, y) \geq \delta/3.$$

As this holds for any distribution σ for which $\mathbb{E}_\sigma[f] = 0$, we can apply Yao's mini-max principle and deduce the following. There exists a distribution ρ over rectangles, such that, for any distribution σ over $X \times Y$ for which $\mathbb{E}_\sigma[f] = 0$, we have that

$$\mathbb{E}_{R\sim\rho}\left[\sum_{(x,y)\in R} \sigma(x, y) f(x, y)\right] \geq \delta/3.$$

Equivalently,

$$\sum_{x\in X, y\in Y} \Pr_{R\sim\rho}[(x, y) \in R] \cdot \sigma(x, y) f(x, y) \geq \delta/3.$$

Fix $(x_1, y_1) \in f^{-1}(1)$ and $(x_2, y_2) \in f^{-1}(-1)$. Let σ be the distribution given by $\sigma(x_1, y_1) = \sigma(x_2, y_2) = 1/2$. As $\mathbb{E}_\sigma[f] = 0$, we have that

$$\Pr_{R\sim\rho}[(x_1, y_1) \in R] - \Pr_{R\sim\rho}[(x_2, y_2) \in R] \geq (2/3)\delta.$$

Let p be the *minimal* probability that $(x_1, y_1) \in R$ over all $(x_1, y_1) \in f^{-1}(1)$, where R is sampled according to ρ; and let q be the *maximal* probability that $(x_2, y_2) \in R$ over all $(x_2, y_2) \in f^{-1}(-1)$. We established that

$$p - q \geq (2/3)\delta.$$

Fix $t \geq 1$ and let $R_1, \ldots, R_t \sim \rho$ be chosen independently, and let $R^* = R_1 \cap \ldots \cap R_t$ be their intersection. We will show that, for an appropriate choice of t, the rectangle R^* satisfies the requirements of the lemma with positive probability (and hence such a rectangle exists). We will use the fact that, for any $x \in X$, $y \in Y$,

$$\Pr[(x, y) \in R^*] = \Pr_{R\sim\rho}[(x, y) \in R]^t.$$

Consider the random variable

$$T = \mu(R^*) - (2/\varepsilon) \cdot \mu(R^* \cap f^{-1}(-1)).$$

By linearity of expectation, we have that

$$\begin{aligned}\mathbb{E}[T] &= \sum_{(x,y)\in f^{-1}(1)} \mu(x, y) \Pr[(x, y) \in R^*] - \sum_{(x,y)\in f^{-1}(-1)} \mu(x, y)((2/\varepsilon) - 1) \Pr[(x, y) \in R^*] \\ &\geq \mu(f^{-1}(1)) \cdot p^t - \mu(f^{-1}(-1)) \cdot q^t \cdot (2/\varepsilon) \\ &\geq 1/2 \cdot (p^t - q^t(2/\varepsilon)),\end{aligned}$$

where we used our initial assumption that $\mathbb{E}_\mu[f] = \mu(f^{-1}(1)) - \mu(f^{-1}(-1)) \geq 0$. We choose $t = O(p/\delta \cdot \log(1/\varepsilon))$ so that

$$q^t/p^t \leq (1 - (2/3)\delta/p)^t \leq \varepsilon/4.$$

For this choice of t, we have that

$$\mathbb{E}[T] \geq p^t/4 = 2^{-O(\delta^{-1}\cdot\log(1/\varepsilon))},$$

where we used the inequality $p^p \geq 1/2$, which holds for all $0 < p \leq 1$. Let R^* be a rectangle that achieves this average, that is,

$$\mu(R^*) - (2/\varepsilon)\cdot\mu(R^* \cap f^{-1}(-1)) \geq 2^{-O(\delta^{-1}\cdot\log(1/\varepsilon))}.$$

In particular, we learn that both $\mu(R^*) \geq 2^{-O(\delta^{-1}\cdot\log(1/\varepsilon))}$ (which satisfies the first requirement) and that $\mu(R^* \cap f^{-1}(-1)) \leq (\varepsilon/2)\cdot\mu(R^*)$, which implies that $\mathbb{E}_\mu[f|R^*] \geq 1-\varepsilon$ (which satisfies the second requirement). □

4. DETERMINISTIC PROTOCOLS FOR LOW-RANK FUNCTIONS

We recall Theorem 1.1 for the convenience of the reader.

