

Improved Quantum Communication Complexity Bounds for Disjointness and Equality

Peter Høyer^{1,*} and Ronald de Wolf^{2,**}

¹ Dept. of Comp. Sci., Univ. of Calgary, AB, Canada. hoyer@cpsc.ucalgary.ca

² UC Berkeley. 583 Soda Hall, Berkeley CA 94720, USA. rdewolf@cs.berkeley.edu

Abstract. We prove new bounds on the quantum communication complexity of the disjointness and equality problems. For the case of exact and non-deterministic protocols we show that these complexities are all equal to $n+1$, the previous best lower bound being $n/2$. We show this by improving a general bound for non-deterministic protocols of de Wolf. We also give an $O(\sqrt{n} \cdot c^{\log^+ n})$ -qubit bounded-error protocol for disjointness, modifying and improving the earlier $O(\sqrt{n} \log n)$ protocol of Buhrman, Cleve, and Wigderson, and prove an $\Omega(\sqrt{n})$ lower bound for a class of protocols that includes the BCW-protocol as well as our new protocol.

1 Introduction

The area of *communication complexity* deals with abstracted models of distributed computing, where one only cares about minimizing the amount of communication between the parties and not about the amount of computation done by the individual parties. The standard setting is the following. Two parties, Alice and Bob, want to compute some function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$. Alice receives input $x \in \{0, 1\}^n$, Bob receives $y \in \{0, 1\}^n$, and they want to compute $f(x, y)$. For example, they may want to find out whether $x = y$ (the *equality* problem) or whether x and y are characteristic vectors of disjoint sets (the *disjointness* problem). A communication protocol is a distributed algorithm where Alice first does some computation on her side, then sends a message to Bob, who does some computation on his side, sends a message back, etc. The *cost* of the protocol is measured by the number of bits (or qubits, in the quantum case) communicated on a worst-case input (x, y) .

As in many other branches of complexity theory, we can distinguish between various different “modes” of computation. Letting $P(x, y)$ denote the acceptance probability of the protocol (the probability of outputting 1), we consider four different kinds of protocols for computing f ,

- An *exact* protocol has $P(x, y) = f(x, y)$, for all x, y

* Supported in part by Canada’s NSERC and the Pacific Institute for the Mathematical Sciences.

** Supported by Talent grant S 62–565 from the Netherlands Organization for Scientific Research. Work conducted while at CWI, Amsterdam, partially supported by EU fifth framework project QAIP, IST–1999–11234.

- A *non-deterministic* protocol has $P(x, y) > 0$ if and only if $f(x, y) = 1$, for all x, y
- A *one-sided error* protocol has $P(x, y) \geq 1/2$ if $f(x, y) = 1$, and $P(x, y) = 0$ if $f(x, y) = 0$
- A *two-sided error* protocol has $|P(x, y) - f(x, y)| \leq 1/3$, for all x, y .

These four modes of computation correspond to those of the computational complexity classes P, NP, RP, and BPP, respectively.

Protocols may be *classical* (send and process classical bits) or *quantum* (send and process quantum bits). Classical communication complexity was introduced by Yao [35], and has been studied extensively. It is well motivated by its intrinsic interest as well as by its applications in lower bounds on circuits, VLSI, data structures, etc. We refer to the book of Kushilevitz and Nisan [26] for definitions and results. We use $D(f)$, $N(f)$, $R_1(f)$, and $R_2(f)$ to denote the minimal cost of classical protocols for f in the exact, non-deterministic, one-sided error, and two-sided error settings, respectively.¹ Note that $R_2(f) \leq R_1(f) \leq D(f) \leq n+1$ and $N(f) \leq R_1(f) \leq D(f) \leq n+1$ for all f . Similarly we define $Q_E(f)$, $NQ(f)$, $Q_1(f)$, and $Q_2(f)$ for the quantum versions of these communication complexities (we will be a bit more precise about the notion of a quantum protocol in the next section). For all of these complexities, we assume Alice and Bob start out without any shared randomness or entanglement.

Quantum communication complexity was introduced by (again) Yao [36] and the first examples of functions where quantum communication complexity is less than classical communication complexity were given in [14, 10, 11, 15, 9]. In particular, Buhrman, Cleve, and Wigderson [9] showed for a specific *promise version* of the equality problem that $Q_E(f) \in O(\log n)$ while $D(f) \in \Omega(n)$. They also showed for the *intersection* problem (the negation of the disjointness problem) that $Q_1(\text{INT}_n) \in O(\sqrt{n} \log n)$, whereas $R_2(\text{INT}_n) \in \Omega(n)$ is a well known and non-trivial result from classical communication complexity [20, 31]. Later, Raz [30] exhibited a promise problem with an exponential quantum-classical gap even in the bounded-error setting: $Q_2(f) \in O(\log n)$ versus $R_2(f) \in \Omega(n^{1/4}/\log n)$. Other results on quantum communication complexity may be found in [25, 2, 28, 13, 21, 34, 24, 23].

