SIGACT News Complexity Theory Column 23

Lane A. Hemaspaandra
Dept. of Computer Science, University of Rochester
Rochester, NY 14627, USA lane@cs.rochester.edu

Introduction to Complexity Theory Column 23

Guest Column: Classical versus Quantum Communication Complexity.

Amnon Ta-Shma[1]

Abstract

Classical Communication complexity has been intensively studied since its conception two decades ago. Recently, its quantum analogue, the quantum communication complexity model, was defined and studied. In this paper we present some of the main results in the area.

1 Introduction

The two party communication complexity model, as introduced by Yao [Yao79], consists of two all-powerful players who have private inputs and wish to compute a known function that depends on both inputs. The players follow a predefined protocol, and exchange communication bits until they are ready to make a decision. A simple example is one where Alice holds $x \in \{0,1\}^n$, Bob holds $y \in \{0,1\}^n$ and the two wish to know whether $x = y$. It turns out that any deterministic protocol requires the two parties to communicate at least n bits. If, however, the players can use probabilistic strategies, a protocol with only $O(\log(n))$ communication bits and a small error is possible. There has been a lot of research in this area in the last two decades, and the interested reader is referred to the excellent book by Kushilevitz and Nisan [KN97].

Yao [Yao93] initiated the study of the quantum communication complexity model where the two players have quantum computers and can communicate qubits. The model is mathematically natural and has some unique features that do not appear in the classical analogue. One very interesting problem is determining the biggest possible gap between the classical and quantum models. While this is still open, many partial interesting answers were recently given. These answers also have some surprising connections to quantum computation and quantum information theory.

We define the communication complexity model and the classical and quantum variants we study in Section 2. In Section 3 we present some of the lower bound techniques used in the classical model. In Section 4 we describe an upper bound technique for the quantum model that uses a connection between quantum communication complexity and quantum black box computation to show a quadratic separation between classical and quantum communication complexity. An

[1]International Computer Science Institute, Berkeley, CA 94704. E-mail: amnon@icsi.berkeley.edu.

exponential gap between classical and quantum communication complexity for a *promise* problem follows in Section 5. Finally, in Section 6 we present several lower bound techniques for quantum computation. We also present a generalization of the low-rank lower bound of classical deterministic communication complexity to the probabilistic and quantum model.

2 The Model

2.1 Deterministic Communication Complexity

We explain the model by example:

Example 1 *(The EQ function)*
Input*: Alice holds $x \in \{0,1\}^n$, Bob holds $y \in \{0,1\}^n$.*
Output*: $f(x,y)$ is true iff $x = y$*
Protocol*: Alice sends x to y. Bob checks whether x equals y, and lets Bob know the answer.*

The above deterministic algorithm communicates $n+1$ bits, and we say it has complexity $n+1$. In general, we denote by $D(f)$, the number of communication bits needed for Alice and Bob to deterministically compute the answer $f(x,y)$. Any function f can be solved by Alice communicating her input to Bob, and thus for any f, $D(f) \leq n+1$. There is no way to solve the EQ function better with a deterministic protocol, i.e., $D(EQ) = n+1$

2.2 Probabilistic Communication Complexity

2.2.1 The Private Coin Model

We now present a probabilistic algorithm, due to Rabin and Yao, for the EQ function. Alice picks a random prime $p \in [1..n^3]$ and sends p and $x \ mod \ p$ to Bob. Bob computes $y \ mod \ p$ and the answer is yes iff $y = x(mod \ p)$. The probabilistic algorithm communicates only $O(\log(n))$ bits, and for any input x, y, the probability (over the random coins p) the algorithm is correct is at least $1 - \frac{1}{n}$. In general, we denote by $R_\epsilon(f)$, the complexity of the best probabilistic protocol that for any input x, y computes the right answer $f(x,y)$ with probability at least $1-\epsilon$. When ϵ is an arbitrary constant we shorten this to $R(f)$. The above example shows that $R(EQ) \leq O(log(n))$ which is in fact tight.

As $D(EQ) = n+1$ and $R(EQ) = \Theta(logn)$ we conclude that there is an exponential gap between deterministic and probabilistic communication complexity.

