Probabilistic Rank and Matrix Rigidity

Josh Alman*
MIT CSAIL & EECS
jalman@mit.edu

Ryan Williams†
MIT CSAIL & EECS
rrw@mit.edu

ABSTRACT

We consider a notion of probabilistic rank and probabilistic sign-rank of a matrix, which measures the extent to which a matrix can be probabilistically represented by low-rank matrices. We demonstrate several connections with matrix rigidity, communication complexity, and circuit lower bounds, such as:

The Walsh-Hadamard Transform is Not Very Rigid. We give surprising *upper bounds* on the rigidity of a family of matrices whose rigidity has been extensively studied, and was conjectured to be highly rigid. For the $2^n \times 2^n$ Walsh-Hadamard transform H_n (a.k.a. Sylvester matrices, a.k.a. the communication matrix of Inner Product modulo 2), we show how to modify only $2^{\epsilon n}$ entries in each row and make the rank of H_n drop below $2^{n(1-\Omega(\epsilon^2/\log(1/\epsilon)))}$, for all small $\epsilon > 0$, over any field. That is, arithmetic circuit lower bounds on Hadamard matrices such as H_n do not follow from L. Valiant’s matrix rigidity approach. We also show non-trivial rigidity upper bounds for H_n with smaller target rank.

Matrix Rigidity and Threshold Circuit Lower Bounds. We give new consequences of rigid matrices for Boolean circuit complexity. First, we show that explicit $n \times n$ Boolean matrices which maintain rank at least $2^{(\log n)^{1-\delta}}$ after $\frac{n^2}{2^{(\log n)^{\delta/2}}}$ modified entries (over *any* field, for any $\delta > 0$) would yield an explicit function that does not have sub-quadratic-size AC^0 circuits with two layers of arbitrary linear threshold gates. Second, we prove that explicit 0/1 matrices over $\mathbb{R}$ which are modestly more rigid than the best known rigidity lower bounds for *sign-rank* would imply *exponential-gate* lower bounds for the infamously difficult class of depth-two linear threshold circuits with arbitrary weights on both layers ($LTF \circ LTF$). In particular, we show that matrices defined by these seemingly-difficult circuit classes actually have low probabilistic rank and sign-rank, respectively.

An Equivalence Among Communication, Probabilistic Rank, and Rigidity. It has been known since Razborov [1989] that explicit rigidity lower bounds would resolve longstanding lower-bound

problems in communication complexity, but it seemed possible that communication lower bounds could be proved without making progress on matrix rigidity. We show that for every function f which is randomly self-reducible in a natural way (the inner product mod 2 is an example), bounding the communication complexity of f (in a precise technical sense) is *equivalent* to bounding the rigidity of the matrix of f , via an equivalence with probabilistic rank.

CCS CONCEPTS

• **Theory of computation** → **Algebraic complexity theory**; Models of computation; Circuit complexity;

KEYWORDS

Matrix Rigidity, Probabilistic Rank, Walsh-Hadamard Transform, Threshold Circuit

ACM Reference format:

Josh Alman and Ryan Williams. 2017. Probabilistic Rank and Matrix Rigidity. In *Proceedings of 49th Annual ACM SIGACT Symposium on the Theory of Computing, Montreal, Canada, June 2017 (STOC’17)*, 12 pages. DOI: 10.1145/3055399.3055484

1 INTRODUCTION

Let R be a ring. In analogy with the notion of a probabilistic polynomial, we define a *probabilistic matrix* over R to be a distribution of matrices $\mathcal{M} \subset R^{n \times n}$. A probabilistic matrix $\mathcal{M}$ *computes* a matrix $A \in R^{n \times n}$ with error $\epsilon > 0$ if for every entry $(i, j) \in [n]^2$,

$$\Pr_{B \sim \mathcal{M}} [A[i, j] = B[i, j]] \geq 1 - \epsilon.$$

In this way, a probabilistic matrix is a *worst-case randomized representation* of a fixed matrix. A probabilistic matrix $\mathcal{M}$ has rank r if the maximum rank of a $M \sim \mathcal{M}$ is r .

We define the ϵ -*probabilistic rank* of a matrix $M \in R^{n \times n}$ to be the minimum rank of a probabilistic matrix computing M with error ϵ . Such probabilistic matrices are of interest and potentially very useful, because some full rank matrices can be represented by probabilistic matrices of rather low rank. For example, every identity matrix has ϵ -probabilistic rank $O(1/\epsilon)$ over any field, by simulating a protocol for EQUALITY using $\log(1/\epsilon) + O(1)$ communication that computes random inner products (cf. Theorem D.2).

We began studying probabilistic rank in the hopes of better understanding the use of probabilistic polynomials in algorithm design. Recent work has shown how substituting low-degree probabilistic polynomials in place of common subroutines can be very useful for speeding up the best known running times for many core problems [Wil14a, Wil14b, AWY15, AW15, ACW16, LPT⁺17]. However, almost every algorithmic application ends up embedding the low-degree polynomial evaluation problem in a *fast multiplication of*

*Supported by NSF DGE-114747. Work performed while the author was a PhD student in the Computer Science Department at Stanford University.

†Supported in part by a Microsoft Research Faculty Fellowship and NSF CCF-1552651 (CAREER). Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. Work performed while the author was affiliated with the Computer Science Department at Stanford University.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC’17, Montreal, Canada

© 2017 ACM. 978-1-4503-4528-6/17/06...\$15.00

DOI: 10.1145/3055399.3055484

two low-rank (rectangular) matrices. That is, this algorithmic work is really using the fact that various circuits and subroutines from core algorithms have *low probabilistic rank*, and is applying low-rank representations to obtain an algorithmic speedup. Because “low probabilistic rank” is potentially a far broader notion than that of “low-degree probabilistic polynomials”, it makes more sense to study probabilistic rank directly, in the hopes of finding stronger algorithmic applications.

In this paper, we consider complexity-theoretic aspects of probabilistic rank. We demonstrate how probabilistic rank can help understand the age-old problem of matrix rigidity, and some models of communication complexity where knowledge is still sparse.

Matrix Rigidity. A central part of our paper connects the probabilistic rank of a matrix to its rigidity. The *rank- r rigidity* of a matrix $A \in R^{n \times n}$, denoted by $\mathcal{R}_A(r)$, is the minimum Hamming distance from A to an $n \times n$ matrix of rank r over R . That is, $\mathcal{R}_A(r)$ is the number of entries of A that must be modified in order for the rank to drop to r . (Sometimes we’ll want to work over a particular field K ; in that case we’ll speak of “ $\mathcal{R}_A(r)$ over K .”) Matrix rigidity was introduced by Leslie Valiant [Val77] in 1977, as a path towards arithmetic circuit lower bounds for linear transformations. Valiant showed that for a field $\mathbb{F}$, and every linear transformation $T : \mathbb{F}^n \rightarrow \mathbb{F}^n$ computable by a circuit of $O(n)$ addition gates of bounded fan-in (with scalar multiplications on the wires) and $O(\log n)$ depth, $\mathcal{R}_T(O(n/\log \log n)) \leq n^{1+\varepsilon}$, for every fixed $\varepsilon > 0$. Thus to prove a circuit lower bound for T , it suffices to lower bound the rigidity of T for rank $O(n/\log \log n)$. Valiant proved that random 0/1 matrices over a field are highly rigid (whp), and strong rigidity lower bounds are known when one allows exponential (or infinite) precision in the matrix entries [Lok06, KLPS14]. However, no explicit rigid matrices T with (say) $\mathcal{R}_T(O(n/\log \log n)) > n^{1.0001}$ are known¹, despite decades of effort (see the surveys [Cod00, Lok09] and the recent work [GT16]). The best known lower bounds for explicit M yield only $\mathcal{R}_M(r) \geq \Omega(\frac{n^2}{r} \cdot \log(n/r))$ [Fri93, SSS97], and Lokam [Lok00] argues that known methods (“untouched minor arguments”) cannot prove rigidity lower bounds larger than this. Very recently, Goldreich and Tal showed an improved “semi-explicit” rigidity lower bound: for random $n \times n$ Toeplitz matrices M , they proved that $\mathcal{R}_M(r) \geq \Omega(\frac{n^3}{r^2 \log n})$ whp, when rank $r \geq \sqrt{n}$ [GT16] (note that such matrices can be generated with $\tilde{O}(n)$ bits of randomness).

In 1989, Razborov [Raz89] (see also [Wun12]) described a connection between matrix rigidity and communication complexity: Letting f be a function in PH^{cc} (the communication complexity equivalent of the polynomial-time hierarchy), the $2^n \times 2^n$ communication matrix M_f of f has $\mathcal{R}_{M_f}(2^{\log^c(n/\varepsilon)}) \leq \varepsilon \cdot 4^n$, where $\varepsilon > 0$ is arbitrary and $c > 0$ is a constant depending only on f , but not n . (Razborov’s proof uses low-degree polynomials which approximate AC^0 functions.) Thus, explicit rigidity lower bounds in the “low” rank and “high” error setting would imply long-open communication lower bounds.

Among the many attempts to prove arithmetic circuit lower bounds via rigidity, perhaps the most commonly studied explicit

matrix has been the Walsh-Hadamard transform [PS88, Alo90, Gri, Nis, KR98, Cod00, Lok01, LTV03, Mid05, dW06, Ras16]:

Definition 1.1. For vectors $x, y \in \mathbb{R}^d$, let $\langle x, y \rangle$ denote their inner product. Let $v_1, \dots, v_{2^n} \in \{0, 1\}^n$ be the enumeration of all n -bit vectors in lexicographical order. The *Walsh-Hadamard matrix* H_n is the $2^n \times 2^n$ matrix defined by $H_n(v_i, v_j) := (-1)^{\langle v_i, v_j \rangle}$.

It was believed that H_n is rigid because its rows are mutually orthogonal (i.e., H_n is Hadamard), so in several of the above references, only that property was assumed of the matrices. The best rigidity lower bounds known for H_n have the form $\mathcal{R}_{H_n}(r) \geq \Omega(4^n/r)$; for the target rank $r = O(2^n/\log n)$ in Valiant’s problem, the lower bound is only $\Omega(2^n \log n)$. It was a folklore theorem that one can modify only $O(n)$ entries of an $n \times n$ Hadamard matrix and make its rank at most $n/2$ [Lok14], but it was believed that for lower rank many more entries would require modification.

Hadamard Ain’t So Rigid. We give a good excuse for the weakness of these lower bounds:

THEOREM 1.2 (NON-RIGIDITY OF HADAMARD MATRICES). *For every field K , for every sufficiently small $\varepsilon > 0$, and for all n , we have $\mathcal{R}_{H_n}(2^{n-f(\varepsilon)n}) \leq 2^{n(1+\varepsilon)}$ over K , for a function f where $f(\varepsilon) = \Theta(\varepsilon^2/\log(1/\varepsilon))$.*

In fact, we show a strong non-rigidity upper bound: by modifying at most $2^{\varepsilon n}$ entries in each row of H_n , the rank of H_n drops to $2^{n-f(\varepsilon)n}$. That is, the usual matrix rigidity approach to arithmetic circuit lower bounds *does not* apply to Hadamard matrices such as the Walsh-Hadamard transform. We would have required lower bounds of the form $\mathcal{R}_{H_n}(2^n/(\log n)) \geq 2^{n(1+\varepsilon)}$ for some $\varepsilon > 0$ to obtain circuit lower bounds; the upper bound of Theorem 1.2 shows this is impossible. The proof is in Section 3. (But please see the conclusion of this paper for a potential loophole, pointed out to us by Pavel Pudlák. Briefly, Valiant’s transformation from arithmetic circuits to non-rigid matrices can yield a somewhat specialized decomposition. Thus it is still possible that one might derive arithmetic circuit lower bounds for Walsh-Hadamard by considering this more stringent rigidity notion.)

