Quantum Query-to-Communication Simulation Needs a Logarithmic Overhead

Sourav Chakraborty[1], Arkadev Chattopadhyay[2], Nikhil S. Mande[3], and Manaswi Paraashar[1]

[1]Indian Statistical Institute, Kolkata
[2]TIFR, Mumbai
[3]Georgetown University

Abstract

Buhrman, Cleve and Wigderson (STOC'98) observed that for every Boolean function $f : \{-1,1\}^n \to \{-1,1\}$ and $\bullet : \{-1,1\}^2 \to \{-1,1\}$ the two-party bounded-error quantum communication complexity of $(f \circ \bullet)$ is $O(Q(f) \log n)$, where $Q(f)$ is the bounded-error quantum query complexity of f. Note that the bounded-error randomized communication complexity of $(f \circ \bullet)$ is bounded by $O(R(f))$, where $R(f)$ denotes the bounded-error randomized query complexity of f. Thus, the BCW simulation has an extra $O(\log n)$ factor appearing that is absent in classical simulation. A natural question is if this factor can be avoided. Høyer and de Wolf (STACS'02) showed that for the Set-Disjointness function, this can be reduced to $c^{\log^* n}$ for some constant c, and subsequently Aaronson and Ambainis (FOCS'03) showed that this factor can be made a constant. That is, the quantum communication complexity of the Set-Disjointness function (which is $\mathsf{NOR}_n \circ \wedge$) is $O(Q(\mathsf{NOR}_n))$.

Perhaps somewhat surprisingly, we show that when $\bullet = \oplus$, then the extra $\log n$ factor in the BCW simulation is unavoidable. In other words, we exhibit a total function $F : \{-1,1\}^n \to \{-1,1\}$ such that $Q^{cc}(F \circ \oplus) = \Theta(Q(F) \log n)$.

To the best of our knowledge, it was not even known prior to this work whether there existed a total function F and 2-bit function $\bullet$, such that $Q^{cc}(F \circ \bullet) = \omega(Q(F))$.

1 Introduction

Classical communication complexity, introduced by Yao [Yao79], is aptly called the 'swiss-army-knife' for understanding, especially the limitations of, classical computing. Quantum communication complexity holds the same promise with regards to quantum computing. Yet, there are many problems that remain open. One broad theme is to understand the fundamental differences between classical randomized and quantum protocols, especially for computing total functions.

Recall a standard way to derive a communication problem from a function $f : \{-1,1\}^n \to \{-1,1\}$. Each input bit of f is *encoded* between the two players Alice and Bob, using an instance of a binary primitive, denoted by $\bullet : \{-1,1\} \times \{-1,1\} \to \{-1,1\}$, giving rise to the communication problem of evaluating $f \circ \bullet$. Each input bit to f is obtained by evaluating $\bullet$ on the relevant bit of Alice and that of Bob, i.e. $(f \circ \bullet)(x,y) = f(\bullet(x_1,y_1), \ldots, \bullet(x_n,y_n))$ and x, y are each n-bit strings given to Alice and Bob respectively. Many well known functions in communication complexity are derived in this way: Set-Disjointness is $\mathsf{NOR} \circ \wedge$, Inner-Product being $\oplus \circ \wedge$, Equality being $\mathsf{NOR} \circ \oplus$. Set-Disjointness is also a standard total function where quantum protocols provably yield a significant cost saving over their classical counterpart.

A natural and well studied question is this regard is what is the relationship between the query complexity of f and the communication problem of $f \circ \bullet$, when the $\bullet$ is $\wedge$ or $\oplus$.[1] This question has been studied for particular interesting functions or special classes of functions. Classically, it is folklore that

$$R^{cc}(f \circ \bullet) = O(R(f)),$$

where $R(f)$ denotes the bounded-error randomized query complexity of f and $R^{cc}(f \circ \bullet)$ denotes the bounded-error randomized communication complexity for computing $(f \circ \bullet)(x, y)$ when Alice and Bob has inputs x and y respectively. However in the quantum world such a relation is not at all obvious. In an influential work, Buhrman, Cleve and Wigderson [BCW98] observed that a general and natural recipe exists for constructing a quantum communication protocol for $(f \circ \bullet)$, using a quantum query algorithm for f as a black-box.

Theorem 1.1 ([BCW98])**.** *For any Boolean function $f : \{-1, 1\}^n \to \{-1, 1\}$, we have*

$$Q^{cc}(f \circ \bullet) = O(Q(f) \cdot \log n),$$

where $\bullet$ is either $\wedge$ or $\oplus$.

Here $Q(f)$ denotes the bounded-error quantum query complexity of f, and $Q^{cc}(f \circ \bullet)$ denotes the bounded-error quantum communication complexity for computing $(f \circ \bullet)(x, y)$. Thus in the quantum world one incurs a logarithmic factor in the BCW simulation while no such factor is needed in the randomized setting. The basic question that arises naturally and which we completely answer in this work, is the following: analogous to the classical model, can this multiplicative $\log n$ blow-up in the communication cost be always avoided by designing quantum communication protocols that more cleverly simulate quantum query algorithms?

A priori, it is not clear what the answer to this question ought to be. For certain special functions and some classes of functions, quantum protocols exist where the $\log n$ factor can be saved. First, Høyer and de Wolf [HdW02] designed a quantum protocol for Set-Disjointness of cost $O(\sqrt{n}c^{\log^* n})$, speeding up the BCW simulation significantly. Later, Aaronson and Ambainis [AA05] gave a more clever protocol that only incurred a constant factor overhead from Grover's search using more involved ideas.