Theorem 1.1 (restated). *Let $f : X \times Y \to \{-1, 1\}$ be a Boolean function with rank r. Then, there exists a deterministic protocol computing f that uses $O(\sqrt{r}\cdot\log r)$ bits of communication.*

We prove Theorem 1.1 in the remainder of this section. Let $f : X \times Y \to \{-1, 1\}$ be a function of rank r. By Theorem 2.1, we have that $\mathrm{disc}(f) \geq 1/8\sqrt{r}$. We apply Lemma 3.1 with μ the uniform distribution and $\varepsilon = 1/2r$ to derive the existence of a rectangle R such that

$$|R| \geq 2^{-O(\sqrt{r}\cdot\log(r))}\cdot|X\times Y|, \qquad \mathbb{E}[f|R] \geq 1 - 1/2r.$$

Next, we apply a claim from Gavinsky and Lovett [2013] that shows that nearly monochromatic rectangles in low-rank matrices contain large monochromatic matrices.

Claim 4.1 ([Gavinsky and Lovett 2013]). *Let $f : X \times Y \to \{-1, 1\}$ be a function with $\mathrm{rank}(f) = r$ and $\mathbb{E}[f|R] \geq 1 - 1/2r$. Then, there exists a rectangle $R' \subset R$ of size $|R'| \geq |R|/4$ such that f is monochromatic on R'.*

For completeness, we include the proof.

Proof. Let $R = A \times B$. Since f is a sign matrix, the condition $\mathbb{E}[f|R] \geq 1 - 1/2r$ implies that $f(x, y) = -1$ for at most $1/4r$ fraction of the inputs in R. Let $A' \subset A$ be the set of rows for which at most $1/2r$ fraction of the elements are -1,

$$A' = \{x \in A : \left|\{y \in B : f(x, y) = -1\}\right| \leq |B|/2r\}.$$

By Markov inequality, $|A'| \geq |A|/2$. Let $x_1, \ldots, x_r \in A'$ be indices so that their rows span $A' \times B$. Let

$$B' = \{y \in B : f(x_1, y) = \cdots = f(x_r, y) = 1\}.$$

Since each of the rows $x_1, \ldots, x_r$ contain at most $1/2r$ fraction of elements that are -1, we have that $|B'| \geq |B|/2$. Consider the restriction of the matrix to $R' = A' \times B'$. It has rank one; as it is spanned by r rows, each of them is all one. Hence, all of its rows are all one, or all minus one. However, by the construction of A', the rows must be all one. □

Hence, we showed that any function $f : X \times Y \to \{-1, 1\}$ of rank r contains a monochromatic rectangle of size $2^{-O(\sqrt{r}\cdot\log(r))}\cdot|X\times Y|$. Applying Theorem 1.5 with $c(r) = O(\sqrt{r}\cdot\log(r))$, we conclude that any such function can be computed by a deterministic protocol that used $O(\sqrt{r}\cdot\log(r))$ bits of communication.

4.1. Proof Sketch of the Nisan-Wigderson Theorem

We recall Theorem 1.5 of Nisan and Wigderson [1994] for the convenience of the reader.

Theorem 1.5 (restated). *Assume that, for any function $f : X \times Y \to \{-1, 1\}$ of $\mathrm{rank}(f) = r$, there exists a monochromatic rectangle of size $|R| \geq 2^{-c(r)}|X \times Y|$. Then, any Boolean function of rank r is computable by a deterministic protocol of complexity $O(\log^2 r + \sum_{i=0}^{\log r} c(r/2^i))$.*

Proof. Let f be a function of rank r, and consider the partition of its corresponding matrix as

$$\begin{pmatrix} R & S \\ P & Q \end{pmatrix}.$$

As R is monochromatic, $\mathrm{rank}(R) = 1$. Hence, $\mathrm{rank}(S) + \mathrm{rank}(P) \leq r + 1$. (To see this, let A be the matrix formed by replacing R with the all-zeros matrix. Then, $\mathrm{rank}(S) + \mathrm{rank}(P) \leq \mathrm{rank}(A) \leq r + 1$). Assume without loss of generality that $\mathrm{rank}(S) \leq r/2 + 1$ (otherwise, exchange the role of the rows and columns player). The row player sends one bit, indicating whether its input x is in the top or bottom half of the matrix. If it is in the top half, the rank decreases to $\leq r/2 + 1$. If it is in the bottom half, the size of the matrix reduces to at most $(1 - 2^{-c(r)})|X \times Y|$. Iterating this process defines a protocol tree. We next bound the number of leaves of the protocol. By standard techniques, any protocol tree can be balanced so that the communication complexity is logarithmic in the number of leaves see Kushilevitz and Nisan [1997, Chapter 2, Lemma 2.8].