We do not (yet) believe that the Walsh-Hadamard transform has $O(2^n)$ -size $O(n)$ -depth circuits; a more appropriate conclusion is that rigidity in itself does not adequately capture the lower bound problem in this case. Having said that, Theorem 1.2 does imply new circuit constructions: it follows that there is a depth-two unbounded fan-in arithmetic circuit for the Walsh-Hadamard transform with $2^{n+O(\varepsilon \log(1/\varepsilon))n} + 2^{2n-\Omega(\varepsilon^2 n)}$ gates; setting $\varepsilon > 0$ appropriately, we have a $4^{\delta n}$ -size circuit for some $\delta < 1$.

We also show non-trivial rigidity upper bounds for H_n in the regime that would be useful for communication complexity, where the rigidity is much closer to 4^n .

THEOREM 1.3 (NON-RIGIDITY OF HADAMARD MATRICES, PART II). *For every integer $r \in [2^{2^n}]$, one can modify at most $2^{2^n}/r$ entries of H_n and obtain a matrix of rank $(n/\ln(r))^{O(\sqrt{n \log(r)})}$.*

See Appendix A for the proof. While the product of rank and rigidity (a natural measure) of H_n is only known to be at least $\Omega(4^n)$, Theorem 1.3 provides an upper bound of $4^n \cdot n^{O(\sqrt{n \log(r)})}/r$, which

¹An infinite family of matrices $\{M_n \mid n \in S\}$ is said to be explicit if there is a polynomial time algorithm A such that $A(1^n)$ prints M_n when $n \in S$.

is not small enough to refute the conjectured rigidity lower bounds required for communication complexity applications. But as we show later, these upper bounds still have non-trivial consequences for the communication complexity of IP2.

New Applications of Explicit Rigid Matrices. Rigidity has been studied primarily for its connections to communication complexity and to lower bounds on arithmetic circuits computing linear transformations. We show new implications of constructing explicit rigid matrices for Boolean circuit complexity.

First, we show how explicit rigidity lower bounds would yield Boolean circuit lower bounds where only somewhat weak results are known:

THEOREM 1.4. *Let K be an arbitrary field, and $\{M_n\}$ be a family of Boolean matrices such that (a) M_n is $n \times n$, (b) there is a $\text{poly}(\log n)$ time algorithm A such that $A(n, i, j)$ prints $M_n(i, j)$, and (c) there is a $\delta > 0$ such that for infinitely many n ,*

$$\mathcal{R}_{M_n} \left(2^{(\log n)^{1-\delta}} \right) \geq \frac{n^2}{2^{(\log n)^{\delta/2}}} \text{ over } K.$$

Then the language $\{(n, i, j) \mid M_n(i, j) = 1\} \in P$ does not have $\text{AC}^0 \circ \text{LTF} \circ \text{AC}^0 \circ \text{LTF}$ circuits of $n^{2-\epsilon}$ -size and $o(\log n / \log \log n)$ -depth, for all $\epsilon > 0$.

The theorem is obtained by giving non-trivial probabilistic rank bounds for such circuits, building on Lokam [Lok01]. Therefore, proving rigidity (or probabilistic rank) lower bounds for explicit 0/1 matrices over a field K would imply nearly-quadratic size lower bounds for $\text{AC}^0 \circ \text{LTF} \circ \text{AC}^0 \circ \text{LTF}$ circuits of unbounded depth, a powerful class of Boolean circuits. (The best known lower bounds are that functions in the huge class E^{NP} do not have such circuits [ACW16].) See Appendix C for these results.

Sign-Rank Rigidity. The sign rank of a $\{-1, 1\}$ matrix M is the lowest rank of a matrix N such that $\text{sign}(M[i, j]) = \text{sign}(N[i, j])$, for all (i, j) . Lower bounds on the sign-rank of matrices were used 15 years ago to prove exponential lower bounds against $\text{LTF} \circ \text{MAJ}$ and $\text{LTF} \circ \text{SYM}$ circuits [For02, FKL⁺01], i.e. restricted versions of depth-two threshold circuits. We extend the sign-rank connection to a circuit class for which strong lower bounds have long been open: explicit matrices with high rigidity under sign-rank would imply strong depth-two threshold circuit lower bounds. (Here, sign-rank rigidity is defined in the natural way, with “rank” replaced with “sign-rank” in the rigidity definition.)

A corollary of a theorem of Razborov and Sherstov [RS10] (see Appendix D) is that for all n , H_n has sign-rank r -rigidity at least $\Omega(4^n/r)$, just as in the case of normal rank rigidity. We show that even a somewhat minor improvement would already imply exponential-size lower bounds for depth-two linear threshold circuits with unbounded weights on both layers, a problem open for decades [HMP⁺93, KW16]:

THEOREM 1.5. *Suppose the sign rank r -rigidity of H_n is $\Omega(4^n/r^{.999})$ for some rank bound $r \geq 2^{\alpha n}$ and some $\alpha > 0$. Then the Inner Product Modulo 2 requires $2^{\Omega(n)}$ -size $\text{LTF} \circ \text{LTF}$ circuits.*

Theorem D.6 gives a more general statement. Under the hood is an upper bound: matrices defined by small $\text{LTF} \circ \text{LTF}$ circuits have low *probabilistic sign-rank*: for every such circuit of s gates,

viewing its truth table as a $2^{n/2} \times 2^{n/2}$ matrix, there is a distribution of $O(s^2 n^2 / \epsilon)$ -rank matrices which sign-represent the truth table in a worst-case probabilistic way with error ϵ .

Rigidity, Communication, and Probabilistic Rank: An Equivalence. Probabilistic rank arises very naturally in studying generalized models of communication complexity. For a Boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, let M_f be the $2^n \times 2^n$ truth table matrix of f with $M_f[x, y] = f(x, y)$ for all x, y . The following correspondence between probabilistic rank and communication complexity is immediate (one could even take the proposition as a definition of $\text{BP} \cdot \text{MOD}_m P$ communication complexity).

PROPOSITION 1.6. *Let $m > 1$ be an integer, let $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, and let M_f be its truth table matrix. The $\text{BP} \cdot \text{MOD}_m P$ communication complexity of f with error ϵ equals the (base-2) logarithm of the ϵ -probabilistic rank of M_f over $\mathbb{Z}_m$ (within additive constants).*

Similarly, AM (Arthur-Merlin communication complexity) is equivalent to probabilistic Boolean rank.

It's easy to see that if a matrix has ϵ -probabilistic rank r , then its rank- r rigidity is at most $\epsilon 2^{2n}$; thus rigidity lower bounds imply communication lower bounds. But conversely, it seems easier to prove lower bounds on probabilistic rank compared to rigidity: with probabilistic rank, we need to rule out a “distribution” of erroneous matrix entries which are required to “spread the errors” around; with rigidity, we have to rule out *any* adversarial choice of bad entries.

We show (in Appendix G) that for every randomly self-reducible function $f : \{0, 1\}^{2n} \rightarrow R$ in which the self-reduction makes k non-adaptive queries, low rigidity implies low probabilistic rank: the ϵ -probabilistic rank of its corresponding matrix is at most $(kr)^k$ if its rank- r rigidity is at most $\epsilon \cdot 4^n$. Thus there is a strong relationship between ϵ -probabilistic rank (and communication complexity, by Proposition 1.6) and the rank for which the rigidity is an ϵ -fraction of the matrix. For the Walsh-Hadamard transform, we prove (in Section 4) that the probabilistic rank of H_n and the rigidity of H_n are equivalent concepts over fields:

THEOREM 1.7. *For every field K and for every n , $\mathcal{R}_{H_n}(r) \leq \epsilon \cdot 4^n$ over K if and only if H_n has ϵ -probabilistic rank r over K .*

The matrices H_n represent the communication matrices of the widely-studied Inner Product Modulo 2 (IP2) function. By Proposition 1.6, the $\text{BP} \cdot \text{MOD}_p P$ communication complexity of IP2 and the rigidity of H_n over $\mathbb{F}_p$ are really equivalent concepts. Applying this theorem, our earlier rigidity upper bounds also imply some modest but interesting improvements on communication complexity protocols. From the rigidity upper bound of Theorem 1.3, we obtain a communication protocol for IP2 with $O(\sqrt{n} \log(1/\epsilon) \log(\frac{n}{\log(1/\epsilon)}))$ bits and error ϵ in the $\text{BP} \cdot \text{MOD}_p P$ communication model, for every prime p . (Aaronson and Wigderson gave an MA protocol for IP with $O(\sqrt{n} \log(n/\epsilon))$ communication complexity and error ϵ [AW09]; ours is more efficient for $\epsilon \ll 1/2^{\sqrt{\log n}}$.) Applying Theorem 1.2 yields an IP2 protocol with $n(1 - \Omega(\epsilon^2/\log(1/\epsilon)))$ communication and only $1/2^{n-\epsilon n}$ error. We are skeptical that our rigidity upper bounds for H_n are tight; we hope these results will aid future

work (to prove rigidity upper bounds, one only has to think about communication protocols for IP2).

1.1 Related Work

Besides the many references already mentioned earlier, there are a few other related works we know of.

Toggle Rank. By Yao's minimax principle [Yao83], $\text{BP} \cdot \text{MOD}_m \text{P}$ communication complexity (randomized communication with "counting modulo m " power) equals worst-case distributional $\text{MOD}_m \text{P}$ communication complexity. In matrix terms, putting an arbitrary distribution $\mathcal{P}$ on the pairs $\{0, 1\}^n \times \{0, 1\}^n$, the worst-case ε -distributional complexity of M is the lowest rank (over $\mathbb{Z}_m$) of a $2^n \times 2^n$ matrix N with error $\|M - N\| \leq \varepsilon$ over $\mathcal{P}$. Wunderlich [Wun12] calls this rank notion the *approximate toggle rank*. Proposition 1.6 shows that probabilistic rank and approximate toggle rank are very closely related, but they are **not** the same as the usual rigidity concept, which corresponds to the uniform distribution on pairs. For structured functions like IP2, we prove (Theorem 1.7) that the uniform distribution is the worst case.

Sign-rank Rigidity and AC0-MOD2 circuits. A tantalizing open problem that has gained popularity in recent years [SV12, ABG⁺14, CGJ⁺16] is whether IP2 has polynomial-size $\text{AC}^0 \circ \text{MOD}_2$ circuits: i.e., circuits of $O(1)$ -depth over AND/OR/NOT, but with a layer of gates computing PARITY at the bottom nearest the inputs. Servedio and Viola [SV12] propose an interesting attack: in our terminology, they note that $\text{AC}^0 \circ \text{MOD}_2$ circuits of size s have $n^{O(\log^{d-1} s) \log(1/\varepsilon)}$ sign-rank rigidity at most $\varepsilon 2^{2n}$ over $\mathbb{R}$, and prove a lower bound on the correlation of signs of sparse polynomials (taken as a proxy for low-rank sign-matrices) with IP2. That is, they prove a weak sign-rank rigidity lower bound (note Razborov and Sherstov prove an analogous lower bound for sign-rank rigidity of IP2; see Appendix D). Our results have two consequences for this sort of approach. First, Theorem D.1 shows that a sign-rank rigidity lower bound would prove something much stronger: a lower bound for *depth-two threshold circuits* computing IP2, a longstanding open problem. Second, our non-trivial upper bounds on the rank rigidity of the IP2 matrix (which is H_n) suggest that IP2 may have much lower sign-rank rigidity than expected.