For partial functions, tightness of the BCW simulation is known in some settings. For example, consider the Deutsch-Jozsa (DJ) problem, where the input is an n-bit string with the promise that its Hamming weight is either 0 or $n/2$, and DJ outputs -1 if the Hamming weight is $n/2$, and 1 otherwise. DJ has quantum query complexity 1 whereas the *exact* quantum communication complexity of $(\mathsf{DJ} \circ \oplus)$ is $\log n$. Note that it is unclear whether the $\log n$ factor loss here is additive or multiplicative[2]. Montanaro, Nishimura and Raymond [MNR11] exhibited a partial function for which the BCW simulation is tight (up to constants) in the *exact* and *non-deterministic* quantum settings. They also observed the existence of a total function for which the BCW simulation is tight (up to constants) in the unbounded-error setting. If we allow multi-output partial functions, then tightness of the BCW simulation is known: Consider the following function: f takes an n-bit string x as input which is promised to be a Hadamard codeword (see Definition 2.3), and outputs a $\log n$ bit string z for which x is its Hadamard codeword. The communication problem $f \circ \wedge$, where the inputs to the two players are promised to be such that their bitwise-AND yields a Hadamard

---

codeword, has bounded-error quantum communication complexity $\log n$, but $Q(f) = O(1)$. Again, it is not clear here whether the $\log n$ factor loss is additive or multiplicative.

As far as we know, there was no (partial or total) Boolean-valued function f known prior to our work for which the *bounded-error* quantum communication complexity of $f \circ \bullet$ (i.e. $Q^{cc}(f \circ \bullet)$) is even $\omega(Q(f))$, where $\bullet$ is either $\wedge$ or $\oplus$.

In this paper, we exhibit the first *total function* witnessing the tightness of the BCW simulation in arguably the most well-known quantum model, which is the bounded-error model.

Theorem 1.2. *There exists a total function $F : \{-1,1\}^n \to \{-1,1\}$ for which,*

$$Q^{cc}(F \circ \oplus) = \Theta(Q(F) \log n). \tag{1}$$

The statement of Theorem 1.2 does not necessarily guarantee that a function exists that both satisfies Equation 1 and has bounded-error quantum query complexity (as a function of n) arbitrarily close to n. We answer this question by proving a more general result, from which Theorem 1.2 immediately follows.

Theorem 1.3 (Main Theorem). *For any constant $0 < \delta < 1$, there exists a total function $F : \{-1,1\}^n \to \{-1,1\}$ for which $Q(F) = \Theta(n^\delta)$ and*

$$Q^{cc}(F \circ \oplus) = \Theta(Q(F) \log n).$$

1.1 Overview of our approach and techniques

To demonstrate the tightness of the BCW simulation for a total function in the quantum bounded-error setting we have to find a function F such that $Q^{cc}(F \circ \bullet) = \Theta(Q(F) \log n)$ for some choice of $\bullet$ (that is, either $\bullet$ is $\wedge$ or $\oplus$). This requires us to prove an upper bound of $Q(F)$ and a lower bound on $Q^{cc}(F \circ \bullet)$. We consider the case when $\bullet$ is the $\oplus$ function.

For the inner function the $\oplus$ function is preferred over the $\wedge$ function for one crucial reason: we have an analytical technique for proving lower bounds on $Q^{cc}(F \circ \oplus)$, due to Lee and Shraibman [LS09]. They reduced the problem of lower bounding the bounded-error quantum communication complexity of $(F \circ \oplus)$ to proving lower bounds on an analytic property of F, called its *approximate spectral norm*. The ϵ-approximate spectral norm of F, denoted by $\|\hat{F}\|_{1,\epsilon}$, is defined to be the minimum ℓ_1-norm of the coefficients of a polynomial that approximates F uniformly to error ϵ (see Definition 2.9). Lee and Shraibman [LS09] showed that $Q^{cc}(F \circ \oplus) = \Omega(\log \|\hat{F}\|_{1,1/3})$. Thus, the lower bound of Theorem 1.3 follows immediately from the result below.

Theorem 1.4. *For any constant $0 < \delta < 1$, there exists a total function $F : \{-1,1\}^n \to \{-1,1\}$ for which $Q(F) = \Theta(n^\delta)$ and*

$$\log \left(\|\hat{F}\|_{1,1/3} \right) = \Theta(Q(F) \log n).$$

There are not many techniques known to bound the approximate spectral norm of a function. This sentiment was expressed both in [LS09] and in the work of Ada, Fawzi and Hatami [AFH12]. On the other hand, classical approximation theory offers tools to prove bounds on a simpler and better known concept called *approximate degree* which has been invaluable particularly for quantum query complexity. The ϵ-approximate degree of f, denoted by $\widetilde{\deg}_\epsilon(f)$, is the minimum degree required by a real polynomial to uniformly approximate f to error ϵ (see Definition 2.8). Recently, two of the authors [CM17] devised a way of lifting approximate degree bounds to approximate spectral norm bounds. We first show here that technique works a bit more

generally, to yield the following: let $\mathsf{ADDR}_{m,t} : \{-1, 1\}^m \to [t]$ be a (possibly partial) addressing function (see Definition 2.4). For any function $f : \{-1, 1\}^n \to \{-1, 1\}$, define the (partial) function $f^{\mathsf{ADDR}_{m,t}} : \{-1, 1\}^{n \times t} \times \{-1, 1\}^{n \times m} \to \{-1, 1\}$ as follows (formally defined in Definition 2.6):

$$f^{\mathsf{ADDR}_{m,t}}(x, y) = f\left(x_{1, \mathsf{ADDR}_{m,t}(y_1)}, x_{2, \mathsf{ADDR}_{m,t}(y_2)}, \ldots, x_{n, \mathsf{ADDR}_{m,t}(y_n)}\right).$$

Our main result on lower bounding the spectral norm is stated below.

Lemma 1.5 (extending [CM17]). *Let $t > 1$ be any integer, $\mathsf{ADDR}_{m,t}$ be any (partial) addressing function and $f : \{-1, 1\}^n \to \{-1, 1\}$ be any function. Then,*

$$\log\left(\|\widehat{f^{\mathsf{ADDR}_{m,t}}}\|_{1,1/3}\right) = \Omega\left(\widetilde{\deg}(f) \log t\right).$$

The functions F constructed for the proof of Theorem 1.4 are completions of instances of $\mathsf{PARITY}^{\mathsf{ADDR}_{\ell,\ell}}$, and hence Lemma 1.5 yields lower bounds on the approximate spectral norm of F in terms of the approximate degree of PARITY (which is known to be maximal).