Consider the protocol that stops once the rank drops to $r/2$. Let $m = |X \times Y|$. The protocol tree, in this case, has at most $O(2^{c(r)} \cdot \log(m))$ leaves, hence can be simulated by a protocol sending only $O(c(r) + \log\log(m))$ bits. Note that, since we can assume that f has no repeated rows or columns, $m \leq 2^{2r}$, hence $\log\log(m) \leq \log(r) + 1$. Next, consider the phase in which the protocol continues until the rank drops to $r/4$. Again, this protocol can be simulated by $O(c(r/2) + \log(r))$ bits of communication. Summing over $r/2^i$ for $i = 0, \ldots, \log(r)$ gives the bound. □

5. A CONJECTURE RELATED TO MATRIX RIGIDITY

The proof of Theorem 1.1 relies on the matrix f being Boolean. However, we conjecture that it can be generalized to show that any low-rank sparse matrix contains a large zero rectangle.

Conjecture 5.1. Let M be an $n \times n$ real matrix with $\mathrm{rank}(M) = r$ and such that $M_{i,j} \neq 0$ for at most εn^2 entries, where $\varepsilon \leq 1/2$. Then, there exists $A, B \subset [n]$ such that

$$M_{a,b} = 0 \qquad \forall a \in A, b \in B$$

such that $|A|, |B| \geq n \cdot 2^{-c\sqrt{\varepsilon r}}$, for some absolute constant $c > 0$.

A related conjecture over $\mathbb{F}_2^n$, called the *approximate duality conjecture*, was studied in Ben-Sasson and Zewi [2011] and Ben-Sasson et al. [2012], with relations to two-source extractors and the log-rank conjecture. Here, we show that Conjecture 5.1, if true, would imply stronger bounds for matrix rigidity than currently known.

The bound in Conjecture 5.1, if true, is the best possible, as the following example shows. Let $M = NN^t$, where N is an $n \times r$ matrix whose rows are all the $\{0, 1\}^r$ vectors of hamming weight $\sqrt{r}/10$ and $n = \binom{r}{\sqrt{r}/10} = r^{\Omega(\sqrt{r})}$. The matrix M is $\varepsilon = 1/100$ sparse, as the probability that two uniformly chosen vectors intersect is at most $1/100$. However, one can verify that the largest subsets $A, B \subset [n]$ such that $M_{a,b} = 0$ for all $a \in A, b \in B$ correspond to choosing A to be all vectors whose support lies in the first half of the coordinates, and B to be all vectors whose support lies in the last half of the coordinate. Furthermore, $|A|, |B| \le n \cdot 2^{-\Omega(\sqrt{r})}$. The bound for general $\varepsilon > 0$ can be similarly obtained by considering all vectors in $\{0,1\}^r$ of hamming weight $\sqrt{\varepsilon r}$.

Matrix Rigidity. A matrix M is called (r, s)-rigid if its rank cannot be made smaller than r by changing at most s entries in M. The problem of explicitly constructing rigid matrices was introduced by Valiant [1977] in the context of arithmetic circuits lower bounds, and was also studied by Razborov [1989] in the context of separation of the analogs of PH and PSPACE in communication complexity. Despite much research, the best results to date are achieved by the so-called "untouched minor" argument, which gives explicit matrices that are (r, s)-rigid with $s = \Omega(\frac{n^2}{r}\log(\frac{n}{r}))$. See, for example, the excellent survey of Lokam [2009] for details. We will prove the following corollary of Conjecture 5.1, which improves previous bounds by a logarithmic factor.

COROLLARY 5.2. *Assuming Conjecture 5.1, for any $r \ge 1$ and $r \le n \le r \cdot 2^{c\sqrt{r}}$, there exists an explicit $n \times n$ real matrix that is (r, s)-rigid for $s = \frac{n^2}{2c^2 r}\log^2\left(\frac{n}{r}\right)$.*