The aim of this paper is to sharpen the bounds on the quantum communication complexities of the equality and disjointness (or intersection) problems, in the four modes we distinguished above. We summarize what was known prior to this paper,

- $n/2 \leq Q_1(\text{EQ}_n), Q_E(\text{EQ}_n) \leq n+1$ [25, 13]
- $n/2 \leq NQ(\text{EQ}_n) \leq n+1$ [34]
- $Q_2(\text{EQ}_n) \in \Theta(\log n)$ [25]
- $n/2 \leq Q_1(\text{DISJ}_n), Q_E(\text{DISJ}_n) \leq n+1$ [25, 13]
- $n/2 \leq NQ(\text{DISJ}_n) \leq n+1$ [34]
- $\log n \leq Q_1(\text{INT}_n), Q_2(\text{DISJ}_n) \in O(\sqrt{n} \log n)$ [9].

¹ Kushilevitz and Nisan [26] use $N^1(f)$ for our $N(f)$, $R^1(f)$ for our $R_1(f)$ and $R(f)$ for our $R_2(f)$.

In Section 3 we first sharpen the non-deterministic bounds, by proving a general algebraic characterization of $NQ(f)$. In [34] it was shown for all functions f that

$$\frac{\log nrank(f)}{2} \leq NQ(f) \leq \log(nrank(f)) + 1,$$

where $nrank(f)$ denotes the rank of a “non-deterministic matrix” for f (to be defined more precisely below). It is interesting to note that in many places in quantum computing one sees factors of $\frac{1}{2}$ appearing that are essential, for example in the query complexity of parity [4, 17], in the bounded-error query complexity of all functions [16], in superdense coding [5], and in lower bounds for entanglement-enhanced quantum communication complexity [13, 28]. In contrast, we show here that the $\frac{1}{2}$ in the above lower bound can be dispensed with, and the upper bound is tight,²

$$NQ(f) = \log(nrank(f)) + 1.$$

Equality and disjointness both have non-deterministic rank 2^n , so their non-deterministic complexities are maximal: $NQ(EQ_n) = NQ(DISJ_n) = n + 1$. (This contrasts with their complements: $NQ(NEQ_n) = 2$ [27] and $NQ(INT_n) \leq N(INT_n) = \log n + 1$.) Since $NQ(f)$ lower bounds $Q_1(f)$ and $Q_E(f)$, we also obtain optimal bounds for the one-sided and exact quantum communication complexities of equality and disjointness. In particular, $Q_E(EQ_n) = n + 1$, which answers a question posed to one of us (RdW) by Gilles Brassard in December 2000.

The two-sided error bound $Q_2(EQ_n) \in \Theta(\log n)$ is easy to show, whereas the two-sided error complexity of disjointness is still wide open. In Section 4 we give a one-sided error protocol for the intersection problem that improves the $O(\sqrt{n} \log n)$ protocol of Buhrman, Cleve, and Wigderson by nearly a log-factor,

$$Q_1(INT_n) \in O(\sqrt{n} \cdot c^{\log^* n}),$$

where c is a (small) constant. The function $\log^* n$ is defined as the minimum number of iterated applications of the logarithm function necessary to obtain a number less than or equal to 1: $\log^* n = \min\{r \geq 0 \mid \log^{(r)} n \leq 1\}$, where $\log^{(0)}$ is the identity function and $\log^{(r)} = \log \circ \log^{(r-1)}$. Even though $c^{\log^* n}$ is exponential in $\log^* n$, it is still very small in n , in particular $c^{\log^* n} \in o(\log^{(r)} n)$ for every constant $r \geq 1$. It should be noted that our protocol is asymptotically somewhat more efficient than the BCW-protocol ($\sqrt{n}c^{\log^* n}$ versus $\sqrt{n} \log n$), but is also more complicated to describe; it is based on a recursive modification of the BCW-protocol, an idea that previously has been used for claw-finding by Buhrman *et al.* [12, Section 5].