2.2.2 The Public Coin Model

If Alice and Bob share a common random string $r \in \{0,1\}^n$, they can do even better. Alice computes the single bit $\Sigma r_i x_i (mod \ 2)$ and sends it Bob. Bob compares it to $\Sigma r_i y_i (mod \ 2)$ and answers "yes" iff both terms equal. If $x = y$ the protocol always gives the right answer "yes", while if $x \neq y$ the protocol errs with probability (over the random coins r) of at most $\frac{1}{2}$. Using c communication bits the error probability can be reduced to 2^{-c}. We denote the communication complexity of a probabilistic protocol in the public coin model by $R_\epsilon^{pub}(f)$. Although the above example shows a gap between the private coin and public coin models, the two models are almost the same. More specifically,

Theorem 1 *[New91]* $R^{pub}_\epsilon(f) \leq R_{\epsilon+\delta}(f) \leq R^{pub}_\epsilon(f) + O(\log(n) + \log(\frac{1}{\delta}))$.

Thus, except for a small $O(log(n))$ *additive* penalty the two models are the same.

2.3 The Quantum Model

We start with a short review of the quantum model.

The qubit : A qubit is a vector $v \in \mathbb{C}^2$ of length 1. We represent the basis vectors $\{(1,0),(0,1)\}$ of $\mathbb{C}^2$ as $|0\rangle \stackrel{\text{def}}{=} (1,0)$ and $|1\rangle \stackrel{\text{def}}{=} (0,1)$. A vector v can be represented as $v = a|0\rangle + b|1\rangle$ with $a, b \in \mathbb{C}$ and $|a|^2 + |b|^2 = 1$.

A pure state : A k qubit vector is described by a vector $\phi \in \mathbb{C}^{2^k}$ of length 1. We also write this as $\phi = \Sigma_{i=00..0}^{11..1} a_i|i\rangle$, with $a_i \in \mathbb{C}$ and $\Sigma_{i=0}^{2^k-1}|a_i|^2 = 1$. We say ϕ is a pure-state.

Evolution : We can apply a unitary transformation to a system. If the system was in a pure-state ϕ and we applied the unitary transformation U then the new state of the system is $U\phi$.

Measurement : A pure state ϕ lives in C^{2^k}. Suppose C^{2^k} is partitioned into two subspaces $C^{2^k} = W \oplus W^\perp$. A measurement of ϕ according to W projects ϕ on W and $W^\perp$. The result is W ($W^\perp$) with probability that is equal to the projection length of ϕ on W ($W^\perp$).

We give two simple examples:

Example 2 *ϕ is initially $|0\rangle$. We first apply the unitary transformation $U = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$, to get the super-position $U|0\rangle = U\begin{pmatrix} 1 \\ 0 \end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ 1 \end{pmatrix} = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. We then measure according to the basis $\{|0\rangle, |1\rangle\}$. The result is $|0\rangle$ with probability half, and $|1\rangle$ with probability half.*

Example 3 *ϕ is an EPR pair $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. We assume Alice holds the first qubit and Bob holds the second. If Alice measures her qubit according to $\{|0\rangle, |1\rangle\}$ then with probability half both Alice and Bob hold $|0\rangle$, and with probability half both Alice and Bob hold $|1\rangle$. As it turns out ϕ can be also represented as $\frac{1}{\sqrt{2}}(|\uparrow\uparrow\rangle + |\downarrow\downarrow\rangle)$, where $|\uparrow\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $|\downarrow\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$. Thus, if Alice measures her qubit according to the basis $\{|\uparrow\rangle, |\downarrow\rangle\}$ then with probability half both Alice and Bob hold $|\uparrow\rangle$, and with probability half both Alice and Bob hold $|\downarrow\rangle$.*

2.3.1 The Private Coin Model

In the quantum communication complexity model [Yao93] Alice and Bob hold qubits. When the game starts Alice holds x and Bob holds y, and so the initial super-position is simply $|x, y\rangle$. The players take turns. Suppose it is Alice's turn to play. Alice can make an arbitrary unitary transformation on her qubits and then send one or more qubits to Bob. Sending qubits does not change the overall superposition, but rather changes the ownership of the qubits, allowing Bob to apply his next unitary transformation on the newly received qubits. Each player can also measure his/her qubits. By the end of the protocol the two players have to decide on a value.