For the upper bound on $Q(F)$ we use two famous query algorithms. These are Grover's search [Gro96] and the Bernstein-Vazirani algorithm [BV97]. The use of these algorithms for upper bounding $Q(F)$ is in the same taste as in the work of Ambainis and de Wolf [AdW14] although their motivation was quite different than ours. Interestingly, Ambainis and de Wolf used their function to pin down the minimal *approximate degree* of a total Boolean function, all of whose input variables are influential.

1.2 Intuition behind the function construction

- From Theorem 1.1 it is known that for all Boolean functions f, $Q^{cc}(f \circ \oplus) \leq O(Q(f) \log n)$.

- In order to prove a matching lower bound, we construct a Boolean function F on n variables such that $\|\widehat{F}\|_{1,1/3} = 2^{\Omega(Q(F) \log n)}$ (Theorem 1.4). From Theorem 2.13, this shows that $Q^{cc}(F \circ \oplus) = \Omega(\log \|\widehat{F}\|_{1,1/3}) = \Omega(Q(F) \log n)$. We want to additionally ensure that $Q(F) = \Theta(n^\delta)$ for a given constant $0 < \delta < 1$. A formal definition of F is given in Figure 1.2, we attempt to provide an overview on how we arrived at this function below.

- Assume δ is a constant that is least $1/2$, else the argument follows along similar lines by ignoring suitably many input variables when defining the function. A natural first attempt is to try to construct a composed function of the form $F = f^{\mathsf{ADDR}}$, for some addressing function ADDR (see Definition 2.4) with $\Omega(n^{1-\delta})$ many target bits, for which $Q(f^{\mathsf{ADDR}}) = \Theta(\widetilde{\deg}(f))$. For the lower bound we use Lemma 1.5 to show that $\log \|\widehat{f^{\mathsf{ADDR}}}\|_{1,1/3} = \Omega(\widetilde{\deg}(f) \log(n^{1-\delta})) = \Omega(\widetilde{\deg}(f) \log n)$.

- Given the upper bound target, we are led to a natural choice of addressing function. Let $\mathsf{HADD}_{n^{1-\delta}}$ be the $(n^{1-\delta}, n^{1-\delta})$-addressing function defined as follows. Fix an arbitrary order on the $n^{1-\delta}$-bit Hadamard codewords (see Definition 2.3) , say $w_1, \ldots, w_{n^{1-\delta}}$. Define g to be the selector function of $\mathsf{HADD}_{n^{1-\delta}}$ such that $g(w_i) = i$ for all $i \in [n^{1-\delta}]$, and $g(x) = \star$ for $x \neq w_i$ for any $i \in [n^{1-\delta}]$.

- For any function f on $n^\delta/2$ bits, the *partial* function $f^{\mathsf{HADD}_{n^{1-\delta}}}$ on n inputs has quantum query complexity $O(Q(f) + n^\delta/2)$, as we sketch in the next step. We select f appropriately such that this is $\Theta(Q(f))$. Finally, we define the *total* function $F = f^{\mathsf{HADD}_{n^{1-\delta}}}$ to be the completion of $f^{\mathsf{HADD}_{n^{1-\delta}}}$ that evaluates to -1 on the non-promised inputs of $f^{\mathsf{HADD}_{n^{1-\delta}}}$.
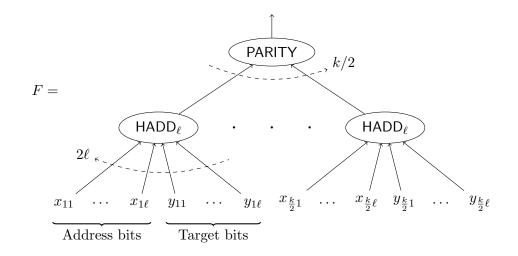
Figure 1: $k = n^\delta, \ell = n^{1-\delta}$. If the address bits of an input to an HADD_ℓ is the j'th Hadamard codeword, then y_j is selected. If on an input, there exists at least one HADD_ℓ for which the address bits do not correspond to a Hadamard codeword, F outputs -1. Else it outputs the Parity of the $k/2$ selected target bits.

- We choose the outer function to be $f = \mathsf{PARITY}_{n^\delta/2}$ to ensure $Q(F) = \Theta(n^\delta)$. To prove the upper bound on $Q(F)$, we crucially use the Bernstein-Vazirani and Grover's search algorithms.

    - Run $n^\delta/2$ instances of the Bernstein-Vazirani algorithm, one on each block. This guarantees that if the address variables were all Hadamard codewords, then we would receive the correct indices of the target variables with probability 1 and just $n^\delta/2$ queries.

    - In the next step, we run Grover's search on two $n/2$-bit strings to test whether the output of the first step was correct. If it was correct, we succeed with probability 1, and proceed to query the $n^\delta/2$ selected target variables and output the parity of them. If it was not correct, Grover's search catches a discrepancy with probability at least $2/3$ and we output -1, succeeding with probability at least $2/3$ in this case.

    - The $n^\delta/2$ invocations of the Bernstein-Vazirani algorithm use a total of $n^\delta/2$ queries, Grover's search uses another $O(\sqrt{n})$ queries, and the final parity (if Grover's search outputs that the strings are equal) uses another $n^\delta/2$ queries, for a cumulative total of $O(n^\delta + \sqrt{n}) = O(n^\delta)$ queries (recall that we assume $\delta \geq 1/2$).

1.3 Other implications of our result

Zhang [Zha09] showed that for all Boolean functions f, there must exists gadgets g_i, each either $\wedge$ or $\vee$, such that $Q^{cc}(f(g_1, \ldots, g_n)) = \Omega(\mathrm{poly}(Q(f)))$. For monotone f, they showed that either $Q^{cc}(f \circ \wedge) = \Omega(\mathrm{poly}(Q(f)))$ or $Q^{cc}(f \circ \mathsf{OR}_2) = \Omega(\mathrm{poly}(Q(f)))$. They also state that it is unclear how tight the BCW simulation is. We show that there exists a function for which it is tight up to constants (on composition with $\oplus$).