Proving good *lower* bounds on the Q_2 -complexity of the disjointness and intersection problems is one of the main open problems in quantum communication complexity. Only logarithmic lower bounds are known so far for general

² Similarly we can improve the query complexity result $ndeg(f)/2 \leq NQ_q(f) \leq ndeg(f)$ of [34] to the optimal $NQ_q(f) = ndeg(f)$.

protocols [25, 2, 13]. A lower bound of $\Omega(n^{1/k})$ is shown in [24] for protocols exchanging at most $k \in O(1)$ messages. In Section 4.1 we prove a nearly tight lower bound of $\Omega(\sqrt{n})$ qubits of communication for all protocols that satisfy the constraint that their acceptance probability is a function of $x \wedge y$ (the n -bit AND of Alice’s x and Bob’s y), rather than of x and y “separately.” Since DISJ_n itself is also a function only of $x \wedge y$, this does not seem to be an extremely strong constraint. The constraint is satisfied by a class of protocols that includes the BCW-protocol and our new protocol. It seems plausible that the general bound is $Q_2(\text{DISJ}_n) \in \Omega(\sqrt{n})$ as well, but we have so far not been able to weaken the constraint that the acceptance probability is a function of $x \wedge y$.

2 Preliminaries

2.1 Quantum Computing

Here we briefly sketch the setting of quantum computation, referring to the book of Nielsen and Chuang [29] for more details. An m -qubit quantum *state* $|\phi\rangle$ is a superposition or linear combination over all classical m -bit states,

$$|\phi\rangle = \sum_{i \in \{0,1\}^m} \alpha_i |i\rangle,$$

with the constraint that $\sum_i |\alpha_i|^2 = 1$. Equivalently, $|\phi\rangle$ is a unit vector in \mathbb{C}^{2^m} . Quantum mechanics allows us to change this state by means of unitary (i.e., norm-preserving) operations: $|\phi_{\text{new}}\rangle = U|\phi\rangle$, where U is a $2^m \times 2^m$ unitary matrix. A *measurement* of $|\phi\rangle$ produces the outcome i with probability $|\alpha_i|^2$, and then leaves the system in the state $|i\rangle$.

The two main examples of quantum algorithms so far, are Shor’s algorithm for factoring n -bit numbers using a polynomial number (in n) of elementary unitary transformations [32] and Grover’s algorithm for searching an unordered n -element space using $O(\sqrt{n})$ “look-ups” or queries in the space [18]. Below we use a technique called *amplitude amplification*, which generalizes Grover’s algorithm.

Theorem 1 (Amplitude amplification [7]). *There exists a quantum algorithm **QSearch** with the following property. Let \mathcal{A} be any quantum algorithm that uses no measurements, and let $\chi : \{1, \dots, n\} \rightarrow \{0, 1\}$ be any Boolean function. Let a denote the initial success probability of \mathcal{A} of finding a solution (i.e., the probability of outputting some $i \in \{1, \dots, n\}$ so that $\chi(i) = 1$). Algorithm **QSearch** finds a solution using an expected number of $O\left(\frac{1}{\sqrt{a}}\right)$ applications of \mathcal{A} , \mathcal{A}^{-1} , and χ if $a > 0$, and it runs forever if $a = 0$.*

Consider the problem of searching an unordered n -element space. An algorithm \mathcal{A} that creates a uniform superposition over all $i \in \{1, \dots, n\}$ has success probability $a \geq 1/n$, so plugging this into the above theorem and terminating after $O(\sqrt{n})$ applications gives us an algorithm that finds a solution with probability at least $1/2$ provided there is one, and otherwise outputs ‘no solution’.

2.2 Communication Complexity

For classical communication protocols we refer to [26]. Here we briefly define quantum communication protocols, referring to the surveys [33, 8, 22, 6] for more details. The space in which the quantum protocol works, consists of three parts: Alice's part, the communication channel, and Bob's part (we do not write the dimensions of these spaces explicitly). Initially these three parts contain only 0-qubits,

$$|0\rangle|0\rangle|0\rangle.$$

We assume Alice starts the protocol. She applies a unitary transformation $U_1^A(x)$ to her part and the channel. This corresponds to her initial computation and her first message. The length of this message is the number of channel qubits affected. The state is now

$$(U_1^A(x) \otimes I^B)|0\rangle|0\rangle|0\rangle,$$

where \otimes denotes tensor product, and I^B denotes the identity transformation on Bob's part. Then Bob applies a unitary transformation $U_2^B(y)$ to his part and the channel. This operation corresponds to Bob reading Alice's message, doing some computation, and putting a return-message on the channel. This process goes back and forth for some k messages, so the final state of the protocol on input (x, y) will be (in case Alice goes last)

$$(U_k^A(x) \otimes I^B)(I^A \otimes U_{k-1}^B(y)) \cdots (I^A \otimes U_2^B(y))(U_1^A(x) \otimes I^B)|0\rangle|0\rangle|0\rangle.$$

The total *cost* of the protocol is the total length of all messages sent, on a worst-case input (x, y) . For technical convenience, we assume that at the end of the protocol the output bit is the first qubit on the channel. Thus the acceptance probability $P(x, y)$ of the protocol is the probability that a measurement of the final state gives a '1' in the first channel-qubit. Note that we do not allow intermediate measurements during the protocol. This is without loss of generality: it is well known that such measurements can be postponed until the end of the protocol at no extra communication cost. As mentioned in the introduction, we use $Q_E(f)$, $NQ(f)$, $Q_1(f)$, and $Q_2(f)$ to denote the cost of optimal exact, non-deterministic, one-sided error, and two-sided error protocols for f , respectively.