The complexity of a quantum protocol is the number of qubits exchanged between the two players. We say a quantum protocol *computes* $f : X \times Y \mapsto \{0,1\}$ with at most ϵ error, if for any input x, y the probability that the protocol computes the result $f(x,y)$ is at least $1-\epsilon$. We denote by $Q_\epsilon(f)$ the complexity of the best quantum protocol that computes f with at most ϵ error.

2.3.2 The Public Coin Model

As with the classical model a "public coin" quantum model can be defined [CvDNT98]. While in the classical case the parties share a common random string, in the quantum case the parties share entangled qubits, i.e., Alice and Bob share infinitely many EPR pairs in which Alice holds one qubit from each pair and Bob holds the second qubit of the pair. If Alice measures her part of the EPR pairs, then Alice and Bob get to see a random classical string, and so the model generalizes the probabilistic public coin model. We denote the communication complexity of the quantum public coin model by $Q_\epsilon^{pub}(f)$. Unlike the classical case, not much is known about the relationship between the quantum public coin and the private coin models.

3 Some Lower Bounds for the Classical Model

3.1 The Deterministic Model

Several lower bound techniques exist for the deterministic model. One notable technique is a simple and elegant theorem discovered by Mehlhorn and Schmidt. Given a function $f : \{0,1\}^n \mapsto \{0,1\}$ we define the $2^n \times 2^n$ matrix $M_f[i,j] = f(i,j)$.

Theorem 2 *[MS82] For any function f, $D(f) \geq \log(Rank(M_f))$.*

E.g., for the EQ function, $M_f = I_{2^n \times 2^n}$ which has full rank and hence $D(EQ) \geq n$. It is not known how tight the log-rank bound is. In particular, the famous "Log-Rank" conjecture states that:

Conjecture 1 $D(f) \leq poly(\log(Rank(M_f)))$.

Raz and Spieker [RS95] were the first to show a super-linear gap between $D(f)$ and $\log(Rank(M_f))$ and the biggest gap known as of today, due to Nisan and Wigderson [NW95], exhibits an f with $D(f) \geq \log(Rank(M_f))^{1.6\cdots}$. It is quite possible, for example, that $D(f) \leq \log(Rank(M_f))^2$. We will say a little bit more about this conjecture in Section 6.

3.2 The Probabilistic Model

Lower bounds for the probabilistic model are harder to achieve. Many of the lower bounds are achieved via a reduction to the deterministic case using a min-max principle exhibited by Yao [Yao83]. We define:

Definition 3.1 *Let $f : X \times Y \mapsto \{0,1\}$, and μ be a distribution over $X \times Y$. $D_\epsilon^\mu(f)$, the* distributional complexity *of f, is the complexity of the best* deterministic *protocol P s.t.*

$$\Pr_{(x,y)\in_\mu X \times Y}(P \text{ is correct for } (x,y)) \geq 1-\epsilon$$

where (x,y) are drawn from $X \times Y$ according to the distribution μ.

We stress that P is a *deterministic* protocol, and the probability is taken over the *inputs*.

Example 4 *(The DISJ function)*
Input*: Alice holds $x \subseteq [1..n]$, Bob holds $y \subseteq [1..n]$.*
Output*: $f(x,y)$ is true iff $x \cap y \neq \emptyset$.*
Protocol*: The answer is always "yes".*

The protocol has complexity 0 (no need for any communication) and is right for almost all inputs x, y. Thus, $D_{0.5}^{uniform}(DISJ) = 0$. There are, however, "hard" distributions μ for which $D_\epsilon^\mu(DISJ) = \Omega(n)$ for some constant ϵ.

The key fact is:

Theorem 3 *[Yao83] For any $f : X \times Y \mapsto \{0,1\}$, $R_\epsilon^{pub}(f) = max_\mu D_\epsilon^\mu(f)$.*

This method was used to get lower bounds on the probabilistic communication complexity of many functions, including linear lower bounds for the inner-product function [Vaz85, CG88], $IP(x_1, \ldots, x_n; y_1, \ldots, y_n) = \bigoplus_{i=1}^n x_i \cdot y_i$, and the disjointness function [BFS86, KS87, Raz92].

4 A Quadratic Gap Using Black-Box Simulations

Buhrman, Cleve and Wigderson [BCW98] show a connection between quantum communication complexity and quantum black box computation, and derive a quadratic separation between classical and quantum communication complexity. We start by describing what a black box computation is.