Another implication of our result is related to the Entropy Influence Conjecture, which is an interesting question in the field of analysis of Boolean functions, posed by Friedgut and Kalai [FK96]. This conjecture is wide open for general functions. A much weaker version of this conjecture is

called the Min-Entropy Influence Conjecture. For the statement of the conjecture we need to consider the Fourier expansion of $f : \{-1,1\}^n \to \{-1,1\}$ as

$$f(\mathbf{x}) = \sum_{S \subseteq [n]} \widehat{f}(S)\chi_S(\mathbf{x}),$$

where $\{\chi_S : S \subseteq [n]\}$ are the *parity* functions ($\chi_S(\mathbf{x}) = \Pi_{i \in S}x_i$, when $\mathbf{x} = (x_1, \ldots, x_n) \in \{-1,1\}^n$) and $\{\widehat{f}(S) : S \subseteq [n]\}$ are the corresponding Fourier coefficients.

Conjecture 1.6. *(Min-Entropy Influence Conjecture) For any Boolean function $f : \{-1,1\}^n \to \{-1,1\}$ there exists a non-zero Fourier coefficient $\widehat{f}(S)$ such that*

$$\log\left(1/|\widehat{f}(S)|\right) = O(I(f)),$$

where $I(f)$ denotes the influence (or average sensitivity) of f ($I(f) = \sum_{S \subseteq [n]} |S|\widehat{f}(S)^2$).

While this conjecture is also wide open, some attempts have been made to prove various implications of this conjecture. One interesting implication of the Min-Entropy Influence Conjecture that is still open is whether the min-entropy of the Fourier spectrum (that is, $\log\left(1/\max_{S \subseteq [n]} |\widehat{f}(S)|\right)$) is less than $O(Q(f))$. In [ACK+18] using a primal-dual technique it was shown that the min-entropy of the Fourier spectrum is less than a constant times $\log(\|\hat{f}\|_{1,\epsilon})$, where the constant depends on ϵ. Thus if it was the case that $\log(\|\hat{f}\|_{1,\epsilon} = O(Q(f)))$, we would have upper bounded the min-entropy of Fourier spectrum by $O(Q(f))$. This was stated in [ACK+18] as a possible approach and was left as an open problem. Our result in this paper nullifies this approach.

2 Preliminaries

For any positive integer n, we denote the set $\{1, \ldots, n\}$ by $[n]$. For $d \leq n$ we use the notation $\binom{n}{\leq d} = \binom{n}{0} + \cdots + \binom{n}{d}$. Note that $\binom{n}{\leq d} < (n+1)^d$.

In this section we review the necessary preliminaries. We first review some basics of Fourier analysis on the Boolean cube. Consider the vector space of functions from $\{-1,1\}^n$ to $\mathbb{R}$, equipped with an inner product defined by

$$\langle f, g \rangle := \mathbb{E}_{x \in \{-1,1\}^n}[f(x)g(x)] = \frac{1}{2^n} \sum_{x \in \{-1,1\}^n} f(x)g(x).$$

for every $f, g : \{-1,1\}^n \to \mathbb{R}$. For any set $S \subseteq [n]$, define the associated *parity* function χ_S by $\chi_S(x) = \prod_{i \in S} x_i$. The set of parity functions $\{\chi_S : S \subseteq [n]\}$ form an orthonormal basis for this vector space. Thus, every function $f : \{-1,1\}^n \to \mathbb{R}$ has a unique multilinear expression as

$$f = \sum_{S \subseteq [n]} \widehat{f}(S)\chi_S.$$

The coefficients $\{\widehat{f}(S) : S \subseteq [n]\}$ are called the *Fourier coefficients* of f.

Fact 2.1 (Parseval's Identity). *For any function $f : \{-1,1\}^n \to \mathbb{R}$,*

$$\sum_{S \subseteq [n]} \widehat{f}(S)^2 = \frac{\sum_{x \in \{-1,1\}^n} f(x)^2}{2^n}.$$

6

Definition 2.2 (Spectral Norm). *For any function $f : \{-1,1\}^n \to \mathbb{R}$, define its spectral norm, which we denote $\|\widehat{f}\|_1$, to be the sum of absolute values of the Fourier coefficients of f. That is,*

$$\|\widehat{f}\|_1 := \sum_{S \subseteq [n]} |\widehat{f}(S)|.$$

Definition 2.3 (Hadamard Codeword). *If a ℓ-bit string $(x_1, \ldots, x_\ell) \in \{-1,1\}^\ell$ is of the form $x_S = \prod_{i \in S} z_i$ for all $S \subseteq [\log \ell]$ for some $z \in \{-1,1\}^{\log \ell}$, then define such an $x = x_1 \ldots x_\ell$ to be the ℓ-bit Hadamard codeword $h(z)$ of the $(\log \ell)$-bit string z.*

2.1 Addressing functions

Definition 2.4 ((m,k)-addressing function). *We define a (partial) function $f : \{-1,1\}^{m+k} \to \{-1, 1, \star\}$ to be an (m,k)-addressing function if there exists $g : \{-1,1\}^m \to \{[k] \cup \star\}$ such that*

- $f(x_1, \ldots, x_m, y_1, \ldots, y_k) = y_{g(x_1, \ldots, x_m)}$ *if $g(x_1, \ldots, x_m) \in [t]$, and $f(x_1, \ldots, x_m, y_1, \ldots, y_k) = \star$ otherwise.*

- *For all $j \in [k]$, there exists $(x_1, \ldots, x_m) \in \{-1,1\}^m$ such that $g(x_1, \ldots, x_m) = j$.*

We call the variables $\{x_1, \ldots, x_m\}$ the address variables *and the variables $\{y_1, \ldots, y_k\}$ the* target variables. *The function g is called the* selector *function of f.*