The following lemma was stated summarily without proof by Yao [36] and in more detail by Kremer [25]. It is key to many of the earlier lower bounds on quantum communication complexity as well as to ours, and is easily proven by induction on ℓ .

Lemma 2 (Yao [36]; Kremer [25]). *The final state of an ℓ -qubit protocol on input (x, y) can be written as*

$$\sum_{i \in \{0,1\}^\ell} |A_i(x)\rangle |i_\ell\rangle |B_i(y)\rangle,$$

where the $A_i(x), B_i(y)$ are vectors (not necessarily of norm 1), and i_ℓ denotes the last bit of the ℓ -bit string i (the output bit).

The acceptance probability $P(x, y)$ of the protocol is the squared norm of the part of the final state that has $i_\ell = 1$. Letting a_{ij} be the 2^n -dimensional complex column vector with the inner products $\langle A_i(x) | A_j(x) \rangle$ as entries, and b_{ij} the 2^n -dimensional column vector with entries $\langle B_i(y) | B_j(y) \rangle$, we can write P (viewed as a $2^n \times 2^n$ matrix) as the sum $\sum_{i,j:i_\ell=j_\ell=1} a_{ij} b_{ij}^T$ of $2^{2\ell-2}$ rank 1 matrices, so the rank of P is at most $2^{2\ell-2}$. For example, for exact protocols this gives immediately that ℓ is lower bounded by $\frac{1}{2}$ times the logarithm of the rank of the communication matrix, and for non-deterministic protocols ℓ is lower bounded by $\frac{1}{2}$ times the logarithm of the non-deterministic rank (defined below). In the next section we show how we can get rid of the factor $\frac{1}{2}$ in the non-deterministic case.

We use $x \wedge y$ to denote the bitwise-AND of n -bit strings x and y , and similarly $x \oplus y$ denotes the bitwise-XOR. Let OR denote the n -bit function which is 1 if at least one of its n input bits is 1, and NOR be its negation. We consider the following three communication complexity problems,

- *Equality*: $\text{EQ}_n(x, y) = \text{NOR}(x \oplus y)$
- *Intersection*: $\text{INT}_n(x, y) = \text{OR}(x \wedge y)$
- *Disjointness*: $\text{DISJ}_n(x, y) = \text{NOR}(x \wedge y)$.

3 Optimal Non-Deterministic Bounds

Let $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$. A $2^n \times 2^n$ complex matrix M is called a *non-deterministic matrix* for f if it has the property that $M_{xy} \neq 0$ if and only if $f(x, y) = 1$ (equivalently, $M_{xy} = 0$ if and only if $f(x, y) = 0$). We use $\text{nrnk}(f)$ to denote the *non-deterministic rank* of f , which is the minimal rank among all non-deterministic matrices for f . In [34] it was shown that

$$\frac{\log \text{nrnk}(f)}{2} \leq \text{NQ}(f) \leq \log(\text{nrnk}(f)) + 1.$$

In this section we show that the upper bound is the true bound. The proof uses the following technical lemma.

Lemma 3. *If there exist two families of vectors $\{A_1(x), \dots, A_m(x)\} \subseteq \mathbb{C}^d$ and $\{B_1(y), \dots, B_m(y)\} \subseteq \mathbb{C}^d$ such that for all $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^n$, we have*

$$\sum_{i=1}^m A_i(x) \otimes B_i(y) = 0 \text{ if and only if } f(x, y) = 0,$$

then $\text{nrnk}(f) \leq m$.

Proof. Assume there exist two such families of vectors. Let $A_i(x)_j$ denote the j th entry of vector $A_i(x)$, and let similarly $B_i(y)_k$ denote the k th entry of vector $B_i(y)$. We use pairs $(j, k) \in \{1, \dots, d\}^2$ to index entries of vectors in the d^2 -dimensional tensor space. Note that

if $f(x, y) = 0$ then $\sum_{i=1}^m A_i(x)_j B_i(y)_k = 0$ for all (j, k) ,
 if $f(x, y) = 1$ then $\sum_{i=1}^m A_i(x)_j B_i(y)_k \neq 0$ for some (j, k) .

As a first step, we want to replace the vectors $A_i(x)$ and $B_i(y)$ by numbers $a_i(x)$ and $b_i(y)$ that have similar properties. We use the probabilistic method [1] to show that this can be done.

