Quantum Communication and Complexity

Ronald de Wolf [a,1]

[a] *University of California, Berkeley. 583 Soda Hall, Berkeley, CA 94720–1776, USA. E-mail:* rdewolf@cs.berkeley.edu

Abstract

In the setting of communication complexity, two distributed parties want to compute a function depending on both their inputs, using as little communication as possible. The required communication can sometimes be significantly lowered if we allow the parties the use of *quantum* communication. We survey the main results of the young area of quantum communication complexity: its relation to teleportation and dense coding, the main examples of fast quantum communication protocols, lower bounds, and some applications.

Keywords: Quantum computing. Communication complexity.

1 Introduction

The area of communication complexity deals with the following type of problem. There are two separated parties, called Alice and Bob. Alice receives some input $x \in X$, Bob receives some $y \in Y$, and together they want to compute some function $f(x, y)$. As the value $f(x, y)$ will generally depend on both x and y, neither Alice nor Bob will have sufficient information to do the computation by themselves, so they will have to communicate in order to achieve their goal. In this model, individual computation is free, but communication is expensive and has to be minimized. How many bits do they need to communicate between them in order to solve this? Clearly, Alice can just send her complete input to Bob, but sometimes more efficient schemes are possible. This model was introduced by Yao [64] and has been studied extensively, both

[1] Supported by Talent grant S 62–565 from the Netherlands Organization for Scientific Research (NWO). Most of this paper was written when the author was a PhD student at CWI and the University of Amsterdam, partially supported by the EU Fifth Framework project QAIP, IST–1999–11234.

for its applications (like lower bounds on VLSI and circuits) and for its own sake. We refer to [45,38] for definitions and results.

An interesting variant of the above is *quantum* communication complexity: suppose that Alice and Bob each have a quantum computer at their disposal and are allowed to exchange quantum bits (qubits) and/or to make use of the quantum correlations given by shared *EPR-pairs* (entangled pairs of qubits named after Einstein, Podolsky, and Rosen [31]). Can Alice and Bob now compute f with less communication than in the classical case? Quantum communication complexity was first considered by Yao [65] for the model with qubit communication and no prior EPR-pairs, and it was shown later that for some problems the amount of communication required in the quantum world is indeed considerably less than the amount of classical communication.

In this survey, we first give brief explanations of quantum computation and communication, and then cover the main results of quantum communication complexity: upper bounds (Section 5), lower bounds (Section 6), and applications (Section 7). We include proofs of some of the central results and references to others. Some other recent surveys of quantum communication complexity are [60,18,41,16], and a more popular account can be found in [59]. Our survey differs from these in being a bit more extensive and up to date.

2 Quantum Computation

In this section we briefly give the relevant background from quantum computation, referring to the book of Nielsen and Chuang [53] for more details.

2.1 States and operations

The classical unit of computation is a *bit*, which can take on the values 0 or 1. In the quantum case, the unit of computation is a *qubit*, which is a linear combination or *superposition* of the two classical values:

$$\alpha_0|0\rangle + \alpha_1|1\rangle.$$

More generally, an m-qubit state $|\phi\rangle$ is a superposition of all 2^m different classical m-bit strings:

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

The classical state $|i\rangle$ is called a *basis* state. The coefficient α_i is a complex number, which is called the *amplitude* of $|i\rangle$. The amplitudes form a 2^m-dimensional complex vector, which we require to have norm 1 (i.e. $\sum_i |\alpha_i|^2 =$

1). If some system is in state $|\phi\rangle$ and some other is in state $|\psi\rangle$, then their joint state is the *tensor product* $|\phi\rangle \otimes |\psi\rangle = |\phi\rangle|\psi\rangle$.

We can basically do two things to a quantum state: measure it or perform a unitary operation to it. If we measure $|\phi\rangle$, then we will see a basis state; we will see $|i\rangle$ with probability $|\alpha_i|^2$. Because $|\phi\rangle$ has norm 1, the probabilities $|\alpha_i|^2$ sum to 1, as they should. A measurement "collapses" the measured state to the measurement outcome: if we see $|i\rangle$, then $|\phi\rangle$ has collapsed to $|i\rangle$, and all other information in $|\phi\rangle$ is gone.