Sign-Rank Rigidity and Margin Complexity. Linial and Shraibman [LS09] prove (in our terminology) that the sign-rank rigidity of an $n \times n$ matrix A is at most εn^2 for target rank $O(mc(A)^2 \log(1/\varepsilon))$, where $mc(A)$ is the "margin complexity" of A . Thus the margin complexity of a matrix can be used to upper bound sign-rank rigidity. They also study a rigidity notion based on mc , conjecture that high mc implies high margin-complexity rigidity, and show that high margin-complexity also implies communication complexity lower bounds (for similar parameters as the usual rank-rigidity).

Approximate Rank. A different "approximating" rank notion has been studied in [BdW01, KS10, ALSV13], with connections to quantum computing and approximation algorithms. The ε -approximate rank of $M \in \mathbb{R}^{n \times n}$ is the lowest rank of a matrix A such that $\|M - A\|_\infty \leq \varepsilon$. That is, we can obtain one matrix from the other by perturbing each entry by at most $\varepsilon > 0$. The appropriate analogy here seems to be that probabilistic polynomials are to probabilistic rank, as ℓ_∞ -approximate polynomials are to approximate rank:

both are natural generalizations of polynomial representations to matrix representations, with different properties.

2 PRELIMINARIES

We assume basic familiarity with complexity theory. For circuit complexity, we use $C \circ \mathcal{D}$ to denote depth-two circuits where the output gate is of type C and the "hidden" layer is of type $\mathcal{D}$, e.g., $\text{LTF} \circ \text{LTF}$ denotes "depth-two linear threshold circuits", $\text{LTF} \circ \text{MOD}_2$ denotes "linear threshold function of parities", etc. For variables $x_1, \dots, x_n$, we use $\vec{x}$ to denote $(x_1, \dots, x_n)$, and for $\vec{x} \in \{0, 1\}^n$, we use $|\vec{x}| = \sum_i x_i$ to denote its Hamming weight. We use the Iverson bracket $[P] : \{0, 1\}^n \rightarrow \{0, 1\}$ to denote the Boolean function which outputs 1 if and only if property P is true of the n inputs. Below we describe some basic properties relating probabilistic polynomials, probabilistic rank, and rigidity.

Definition 2.1. Let R be any ring, and $f : \{0, 1\}^{2n} \rightarrow R$ be any function on $2n$ Boolean variables. The *truth table matrix* M_f of f is the $2^n \times 2^n$ matrix given by

$$M_f(v_i, v_j) = f(v_i, v_j),$$

where $v_1, \dots, v_{2^n} \in \{0, 1\}^n$ is the enumeration of all n -bit vectors in lexicographical order.

Given the above definition, it is natural to define the probabilistic rank of a function:

Definition 2.2. The ε -probabilistic rank of a Boolean-input function $f : \{0, 1\}^{2n} \rightarrow R$ is the ε -probabilistic rank of its truth table matrix M_f . The rank of f and the rigidity of f are defined similarly.

Definition 2.3 (Razborov [Raz87], Smolensky [Smo87]). Let R be a ring, and let $f : \{0, 1\}^n \rightarrow R$. A *probabilistic polynomial* for f with error ε and degree d is a distribution $\mathcal{P}$ on polynomials $p : \{0, 1\}^n \rightarrow R$ of degree at most d such that for every $x \in \{0, 1\}^n$, $\Pr_{p \sim \mathcal{P}}[p(x) = f(x)] \geq 1 - \varepsilon$. We may similarly refer to a probabilistic polynomial with m monomials.

The following simple mapping from sparse polynomials to low-rank matrices is very useful:

LEMMA 2.4. Let R be any ring, and $f : \{0, 1\}^{2n} \rightarrow R$. Let $p : R^{2n} \rightarrow R$ be a polynomial with m monomials such that $p(x, y) = f(x, y)$ for any $x, y \in \{0, 1\}^n$. Then the rank of f is at most m .

PROOF. Let $a_1, \dots, a_m, b_1, \dots, b_m : R^n \rightarrow R$ be monomials such that $p(x, y) = \sum_{i=1}^m a_i(x) \cdot b_i(y)$ is the monomial expansion of p . For $1 \leq i \leq m$, define vectors $\vec{a}_i, \vec{b}_i \in R^{2^n}$ by $\vec{a}_i[x] = a_i(x)$ and $\vec{b}_i[y] = b_i(y)$ for each $x, y \in \{0, 1\}^n$. Then $M_f = \sum_{i=1}^m \vec{a}_i \otimes \vec{b}_i$, where $\otimes$ denotes the outer product of vectors. Thus $\text{rank}(M_f) \leq m$. $\square$

As a corollary, the probabilistic rank of f is at most the sparsity of a probabilistic polynomial for f :

COROLLARY 2.5. Let R be any ring, and $f : \{0, 1\}^{2n} \rightarrow R$. If f has a probabilistic polynomial $\mathcal{P}$ with at most m monomials and error ε , then the ε -probabilistic rank of f is at most m .

PROOF. Let p be a polynomial in the support of the distribution $\mathcal{P}$. Since p has at most m monomials, by Lemma 2.4 the truth table matrix M_p of p (restricted to the domain $\{0, 1\}^{2n}$) has rank at most

m . The distribution of M_p over p drawn from $\mathcal{P}$ is therefore an ε -probabilistic rank- m distribution for M_f , since $M_f(x, y) = M_p(x, y)$ if and only if $f(x, y) = p(x, y)$. $\square$

By drawing a ‘typical’ matrix from the probabilistic rank distribution, we can always obtain a matrix rigidity upper bound from a sparse probabilistic polynomial.

COROLLARY 2.6. *Let R be any ring, and $f : \{0, 1\}^{2n} \rightarrow R$ be any function on $2n$ Boolean variables. If f has a probabilistic polynomial P with at most m monomials and error ε , then one can modify $\varepsilon 2^{2n}$ entries of the truth table matrix M_f and obtain a matrix of rank at most m .*

3 NON-RIGIDITY OF WALSH-HADAMARD

Now we turn to proving that the Walsh-Hadamard matrices are not rigid.

Reminder of Theorem 1.2 *For every field K , for every sufficiently small $\varepsilon > 0$, and for all n , we have $\mathcal{R}_{H_n}(2^{n-f(\varepsilon)n}) \leq 2^{n(1+\varepsilon)}$ over K , for a function f where $f(\varepsilon) = \Theta(\varepsilon^2/\log(1/\varepsilon))$.*

For a vector $v \in \{0, 1\}^n$, let $|v|$ be the number of ones in v . Let $H : [0, 1] \rightarrow [0, 1]$ denote the binary entropy function

$$H(p) = -p \log_2 p - (1-p) \log_2 (1-p).$$

We will need some estimates of binomial coefficients. For $\varepsilon \in (0, 1/2)$:

$$\binom{n}{\varepsilon n} \leq n \cdot 2^{H(\varepsilon)n}, \text{ and} \quad (1)$$

$$2^{n-O(\varepsilon^2 n)} \leq \binom{n}{(1/2-\varepsilon)n} \leq 2^{n-\Omega(\varepsilon^2 n)}. \quad (2)$$

Equation (1) is standard; equation (2) follows from standard tail bounds on the binomial distribution. In particular, the probability that a uniform random bit string has at most $(1/2 - \varepsilon)n$ ones is at most $2^{-c_1 \varepsilon^2 n}$ and at least $2^{-c_2 \varepsilon^2 n}$, for universal constants $c_1, c_2 > 0$.

Our first (simple) lemma uses a polynomial to compute a large fraction of H_n 's entries with a low-rank matrix. However, this fraction won't be high enough; we'll need another idea to ‘correct’ many entries later.

LEMMA 3.1. *For every field K , and for every $\varepsilon \in (0, 1/2)$, there is a multilinear polynomial $p(x_1, \dots, x_n, y_1, \dots, y_n)$ over K with at most $2^{n-\Omega(\varepsilon^2 n)}$ monomials, such that for all $\vec{x}, \vec{y} \in \{0, 1\}^n$ with $\langle \vec{x}, \vec{y} \rangle \in [2\varepsilon n, (1/2 + \varepsilon)n]$,*

$$p(\vec{x}, \vec{y}) = (-1)^{\langle \vec{x}, \vec{y} \rangle}.$$

The proof uses properties of multivariate polynomial interpolation over the integers. To be concrete, we will apply the following lemma from one of our previous papers:

LEMMA 3.2 ([AW15], LEMMA 3.1). *For any integers n, r, k with $n \geq r + k$ and any integers $c_1, \dots, c_r$, there is a multivariate polynomial $p : \{0, 1\}^n \rightarrow \mathbb{Z}$ of degree $r - 1$ with integer coefficients such that $p(\vec{z}) = c_i$ for all $\vec{z} \in \{0, 1\}^n$ with Hamming weight $|\vec{z}| = k + i$.*

Intuitively, Lemma 3.2 is true because the dimension of the space of degree- $(r - 1)$ polynomials in n variables is large enough that we can always construct a polynomial with the desired constraints.

PROOF OF LEMMA 3.1. By Lemma 3.2 with $k = 2\varepsilon n - 1$, $r = (1/2 - \varepsilon)n + 1$, and $c_i = (-1)^{k+i}$, one can construct a multivariate polynomial $q : \{0, 1\}^n \rightarrow \mathbb{Z}$ with integer coefficients, of degree $(1/2 - \varepsilon)n$, such that for all $\vec{z} \in \{0, 1\}^n$ with $|\vec{z}| \in [2\varepsilon n, (1/2 + \varepsilon)n]$, we have $q(\vec{z}) = (-1)^{|\vec{z}|}$. Since the prime subfield of every field K is either $\mathbb{Q}$ or $\mathbb{F}_m$ for some prime m , and q has integer coefficients, q can be viewed as a polynomial over K (by taking the coefficients modulo m if appropriate). Then our desired polynomial is

$$p(x_1, \dots, x_n, y_1, \dots, y_n) = q(x_1 y_1, x_2 y_2, \dots, x_n y_n).$$

We can upper-bound the number of monomials in p as follows. First, since we only care about the value of p on $\{0, 1\}^{2n}$, we can make p multilinear by applying the equation $v^2 = v$ to all variables. Second, observe that for all $i = 1, \dots, n$, x_i and y_i appear in exactly the same monomials. So if we introduce a variable z_i in place of each $x_i \cdot y_i$ in p , the number of monomials in our new n -variate polynomial p' equals the number of monomials in p .

Since p' is multilinear and degree $(1/2 - \varepsilon)n + 1$, the number of monomials is at most $n \binom{n}{(1/2-\varepsilon)n+1}$, which by (2) is at most $2^{n-c_2 \varepsilon^2 n}$ for some constant $c_2 > 0$. $\square$

Our second lemma says: fixing a vector x with about $1/2$ ones, there is a strong upper bound the number of vectors which has about $1/2$ ones but has small (integer) inner product with x ; we'll use this to upper bound the number of erroneous entries at the very end.