Definition 2.5 (Indexing Function). *The Indexing function, which we denote by IND_k, is a $(k, 2^k)$-addressing function defined by $\mathsf{IND}(x_1, \ldots, x_k, y_1, \ldots, y_{2^k}) = y_{\mathrm{bin}(x)}$, where $\mathrm{bin}(x)$ denotes the integer represented by the binary string $x_1, \ldots, x_k$.*

Definition 2.6 (Composition with addressing functions). *For any function $f : \{-1,1\}^n \to \{-1,1\}$ and an (m,k)-addressing function ADDR, define the (partial) function $f^{\mathsf{ADDR}} : \{-1,1\}^{n(m+k)} \to \{-1, 1, \star\}$ by*

$$f^{\mathsf{ADDR}}(x_1, y_1, \ldots, x_n, y_n) = \begin{cases} f(\mathsf{ADDR}(x_1, y_1), \ldots, \mathsf{ADDR}(x_n, y_n)) & \text{if } \mathsf{ADDR}(x_i, y_i) \in \{-1,1\} \text{ for all } i \in [n] \\ \star & \text{otherwise.} \end{cases}$$

where $x_i \in \{-1,1\}^m$ and $y_i \in \{-1,1\}^k$ for all $i \in [n]$.

Definition 2.7 (Hadamard Addressing Function). *We define the* Hadamard addressing function, *which we denote $\mathsf{HADD}_\ell : \{-1,1\}^{2\ell} \to \{-1, 1, \star\}$, as follows. Fix an arbitrary order on the ℓ-many Hadamard codewords of $(\log \ell)$-bit strings, say $w_1, \ldots, w_\ell$. Define the selector function of HADD_ℓ by*

$$g(x) = \begin{cases} i & \text{if } x = w_i \text{ for some } i \in [\ell] \\ \star & \text{otherwise.} \end{cases}$$

Note that HADD_ℓ is an (ℓ, ℓ)-addressing function.

2.2 Polynomial approximation

Definition 2.8 (Approximate Degree)**.** *The ϵ-approximate degree of a function $f : \{-1,1\}^n \to \{-1,1,\star\}$, denoted by $\widetilde{\deg}_\epsilon(f)$ is defined to be the minimum degree of a real polynomial $p : \{-1,1\}^n \to \mathbb{R}$ that satisfies $\left|p(x) - f(x)\right| \leq \epsilon$ for all $x \in \{-1,1\}^n$ for which $f(x) \in \{-1,1\}$.*[3] *That is,*

$$\widetilde{\deg}_\epsilon(f) := \min\{d : \deg(p) \leq d, \left|p(x) - f(x)\right| \leq \epsilon \text{ for all } x \in \{-1,1\}^n \text{ for which } f(x) \in \{-1,1\}\}.$$

Henceforth, we will use the notation $\widetilde{\deg}(f)$ to denote $\widetilde{\deg}_{1/3}(f)$.

Definition 2.9 (Approximate Spectral Norm)**.** *The* approximate spectral norm *of a function $f : \{-1,1\}^n \to \{-1,1,\star\}$, denoted by $\|\widehat{f}\|_{1,\epsilon}(f)$ is defined to be the minimum spectral norm of a real polynomial $p : \{-1,1\}^n \to \mathbb{R}$ that satisfies $\left|p(x) - f(x)\right| \leq \epsilon$ for all $x \in \{-1,1\}^n$ for which $f(x) \in \{-1,1\}$.*

$$\|\widehat{f}\|_{1,\epsilon}(f) := \min\{\|\widehat{p}\|_1 : \left|p(x) - f(x)\right| \leq \epsilon \text{ for all } x \in \{-1,1\}^n \text{ for which } f(x) \in \{-1,1\}\}.$$

Lemma 2.10 ([BNRdW07])**.** *Let $f : \{-1,1\}^n \to \{-1,1\}$ be a total function. Then for all constants $0 < \delta, \epsilon < 1$ we have $\widetilde{\deg}_\epsilon(f) = \Theta(\widetilde{\deg}_\delta(f))$.*

The following is a standard upper bound on the approximate spectral norm of a Boolean function in terms of its approximate degree.

Claim 2.11. *For all total functions $f : \{-1,1\}^n \to \{-1,1\}$, we have*

$$\log\|\widehat{f}\|_{1,1/3}(f) = O(\widetilde{\deg}(f) \log n).$$

Proof. Let d denote the approximate degree of f. Take any $1/3$-approximating polynomial of degree d, say p, to f. Then,

$$\sum_{S \subseteq [n]} \left|\widehat{p}(S)\right| \leq \sqrt{\binom{n}{\leq d}} \cdot \sqrt{\sum_{S:|S|\leq d} \widehat{p}(S)^2} \leq 4/3 \cdot (n+1)^{d/2} = 2^{O(d \log n)},$$

where the first inequality follows by the Cauchy-Schwarz inequality, the second inequality follows by Parseval's identity (Fact 2.1) and the fact that the absolute value of p is at most $4/3$ for any input $x \in \{-1,1\}^n$. $\square$

It is easy to exhibit functions $f : \{-1,1\}^n \to \{-1,1\}$ such that $\log\|\widehat{f}\|_{1,1/3}(f) = \Omega(\widetilde{\deg}(f))$. Bent functions satisfy this bound, for example.

Building upon ideas in [KP97], the approximate spectral norm of $f \circ \mathsf{IND}_1$ was shown to be bounded below by $2^{\Omega(\widetilde{\deg}(f))}$ in [CM17].

Theorem 2.12 ([CM17])**.** *Let $f : \{-1,1\}^n \to \{-1,1\}$ be any function. Then $\|\widehat{f}\|_{1,1/3}(f \circ \mathsf{IND}_1) \geq 2^{c \cdot \widetilde{\deg}_{2/3}(f)}$ for any constant $c < 1 - 3/\widetilde{\deg}_{2/3}(f)$.*

---

[3]When dealing with partial functions, another notion of approximation is sometimes considered, where the approximating polynomial p is required to have bounded values even on the non-promise inputs of f. For the purpose of this paper, we do not require this constraint.