Let I be an arbitrary set of 2^{2n+1} numbers. Choose coefficients $\alpha_1, \dots, \alpha_d$ and β_1, \dots, β_d , each coefficient picked uniformly at random from I . For every x , define $a_i(x) = \sum_{j=1}^d \alpha_j A_i(x)_j$, and for every y define $b_i(y) = \sum_{k=1}^d \beta_k B_i(y)_k$. Consider the number

$$v(x, y) = \sum_{i=1}^m a_i(x) b_i(y) = \sum_{j,k=1}^d \alpha_j \beta_k \left(\sum_{i=1}^m A_i(x)_j B_i(y)_k \right).$$

If $f(x, y) = 0$, then $v(x, y) = 0$ for all choices of the α_j, β_k .

Now consider some (x, y) with $f(x, y) = 1$. There is a pair (j', k') for which $\sum_{i=1}^m A_i(x)_{j'} B_i(y)_{k'} \neq 0$. We want to prove that $v(x, y) = 0$ happens only with very small probability. In order to do this, fix the random choices of all α_j , $j \neq j'$, and β_k , $k \neq k'$, and view $v(x, y)$ as a function of the two remaining not-yet-chosen coefficients $\alpha = \alpha_{j'}$ and $\beta = \beta_{k'}$,

$$v(x, y) = c_0 \alpha \beta + c_1 \alpha + c_2 \beta + c_3.$$

Here we know that $c_0 = \sum_{i=1}^m A_i(x)_{j'} B_i(y)_{k'} \neq 0$. There is at most one value of α for which $c_0 \alpha + c_2 = 0$. All other values of α turn $v(x, y)$ into a linear equation in β , so for those α there is at most one choice of β that gives $v(x, y) = 0$. Hence out of the $(2^{2n+1})^2$ different ways of choosing (α, β) , at most $2^{2n+1} + (2^{2n+1} - 1) \cdot 1 < 2^{2n+2}$ choices give $v(x, y) = 0$. Therefore

$$\Pr[v(x, y) = 0] < \frac{2^{2n+2}}{(2^{2n+1})^2} = 2^{-2n}.$$

Using the union bound, we now have

$$\begin{aligned} & \Pr \left[\text{there is an } (x, y) \in f^{-1}(1) \text{ for which } v(x, y) = 0 \right] \\ & \leq \sum_{(x,y) \in f^{-1}(1)} \Pr[v(x, y) = 0] < 2^{2n} \cdot 2^{-2n} = 1. \end{aligned}$$

This probability is strictly less than 1, so there exist sets $\{a_1(x), \dots, a_m(x)\}$ and $\{b_1(y), \dots, b_m(y)\}$ that make $v(x, y) \neq 0$ for every $(x, y) \in f^{-1}(1)$. We thus have that

$$\sum_{i=1}^m a_i(x) b_i(y) = 0 \text{ if and only if } f(x, y) = 0.$$

View the a_i and b_i as 2^n -dimensional vectors, let A be the $2^n \times m$ matrix having the a_i as columns, and B be the $m \times 2^n$ matrix having the b_i as rows. Then $(AB)_{xy} = \sum_{i=1}^m a_i(x) b_i(y)$, which is 0 if and only if $f(x, y) = 0$. Thus AB is a non-deterministic matrix for f , and $nrank(f) \leq rank(AB) \leq rank(A) \leq m$. \square

Lemma 3 allows us to prove tight bounds for non-deterministic quantum protocols.

Theorem 4. $NQ(f) = \log(\text{nrank}(f)) + 1$.

Proof. The upper bound $NQ(f) \leq \log(\text{nrank}(f)) + 1$ was shown in [34] (actually, the upper bound shown there was $\log(\text{nrank}(f))$ for protocols where only Bob has to know the output value). For the sake of completeness we repeat that proof here. Let $r = \text{nrank}(f)$ and M be a rank- r non-deterministic matrix for f . Let $M^T = U\Sigma V$ be the singular value decomposition of the transpose of M [19], so U and V are unitary, and Σ is a diagonal matrix whose first r diagonal entries are positive real numbers and whose other diagonal entries are 0. Below we describe a one-round non-deterministic protocol for f , using $\log(r) + 1$ qubits. First Alice prepares the state $|\phi_x\rangle = c_x \Sigma V|x\rangle$, where $c_x > 0$ is a normalizing real number that depends on x . Because only the first r diagonal entries of Σ are non-zero, only the first r amplitudes of $|\phi_x\rangle$ are non-zero, so $|\phi_x\rangle$ can be compressed into $\log r$ qubits. Alice sends these qubits to Bob. Bob then applies U to $|\phi_x\rangle$ and measures the resulting state. If he observes $|y\rangle$, then he puts 1 on the channel and otherwise he puts 0 there. The acceptance probability of this protocol is

$$P(x, y) = |\langle y|U|\phi_x\rangle|^2 = c_x^2 |\langle y|U\Sigma V|x\rangle|^2 = c_x^2 |M_{yx}^T|^2 = c_x^2 |M_{xy}|^2.$$

Since M_{xy} is non-zero if and only if $f(x, y) = 1$, $P(x, y)$ will be positive if and only if $f(x, y) = 1$. Thus we have a non-deterministic quantum protocol for f with $\log(r) + 1$ qubits of communication.