LEMMA 3.3. *For every vector $x \in \{0, 1\}^n$ with $|x| \in [(1/2 - a)n, (1/2 + a)n]$, and any parameters $a, b \in (0, 1/5)$, the probability that a uniformly random vector y from $\{0, 1\}^n$ satisfies both*

- $|y| \in [(1/2 - a)n, (1/2 + a)n]$, and
- $\sum_{k=1}^n x_k y_k \leq bn$

is at most $(2an + 1)(bn + 1) \cdot 2^{f(a,b)-1)n}$, where f is a function such that $f(a, b) \rightarrow 0$ as $a, b \rightarrow 0$.

The usual toolbox of small-deviation estimates does not seem to yield the lemma; we give a direct proof.

PROOF. For all x of the above form, every $k \in [(1/2 - a)n, (1/2 + a)n]$, and every $s \leq bn$, we count the number of $y \in \{0, 1\}^n$ with $|y| = k$ and $\sum_{k=1}^n x_k y_k = s$. A vector y satisfies these properties if and only if:

- there are exactly s integers i with $y[i] = 1$ and $x[i] = 1$, and
- there are exactly $k - s$ integers i with $y[i] = 1$ and $x[i] = 0$.

So there are $\binom{|x|}{s} \binom{n-|x|}{k-s}$ such choices of y . The total probability is hence

$$\begin{aligned} & \frac{1}{2^n} \sum_{k=(1/2-a)n}^{(1/2+a)n} \sum_{s=0}^{bn} \binom{|x|}{s} \binom{n-|x|}{k-s} \\ &= \frac{1}{2^n} \sum_{k=(1/2-a)n}^{(1/2+a)n} \sum_{s=0}^{bn} \binom{|x|}{s} \binom{n-|x|}{k-s} \\ &\leq \frac{1}{2^n} \sum_{k=(1/2-a)n}^{(1/2+a)n} \sum_{s=0}^{bn} \binom{(1/2+a)n}{s} \binom{(1/2+a)n}{k-s} \\ &\leq \frac{1}{2^n} \sum_{k=(1/2-a)n}^{(1/2+a)n} (bn+1) \cdot \binom{(1/2+a)n}{bn} \binom{(1/2+a)n}{(1/2-a-b)n}. \quad (3) \end{aligned}$$

Recall that if $k_1 < k_2 < n/3$ then $\binom{n}{k_1} < \binom{n}{k_2}$, and if $k_3 > k_4 > n/2$, then $\binom{n}{k_3} < \binom{n}{k_4}$. Inequality (3) therefore follows since $s \leq bn < \frac{1}{2}(1/2+a)n$ and $k-s \geq (1/2-a-b)n > \frac{1}{2}(1/2+a)n$ whenever $0 < a, b < 1/5$. Let $g(n) = (2an+1) \cdot (bn+1)$. Simplifying further, the above expression is at most

$$\begin{aligned} & \frac{g(n)}{2^n} \binom{(1/2+a)n}{bn} \binom{(1/2+a)n}{(2a+b)n} \\ &\leq \frac{g(n)}{2^n} \cdot 2^{(1/2+a)n \cdot H(b/(1/2+a))} 2^{(1/2+a)n \cdot H((2a+b)/(1/2+a))} \quad (\text{by (1)}) \\ &\leq \frac{g(n)}{2^n} \cdot 2^{(1/2+a)n \cdot H(2b)} 2^{(1/2+a)n \cdot H(4a+2b)} \\ &\leq \frac{g(n)}{4^n} \cdot 2^{(1/2+a)n \cdot 2 \cdot 2b \cdot \log(1/2b)} 2^{(1/2+a)n \cdot 2 \cdot (4a+2b) \cdot \log(1/(4a+2b))} \\ & \quad (\text{since } H(\epsilon) \leq 2\epsilon \log_2(1/\epsilon) \text{ for } \epsilon < 1/2) \\ &\leq \frac{g(n)}{2^n} \cdot 2^{f(a,b)n}, \end{aligned}$$

where $f(a, b) = (1/2+a)(4b \log(1/2b) + (8a+4b) \log(1/(4a+2b)))$. $\square$

Our third lemma is a simple linear-algebraic observation: given a low-rank matrix M that computes another matrix N on all but a small number of rows and columns, N must also have relatively low rank.

LEMMA 3.4. *Let M' be a matrix of rank r which is equal to M except in at most k columns and ℓ rows. Then the rank of M is at most $r + k + \ell$.*

PROOF. We will start with M' , and add at most $k + \ell$ rank-one matrices to M' so that it equals M .

Consider a column c on which M does not equal M' . We can add to M' a correction matrix C_c given by

$$C_c(i, j) = \begin{cases} M(i, v) - M'(i, v) & \text{if } j = v, \\ 0 & \text{otherwise.} \end{cases}$$

Then, $M' + C_c$ equals M on column c , and is unchanged in any other column. Moreover, since C_c is only nonzero on a single column, it has rank one. So all we have to do is add the correction matrix C_c for each column c on which M and M' differ. The rows of M' can be corrected analogously. $\square$

COROLLARY 3.5. *Let T be any $2^n \times 2^n$ matrix. Let $a \in (0, 1/2)$, and let M be a $2^n \times 2^n$ matrix of rank r , indexed by n -bit vectors. There is a $2^n \times 2^n$ matrix M' of rank at most $r + 4 \cdot n \cdot 2^{n-\Omega(a^2n)}$ such that $M'(v_i, v_j) = T(v_i, v_j)$ on all $v_i, v_j \in \{0, 1\}^n$ where at least one of the following holds:*

- $|v_i| \notin [(1/2-a)n, (1/2+a)n]$,
- $|v_j| \notin [(1/2-a)n, (1/2+a)n]$, or,
- $M(v_i, v_j) = T(v_i, v_j)$.

PROOF. The number of $v_i \in \{0, 1\}^n$ with $|v_i| \notin [(1/2-a)n, (1/2+a)n]$ is at most

$$\sum_{i=0}^{(1/2-a)n} \binom{n}{i} + \sum_{i=(1/2+a)n}^n \binom{n}{i} = 2 \sum_{i=0}^{(1/2-a)n} \binom{n}{i} \leq n \cdot 2^{n-\Omega(a^2n)},$$

by (2). Applying Lemma 3.4 to M and M' with k and ℓ set to $2 \cdot n \cdot 2^{n-\Omega(a^2n)}$, the result follows. $\square$

Let us outline how we'll use all of the above. First, we construct a matrix M of rank about $2^{n-\Omega(\epsilon^2n)}$ approximating H_n , using the polynomial from Lemma 3.1 in a straightforward way. This matrix M has far more erroneous entries than what we desire. But by Lemma 3.3, we can infer that the errors in M are highly concentrated on a relatively small number of rows and columns. Applying Corollary 3.5, the rows and columns can be "corrected" in a way that increases the rank of M by only $2^{n-\Omega(\epsilon^2n)}$. By Lemma 3.3, each row of the matrix left over will have $2^{O(\epsilon \log(1/\epsilon)n)}$ erroneous entries.

PROOF OF THEOREM 1.2. In fact, we prove that one only has to modify $2^{O(\epsilon \log(1/\epsilon)n)}$ entries in each row of H_n , to obtain the desired rank.

Let $\epsilon > 0$ be given. By Lemma 3.1, there is a polynomial $p(x, y)$ in $2n$ variables with $m = 2^{n-\Omega(\epsilon^2n)}$ monomials which computes $(-1)^{\langle x, y \rangle}$ correctly, on all $(x, y) \in \{0, 1\}^{2n}$ such that $\langle x, y \rangle \in [2\epsilon n, (1/2 + \epsilon)n]$.

Construct a $2^n \times 2^n$ matrix M of rank m as in Corollary 2.4, so that $M(x, y) = p(x, y)$. By definition, M equals H_n on all $(x, y) \in \{0, 1\}^{2n}$ satisfying $\langle x, y \rangle \in [2\epsilon n, (1/2 + \epsilon)n]$.

Applying Corollary 3.5 to M with $T = H_n$ and $a = \epsilon$, we obtain a matrix M' of rank $m + 4 \cdot n \cdot 2^{n-\Omega(\epsilon^2n)}$ which is correct on all (x, y) where either $|x| \notin [(1/2-\epsilon)n, (1/2+\epsilon)n]$, $|y| \notin [(1/2-\epsilon)n, (1/2+\epsilon)n]$, or $\langle x, y \rangle \in [2\epsilon n, (1/2 + \epsilon)n]$.

Fix a row of H_n indexed by $x \in \{0, 1\}^n$ with $|x| \in [(1/2 - \epsilon)n, (1/2 + \epsilon)n]$ (note the other rows are already correct). To show that M' differs from H_n on a small number of entries, we need to bound the number of y such that none of the above conditions hold, i.e.,

- (1) $|y| \in [(1/2 - \epsilon)n, (1/2 + \epsilon)n]$ and
- (2) $\langle x, y \rangle \notin [2\epsilon n, (1/2 + \epsilon)n]$.

Note for our given x , it is never true that $\langle x, y \rangle > (1/2 + \epsilon)n$. Therefore we only need to bound the number N of y such that $|y| \in [(1/2-a)n, (1/2+a)n]$ and yet $\langle x, y \rangle < 2\epsilon n$. By Lemma 3.3 with $a = \epsilon$ and $b = \epsilon$, the probability that a random y satisfies $\langle x, y \rangle < 2\epsilon n$ and $|y| \in [(1/2 - \epsilon)n, (1/2 + \epsilon)n]$, is at most $O(n^2) \cdot 2^{f(\epsilon, \epsilon)-1)n}$, where $f \rightarrow 0$ as $\epsilon \rightarrow 0$. Therefore $N \leq 2^n \cdot O(n^2) \cdot 2^{f(\epsilon, \epsilon)-1)n} \leq O(n^2) \cdot 2^{f(\epsilon, \epsilon)n}$.

Now for sufficiently large n and $\varepsilon \in (0, 1/2)$, M' has rank at most $m + 4 \cdot n \cdot 2^{n-\Omega(\varepsilon^2 n)} \leq 5n \cdot 2^{n-\Omega(\varepsilon^2 n)}$. Furthermore, on every row, M' differs from H_n in at most $n^2 \cdot 2^{f(\varepsilon, \varepsilon)} \leq 2^{O(\varepsilon \log(1/\varepsilon)n)}$ entries. $\square$

Other rigidity upper bounds for H_n are described in Appendix A.

4 PROBABILISTIC RANK AND RIGIDITY: AN EQUIVALENCE

In this section, we show that the probabilistic rank of H_n and the rigidity of H_n are the *same* concept over fields. It is easy to see that if ε -probabilistic rank of H_n is k over a field K , then the rank- k rigidity of H_n is at most $\varepsilon 2^{2n}$ over K . Exploiting the random self-reducibility of the H_n function, we can show a converse: lower bounds on probabilistic rank imply proportionate rigidity lower bounds. This is of interest because probabilistic rank lower bounds appear to be fundamentally easier to prove than rigidity lower bounds.