Apart from measuring, we can also transform the state, i.e., change the amplitudes. Quantum mechanics stipulates that this transformation U must be a *linear* transformation on the 2^m-dimensional vector of amplitudes:

$$U \begin{pmatrix} \alpha_{0\ldots0} \\ \vdots \\ \alpha_{1\ldots1} \end{pmatrix} = \begin{pmatrix} \beta_{0\ldots0} \\ \vdots \\ \beta_{1\ldots1} \end{pmatrix}.$$

Since the new vector of amplitudes β_i must also have norm 1, it follows that the linear transformation U must be norm-preserving and hence *unitary*.[2] This in turn implies that U has an inverse (in fact equal to its conjugate transpose U^*), hence non-measuring quantum operations are reversible.

2.2 Quantum algorithms

We describe quantum algorithms in the quantum circuit model [29,65], rather than the somewhat more cumbersome quantum Turing machine model [28,14]. A classical Boolean circuit is a directed acyclic graph of elementary Boolean gates (usually AND, OR, and NOT), only acting on one or two bits at a time. It transforms an initial vector of bits (containing the input) into the output. A quantum circuit is similar, except that the classical Boolean gates now become elementary *quantum* gates. Such a gate is a unitary transformation acting only on one or two qubits, and implicitly acting as the identity on the other qubits of the state. A simple example of a 1-qubit gate is the *Hadamard transform*, which maps basis state $|b\rangle$ to $\frac{1}{\sqrt{2}}(|0\rangle + (-1)^b|1\rangle)$. In matrix form, this is

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

[2] Both quantum measurements and quantum operations allow for a somewhat more general description than given here (POVMs and superoperators, respectively, see [53]), but the above definitions suffice for our purposes.

An example of a 2-qubit gate is the controlled-NOT (CNOT) gate, which negates the second bit of the state depending on the first bit: $|c, b\rangle \to |c, b \oplus c\rangle$. In matrix form, this is

$$C = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

It is known that the set of gates consisting of CNOT and all 1-qubit gates is *universal*, meaning that any other unitary transformation can be written as a product of gates from this set. We refer to [6,53] for more details.

The product of all elementary gates in a quantum circuit is a big unitary transformation that transforms the initial state (usually a classical bitstring containing the input x) into a final superposition. The *output* of the circuit is then the outcome of measuring some dedicated part of the final state. We say that a quantum circuit computes some function $f : \{0,1\}^n \to Z$ *exactly* if it always outputs the correct value $f(x)$ on input x. The circuit computes f *with bounded error* if it outputs $f(x)$ with probability at least $2/3$, for all x. Notice that a quantum circuit involves only one measurement; this is without loss of generality, since it is known that measurements can always be pushed to the end at the cost of a moderate amount of extra memory.

The *complexity* of a quantum circuit is usually measured by the number of elementary gates it contains. A circuit is deemed *efficient* if its complexity is at most polynomial in the length n of the input. The most spectacular instance of an efficient quantum circuit (rather, a uniform family of such circuits, one for each n) is still Shor's 1994 efficient algorithm for finding factors of large integers. It finds a factor of arbitrary n-bit numbers with high probability using only $n^2 \text{polylog}(n)$ elementary gates. This compromises the security of modern public-key cryptographic systems like RSA, which are based on the assumed hardness of factoring.

2.3 Query algorithms

A type of quantum algorithms that we will refer to later are the *query algorithms*. In fact, most existing quantum algorithms are of this type. Here the input is not part of the initial state, but encoded in a special "black box" quantum gate. The black box maps basis state $|i, b\rangle$ to $|i, b \oplus x_i\rangle$, thus giving access to the bits x_i of the input. Note that a quantum algorithm can run the black box on a superposition of basis states, gaining access to several input bits x_i at the same time. One such application of the black box is called a

query. The complexity of a quantum circuit for computing some function f is now the number of queries we need on the worst-case input; we don't count the complexity of other operations in this model. In the classical world, this query complexity is known as the *decision tree* complexity of f.