PROOF. Let M be an $n \times n$ matrix of rank r, such that all $r \times r$ minors of M have full rank. For example, such a matrix may be constructed as $M = NN^t$, where N is an $n \times r$ matrix such that any r rows of N are linearly independent. Assume that M is not (r, s)-rigid. Then, we can decompose

$$M = L + S, \quad \mathrm{rank}(L) < r, \quad S \text{ is s-sparse.}$$

Let $\varepsilon = \frac{\log^2(n/r)}{2c^2 r}$ so that the matrix S is $s = \varepsilon n^2$ sparse. Moreover, S is low rank, as $\mathrm{rank}(S) \le \mathrm{rank}(M) + \mathrm{rank}(L) < 2r$. Hence, by Conjecture 5.1, there exist $A, B \subset [n]$ of size $|A|, |B| \ge n \cdot 2^{-c\sqrt{2\varepsilon r}} = r$ such that $S_{a,b} = 0$ for all $a \in A, b \in B$. Hence, $M_{a,b} = L_{a,b}$. However, as by construction the rank of M on the minor $A \times B$ is at least r, the same holds for L. Hence, $\mathrm{rank}(L) \ge r$, and we have reached a contradiction. To conclude, note that the upper bound on n follows from requiring that $\varepsilon \le 1/2$. □

6. FURTHER RESEARCH

We proved a bound on the communication complexity that is near linear in the discrepancy. This seem to be tight for our proof technique. The dependence of the discrepancy on the rank, $\mathrm{disc}(f) \ge \Omega(1/\sqrt{\mathrm{rank}(f)})$, is tight, in general, as can be seen, for example, by taking f to be the inner product function. However, it may be that further assuming that the rank of f is much smaller than its size might allow the proof of better bounds. Another interesting direction is to combine our current approach with the additive combinatorics approach of Ben-Sasson et al. [2012]. Finally, we note that it may be possible to generalize the techniques developed here in order to relate the approximate rank of a function and its randomized or quantum communication complexity. However, there seem to be some technical challenges in implementing this.

ACKNOWLEDGMENTS

I thank Dmitry Gavinsky, Pooya Hatami, Russell Impagliazzo, and Adi Shraibman for helpful discussions, and Salil Vadhan for allowing me to present his simplified proof of Lemma 3.1. I thank the anonymous reviewers for their helpful suggestions and corrections.

REFERENCES

Eli Ben-Sasson, Shachar Lovett, and Noga Ron-Zewi. 2012. An additive combinatorics approach relating rank to communication complexity. In *Proceedings of the 53rd Annual Symposium on Foundations of Computer Science,* 177–186.

Eli Ben-Sasson and Noga Zewi. 2011. From affine to two-source extractors via approximate duality. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*. ACM, New York, NY, 177–186.

D. Gavinsky and S. Lovett. 2013. En route to the log-rank conjecture: New reductions and equivalent formulations. *Electronic Colloquium on Computational Complexity (ECCC'13),* 20, 80.

Rahul Jain and Hartmut Klauck. 2010. The partition bound for classical communication complexity and query complexity. In *IEEE 25th Annual Conference on Computational Complexity (CCC'10)*. IEEE, 247–258.

Hartmut Klauck. 2003. Rectangle size bounds and threshold covers in communication complexity. In *Proceedings of the 18th IEEE Annual Conference on Computational Complexity.* IEEE, 118–134.

A. Kotlov. 1997. Rank and chromatic number of a graph. *Journal of Graph Theory* 26, 1, 1–8.

E. Kushilevitz and N. Nisan. 1997. Communication complexity. Cambridge University Press, New York, NY.

N. Linial, S. Mendelson, G. Schechtman, and A. Shraibman. 2007. Complexity measures of sign matrices. *Combinatorica* 27, 4, 439–463.

N. Linial and A. Shraibman. 2009. Learning complexity vs. communication complexity. *Combinatorics, Probability & Computing* 18, 1–2, 227–245.

Satyanarayana V. Lokam. 2009. *Complexity Lower Bounds Using Linear Algebra*. Now Publishers Inc., Delft, The Netherlands.

L. Lovász and M. Saks. 1988. Lattices, Möbius functions and communication complexity. *Annual Symposium on Foundations of Computer Science,* 81–90.

N. Nisan and A. Wigderson. 1994. On rank vs. communication complexity. *Proceedings of the 35rd Annual Symposium on Foundations of Computer Science,* 831–836.

Alexander Razborov. 1989. On rigid matrices (in Russian). *Technical Report, Steklov Mathematical Institute.*

Hing Yin Tsang, Chung Hoi Wong, Ning Xie, and Shengyu Zhang. 2013. Fourier sparsity, spectral norm, and the log-rank conjecture. In *IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS'13)*. IEEE, 658–667.

Leslie G. Valiant. 1977. Graph-theoretic arguments in low-level complexity. In *Mathematical Foundations of Computer Science 1977*, Jozef Gruska (Ed.). Lecture Notes in Computer Science, Vol. 53. Springer, Berlin, 162–176. DOI:http://dx.doi.org/10.1007/3-540-08353-7_135

Received October 2014; revised January 2015; accepted February 2015