Reminder of Theorem 1.7 *For every field K and for every n , $\mathcal{R}_{H_n}(r) \leq \varepsilon 2^{2n}$ over K if and only if H_n has ε -probabilistic rank r over K .*

First let us give some definitions. Let $\otimes$ denote the outer product of vectors. For vectors $a \in K^{2^n}$ whose entries are indexed by $v_1, \dots, v_{2^n} \in \{0, 1\}^n$, and $x, y \in \{0, 1\}^n$, let $a^{(x, y)}$ denote the vector in K^{2^n} given by

$$a^{(x, y)}[v_i] = (-1)^{\langle v_i, y \rangle} a[v_i \oplus x].$$

This permutes the entries of a , then negates half of the entries.

PROOF. One direction is trivial: low probabilistic rank implies low rigidity, by simply drawing a “typical” matrix from the distribution. For the other direction, suppose $a_1, \dots, a_r$ and $b_1, \dots, b_r$ are vectors in K^{2^n} such that the $2^n \times 2^n$ matrix

$$M := \sum_{k=1}^r a_k \otimes b_k \quad (4)$$

differs from H_n in at most $\varepsilon 2^{2n}$ entries. Pick vectors $x, y \in \{0, 1\}^n$ uniformly at random, and consider the $2^n \times 2^n$ matrix

$$M' = (-1)^{\langle x, y \rangle} \sum_{k=1}^r a_k^{(x, y)} \otimes b_k^{(y, x)}. \quad (5)$$

In this form it is clear that M' has rank at most r . We claim that each entry of M' is equal to the corresponding entry of H_n with probability at least $1 - \varepsilon$, over the choice of x and y , which will complete the proof.

Consider a given entry $M'(v_i, v_j)$. It is sufficient to show that if $M(v_i \oplus x, v_j \oplus y) = H_n(v_i \oplus x, v_j \oplus y)$ then $M'(v_i, v_j) = H_n(v_i, v_j)$, since $(v_i \oplus x, v_j \oplus y)$ is a uniformly random pair of vectors in $\{0, 1\}^n$. Suppose this is the case, meaning $M(v_i \oplus x, v_j \oplus y) =$

$$\begin{aligned} & (-1)^{\langle v_i \oplus x, v_j \oplus y \rangle}. \text{ Applying definition (4) and then (5) we see that} \\ & (-1)^{\langle v_i \oplus x, v_j \oplus y \rangle} \\ &= \sum_{k=1}^r a_k[v_i \oplus x] \cdot b_k[v_j \oplus y] \\ &= (-1)^{\langle v_i, y \rangle + \langle v_j, x \rangle} \sum_{k=1}^r (-1)^{\langle v_i, y \rangle} a_k[v_i \oplus x] \cdot (-1)^{\langle v_j, x \rangle} b_k[v_j \oplus y] \\ &= (-1)^{\langle v_i, y \rangle + \langle v_j, x \rangle} \sum_{k=1}^r a_k^{(x, y)}[v_i] \cdot b_k^{(y, x)}[v_j] \\ &= (-1)^{\langle v_i, y \rangle + \langle v_j, x \rangle} \cdot (-1)^{\langle x, y \rangle} \cdot M'(v_i, v_j). \end{aligned}$$

Rearranging, we see as desired that

$$M'(v_i, v_j) = (-1)^{\langle v_i \oplus x, v_j \oplus y \rangle + \langle v_i, y \rangle + \langle v_j, x \rangle + \langle x, y \rangle} = (-1)^{\langle v_i, v_j \rangle},$$

where the last step follows from the bilinearity of the inner product $\langle \cdot, \cdot \rangle$. $\square$

Therefore, proving communication lower bounds for the IP2 function against (for example) the class $\text{BP} \cdot \text{MOD}_m \text{P}$ is *equivalent* to proving rigidity lower bounds for H_n over the ring $\mathbb{Z}_m$. Applying our rigidity upper bounds for H_n (Theorems 1.2 and 1.3), we obtain surprisingly low probabilistic rank bounds for H_n (and therefore communication-efficient protocols as well):

COROLLARY 4.1. *For every field K , every sufficiently small $\varepsilon > 0$, and all n , H_n has $1/2^{n(1-\varepsilon)}$ -probabilistic rank at most $2^{n-\Omega(\varepsilon^2/\log(1/\varepsilon))n}$ over K , and ε -probabilistic rank at most $(1/\varepsilon)^{O(\sqrt{n} \log n)}$.*

Our reduction from rigidity to probabilistic rank in fact works for any (non-adaptive) random self-reducible function [FF93] that makes a small number of oracle calls. See Appendix G.

5 CONCLUSION

Our most significant finding is that Hadamard matrices are not as rigid as previously believed: for every $\varepsilon > 0$, there are infinitely many N and $N \times N$ Hadamard matrices whose rank drops below $N^{1-\Omega(\varepsilon^2/\log(1/\varepsilon))}$ after modifying only N^ε entries in each row. This rules out a proof of arithmetic circuit lower bounds for the DFT over $\mathbb{Z}_2^n$ via matrix rigidity. Our proof shows precisely how low rank-rigidity can be more powerful than low-sparsity polynomial approximations: we start with a sparse polynomial that has errors concentrated on negligibly many rows and columns, and use a simple lemma to correct most erroneous rows and columns.

Recently, Pavel Pudlák (private communication) observed that Valiant’s transformation from $O(n)$ -size $O(\log n)$ -depth arithmetic circuits to non-rigid matrices can yield a somewhat stronger property: if the linear transformation $T : \mathbb{F}^n \rightarrow \mathbb{F}^n$ is computed by such circuits, we can find $A, B \in \mathbb{F}^{n \times O(n/\log \log n)}$ and $n^{1+\varepsilon}$ -sparse $C \in \mathbb{F}^{n \times n}$ such that $T = (A \cdot B) + C$ but also A can be made $n^{1+\varepsilon}$ -sparse. Curiously, our matrix decomposition based on polynomials does not appear to extend to this more stringent decomposition. Therefore it is still possible that one might prove $O(n)$ -size $O(\log n)$ -depth arithmetic circuit lower bounds for the Walsh-Hadamard transform by considering a stronger form of matrix non-rigidity.

Are there other conjectured-to-be-rigid matrices which are not? One candidate would be the generating matrix of a good linear code

over $\mathbb{F}_2$. Very recently, Goldreich [Dvi16] has reported a distribution of matrices in which most of them are the generating matrix of a good linear code that is not rigid, found by Dvir. It would be very interesting to find an explicit code with this property. Another next natural target would be Vandermonde matrices. Given a field $\mathbb{F}$ of order n , and letting g be a generator of the multiplicative group $\mathbb{F}^\times$, the $n \times n$ matrix $V[i, j] := g^{(i-1) \cdot (j-1)}$ also has structure that may be similarly exploitable.

Our proof that functions with small LTF $\circ$ LTF circuits have low probabilistic sign-rank (Theorem D.1) effectively shows how to randomly reduce an “inner product defined by a LTF $\circ$ LTF” to an “inner product defined by a LTF $\circ$ XOR.” It seems likely that this result could have further applications (beyond what we showed). The theorem suggests the research question: is it possible to write a LTF $\circ$ LTF circuit as an small “approximate-MAJORITY” of LTF $\circ$ XOR circuits, i.e. a probabilistic PTF, in the sense of [ACW16]? This would be an intriguing simulation of depth-two threshold circuits.

Another significant theme in this paper is the close relationship between probabilistic rank and threshold circuits, as well as rigidity. It seems likely that more algorithmic applications will be found by further study of probabilistic rank of matrices; perhaps some lower bounds can also be proved via these connections.

A RIGIDITY UPPER BOUNDS FOR HIGH ERROR

In this section, we prove upper bounds on the rigidity of the Walsh-Hadamard transform in the regime where the error is constant, or much larger than $1/2^n$; this setting is of interest for communication complexity lower bounds.

Reminder of Theorem 1.3 *For every integer $r \in [2^{2n}]$, one can modify at most $2^{2n}/r$ entries of H_n and obtain a matrix of rank at most $(n/\log(r))^{O(\sqrt{n \log(r)})}$.*

The proof follows from applying an optimal-degree probabilistic polynomial for symmetric functions:

THEOREM A.1 ([AW15]). *There is a probabilistic polynomial over any field, or the integers, for any symmetric Boolean function on n variables, with error ϵ and degree $O(\sqrt{n \log(1/\epsilon)})$.*

Proof of Theorem 1.3. Set $\epsilon = 1/r$, and define the Boolean function $IP2 : \{0, 1\}^{2n} \rightarrow \{-1, 1\}$ by $IP2(x, y) = (-1)^{\langle x, y \rangle}$ for all $x, y \in \{0, 1\}^n$. We can see that H_n is the truth table matrix M_{IP2} . By Corollary 2.6, it is sufficient to construct a probabilistic polynomial for $IP2$ with error ϵ and $(n/\ln(1/\epsilon))^{O(\sqrt{n \log(1/\epsilon)})}$ monomials. Consider the Boolean function $PARITY(z_1, \dots, z_n) = (-1)^{z_1 + \dots + z_n}$ for all $z \in \{0, 1\}^n$, and note that $IP2(x_1, \dots, x_n, y_1, \dots, y_n) = PARITY(x_1 y_1, x_2 y_2, \dots, x_n y_n)$. Since $PARITY$ is symmetric, by Theorem A.1 it has a probabilistic polynomial P of error ϵ and degree $d = O(\sqrt{n \log(1/\epsilon)})$. Hence, the distribution of $p(x_1 y_1, \dots, x_n y_n)$ over p drawn from P is a probabilistic polynomial for $IP2$. Since we are only interested in the value of $p(z)$ when $z \in \{0, 1\}^n$, we can make p multilinear by applying the equation $v^2 = v$ to all variables. Then the number of monomials of p is at most $\sum_{i=0}^{O(\sqrt{n \log(1/\epsilon)})} \binom{n}{i} \leq (n/\ln(1/\epsilon))^{O(\sqrt{n \log(1/\epsilon)})}$. Since in $p(x_1 y_1, \dots, x_n y_n)$ we are substituting in a monomial for each variable, its expansion has the same number of monomials as p , as desired. $\square$

B RIGIDITY UPPER BOUND FOR SYM-AND CIRCUITS

Here we generalize Theorems 1.2 and 1.3 to $SYM \circ AND$ circuits. In the proof of Theorem 1.2, the key property of the $IP2$ function required is that has the form

$$IP2(x_1, \dots, x_n, y_1, \dots, y_n) = f(x_1 \wedge y_1, \dots, x_n \wedge y_n),$$

where f is a symmetric Boolean function (in our case, f computes parity). The same proof yields the following generalization:

THEOREM B.1. *For every symmetric function $f : \{0, 1\}^n \rightarrow R$, define the function $IP_f : \{0, 1\}^{2n} \rightarrow R$ by $IP_f(x, y) = f(x_1 \wedge y_1, \dots, x_n \wedge y_n)$. For all sufficiently small ϵ , there is a $\delta < 1$ and a matrix of rank $2^{\delta n}$ which differs from the truth table matrix M_{IP_f} in at most $2^{(1+\epsilon)n}$ entries.*

The proof of Theorem 1.3 only requires a probabilistic polynomial construction in Corollary 2.6. Our probabilistic matrix distribution simply substitutes monomials into the probabilistic polynomial of Theorem A.1 for any symmetric function. Since each monomial can be viewed as an AND, the same argument will work for any $SYM \circ AND$ circuit.