For the lower bound, consider a non-deterministic ℓ -qubit protocol for f . By Lemma 2, its final state on input (x, y) can be written as

$$\sum_{i \in \{0,1\}^\ell} |A_i(x)\rangle |i_\ell\rangle |B_i(y)\rangle.$$

Without loss of generality we assume the vectors $A_i(x)$ and $B_i(y)$ all have the same dimension d . Let $S = \{i \in \{0,1\}^\ell \mid i_\ell = 1\}$ and consider the part of the state that corresponds to output 1 (we drop the $i_\ell = 1$ and the $|\cdot\rangle$ -notation here),

$$\phi(x, y) = \sum_{i \in S} A_i(x) \otimes B_i(y).$$

Because the protocol has acceptance probability 0 if and only if $f(x, y) = 0$, this vector $\phi(x, y)$ will be the zero vector if and only if $f(x, y) = 0$. The previous lemma gives $\text{nrank}(f) \leq |S| = 2^{\ell-1}$, and hence that $\log(\text{nrank}(f)) + 1 \leq NQ(f)$. \square

Note that any non-deterministic matrix for the equality function has non-zeroes on its diagonal and zeroes off-diagonal, and hence has full rank. Thus $NQ(\text{EQ}_n) = n + 1$, which contrasts sharply with the non-deterministic complexity of its complement (inequality), which is only 2 [27]. Similarly, a non-deterministic matrix for disjointness has full rank, because reversing the ordering of the columns gives an upper triangular matrix with non-zero elements on

the diagonal. This gives tight bounds for the exact, one-sided error, and non-deterministic settings.

Corollary 5. *We have that $Q_E(EQ_n) = Q_1(EQ_n) = NQ(EQ_n) = n + 1$ and that $Q_E(DISJ_n) = Q_1(DISJ_n) = NQ(DISJ_n) = n + 1$.*

4 On the Bounded-Error Complexity of Disjointness

4.1 Improved Upper Bound

Here we show that we can take off most of the $\log n$ factor from the $O(\sqrt{n} \log n)$ protocol for the intersection problem (the complement of disjointness) that was given by Buhrman, Cleve, and Wigderson in [9].

Theorem 6. *There exists a constant c such that $Q_1(INT_n) \in O(\sqrt{n} \cdot c^{\log^* n})$.*

Proof. We recursively build a one-sided error protocol that can find an index i such that $x_i = y_i = 1$, provided such an i exists (call such an i a ‘solution’). Clearly this suffices for computing $INT_n(x, y)$. Let C_n denote the cost of our protocol on n -bit inputs.

Alice and Bob divide the n indices $\{1, \dots, n\}$ into $n/(\log n)^2$ blocks of $(\log n)^2$ indices each. Alice picks a random number $j \in \{1, \dots, n/(\log n)^2\}$ and sends the number j to Bob. Now they recursively run our protocol on the j th block, at a cost of $C_{(\log n)^2}$ qubits of communication. Alice then measures her part of the state, and they verify whether the measured i is indeed a solution. If there is a solution in the j th block, then Alice finds one with probability at least $1/2$, so the overall probability of finding a solution (if there is one) is at least $(\log n)^2/2n$. By using a superposition over all j we can push all intermediate measurements to the end without affecting the success probability. Therefore, applying $O(\sqrt{n}/\log n)$ rounds of amplitude amplification (Theorem 1) boosts this protocol to having error probability at most $1/2$. We thus have the recurrence

$$C_n \leq O(1) \frac{\sqrt{n}}{\log n} (C_{(\log n)^2} + O(\log n)).$$

Since $C_1 = 2$, this recursion unfolds to the bound $C_n \in O(\sqrt{n} \cdot c^{\log^* n})$ for some constant c . Careful inspection of the protocol gives that the constant c is reasonably small. \square

4.2 Lower Bound for a Specific Class of Protocols

We give a lower bound for two-sided error quantum protocols for disjointness. The lower bound applies to all protocols whose acceptance probability $P(x, y)$ is a function just of $x \wedge y$, rather than of x and y ‘separately.’ In particular, the protocols of [9] and of our previous section fall in this class.