2.3 Communication complexity

The classical model of communication complexity was introduced by Yao in [Yao79]. In this model two parties, say Alice and Bob, wish to compute a function whose output depends on both their inputs. Alice is given an input $x \in \mathcal{X}$, Bob is given $y \in \mathcal{Y}$, and they want to jointly compute the value of a given function $F(x, y)$ by communicating with each other. Alice and Bob individually have unbounded computational power and the number of bits communicated is the resource we wish to minimize. Alice and Bob communicate using a *protocol* that is agreed upon in advance. In the randomized model, Alice and Bob have access to unlimited public random bits and the goal is to compute the correct value of $F(x, y)$ with probability at least 2/3 for all inputs $(x, y) \in \mathcal{X} \times \mathcal{Y}$. The *bounded-error randomized communication complexity* of a function F, denoted $R^{cc}(F)$, is the number of bits that must be communicated in the worst case by any randomized protocol to compute the correct value of the function $F(x, y)$, with probability at least 2/3, for every $(x, y) \in \mathcal{X} \times \mathcal{Y}$.

The quantum model of communication complexity was introduced by Yao in [Yao93]. We refer the reader to the survey [dW02] for details. The *bounded-error quantum communication complexity* of a function F, denoted $Q^{cc}(F)$ is the number of bits that must be communicated by any quantum communication protocol in the worst case to compute the correct value of the function $F(x, y)$, with probability at least 2/3, for every (x, y) in domain of F. Buhrman, Cleve and Wigderson [BCW98] observed a quantum simulation theorem, which gives an upper bound on the bounded-error quantum communication complexity of a composed function of the form $f \circ \wedge$ or $f \circ \oplus$ in terms of the bounded-error quantum query complexity of f (see Theorem 1.1).

Lee and Shraibman [LS09] showed that the bounded-error quantum communication complexity of $f \circ \oplus$ is bounded below by the logarithm of the approximate spectral norm of f. Also see [CM17] for an alternate proof.

Theorem 2.13 ([LS09]). *For any Boolean function $f : \{-1, 1\}^n \to \{-1, 1\}$,*

$$Q^{cc}(f \circ \oplus) = \Omega(\log \|\widehat{f}\|_{1,1/3}).$$

3 Proof of Theorem 1.3

In this section, we prove Theorem 1.3. We first formally define the function we use.

3.1 Definition of the function

If $\delta < 1/2$, then ignore the last $n - 2n^{2\delta}$ bits of the input, and define the following function on the first $2n^\delta$ bits of the input. The same argument as in Sections 3.2 and 3.3 give the required bounds for Theorem 1.3 and Theorem 1.4. Hence, we may assume without loss of generality that $\delta \geq 1/2$.

Define the partial function $f : \{-1, 1\}^n \to \{-1, 1, \star\}$ by $f = \mathsf{PARITY}_{n^\delta/2}^{\mathsf{HADD}_{n^{1-\delta}}}$. Define F to be the completion of f that evaluates to -1 on the non-promise domain of f (see Figure 1.2).

3.2 Upper bound

In this section, we prove the following.

Claim 3.1. *For $F : \{-1, 1\}^n \to \{-1, 1\}$ defined as in Section 3.1, we have*

$$Q(F) = \Theta(n^\delta).$$

The upper bound follows along the lines of a proof in [AdW14], and the lower bound just uses the fact that F is at least as hard as $\mathsf{PARITY}_{n^\delta/2}$.

Proof. The following is an $O(n^\delta)$-query quantum algorithm computing F. For convenience, set $\ell = n^{1-\delta}$ and $k = n^\delta$. Note that since $\delta \geq 1/2$, we have $k = \Omega(\ell)$.

1. Run $k/2$ instances of Bernstein-Vazirani algorithm on inputs $(x_{11}, \ldots, x_{1\ell}), \ldots, (x_{\frac{k}{2}1}, \ldots, x_{\frac{k}{2}\ell})$ to obtain $k/2$ strings $z_{i_1}, \ldots, z_{i_{k/2}}$.

2. Run Grover's search [Gro96, BHMT02] to check equality of the two strings: $h(z_{i_1}), \ldots, h(z_{i_\ell})$ and $x_{11}, \ldots, x_{1\ell}, \ldots, x_{\frac{k}{2}1}, \ldots, x_{\frac{k}{2}\ell}$, i.e. to check whether the addressing bits of the input are indeed all Hadamard codewords which are output by the first step.

3. If the step above outputs that the strings are equal, then query the $k/2$ selected variables and output their parity. Else, output -1.

   - If the input was indeed of the form as claimed in the first step, then Bernstein-Vazirani outputs the correct $z_{i_1}, \ldots, z_{i_\ell}$ with probability 1, and Grover's search verifies that the strings are equal with probability 1. Hence the algorithm is correct with probability 1 in this case.

   - If the input was not of the claimed form, then the two strings for which equality is to be checked in the second step are not equal. Grover's search catches a discrepancy with probability at least $2/3$. Hence, the algorithm is correct with probability at least $2/3$ in this case.

The correctness of the algorithm is argued above, and the cost is $k/2$ queries for the first step, $O(\sqrt{k\ell})$ queries for the second step, and at most $k/2$ for the third step. Thus, we have

$$Q(F) = O(k + \sqrt{k\ell}) = O(k),$$

since $k = \Omega(\ell)$. The upper bound in the lemma follows.

For the lower bound, we argue that F is at least as hard as Parity on $k/2$ inputs. To see this formally, set all the address variables such that the selected target variables are the first target variable in each block. Under this restriction, F equals $\mathsf{PARITY}(y_{11}, \ldots, y_{\frac{k}{2}1})$. Thus any quantum query algorithm computing F must be able to compute $\mathsf{PARITY}_{k/2}$, and thus $Q(F) \geq k/2$. □

Remark 3.2. *The same argument as above works when the function f is defined to be g^{HADD_ℓ} for any $g : \{-1, 1\}^{n^\delta} \to \{-1, 1\}$ satisfying $\widetilde{\deg}(g) = \Omega(n^\delta)$, and F is the completion of f that evaluates to -1 on all non-promise inputs. The same proof of Theorem 1.3 also goes through, but we fix $g = \mathsf{PARITY}_{n^\delta/2}$ for convenience.*

3.3 Lower bound

In this section, we first prove Lemma 1.5. We require the following observation.