THEOREM B.2. *For any Boolean function $f : \{0, 1\}^{2n} \rightarrow R$ which can be written as a $SYM \circ AND$ circuit with s AND gates, and for every integer $r \in [2^{2n}]$, one can modify $2^{2n}/r$ entries of the truth table matrix M_f and obtain a matrix of rank at most $(s/\log r)^{O(\sqrt{s \log r})}$.*

C EXPLICIT RIGID MATRICES AND THRESHOLD CIRCUITS

In this section, we show how explicit rigidity lower bounds would also imply circuit lower bounds where we currently only know weak results (e.g., we know that some functions in E^{NP} do not have such circuits).

THEOREM C.1. *For every $\delta > 0$ and every $AC^0 \circ LTF \circ AC^0 \circ LTF$ circuit C of size $s = n^{2-\delta}$ and depth $d = o(\log(n)/\log \log(n/\epsilon))$, there is a $\gamma > 0$ such that the truth table of C as a $2^{n/2} \times 2^{n/2}$ matrix M_C has rigidity $\mathcal{R}_{M_C}(2^{n^{1-\gamma} \log(1/\epsilon)}) \leq \epsilon 2^n$, for all $\epsilon \in (1/2^n, 1)$, over any field.*

Our proof will use a technique by Maciel and Therien for converting each middle layer LTF gate into an equivalent $AC^0 \circ MAJ$ circuit:

THEOREM C.2 ([MT98] THEOREM 3.3, [ACW16] THEOREM 7.1). *For every $\alpha > 0$, every LTF on n inputs can be computed by a polynomial-size $AC^0 \circ MAJ$ circuit where the fan-in of each MAJ gate is $n^{1+\alpha}$ and the circuit has depth $O(\log(1/\alpha))$.*

We will also use Tarui’s probabilistic polynomial for AC^0 :

THEOREM C.3 ([Tar93] THEOREM 3.6). *Every circuit in AC^0 with depth d has a probabilistic polynomial over $\mathbb{Z}$ of degree $O(\log^d(n))$ and error $1/2^{\log^{O(1)}(n)}$.*

Proof of Theorem C.1. By Lemma D.4, each LTF gate in the bottom layer has ϵ/s -probabilistic rank $O(n^2 s/\epsilon)$. We will design a probabilistic polynomial for the upper $AC^0 \circ LTF \circ AC^0$ circuitry,

which will give the desired result when composed with this probabilistic rank expression.

First, each LTF gate in the middle layer has fan-in at most $s = n^{2-\delta}$. Applying Theorem C.2 with $\alpha = \delta/2$ to each, the upper $AC^0 \circ LTF \circ AC^0$ circuit becomes a $AC^0 \circ MAJ \circ AC^0$ where each MAJ gate has fan-in at most $n^{(2-\delta)(1+\delta/2)} = n^{2-\delta^2/2}$, and the depth is still $O(d)$.

We can now apply the probabilistic polynomial for AC^0 from Theorem C.3 with degree $O(\log^d(n))$ error $1/2^{\log^{O(1)}(n)}$ to the AC^0 circuits, and the probabilistic polynomial for symmetric functions on $n^{2-\delta^2/2}$ bits from Theorem A.1 with error ε/s and degree $O(n^{1-\delta^2/4} \log(s/\varepsilon))$ to the MAJ gates in the middle. This results in a probabilistic polynomial of degree $O(n^{1-\delta^2/4} \log^{O(d)}(n/\varepsilon))$. For $d = o(\log(n)/\log \log(n/\varepsilon))$, this is $O(n^{1-\beta})$ for any $\beta \in (0, \delta^2/4)$.

We can view the terms in the probabilistic rank expression for the LTF gates in the bottom layer as variables that we substitute into this probabilistic polynomial; the number of monomials in this expansion will upper bound the rank, as in Lemma 2.4. Since there are at most s such gates, and each probabilistic rank expression has $O(n^2 s/\varepsilon)$ terms, we are substituting $O(n^2 s^2/\varepsilon)$ terms into our polynomial. Hence, the number of monomials will be upper bounded by

$$(n^2 s^2/\varepsilon)^{O(n^{1-\beta})} = 2^{O(n^{1-\gamma} \log(1/\varepsilon))},$$

for any $\gamma < \beta$. This is of the desired form, where we can pick any positive value $\gamma < \delta^2/4$. The correctness follows by union bounding over all s probabilistic substitutions we make, each of which has error probability at most ε/s . $\square$

From the above theorem, setting the error ε appropriately, we infer a new consequence of explicit rigid matrices:

Reminder of Theorem 1.4 *Let K be an arbitrary field, and $\{M_n\}$ be a family of Boolean matrices such that (a) M_n is $n \times n$, (b) there is a $\text{poly}(\log n)$ time algorithm A such that $A(n, i, j)$ prints $M_n(i, j)$, and (c) there is a $\delta > 0$ such that for infinitely many n ,*

$$\mathcal{R}_{M_n} \left(2^{(\log n)^{1-\delta}} \right) \geq \frac{n^2}{2^{(\log n)^{\delta/2}}} \text{ over } K.$$

Then the language $\{(n, i, j) \mid M_n(i, j) = 1\} \in P$ does not have $AC^0 \circ LTF \circ AC^0$ of $n^{2-\varepsilon}$ -size and $o(\log n / \log \log n)$ -depth, for all $\varepsilon > 0$.

Therefore, proving strong rigidity lower bounds for explicit matrices over $\mathbb{R}$ has consequences for Boolean circuit complexity as well. Indeed, the desired circuit lower bounds could be derived from lower-bounding probabilistic rank.

D SIGN-RANK RIGIDITY AND DEPTH-TWO THRESHOLD CIRCUITS

Given a matrix $A \in \mathbb{R}^{n \times n}$, its sign rank is the minimum rank of any $B \in \{-1, 1\}^{n \times n}$ such that $\text{sign}(A[i, j]) = \text{sign}(B[i, j])$ for all $i, j \in [n]$. The ε -probabilistic sign-rank of A is defined analogously as with probabilistic rank. We say A has *sign rank r -rigidity t* if a minimum of t entries of A need to be modified in order for A to have sign rank at most r .

First, we observe (in Appendix E) that in the sign-rank setting, random $-1/1$ matrices are still rigid: for example, with high probability, a random $-1/1$ matrix has sign-rank- $(n/\log^2 n)$ rigidity at

least $\Omega(n^2)$. Even though most $-1/1$ matrices have high sign-rank rigidity, we show that the truth table of a small $LTF \circ LTF$ circuit is always close to a matrix of low sign-rank. For even n , we say a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ has ε -probabilistic sign rank r if the truth table of C construed as a $2^{n/2} \times 2^{n/2}$ matrix has ε -probabilistic sign-rank r .

THEOREM D.1. *For every function f with a $LTF \circ LTF$ circuit of size s , and every $\varepsilon > 0$, the ε -probabilistic sign-rank of f is $O(s^2 n^2/\varepsilon)$. Moreover, we can sample a low-rank matrix from the distribution of matrices in $2^{n/2} \cdot \text{poly}(s, n)$ time.*

We will prove this theorem in a few steps. Let $EQ_n : \{0, 1\}^{2n} \rightarrow \{0, 1\}$ be the equality function, i.e., $EQ_n(x, x) = [x = y]$ (using Iverson bracket notation). Similarly, let $LEQ_n : \{0, 1\}^{2n} \rightarrow \{0, 1\}$ be the function $LEQ_n(x, x) = [x \leq y]$ where x and y are interpreted as integers in $\{0, \dots, 2^n - 1\}$.

LEMMA D.2. *For every n , EQ_n has ε -probabilistic rank at most $O(1/\varepsilon)$ over any field.*

PROOF. We mimic a well-known randomized communication protocol for EQ_n . Pick $k = \lceil \log_2(1/\varepsilon) \rceil$ uniformly random subsets $S_1, \dots, S_k \subseteq \{0, 1\}^n$, and define the hash functions $h_1, \dots, h_k : \{0, 1\}^n \rightarrow \{0, 1\}$ by $h_i(x) = \bigoplus_{j \in S_i} x_j$. Note that $h_i(x) \neq h_i(y)$ with $1/2$ chance if $x \neq y$. Hence, the following expression equals $EQ(x, y)$ with error probability at most ε :

$$\prod_{i=1}^k (h_i(x)h_i(y) + (1 - h_i(x))(1 - h_i(y))). \quad (6)$$

When expanded out, (6) is a sum of $2^k = O(1/\varepsilon)$ terms of the form $f(x) \cdot g(y)$ for some functions f and g , each of which has rank one. $\square$

LEMMA D.3. *For every n , LEQ_n has ε -probabilistic rank at most $O(n^2/\varepsilon)$ over any field.*

PROOF. We express LEQ_n in terms of EQ predicates which check for the first bit in which x and y differ, as

$$\begin{aligned} &LEQ_n(x_1, \dots, x_n, y_1, \dots, y_n) \\ &= \sum_{i=1}^n (1 - x_i) \cdot y_i \cdot EQ_{i-1}(x_1, \dots, x_{i-1}, y_1, \dots, y_{i-1}). \end{aligned} \quad (7)$$

We then get the desired rank bound by replacing each EQ with the probabilistic rank expression from Lemma D.2 with error ε/n . By the union bound, all n of the EQ predicates will be correct with probability at least $1 - \varepsilon$, and hence we will correctly compute LEQ_n . $\square$

LEMMA D.4. *For all n , every linear threshold function $f : \{0, 1\}^{2n} \rightarrow \{0, 1\}$ has ε -probabilistic rank $O(n^2/\varepsilon)$.*

PROOF. Since f is a linear threshold function, we can write it as $f(x_1, \dots, x_n, y_1, \dots, y_n) = [\sum_i v_i x_i + \sum_i w_i y_i \geq k]$, where all v_i 's, w_i 's, and k are reals. We want to show that the $2^n \times 2^n$ matrix indexed by x_i -assignments on the rows and y_i -assignments on the columns has low probabilistic rank. We will exploit the fact that the linear forms on x_i 's and y_i 's can be preprocessed separately in a rank decomposition.