4.1 Black-Box Computation

In a classical black box computation we are given an oracle to access an unknown function $f : \{0,1\}^n \mapsto \{0,1\}$. Our goal is to "learn" the function f. More specifically, we want to compute a property G of f. Our only access to the function is by querying it at specific points. Thus, suppose we want to learn the property

$$OR(f) = \vee_{x \in \{0,1\}^n} f(x)$$

A deterministic black box algorithm has no other way than querying the oracle f at all the 2^n possible inputs. A probabilistic black box algorithm can choose a probabilistic distribution over the queries, and is allowed to make an error with probability at most ϵ. However, even this does not help much for the OR function, and for a constant ϵ error probability one needs to make $\Theta(2^n)$ queries.

A quantum black box computation starts in a fixed state (say the $|00\ldots0\rangle$ super-position) and runs a sequence of unitary transformations and f–gates. An f gate is a unitary transformation that takes $|x\rangle|y\rangle$ into $|x\rangle|f(x) \oplus y\rangle$. Finally, it measures the output qubits to get the result. A quantum black box algorithm computes property G with probability p, if for any function f, the probability (over the protocol) that the output is $G(f)$ is at least p. As an example, a computation may begin with $|00\ldots0\rangle$, get into the super-position $\Sigma_x|x,0\rangle$ apply an f gate to get $\Sigma_x|x, f(x)\rangle$, and so on.

Quantum black box computation can be surprisingly efficient. Grover [Gro96] showed the following remarkable upper bound which was later given a matching lower bound in [BBBV97]:

Theorem 4 *[Gro96, BBBV97] The OR property can computed by a quantum black box algorithm with a small error and using only $\Theta(\sqrt{n})$ queries.*

This shows a quadratic gap between classical and quantum black box complexities. Unfortunately, the gap between classical and quantum black box computation can not be much larger than that.

Theorem 5 *[BBC$^+$98] If there is a quantum black box algorithm for $G(f)$ using t queries, then there is a classical black box algorithm for $G(f)$ using at most $O(t^6)$ queries.*

The picture looks completely different when one wants to learn a property G of a function f on which some partial information is known.

Example 5 *(half or none, [DJ92]) $f : \{0,1\}^n \mapsto \{0,1\}$ is s.t. either*

- *For every x, $f(x) = 0$, or*
- *For exactly half of the inputs x, $f(x) = 0$.*

$G(f) = 0$ if f is constant 0, and 1 otherwise.

Deutsch and Jozsa [DJ92] showed that a quantum black box algorithm using only one (!) query can solve the problem with no error. A deterministic algorithm, however, requires $\Theta(2^n)$ queries. We call such a problem a *promise* problem, because we have some partial information about the input that is promised to us. The [DJ92] functions shows an exponential gap between quantum and classical black box complexity for a promise problem and the no-error case.

4.2 Translating Black-Box Gaps to Communication Complexity Gaps

Buhrman, Cleve and Wigderson [BCW98] showed how to translate any black box quantum computation into a quantum communication protocol. We demonstrate this with the $DISJ$ function. We remind the reader that

$$DISJ(x_1, \ldots, x_n; y_1, \ldots, y_n) = \ OR(x_1 \cdot y_1, \ldots, x_n \cdot y_n)$$

We define a function $f : \{0,1\}^{log(n)} \mapsto \{0,1\}$ by $f(i) = x_i \cdot y_i$, and so $DISJ(x;y) = OR(f)$. Alice simulates the black-box algorithm that computes $OR(f)$ using $O(\sqrt{n})$ calls to f.

For the simulation Alice needs to simulate an f gate. Alice can not compute an f gate by herself, as $f(i) = x_i \cdot y_i$ requires the knowledge of y_i. However, a simple argument shows that by communicating with Bob, Alice can simulate an f gate using only $O(\log(n))$ communication qubits. Using this, Alice can simulate the whole black-box algorithm, and compute $OR(f) = DISJ(x;y)$ with a large probability of success and only $O(\sqrt{n}\log(n))$ communication qubits. This shows that:

Theorem 6 *[BCW98] $Q(DISJ) = O(\sqrt{n}\log(n))$.*

As we mentioned before, it is known that $R(DISJ) = \Theta(n)$. Thus, this shows a quadratic gap between probabilistic and quantum communication complexity. We remind the reader, though, that in light of Theorem 5, the best one can hope to achieve using this approach is a polynomial gap between quantum and classical communication complexity.