A simple but illustrative example is the Deutsch-Jozsa algorithm [30,27]: suppose that n is a power of 2, and we get the promise that the input $x \in \{0,1\}^n$ is either $0 \ldots 0$ ("constant") or has exactly $n/2$ 0s and $n/2$ 1s ("balanced"). Define $\mathrm{DeJo}(x) = 1$ in the first case and $\mathrm{DeJo}(x) = 0$ in the second. It is easy to see that a deterministic classical computer needs $n/2 + 1$ queries for this (if the computer has queried $n/2$ bits and they are all 0, then the function value is still undetermined). On the other hand, here is a 1-query quantum algorithm for this problem:

(1) Start in a basis state $|0 \ldots 01\rangle$ of $\log n$ zeroes followed by a 1
(2) Apply a Hadamard transform to each of the $\log n + 1$ qubits
(3) Query the black box once
(4) Apply a Hadamard transform to the first $\log n$ qubits
(5) Measure the first $\log n$ qubits, output 1 if the observed state is $|0 \ldots 0\rangle$ and output 0 otherwise

By following the state through these steps, it may be verified that the algorithm always outputs 1 if the input x is constant, and 0 if it is balanced.

Another important quantum query algorithm is Grover's search algorithm [35], which finds an i such that $x_i = 1$ if such an i exists in the n-bit input. It has error probability $\leq 1/3$ on each input and uses $O(\sqrt{n})$ queries, which is optimal [12,15,66]. Note that the algorithm can also be viewed as computing the OR-function: it can determine whether at least one of the input bits is 1.

3 Quantum Communication

The area of quantum information theory deals with the properties of quantum information and its communication between different parties. We refer to [13,53] for general surveys, and will here restrict ourselves to explaining two important primitives: *teleportation* [10] and *superdense coding* [11]. These pre-date quantum communication complexity and show some of the power of quantum communication.

We first show how teleporting a qubit works. Alice has a qubit $\alpha_0|0\rangle + \alpha_1|1\rangle$ that she wants to send to Bob via a *classical* channel. Without further resources this would be impossible, but Alice also shares an EPR-pair $\frac{1}{\sqrt{2}}(|00\rangle +$

$|11\rangle$) with Bob. Initially, their joint state is

$$(\alpha_0|0\rangle + \alpha_1|1\rangle) \otimes \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle).$$

The first two qubits belong to Alice, the third to Bob. Alice performs a CNOT on her two qubits and then a Hadamard transform on her first qubit. Their joint state can now be written as

$$\begin{aligned}
&\tfrac{1}{2}\ |00\rangle(\alpha_0|0\rangle + \alpha_1|1\rangle) + \\
&\tfrac{1}{2}\ |01\rangle(\alpha_0|1\rangle + \alpha_1|0\rangle) + \\
&\tfrac{1}{2}\ |10\rangle(\alpha_0|0\rangle - \alpha_1|1\rangle) + \\
&\tfrac{1}{2}\ \underbrace{|11\rangle}_{Alice}\underbrace{(\alpha_0|1\rangle - \alpha_1|0\rangle)}_{Bob}.
\end{aligned}$$

Alice then measures her two qubits and sends the result (2 random classical bits) to Bob, who now knows which transformation he must do on his qubit in order to regain the qubit $\alpha_0|0\rangle + \alpha_1|1\rangle$. For instance, if Alice sent 11 then Bob knows that his qubit is $\alpha_0|1\rangle - \alpha_1|0\rangle$. A bit-flip ($|b\rangle \to |1-b\rangle$) followed by a phase-flip ($|b\rangle \to (-1)^b|b\rangle$) will give him Alice's original qubit $\alpha_0|0\rangle + \alpha_1|1\rangle$. In fact, if Alice's qubit had been entangled with other qubits, then teleportation preserves this entanglement: Bob then receives a qubit that is entangled in the same way as Alice's original qubit was.