The lower bound basically follows by combining various results from [13].

Theorem 7. *Any two-sided error quantum protocol for DISJ_n whose acceptance probability is a function of $x \wedge y$, has to communicate $\Omega(\sqrt{n})$ qubits.*

Proof. Consider an ℓ -qubit protocol with error probability at most $1/3$. By the comment following Lemma 2, we can write its acceptance probability $P(x, y)$ as a $2^n \times 2^n$ matrix P of rank $r \leq 2^{2\ell-2}$.

We now invoke a relation between the rank of the matrix P and properties of the $2n$ -variate multilinear polynomial that equals $P(x, y)$.³ There is an n -variate function g such that $P(x, y) = g(x \wedge y)$. Let $g(z) = \sum_S a_S z_S$ be the polynomial representation of g . Then $P(x, y) = g(x \wedge y) = \sum_S a_S (x \wedge y)_S = \sum_S a_S x_S y_S$, so the $2n$ -variate multilinear polynomial P only contains monomials in which the set of x -variables is the same as the set of y -variables. For polynomials of this form (called “even”), [13, Lemmas 2 and 3] imply that the number of monomials in $P(x, y)$ equals the rank r of the matrix P .

Setting $y = x$ in $P(x, y)$ gives a polynomial $p(x) = \sum_S a_S x_S$ that has r monomials and that approximates the n -bit function NOR, since $|p(x) - \text{NOR}(x)| = |P(x, x) - \text{DISJ}(x, x)| \leq 1/3$. But [13, Theorem 8] implies that every polynomial that approximates NOR must have at least $2\sqrt{n/12}$ monomials. Hence $2\sqrt{n/12} \leq r \leq 2^{2\ell-2}$, which gives $\ell \geq \sqrt{n/48} + 1$. \square

5 Open Problems

This paper fits in a sequence of papers that (slowly) extend what is known for quantum communication complexity, e.g. [9, 2, 30, 13, 21, 34, 24, 23]. The main open question is still the bounded-error complexity of disjointness. Of interest is whether it is possible to prove an $O(\sqrt{n})$ upper bound for disjointness, thus getting rid of the factor of $c^{\log^* n}$ in our upper bound of Theorem 6, and whether it is possible to extend the lower bound of Theorem 7 to broader classes of protocols. Since disjointness is coNP-complete for communication complexity problems [3], strong lower bounds on the disjointness problem imply a host of other lower bounds.

A second question is whether qubit communication can be significantly reduced in case Alice and Bob can make use of prior entanglement (shared EPR-pairs). Giving Alice and Bob n shared EPR-pairs trivializes the non-deterministic complexity (use the EPR-pairs as a public coin to randomly guess some n -bit z , Alice then sends Bob 1 bit indicating whether $x = z$, if $x = z$ then Bob can compute the answer $f(x, y)$ and send it to Alice, if $x \neq z$ then they output 0), but for the exact and bounded-error models it is open whether prior entanglement can make a significant difference.

³ For $S \subseteq [n] = \{1, \dots, n\}$, we use x_S for the monomial $\prod_{i \in S} x_i$. An n -variate multilinear polynomial $p(x) = \sum_{S \subseteq [n]} a_S x_S$, $a_S \in \mathbb{R}$, is a weighted sum of such monomials. The number of monomials in p is the number of S for which $a_S \neq 0$. One can show that for every function $g : \{0, 1\}^n \rightarrow \mathbb{R}$ there is a unique n -variate multilinear polynomial p such that $g(x) = p(x)$ for all $x \in \{0, 1\}^n$.

Acknowledgments

We thank Harry Buhrman and Hartmut Klauck for helpful discussions concerning the proof of Lemma 3, and Richard Cleve for helpful discussions on the protocol for disjointness.