5 An Exponential Gap for a Promise Problem

Applying the [BCW98] black-box approach to the Duetch-Jozsa function gives an exponential gap between the *zero-error* quantum communication complexity and *zero-error* classical communication complexity for a *promise* problem. Thus, the separation suffers from two "drawbacks"; it is only for the zero-sided error case, and only for a promise function. Ran Raz [Raz99] recently improved on that. He studied a problem similar to one suggested in [Kre95].

Example 6 *(The problem P_1)*
Alice*: holds $x \subseteq R^n$, and two orthogonal subspaces $M_0, M_1 \subseteq R^n$.*
Bob*: holds a unitary transformation T on R^n.*
Promise*: Tx is either close to M_0 or M_1*
Output*: i iff Tx is close to M_i.*
When x, M_i and T are described by $poly(n)$ bits using sufficient accuracy.

It is not hard to see that the above problem has an efficient quantum communication protocol requiring only $O(\log(n))$ communication qubits. Showing a lower bound for the probabilistic model is, however, a much harder task. Ran Raz shows that:

Theorem 7 $Q(P_1) = \Theta(\log(n))$ *and* $R(P_1) = \Omega(\sqrt{n})$

establishing an exponential gap for the two-sided error case. Yet, the gap applies only for *promise* functions.

6 Lower Bounds for The Quantum Model

Up to now we saw upper bounds showing that certain problems have good quantum protocols. We now turn our attention to lower bounds. We first note that

Theorem 8 *[Kre95] For any f, $Q(f) \geq \Omega(\log(D(f)))$*

Another result that appears in [Kre95] is that $Q(IP) = \Theta(n)$, and was the first linear lower bound on the quantum communication complexity of any function. This result was improved in [CvDNT98] where it was shown that

Theorem 9 *[CvDNT98]* $Q^{pub}(IP) = \Theta(n)$

We now elaborate a bit on the proof technique of Theorem 9. We start by presenting Holevo's theorem:

Theorem 10 *[Hol73] If Alice and Bob communicate l qubits then the mutual information between Alice and Bob is at most $2l$.*

[CvDNT98] generalize this result to the case where the two parties share EPR pairs, and are allowed many communication rounds. Now, suppose a quantum communication complexity protocol computes the IP function with l communication qubits and small error. Then, they show that given an input x to Alice, Bob can approximate the super-position $\Sigma_y (-1)^{IP(x,y)}|y\rangle$ using only $2l$ communication qubits. However, by applying the Hadamard transformation on $\Sigma_y (-1)^{IP(x,y)}|y\rangle$ Bob gets $|x\rangle$ itself. Hence Bob can approximate $|x\rangle$, and Alice and Bob can communicate $2l$ qubits and get mutual information almost n. This implies, by Holevo's theorem, that $l = \Omega(n)$.

As [CvDNT98] note, this also implies that $R(IP) \geq Q(IP) \geq Q^{pub}(IP) = \Omega(n)$, giving us the known classical result $R(IP) = \Theta(n)$ but using a new technique: proving a lower bound by reducing an information theoretic problem to a communication complexity problem.

6.1 The Quantum Communication Complexity of Constructing a Function

We now continue the above line of thought and ask how many communication qubits does it take for Alice and Bob to enter the super-position $\Sigma_{x,y}(-1)^{f(x,y)}|x,y\rangle$. We define:

Definition 6.1 *(q–sampling) [ASTS+98] We say a quantum protocol q–samples a super-position $\psi = \Sigma_{x\in X, y\in Y} a_{x,y}|x,y\rangle$ to within ϵ error if it starts with no inputs to the two players, and by the end of the protocol the two players compute a superposition ϕ where ϕ_{Alice} has support in X, ϕ_{Bob} has support in Y and $|\langle\phi|\psi\rangle| \geq 1-\epsilon$.*

We say a quantum protocol q–samples f according to the uniform distribution with ϵ error if it q–samples the super-position $\Sigma_{x,y}\mu_{x,y}(-1)^{f(x,y)}|x,y\rangle$ to within ϵ error, where $\mu_{x,y} = \frac{1}{\sqrt{|X|\cdot|Y|}}$ is equally distributed over $X \times Y$. We denote the number of communication qubits needed for this by $\overset{\bullet}{Q}_\epsilon(f)$. In this terminology [CvDNT98] show that if $Q(IP) = l$ then $\overset{\bullet}{Q}(IP) \leq 2l$. In fact, using the same technique it can be shown that in general

Theorem 11 $\overset{\bullet}{Q}_\epsilon(f) \leq 2Q_\epsilon(f)$.