Observation 3.3. *For any monomial χ_S, $S \subseteq [n]$ and any $j \in S$, we have*

$$\mathop{\mathbb{E}}_{x_j \sim \{-1,1\}} [\chi_S(x)] = 0,$$

where x_j is distributed uniformly over $\{-1, 1\}$.

Proof of Lemma 1.5. Let $F = f^{\mathsf{ADDR}_{m,t}}$. Recall that our goal is to show that $\log\|\widehat{F}\|_{1,1/3} = \Omega(\widetilde{\deg}(f)\log t)$. We may assume $\widetilde{\deg}(f) \geq 1$, because the lemma is trivially true otherwise.

Towards a contradiction, suppose there exists a polynomial P of spectral norm strictly less than $\left(\frac{1}{10}\widetilde{\deg}_{0.99}(f)\log t\right)$ uniformly approximating F to error $1/3$ on the promise inputs (recall that from Lemma 2.10, we have $\widetilde{\deg}(f) = \Theta(\widetilde{\deg}_{0.99}(f))$).

Let ν be a distribution on the address bits of $\mathsf{ADDR}_{m,t}$ such that ν is supported only on assignments to the address variables that do not select $\star$, and is the uniform distribution over these assignments. Let $\mu = \nu^n$ be the product distribution over the address bits of the addressing functions in F.

- For any assignment z of the address variables from the support of μ, define a relevant (target) variable to be one that is selected z. Analogously, define a target variable to be irrelevant if it is not selected by z. Define a monomial to be relevant if it does not contain irrelevant variables, and irrelevant otherwise.

- Note that for any target variable, the probability with which it is selected is exactly $1/t$.

- Thus under any assignment z drawn from μ, for any monomial of the restricted function $P|_z$ of degree $t \geq \widetilde{\deg}_{0.99}(f)$, the probability that it is relevant is at most $1/t^{\widetilde{\deg}_{0.99}(f)}$.

- Hence,

$$\mathbb{E}_{z\sim\mu}[\ell_1\text{-norm of relevant monomials in } P|_z \text{ of degree } \geq \widetilde{\deg}_{0.99}(f)]$$

$$= \sum_{|S|\geq\widetilde{\deg}_{0.99}(f)} |w_S| \Pr_{z\sim\mu}[\chi_S \text{ is relevant}]$$

$$\leq \Pr_{z\sim\mu}[\chi_S \text{ is relevant}] \cdot \|\widehat{P}\|_1$$

$$< \frac{1}{t^{\widetilde{\deg}_{0.99}(f)}} \cdot 2^{\frac{1}{10}\widetilde{\deg}_{0.99}(f)\log t}$$

$$= 2^{(-\frac{9}{10})\widetilde{\deg}_{0.99}(f)\log t} < \frac{3}{5},$$

where the last inequality holds because $t \geq 2$ and $\widetilde{\deg}_{0.99}(f) \geq 1$.

- Fix an assignment to the address variables from the support of μ such that under this assignment, the ℓ_1-norm of the relevant monomials in P of degree $\geq \widetilde{\deg}_{0.99}(f)$ is less than $3/5$.

- Note that under this assignment (in fact under any assignment in the support of μ), the restricted F is just the function f on the n variables selected by the addressing functions. Denote by P_1 the polynomial on the target variables obtained from P by fixing address variables as per this assignment.

- Drop the relevant monomials of degree $\geq \widetilde{\deg}_{0.99}(f)$ from P_1 to get a polynomial P_2, which uniformly approximates the restricted F (which is f on n variables) to error $1/3+3/5 < 0.99$.

- Take expectation over irrelevant variables (from the distribution where each irrelevant variable independently takes values uniformly from $\{-1,1\}$). Under this expectation, the value of F does not change (since irrelevant variables do not affect F's output by definition), and all

11

irrelevant monomials of P_2 become 0 (using Observation 3.3 and linearity of expectation). Hence, under this expectation we have $\mathbb{E}[P_2] = P_3$, where P_3 is a polynomial of degree strictly less than $\widetilde{\deg}_{0.99}(f)$. Furthermore, P_3 uniformly approximates f to error less than 0.99 which is a contradiction.

$\square$

As a corollary of Lemma 1.5, we obtain a lower bound on the approximate spectral norm of F, where F is defined as in Section 3.1. This yields a proof of Theorem 1.4.

Proof of Theorem 1.4. Given $0 < \delta < 1$, construct F as in Section 3.1. Claim 3.1 implies

$$Q(F) = \Theta(n^\delta).$$

Let $f = \mathsf{PARITY}_{n^\delta/2}^{\mathsf{HADD}_{n^{1-\delta}}}$. Lemma 1.5 implies that

$$\|\widehat{f}\|_{1,1/3} = \Omega(n^\delta \log n).$$

Since F is a completion of f, we have

$$\|\widehat{F}\|_{1,1/3} = \Omega(n^\delta \log n),$$

which proves the lower bound in Theorem 1.4. The upper bound follows from Theorem 1.1. $\square$

We are now ready to prove our main theorem.

Proof of Theorem 1.3. It immediately follows from Theorem 1.4 and Theorem 2.13. $\square$

4 Conclusions

We conclude with the following points: first, we find our main result somewhat surprising that simulating a query algorithm by a communication protocol in the quantum context has a larger overhead than in the classical context. Second, it is remarkable that this relatively fine overhead of $\log n$ can be detected using analytic techniques that are an adaptation of the generalized discrepancy method. Third, the function that we used in this work is an XOR function. Study of this class of functions is proving to be very insightful. A recent example is the refutation of the log-approximate-rank conjecture [CMS19] and even its quantum version [ABT18, SdW18]. Our work further advocates the study of this rich class.