References

1. N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley-Interscience, 1992.
2. A. Ambainis, L. Schulman, A. Ta-Shma, U. Vazirani, and A. Wigderson. The quantum communication complexity of sampling. In *Proceedings of 39th IEEE FOCS*, pages 342–351, 1998.
3. L. Babai, P. Frankl, and J. Simon. Complexity classes in communication complexity theory. In *Proceedings of 27th IEEE FOCS*, pages 337–347, 1986.
4. R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf. Quantum lower bounds by polynomials. In *Proceedings of 39th IEEE FOCS*, pages 352–361, 1998. quant-ph/9802049.
5. C. Bennett and S. Wiesner. Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states. *Physical Review Letters*, 69:2881–2884, 1992.
6. G. Brassard. Quantum communication complexity (a survey). quant-ph/0101005, 1 Jan 2001.
7. G. Brassard, P. Høyer, M. Mosca, and A. Tapp. Quantum amplitude amplification and estimation. quant-ph/0005055. To appear in *Quantum Computation and Quantum Information: A Millennium Volume*, AMS Contemporary Mathematics Series, 15 May 2000.
8. H. Buhrman. Quantum computing and communication complexity. *EATCS Bulletin*, pages 131–141, February 2000.
9. H. Buhrman, R. Cleve, and A. Wigderson. Quantum vs. classical communication and computation. In *Proceedings of 30th ACM STOC*, pages 63–68, 1998. quant-ph/9802040.
10. H. Buhrman, W. van Dam, and R. Cleve. Quantum entanglement and communication complexity. *SIAM Journal on Computing*, 30(6):1829–1841, 2001. quant-ph/9705033.
11. H. Buhrman, W. van Dam, P. Høyer, and A. Tapp. Multiparty quantum communication complexity. *Physical Review A*, 60(4):2737–2741, 1999. quant-ph/9710054.
12. H. Buhrman, Ch. Dürr, M. Heiligman, P. Høyer, F. Magniez, M. Santha, and R. de Wolf. Quantum algorithms for element distinctness. In *Proceedings of 16th IEEE Conference on Computational Complexity*, pages 131–137, 2001. quant-ph/0007016.
13. H. Buhrman and R. de Wolf. Communication complexity lower bounds by polynomials. In *Proceedings of 16th IEEE Conference on Computational Complexity*, pages 120–130, 2001. cs.CC/9910010.
14. R. Cleve and H. Buhrman. Substituting quantum entanglement for communication. *Physical Review A*, 56(2):1201–1204, 1997. quant-ph/9704026.
15. R. Cleve, W. van Dam, M. Nielsen, and A. Tapp. Quantum entanglement and the communication complexity of the inner product function. In *Proceedings of 1st NASA QCQC conference*, volume 1509 of *Lecture Notes in Computer Science*, pages 61–74. Springer, 1998. quant-ph/9708019.

16. W. van Dam. Quantum oracle interrogation: Getting all information for almost half the price. In *Proceedings of 39th IEEE FOCS*, pages 362–367, 1998. quant-ph/9805006.
17. E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser. A limit on the speed of quantum computation in determining parity. *Physical Review Letters*, 81:5442–5444, 1998. quant-ph/9802045.
18. L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of 28th ACM STOC*, pages 212–219, 1996. quant-ph/9605043.
19. R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.
20. B. Kalyanasundaram and G. Schnitger. The probabilistic communication complexity of set intersection. *SIAM Journal on Computing*, 5(4):545–557, 1992.
21. H. Klauck. On quantum and probabilistic communication: Las Vegas and one-way protocols. In *Proceedings of 32nd ACM STOC*, pages 644–651, 2000.
22. H. Klauck. Quantum communication complexity. In *Proceedings of Workshop on Boolean Functions and Applications at 27th ICALP*, pages 241–252, 2000. quant-ph/0005032.
23. H. Klauck. Lower bounds for quantum communication complexity. In *Proceedings of 42nd IEEE FOCS*, pages 288–297, 2001. quant-ph/0106160.
24. H. Klauck, A. Nayak, A. Ta-Shma, and D. Zuckerman. Interaction in quantum communication and the complexity of set disjointness. In *Proceedings of 33rd ACM STOC*, pages 124–133, 2001.
25. I. Kremer. Quantum communication. Master’s thesis, Hebrew University, Computer Science Department, 1995.
26. E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
27. S. Massar, D. Bacon, N. Cerf, and R. Cleve. Classical simulation of quantum entanglement without local hidden variables. *Physical Review A*, 63(5):052305, 2001. quant-ph/0009088.
28. M. A. Nielsen. *Quantum Information Theory*. PhD thesis, University of New Mexico, Albuquerque, 1998. quant-ph/0011036.
29. M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
30. R. Raz. Exponential separation of quantum and classical communication complexity. In *Proceedings of 31st ACM STOC*, pages 358–367, 1999.
31. A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385–390, 1992.
32. P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997. quant-ph/9508027.
33. A. Ta-Shma. Classical versus quantum communication complexity. *ACM SIGACT News (Complexity Column 23)*, pages 25–34, 1999.
34. R. de Wolf. Characterization of non-deterministic quantum query and quantum communication complexity. In *Proceedings of 15th IEEE Conference on Computational Complexity*, pages 271–278, 2000. cs.CC/0001014.
35. A. C.-C. Yao. Some complexity questions related to distributive computing. In *Proceedings of 11th ACM STOC*, pages 209–213, 1979.
36. A. C.-C. Yao. Quantum circuit complexity. In *Proceedings of 34th IEEE FOCS*, pages 352–360, 1993.