Thus, a lower bound on the complexity of q–sampling, gives a lower bound on probabilistic communication complexity. The nice thing is that we have a complete characterization of the complexity of q–sampling. Given a super-position $\psi = \Sigma_{x,y}a_{x,y}|x,y\rangle$ define M_ψ to be the $2^n \times 2^n$ matrix $M = (a_{x,y})$. Define $M = M_\psi M_\psi^\dagger$. [2] Let λ_i be the i'th largest eigenvalue of M. It turns out that all the λ_i are non-negative, real numbers and $\Sigma_i\lambda_i = 1$. Define K_ϵ be the first integer K s.t. $\Sigma_{i=1}^{K}\lambda_i \geq 1-\epsilon$. Then:

Theorem 12 *[ASTS+98] For any pure state ψ and $0 \leq \epsilon \leq \frac{1}{2}$, $\lceil\log K_{2\epsilon}\rceil \leq \overset{\bullet}{Q}_\epsilon(\psi) \leq \lceil\log K_\epsilon\rceil$.*

We note that unlike the results of [CvDNT98], Theorem 12 applies only to the private key model and not to the public key model. We also note that the quantum protocol that approximates ψ runs in one message and thus extra rounds do not help, which is in sharp contrast to classical communication complexity where any extra round can give exponential speedup.

Theorem 12 is a generalization of the log-rank bound. When $\epsilon = 0$, K_ϵ is just the number of non-zero singular values, i.e., the rank of M_ψ, which is almost the same as the rank of M_f. Thus, $\overset{\bullet}{Q}_0(f) = \Theta(\log(Rank(M_f)))$. In particular, the case $\epsilon = 0$, shows that $D(f) = R_0(f) \geq Q_0(f) \geq \overset{\bullet}{Q}_0(f) = \Omega(\log(Rank(M_f)))$ which is the log rank lower bound for deterministic computation of Theorem 2.

Theorem 12, however, is much more general than the log-rank lower bound and applies also, e.g., for probabilistic computation. It can be applied, e.g., to the IP function, where all eigenvalues are the same, to prove that $R(IP) \geq Q(IP) \geq \Omega(n)$, and in fact gives a new general technique for proving lower bounds on probabilistic and quantum communication complexity.

Furthermore, it is the first time $\log(rank(M_f))$ is *characterized* by a communication complexity task, and so we can now restate the log-rank conjecture as comparing zero-error classical computation with zero-error q–sampling. I.e., the following is equivalent to the log-rank conjecture:

Conjecture 2 $D(f) \leq poly(\overset{\bullet}{Q}_0(f))$.

[2] One can also define M as Alice's reduced density matrix of the super-position ψ.

6.2 The $DISJ_k$ Function

We conclude with the $DISJ_k$ function:

Example 7 *(The $DISJ_k$ function)*
Alice: *holds $x \subseteq [1..n]$ of cardinality k.*
Bob: *holds $y \subseteq [1..n]$ of cardinality k.*
Output: *$f(x,y)$ is true iff $x \cap y \neq \emptyset$.*

The $DISJ_k$ is classically a hard function and [BFS86] prove that $R(DISJ_{\sqrt{n}}) \geq \Omega(\sqrt{n})$. Applying the technique from Theorem 12 for the $DISJ_k$ function one gets that $\overset{\bullet}{Q}_\epsilon\ (DISJ_{\sqrt{n}}) = O(\log(n)\log(\frac{1}{\epsilon}))$. Thus, there is a quantum protocol that with only $O(\log(n))$ qubits approximates the super-position $\Sigma_{x,y}(-1)^{f(x,y)}|x,y\rangle$ for $f = DISJ_{\sqrt{n}}$.