Note that the qubit on Alice's side has been destroyed: teleporting moves a qubit from A to B, rather than copying it. In fact, copying an unknown qubit is impossible [62], which can be seen as follows. Suppose C were a 1-qubit copier, i.e. $C|\phi\rangle|0\rangle = |\phi\rangle|\phi\rangle$ for every qubit $|\phi\rangle$. In particular $C|0\rangle|0\rangle = |0\rangle|0\rangle$ and $C|1\rangle|0\rangle = |1\rangle|1\rangle$. But then C would not copy $|\phi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ correctly, since by linearity $C|\phi\rangle|0\rangle = \frac{1}{\sqrt{2}}(C|0\rangle|0\rangle + C|1\rangle|0\rangle) = \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|1\rangle) \neq |\phi\rangle|\phi\rangle$.

In teleportation, Alice uses 2 classical bits and 1 EPR-pair to send 1 qubit to Bob. *Superdense coding* achieves the opposite: using 1 qubit and 1 EPR-pair, Alice can send 2 classical bits b_1, b_2 to Bob. It works as follows. Initially they share an EPR-pair $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. First, if $b_1 = 1$ then Alice applies a phase-flip to her half of the pair. Second, if $b_2 = 1$ she applies a bit-flip. Third, she sends her half of the EPR-pair to Bob, who now has one of 4 states $|\phi_{b_1 b_2}\rangle$:

$$\begin{aligned}
|\phi_{00}\rangle &= \tfrac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \\
|\phi_{01}\rangle &= \tfrac{1}{\sqrt{2}}(|10\rangle + |01\rangle) \\
|\phi_{10}\rangle &= \tfrac{1}{\sqrt{2}}(|00\rangle - |11\rangle) \\
|\phi_{11}\rangle &= \tfrac{1}{\sqrt{2}}(|10\rangle - |01\rangle)
\end{aligned}$$

Since these states are orthogonal, Bob can apply a unitary transformation

that maps $|\phi_{b_1 b_2}\rangle \rightarrow |b_1 b_2\rangle$ and thus learn b_1 and b_2.

Suppose Alice wants to send n classical bits of information to Bob and they do not share any prior entanglement. Alice can just send her n bits to Bob, but, alternatively, Bob can also first send $n/2$ halves of EPR-pairs to Alice and then Alice can send n bits in $n/2$ qubits using dense coding. In either case, n qubits are exchanged between them. If Alice and Bob already share $n/2$ prior EPR-pairs, then $n/2$ qubits suffice by superdense coding. The following result shows that this is optimal. We will refer to it as *Holevo's theorem*, because the first part is an immediate consequence of a result of [36] (the second part was derived in [26]).

Theorem 1 (Holevo [36]) *If Alice wants to send n bits of information to Bob via a qubit channel, and they don't share prior entanglement, then they have to exchange at least n qubits. If they do share unlimited prior entanglement, then Alice has to send at least $n/2$ qubits to Bob, no matter how many qubits Bob sends to Alice.*

A somewhat stronger and more subtle variant of this lower bound was derived by Nayak [48], improving upon [2]. Suppose that Alice doesn't want to send Bob all of her n bits, but just wants to send a message that allows Bob to learn *one* of her bits x_i, where Bob can choose i after the message has been sent. Even for this weaker form of communication, Alice has to send an $\Omega(n)$-qubit message.

4 Quantum Communication Complexity: The Model

First we sketch the setting for classical communication complexity, referring to [45,38] for more details. Alice and Bob want to compute some function $f : \mathcal{D} \rightarrow \{0,1\}$, where $\mathcal{D} \subseteq X \times Y$. If the domain $\mathcal{D}$ equals $X \times Y$ then f is called a *total* function, otherwise it is a *promise* function. Alice receives input $x \in X$, Bob receives input $y \in Y$, with $(x,y) \in \mathcal{D}$. As the value $f(x,y)$ will generally depend on both x and y, some communication between Alice and Bob is required in order for them to be able to compute $f(x,y)$. We are interested in the *minimal* amount of communication they need.