Define $a : \{0, 1\}^n \rightarrow \mathbb{R}$ by $a(x_1, \dots, x_n) = \sum_{i=1}^n v_i x_i$, and $b : \{0, 1\}^n \rightarrow \mathbb{R}$ by $b(y_1, \dots, y_n) = k - \sum_{j=1}^n w_j y_j$. Hence

$$f(x, y) = [a(x) \leq b(y)].$$

Let L be the list, sorted in increasing order, of all values of $a(x)$ and $b(y)$, for all $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^n$. Then define the function $\alpha : \{0, 1\}^n \rightarrow \{0, 1\}^{n+1}$ where $\alpha(x)$ equals the earliest index of $a(x)$ in the sorted list L , interpreted as a $n + 1$ bit number. Define $\beta : \{0, 1\}^n \rightarrow \{0, 1\}^{n+1}$ similarly. Then

$$f(x, y) = LEQ_{n+1}(\alpha(x), \beta(y)).$$

So the ε -probabilistic rank of M_f is at most that of MEQ_{n+1} , which we upper-bound in Lemma D.3. $\square$

Now we are ready to upper-bound the probabilistic sign-rank of depth-two threshold circuits:

Proof of Theorem D.1. We interpret our LTF $\circ$ LTF circuit C as a function on two groups of $n/2$ bits, $x_1, \dots, x_{n/2}$, and $y_1, \dots, y_{n/2}$. Let $w_1, \dots, w_s, k \in \mathbb{R}$ be the weights of the output gate, so that

$$C(x_1, \dots, x_n, y_1, \dots, y_n) = \left[\sum_{i=1}^s w_i \cdot f_i(x_1, \dots, x_n, y_1, \dots, y_n) \leq k \right],$$

for s different LTF functions f_i . By Lemma D.4, the truth table matrix M_{f_i} of each f_i has (ε/s) -probabilistic rank $r = O(n^2 s / \varepsilon)$. Our probabilistic distribution of matrices for M_C can be constructed as follows: for all $i = 1, \dots, s$, draw a random rank- r matrix P_i from the distribution for M_{f_i} , and set

$$Q_C = (-k \cdot I) + \sum_i (w_i \cdot P_i).$$

Q_C has rank at most $sr + 1 \leq O(n^2 s^2 / \varepsilon)$ and for all $(\vec{x}, \vec{y})$, we have $\Pr[\text{sign}(Q_C[\vec{x}, \vec{y}]) \neq C(\vec{x}, \vec{y})] \leq \varepsilon$. $\square$

Are there explicit matrices with non-trivial sign-rank rigidity? We observe that the best-known rank rigidity lower bounds for H_n extend to sign-rank rigidity:

THEOREM D.5 (FOLLOWS FROM RAZBOROV AND SHERSTOV [RS10]). *For all n , and $r \in [2^{n/2}, 2^n]$, the sign-rank- r rigidity of H_n is at least $\Omega(4^n/r)$.*

PROOF. Theorem 5.1 of [RS10] gives the following lower bound on sign-rank: given any matrix $A \in \{-1, 1\}^{n \times n}$, suppose that all but h entries of matrix A have absolute value at least γ . Then

$$\text{sign-rank}(A) \geq \frac{\gamma n^2}{\|A\|n + \gamma h},$$

where $\|A\|$ is the spectral norm of A . For the case of H_n , if we modify $h := 4^n/r$ entries arbitrarily, all but h entries have absolute value equal to 1. Thus

$$\text{sign-rank}(H_n) \geq \frac{4^n}{\|H_n\|n + 4^n/r}.$$

As $\|H_n\| \leq O(2^{n/2})$ [For02], we have $\text{sign-rank}(H_n) \geq \Omega(4^n/(2^{3n/2} + 4^n/r)) \geq \Omega(2^{n/2} + r) \geq \Omega(r)$. $\square$

Can the above lower bound be improved slightly? Combining the previous two theorems, it follows that any minor improvement in the above rank/rigidity trade-off would begin to imply lower bounds for LTF $\circ$ LTF:

THEOREM D.6. *Suppose there is an $\alpha > 0$ such that for infinitely many n , the sign rank r -rigidity of H_n is $\Omega(4^n/r^{1-\alpha})$, for some $r \geq \omega(n^{2/\alpha} s(n)^{2/\alpha})$. Then the Inner Product Modulo 2 does not have LTF $\circ$ LTF circuits of $s(n)$ gates.*

PROOF. Suppose the sign rank r -rigidity of H_m is $\Omega(4^n/r^{1-\alpha})$. Let $\varepsilon = 1/r^{1-\alpha}$. It follows that the ε -probabilistic sign-rank of H_n is greater than r . But for a LTF $\circ$ LTF function with s gates, its matrix always has $\Omega(\varepsilon)$ -probabilistic rank $O(s^2 n^2 / \varepsilon) = O(s^2 n^2 r^{1-\alpha})$, by Theorem D.1. Thus we have a contradiction when $O(s^2 n^2 r^{1-\alpha})$ is asymptotically less than r , i.e.,

$$r = \omega(n^{2/\alpha} s^{2/\alpha}),$$

corresponding to an s -gate lower bound against LTF $\circ$ LTF circuits. Since H_n is just a linear translation of the matrix for Inner Product Modulo 2, the proof is complete. $\square$

For instance, proving the sign-rank $2^{\alpha n}$ -rigidity of H_n is at least $4^n/2^{.999\alpha n}$ for some $\alpha > 0$ would imply exponential-gate lower bounds for depth-two threshold circuits computing IP2.

E RANDOM MATRICES ARE SIGN-RANK RIGID

The proof that random $-1/1$ matrices have high sign-rank rigidity follows readily from recent work:

THEOREM E.1 (FOLLOWS FROM [AMY16]). *Let $r(n) = o(n/\log n)$. For all sufficiently large n , a random $n \times n$ matrix with $-1/1$ entries has sign-rank- $r(n)$ rigidity at least $\Omega(n^2)$, with high probability.*

PROOF. There are 2^{n^2} matrices over $\{-1, 1\}$. The number of distinct matrices with sign rank at most r is at most $2^{O(rn \log n)}$ [AMY16]. For a fixed matrix M , the number of matrices within Hamming distance d of M is at most $O(\binom{n^2}{d})$. Thus the number of matrices for which up to t entries can be changed to obtain a matrix of sign rank at most r , is upper-bounded by

$$2^{O(rn \log n)} \cdot \binom{n^2}{t} \leq n^{O(rn)} \cdot (en^2/t)^t.$$

Suppose we set $t = \varepsilon n^2$. Then the above quantity is at most

$$n^{O(rn)} \cdot (e/\varepsilon)^{\varepsilon n^2}.$$

For $r = o(n/\log n)$ and $\varepsilon \log_2(e/\varepsilon) < 1$, a random matrix is not among these matrices with high probability. Therefore a random matrix has sign rank- $o(n/\log n)$ rigidity $\Omega(n^2)$ with high probability. $\square$

F EQUIVALENCE BETWEEN PROBABILISTIC RANK MODULO m AND BP-MOD m COMMUNICATION COMPLEXITY

Here we sketch how probabilistic rank over $\mathbb{Z}_m$ is equivalent to BP $\cdot$ MOD m P communication complexity:

PROPOSITION F.1. *Let $m > 1$ be an integer, let $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, and let M_f be its truth table matrix. Let $C_\varepsilon(f)$ be the BP $\cdot$ MOD m P communication complexity of f with error ε , and let $\varepsilon\text{-rank}_{\mathbb{Z}_m}(M_f)$ be the ε -probabilistic rank of M_f over $\mathbb{Z}_m$. Then $C_\varepsilon(f) \leq \log_2(\varepsilon\text{-rank}_{\mathbb{Z}_m}(M_f) + 1) \leq 2C_\varepsilon(f)$.*

This proposition is different from the one quoted in the introduction (giving constant-factor equivalences between the log of the rank and the communication complexity) because we are assuming a more stringent communication model here. However, the more general model is often taken as the definition, in which case the probabilistic rank and communication complexity truly coincide.

First, given a distribution of low-rank matrices for M_f , it is easy to construct a protocol for f : Alice and Bob publicly randomly sample a matrix from the distribution, which is a product of two matrices A and B . Alice takes the row of A of length r corresponding to her input, Bob takes the column of B of length r corresponding to his, and they then compute the inner product of these two vectors over $\mathbb{Z}_m$ with $\lceil \log_2(r+1) \rceil$ communication in the MOD_mP model.

To construct a distribution of matrices from communication protocols, we do a simple modification of the $\text{BPP}^{\oplus\text{P}}$ communication model. In fact, sometimes the literature defines the $\text{BPP}^{\oplus\text{P}}$ communication model in this modified way [GPW16]. After the public randomness is chosen, Alice and Bob can, along with their c nondeterministic bits, also sum over all possible transcripts of at most c bits between them. For each choice of randomness and nondeterminism there is a unique accepting transcript, so this extra choice does not alter the number of accepting communication patterns. But in this modified version, now Alice and Bob do not even have to communicate: they only have to send a single bit indicating whether they would accept or not, given the transcript and the nondeterminism. From such a protocol, it is straightforward to construct a $2^n \times 2^{2c}$ matrix A representing Alice's protocol and a $2^{2c} \times 2^n$ matrix B representing Bob, for any given string of public randomness.

G RANDOM SELF-REDUCIBILITY, RIGIDITY, AND PROBABILISTIC RANK

The reduction from rigidity to probabilistic rank works for any (non-adaptive) random self-reducible function [FF93] that makes a small number of oracle calls. Our notion of random self-reducibility is adapted for the communication complexity setting (for example, we do not care about the feasibility of the reduction).

Definition G.1. A function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ is k -random self-reducible if there are random sampling procedures S_1, S_2 and a function $g : \{0, 1\}^k \rightarrow \{0, 1\}$ such that:

- S_1 takes $x \in \{0, 1\}^n$ and a random string r and outputs $x_1, \dots, x_k \in \{0, 1\}^n$ such that for all n -bit strings z , $\Pr_r[x_i = z] = 1/2^n$ for all i ,
- S_2 takes $y \in \{0, 1\}^n$ and a random string s and outputs $y_1, \dots, y_k \in \{0, 1\}^n$ such that for all n -bit strings z , $\Pr_s[y_i = z] = 1/2^n$ for all i , and
- $f(x, y) = g(f(x_1, y_1), \dots, f(x_k, y_k))$.

Requirements (a) and (b) in the definition ensures that each x_i and y_i are uniform random variables; requirement (c) says that we can reconstruct $f(x, y)$ from the values $f(x_1, y_1), \dots, f(x_k, y_k)$.

THEOREM G.2. Let $r, n \in \mathbb{N}$ and $\varepsilon \in (0, 1)$, and let $f : \{0, 1\}^{2n} \rightarrow \{0, 1\}$ be k -random self-reducible. Suppose M_f has rank- r rigidity at most $\varepsilon 4^n$ over K . Then the $(k\varepsilon)$ -probabilistic rank of M_f over K is at most $O(\binom{kr}{k})$.

PROOF. Suppose there is an $2^n \times r$ matrix A and $r \times 2^n$ matrix B , such that M_f and $A \cdot B$ differ in at most $\varepsilon 4^n$ entries. We construct a distribution of low-rank matrices for M_f as follows.

Let $P(z_1, \dots, z_k)$ be the unique multilinear polynomial over K that represents the function g from the random self-reduction for f . Given k rows $X_1, \dots, X_k \in K^r$ of A , and k columns $Y_1, \dots, Y_k \in K^r$ of B , define a polynomial in $2kr$ variables:

$$Q(X_1, \dots, X_k, Y_1, \dots, Y_k) = P(\langle X_1, Y_1 \rangle, \dots, \langle X_k, Y_k \rangle).$$

Treating each term of the form $X_i[j] \cdot Y_i[j]$ as a variable, Q can be written as a sum of $t \leq \binom{kr}{k}$ total terms. Call the terms $m_1, \dots, m_t$, each of which are over $2kr$ variables.