The ability to efficiently construct $\Sigma_{x,y}(-1)^{f(x,y)}|x,y\rangle$ also implies that Alice and Bob can efficiently sample f in the following sense. Alice and Bob start with no input. They communicate and finally end up with Alice holding some random $x \in X$, Bob holding some random $y \in Y$ and both players knowing the value $f(x,y)$. That is, while communication complexity asks for $f(x,y)$ on given inputs x and y, sampling only asks for a random sample of x, y along with the right answer $f(x,y)$. [ASTS+98] prove:

Theorem 13 *Sampling $DISJ_{\sqrt{n}}$ classically requires $\Omega(\sqrt{n})$ communication bits, while sampling it quantumly requires only $O(\log(n))$ qubits.*

Thus, an exponential gap exists between classical and quantum sampling. The general problem of determining the quantum communication complexity of the $DISJ_k$ function, and determining the largest possible gap between quantum and classical communication complexity is still wide open.

References

[ASTS+98] Ambainis, Schulman, Ta-Shma, Vazirani, and Wigderson. The quantum communication complexity of sampling. In *FOCS: IEEE Symposium on Foundations of Computer Science (FOCS)*, 1998.

[BBBV97] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26(5):1510–1523, October 1997.

[BBC+98] Beals, Buhrman, Cleve, Mosca, and de Wolf. Quantum lower bounds by polynomials. In *FOCS: IEEE Symposium on Foundations of Computer Science (FOCS)*, 1998.

[BCW98] H. Buhrman, R. Cleve, and A. Wigderson. Quantum vs. classical communication and computation. In *ACM Symposium on Theory of Computing (STOC)*, May 1998.

[BFS86] L. Babai, P. Frankl, and J. Simon. Complexity classes in communication complexity theory. In *27th Annual Symposium on Foundations of Computer Science*, pages 337–347, Los Angeles, Ca., USA, October 1986. IEEE Computer Society Press.

[CG88] B. Chor and O. Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17:230–261, 1988.

[CvDNT98] R. Cleve, W. van Dam, M. Nielsen, and A. Tapp. Quantum entanglement and the communication complexity of the inner product function. In *Proceedings of the 1st NASA International Conference on Quantum Computing and Quantum Communications (Springer-Verlag)*, 1998.

[DJ92] Deutsch and Jozsa. Rapid solution of problems by quantum computation. In *Proceedings of the Royal Society of London*, volume A439, pages 553–558, 1992.

[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*, pages 212–219, Philadelphia, Pennsylvania, 22–24 May 1996.

[Hol73] A. Holevo. Bounds for the quantity of information transmitted by a quantum communication channel. In *Problemy Peredachi Informatsii*, volume 9 (3), pages 3–11, 1973. English translation Problems of Information Transmission, vol 9, 1973, pp. 177-183.

[KN97] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.

[Kre95] I. Kremer. Quantum communication. Master's thesis, The Hebrew University of Jerusalem, 1995.

[KS87] Balasubramanian Kalyanasundaram and Georg Schnitger. The probabilistic communication complexity of set intersection (preliminary version). In *Proceedings, Structure in Complexity Theory, Second Annual Conference*, pages 41–47, Cornell University, Ithca, NY, 16–19 June 1987. IEEE Computer Society Press.

[MS82] Kurt Mehlhorn and Erik M. Schmidt. Las Vegas is better than determinism in VLSI and distributed computing (extended abstract). In *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing*, pages 330–337, San Francisco, California, 5–7 May 1982.

[New91] Ilan Newman. Private vs. common random bits in communication complexity. *Information Processing Letters*, 39(2):67–71, July 1991.

[NW95] Nisan and Wigderson. On rank vs. communication complexity. *Combinatorica*, 15:557–566, 1995.

[Raz92] A. A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385–390, December 1992.

[Raz99] R. Raz. Exponential separation of quantum and classical communication complexity. In *STOC*, 1999.

[RS95] Raz and Spieker. On the "log rank"-conjecture in communication complexity. *Combinatorica*, 15, 1995.

[Vaz85] Umesh V. Vazirani. Towards a strong communication complexity theory or generating quasi-random sequences from two communicating slightly-random sources (extended abstract). In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, pages 366–378, Providence, Rhode Island, 6–8 May 1985.

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

[Yao83] A. C. Yao. Lower bounds by probabilistic arguments (extended abstract). In *Proc. 24th Ann. IEEE Symp. on Foundations of Computer Science*, pages 420–428, 1983.

[Yao93] A. Yao. Quantum circuit complexity. In *Proceedings of the 34th Annual Symposium on Foundations of Computer Science*, pages 352–361, 1993.