An open question that remains is whether there exists a Boolean function $F : \{-1, 1\}^n \to \{-1, 1\}$ such that $Q^{cc}(F \circ \wedge) = \Omega(Q(F) \log n)$. Or does there exist a better quantum communication protocol for $(F \circ \wedge)$ that does not incur the logarithmic factor loss?

Along with the fact that $\widetilde{\deg}(f) \leq 2Q(f)$ [BBC$^+$01], Theorem 1.4 yields the following corollary.

Corollary 4.1. *For any constant $0 < \delta < 1$, there exists a total function $F : \{-1, 1\}^n \to \{-1, 1\}$ for which $\widetilde{\deg}(F) = O(n^\delta)$ and*

$$\log \|\widehat{F}\|_{1,1/3} = \Omega(\widetilde{\deg}(F) \log n).$$

It is easy to verify that the constructions of F that yield Theorem 1.4 for any fixed constant $0 < \delta < 1$, also satisfy $\widetilde{\deg}(F) = \Theta(n^\delta)$. Thus, Corollary 4.1 also gives a negative answer to Open Problem 2 in [ACK$^+$18, Section 6], where it was asked if any degree-d approximating polynomial to a Boolean function of approximate degree d has weight at most $2^{O(d)}$. Thus to prove min-entropy of the Fourier spectrum of a Boolean function is upper bounded by approximate degree, it cannot follow from their observation that min-entropy is upper bounded by the logarithm of the approximate spectral norm. It remains an interesting and important open problem: (how) can one prove that the min-entropy of the Fourier spectrum of a Boolean function is upper bounded by a constant multiple of its approximate degree? This inequality is an implication of the Fourier Entropy Influence (FEI) Conjecture.

5 Acknowledgements

References

[AA05]   Scott Aaronson and Andris Ambainis. Quantum search of spatial regions. *Theory of Computing*, 1(1):47–79, 2005.

[ABT18]   Anurag Anshu, Naresh Goud Boddu, and Dave Touchette. Quantum log-approximate-rank conjecture is also false. *arXiv preprint arXiv:1811.10525,* to appear in *60th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2019*, 2018.

[ACK$^+$18]   Srinivasan Arunachalam, Sourav Chakraborty, Michal Koucký, Nitin Saurabh, and Ronald de Wolf. Improved bounds on fourier entropy and min-entropy. *CoRR*, abs/1809.09819, 2018.

[AdW14]   Andris Ambainis and Ronald de Wolf. How low can approximate degree and quantum query complexity be for total boolean functions? *Computational Complexity*, 23(2):305–322, 2014.

[AFH12]   Anil Ada, Omar Fawzi, and Hamed Hatami. Spectral norm of symmetric functions. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 15th International Workshop, APPROX 2012, and 16th International Workshop, RANDOM 2012, Cambridge, MA, USA, August 15-17, 2012. Proceedings*, pages 338–349, 2012.

[BBC$^+$01]   Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001.

[BCW98]   Harry Buhrman, Richard Cleve, and Avi Wigderson. Quantum vs. classical communication and computation. In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 63–68, 1998.

[BHMT02]   Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. *Contemporary Mathematics*, 305:53–74, 2002.

[BNRdW07] Harry Buhrman, Ilan Newman, Hein Röhrig, and Ronald de Wolf. Robust polynomials and quantum algorithms. *Theory Comput. Syst.*, 40(4):379–395, 2007.

[BV97] Ethan Bernstein and Umesh V. Vazirani. Quantum complexity theory. *SIAM J. Comput.*, 26(5):1411–1473, 1997.

[CM17] Arkadev Chattopadhyay and Nikhil S. Mande. A lifting theorem with applications to symmetric functions. In *37th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2017, December 11-15, 2017, Kanpur, India*, pages 23:1–23:14, 2017.

[CMS19] Arkadev Chattopadhyay, Nikhil S. Mande, and Suhail Sherif. The log-approximate-rank conjecture is false. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019.*, pages 42–53, 2019.

[dW02] Ronald de Wolf. Quantum communication and complexity. *Theor. Comput. Sci.*, 287(1):337–353, 2002.

[FK96] Ehud Friedgut and Gil Kalai. Every monotone graph property has a sharp threshold. *Proceedings of the American Mathematical Society*, 124(10):2993–3002, 1996.

[Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 212–219, 1996.

[HdW02] Peter Høyer and Ronald de Wolf. Improved quantum communication complexity bounds for disjointness and equality. In *STACS 2002, 19th Annual Symposium on Theoretical Aspects of Computer Science, Antibes - Juan les Pins, France, March 14-16, 2002, Proceedings*, pages 299–310, 2002.

[KP97] Matthias Krause and Pavel Pudlák. On the computational power of depth-2 circuits with threshold and modulo gates. *Theor. Comput. Sci.*, 174(1-2):137–156, 1997.

[LS09] Troy Lee and Adi Shraibman. Lower bounds in communication complexity. *Foundations and Trends in Theoretical Computer Science*, 3(4):263–398, 2009.

[MNR11] Ashley Montanaro, Harumichi Nishimura, and Rudy Raymond. Unbounded-error quantum query complexity. *Theor. Comput. Sci.*, 412(35):4619–4628, 2011.

[SdW18] Makrand Sinha and Ronald de Wolf. Exponential separation between quantum communication and logarithm of approximate rank. *arXiv preprint arXiv:1811.10090,* to appear in *60th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2019*, 2018.

[Yao79] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *Proceedings of the 11h Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1979, Atlanta, Georgia, USA*, pages 209–213, 1979.

[Yao93] Andrew Chi-Chih Yao. Quantum circuit complexity. In *34th Annual Symposium on Foundations of Computer Science, Palo Alto, California, USA, 3-5 November 1993*, pages 352–361, 1993.

[Zha09]   Shengyu Zhang. On the tightness of the Buhrman-Cleve-Wigderson simulation. In *Algorithms and Computation, 20th International Symposium, ISAAC 2009, Honolulu, Hawaii, USA, December 16-18, 2009. Proceedings*, pages 434–440, 2009.