Let r be a random string for S_1 and s be a random string for S_2 . For $x \in \{0, 1\}^n$, let $x_1, \dots, x_k \in \{0, 1\}^n$ be the outputs of $S_1(x)$ with randomness r . We define the x th row of a new $2^n \times t$ matrix A_r to be

$$[m_1(A[x_1, :], \dots, A[x_k, :]), \vec{1}, \dots, \vec{1}), \dots, m_t(A[x_1, :], \dots, A[x_k, :]), \vec{1}, \dots, \vec{1})].$$

For $y \in \{0, 1\}^n$, let $y_1, \dots, y_k$ be the outputs of $S_2(y)$ with randomness s . Define the y th column of a new $t \times 2^n$ matrix B_s to be

$$[m_1(\vec{1}, \dots, \vec{1}, B[:, y_1], \dots, B[:, y_k]), \dots, m_t(\vec{1}, \dots, \vec{1}, B[:, y_1], \dots, B[:, y_k])]^T.$$

Then, for all $(x, y) \in \{0, 1\}^n \times \{0, 1\}^n$, the inner product of the x th row of A_r and the y th column of B_s is

$$\begin{aligned} & \sum_i m_i(A[x_1, :], \dots, A[x_k, :], B[:, y_1], \dots, B[:, y_k]) \\ &= Q(A[x_1, :], \dots, A[x_k, :], B[:, y_1], \dots, B[:, y_k]) \\ &= P(\langle A[x_1, :], B[:, y_1] \rangle, \dots, \langle A[x_k, :], B[:, y_k] \rangle). \end{aligned}$$

Since $(A \cdot B)$ differs from M_f on an ε -fraction of entries, for uniform random $x_i, y_j \in \{0, 1\}^n$ we have $\langle A[x_i, :], B[:, y_j] \rangle \neq f(x_i, y_j)$ with probability at most ε . So with probability at least $1 - k\varepsilon$, $f(x_i, y_j) = \langle A[x_i, :], B[:, y_j] \rangle$ for all $i = 1, \dots, k$. Thus the polynomial $P(\langle A[x_1, :], B[:, y_1] \rangle, \dots, \langle A[x_k, :], B[:, y_k] \rangle)$ being implemented by $A_r \cdot B_s$ outputs $f(x, y)$ with probability at least $1 - k\varepsilon$. Hence all matrices $C_{r,s} = A_r \cdot B_s$ in our defined distribution have rank at most $O(\binom{kr}{k})$, and for every $(x, y) \in \{0, 1\}^n \times \{0, 1\}^n$, $\Pr_{r,s}[C_{r,s}[x, y] = f(x, y)] \geq 1 - k\varepsilon$. $\square$

ACKNOWLEDGMENTS

We thank Mrinal Kumar for helpful comments, Oded Goldreich, Michael Forbes and Satya Lokam for helpful discussions, and Amir Shpilka and Avishay Tal for patiently listening to R.W.'s conjectures and results on non-rigidity at Banff (BIRS) in August 2016.

REFERENCES

- [ABG⁺14] Adi Akavia, Andrej Bogdanov, Siyao Guo, Akshay Kamath, and Alon Rosen. Candidate weak pseudorandom functions in AC^0 o MOD_2 . In *ITCS*, pages 251–260, 2014.
- [ACW16] Josh Alman, Timothy Chan, and Ryan Williams. Polynomial representations of threshold functions and algorithmic applications. In *FOCS*, pages 467–476, 2016.
- [Alo90] Noga Alon. On the rigidity of an Hadamard matrix. Manuscript. See [Juk01, Section 15.1.2], 1990.

- [ALSV13] Noga Alon, Troy Lee, Adi Shraibman, and Santosh Vempala. The approximate rank of a matrix and its algorithmic applications. In *STOC*, pages 675–684. ACM, 2013.
- [AMY16] Noga Alon, Shay Moran, and Amir Yehudayoff. Sign rank versus VC dimension. In *COLT*, pages 47–80, 2016.
- [AW09] Scott Aaronson and Avi Wigderson. Algebrization: A new barrier in complexity theory. *ACM Transactions on Computation Theory (TOCT)*, 1(1):2, 2009.
- [AW15] Josh Alman and Ryan Williams. Probabilistic polynomials and hamming nearest neighbors. In *FOCS*, pages 136–150. IEEE, 2015.
- [AWY15] Amir Abboud, Ryan Williams, and Huacheng Yu. More applications of the polynomial method to algorithm design. In *SODA*, pages 218–230, 2015.
- [BdW01] Harry Buhrman and Ronald de Wolf. Communication complexity lower bounds by polynomials. In *CCC*, pages 120–130. IEEE, 2001.
- [CGJ⁺16] Mahdi Cheraghchi, Elena Grigorescu, Brendan Juba, Karl Wimmer, and Ning Xie. AC^0 modulo 2 lower bounds for the boolean inner product. In *ICALP*, pages 35:1–35:14, 2016.
- [Cod00] Bruno Codenotti. Matrix rigidity. *Linear Algebra and its Applications*, 304(1&2&3):181 – 192, 2000.
- [Dvi16] Zeev Dvir. On the non-rigidity of generating matrices of good codes. Writeup by Oded Goldreich. Available at <http://www.wisdom.weizmann.ac.il/~oded/MC/209.html>, October 30 2016.
- [dW06] Ronald de Wolf. Lower bounds on matrix rigidity via a quantum argument. In *Proc. ICALP*, volume 4051, pages 62–71, 2006.
- [FF93] Joan Feigenbaum and Lance Fortnow. Random-self-reducibility of complete sets. *SIAM J. Comput.*, 22(5):994–1005, 1993.
- [FKL⁺01] Jürgen Forster, Matthias Krause, Satyanarayana V. Lokam, Rustam Mubarakzjanov, Niels Schmitt, and Hans Ulrich Simon. Relations between communication complexity, linear arrangements, and computational complexity. In *FSTTCS*, pages 171–182, 2001.
- [For02] Jürgen Forster. A linear lower bound on the unbounded error probabilistic communication complexity. *J. Comput. System. Sci.*, 65:612–625, 2002.
- [Fri93] Joel Friedman. A note on matrix rigidity. *Combinatorica*, 13(2):235–239, 1993.
- [GPW16] Mika Göös, Toniann Pitassi, and Thomas Watson. Zero-information protocols and unambiguity in arthur-merlin communication. *Algorithmica*, 76(3):684–719, 2016.
- [Gri] D. Yu. Grigor'ev. Unpublished work. Cited in [KR98].
- [GT16] Oded Goldreich and Avishay Tal. Matrix rigidity of random toeplitz matrices. In *STOC*, pages 91–104, 2016.
- [HMP⁺93] András Hajnal, Wolfgang Maass, Pavel Pudlák, Mario Szegedy, and György Turán. Threshold circuits of bounded depth. *J. Comput. Syst. Sci.*, 46(2):129–154, 1993.
- [Juk01] Stasys Jukna. *Extremal Combinatorics, With Applications in Computer Science*. EATCS Series. Springer, 2001.
- [KLPS14] Abhinav Kumar, Satyanarayana V. Lokam, Vijay M. Patankar, and M. N. Jayalal Sarma. Using elimination theory to construct rigid matrices. *computational complexity*, 23(4):531–563, 2014.
- [KR98] B. S. Kashin and A. A. Razborov. Improved lower bounds on the rigidity of Hadamard matrices. *Matematicheskie Zametki*, 63(4):535–540, 1998. (in Russian).
- [KS10] Adam R Klivans and Alexander A Sherstov. Lower bounds for agnostic learning via approximate rank. *Computational Complexity*, 19(4):581–604, 2010.
- [KW16] Daniel M. Kane and Ryan Williams. Super-linear gate and super-quadratic wire lower bounds for depth-two and depth-three threshold circuits. In *STOC*, pages 633–643, 2016.
- [Lok00] Satyanarayana V. Lokam. On the rigidity of vandermonde matrices. *Theoretical Computer Science*, 237(1&2&3):477 – 483, 2000.
- [Lok01] Satyanarayana V. Lokam. Spectral methods for matrix rigidity with applications to size-depth tradeoffs and communication complexity. *Journal of Computer and System Sciences*, 63:449–473, 2001.
- [Lok06] Satyanarayana V. Lokam. Quadratic lower bounds on matrix rigidity. In *Proc. TAMC*, volume 3959, pages 295–307. Springer, 2006.
- [Lok09] Satyanarayana V. Lokam. Complexity lower bounds using linear algebra. *Foundations and Trends in Theoretical Computer Science*, 4(1-2):1–155, 2009.
- [Lok14] Satyanarayana V. Lokam. Exercises on matrix rigidity. Simons Institute for Theory of Computing. Available at <https://simons.berkeley.edu/sites/default/files/docs/1738/exercises.pdf>, 2014.
- [LPT⁺17] Daniel Lokshtanov, Ramamohan Paturi, Suguru Tamaki, Ryan Williams, and Huacheng Yu. Beating brute force for systems of polynomial equations over finite fields. In *SODA*, page to appear, 2017.
- [LS09] Nathan Linial and Adi Shraibman. Learning complexity vs communication complexity. *Combinatorics, Probability & Computing*, 18(1-2):227–245, 2009.
- [LTV03] JM Landsberg, J. Taylor, and N.K. Vishnoi. The geometry of matrix rigidity. preprint available at <https://smartech.gatech.edu/handle/1853/6514>, 2003.
- [Mid05] Gatis Midrijanis. Three lines proof of the lower bound for the matrix rigidity. Technical report, Arxiv e-print, June 2005.
- [MT98] Alexis Maciel and Denis Thérien. Threshold circuits of small majority-depth. *Information and Computation*, 146(1):55–83, 1998.
- [Nis] Noam Nisan. Unpublished work. Cited in [KR98].
- [PS88] P. Pudlak and P. Savicky. Private communication, cited in [Raz89], 1988.
- [Ras16] Cyrus Rashtchian. Bounded matrix rigidity and John's theorem. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:93, 2016.
- [Raz87] A. A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematical Notes of the Academy of Sciences of the USSR*, 41(4):333–338, 1987.
- [Raz89] A. A. Razborov. On rigid matrices (in Russian). Manuscript can be found at <http://people.cs.uchicago.edu/~razborov/files/rigid.pdf>, 1989.
- [RS10] Alexander A. Razborov and Alexander A. Sherstov. The sign-rank of AC^0 . *SIAM J. Comput.*, 39(5):1833–1855, 2010.
- [Smo87] Roman Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *STOC*, pages 77–82, 1987.
- [SSS97] M.A. Shokrollahi, D.A. Spielman, and V. Stemann. A remark on matrix rigidity. *Information Processing Letters*, 64(6):283 – 285, 1997.
- [SV12] Rocco A. Servedio and Emanuele Viola. On a special case of rigidity. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:144, 2012.
- [Tar93] Jun Tarui. Probabilistic polynomials, AC^0 functions and the polynomial-time hierarchy. *Theor. Comput. Sci.*, 113(1):167–183, 1993.
- [Val77] Leslie G. Valiant. Graph-theoretic arguments in low-level complexity. In *Mathematical Foundations of Computer Science (MFCS)*, pages 162–176, Berlin, Heidelberg, 1977. Springer Berlin Heidelberg.
- [Wil14a] Ryan Williams. Faster all-pairs shortest paths via circuit complexity. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 664–673, 2014.
- [Wil14b] Ryan Williams. The polynomial method in circuit complexity applied to algorithm design (invited talk). In *FSTTCS*, pages 47–60, 2014.
- [Wun12] Henning Wunderlich. On a theorem of razborov. *Computational Complexity*, 21(3):431–477, 2012.
- [Yao83] Andrew C. Yao. Lower bounds by probabilistic arguments (extended abstract). In *FOCS*, pages 420–428. IEEE, 1983.