A communication *protocol* is a distributed algorithm where first Alice does some individual computation, and then sends a message (of one or more bits) to Bob, then Bob does some computation and sends a message to Alice, etc. Each message is called a *round*. After one or more rounds the protocol terminates and outputs some value, which must be known to both players. The *cost* of a protocol is the total number of bits communicated on the worst-case input. A *deterministic* protocol for f always has to output the right value

$f(x,y)$ for all $(x,y) \in \mathcal{D}$. In a *bounded-error* protocol, Alice and Bob may flip coins and the protocol has to output the right value $f(x,y)$ with probability $\geq 2/3$ for all $(x,y) \in \mathcal{D}$. We use $D(f)$ and $R_2(f)$ to denote the minimal cost of deterministic and bounded-error protocols for f, respectively. The subscript '2' in $R_2(f)$ stands for 2-sided bounded error. For $R_2(f)$ we can either allow Alice and Bob to toss coins individually (private coin) or jointly (public coin). This makes not much difference: a public coin can save at most $O(\log n)$ bits of communication [50], compared to a protocol with a private coin.

Some often studied total functions where $X = Y = \{0,1\}^n$:

- *Equality:* $\mathrm{EQ}(x,y) = 1$ iff $x = y$
- *Inner product:* $\mathrm{IP}(x,y) = \mathrm{PARITY}(x \wedge y) = \sum_i x_i y_i \pmod 2$
  (for $x, y \in \{0,1\}^n$, x_i is the ith bit of x and $x \wedge y \in \{0,1\}^n$ is the bit-wise AND of x and y)
- *Disjointness:* $\mathrm{DISJ}(x,y) = \mathrm{NOR}(x \wedge y)$. This function is 1 iff there is no i where $x_i = y_i = 1$ (viewing x and y as characteristic vectors of sets, the sets are disjoint)

It is known that $D(\mathrm{EQ}) = D(\mathrm{IP}) = D(\mathrm{DISJ}) = n+1$, $R_2(\mathrm{IP}) = R_2(\mathrm{DISJ}) = \Omega(n)$. However, $R_2(\mathrm{EQ})$ is only $O(1)$, as follows. Alice and Bob jointly toss a random string $r \in \{0,1\}^n$. Alice sends the bit $a = x \cdot r$ to Bob (where '$\cdot$' is inner product mod 2). Bob computes $b = y \cdot r$ and compares this with a. If $x = y$ then $a = b$, but if $x \neq y$ then $a \neq b$ with probability $1/2$. Thus Alice and Bob can decide equality with small error using $O(n)$ public coin flips and $O(1)$ communication. Since public coin and private coin protocols are close, this also implies that $R_2(\mathrm{EQ}) \in O(\log n)$ with a private coin.

Now what happens if we give Alice and Bob a quantum computer and allow them to send each other qubits and/or to make use of EPR-pairs that they share at the start of the protocol? Formally speaking, we can model a quantum protocol as follows. The total state consists of 3 parts: Alice's private space, the channel, and Bob's private space. The starting state is $|x\rangle|0\rangle|y\rangle$: Alice gets x, the channel is initially empty, and Bob gets y. Now Alice applies a unitary transformation to her space and the channel. This corresponds to her private computation as well as to putting a message on the channel (the length of this message is the number of channel-qubits affected by Alice's operation). Then Bob applies a unitary transformation to his space and the channel, etc. At the end of the protocol Alice or Bob makes a measurement to determine the output of the protocol. We use $Q(f)$ to denote the minimal communication cost of a quantum protocol that computes $f(x,y)$ exactly (= with error probability 0). This model was introduced by Yao [65]. In the second model, introduced by Cleve and Buhrman [25], Alice and Bob share an unlimited number of EPR-pairs at the start of the protocol, but now they communicate via a *classical* channel: the channel has to be in a classical state

throughout the protocol. We use $C^*(f)$ for the minimal complexity of an exact protocol for f in this model. Note that we only count the communication, not the number of EPR-pairs used. The third variant combines the strengths of the other two: here Alice and Bob start out with an unlimited number of shared EPR-pairs *and* they are allowed to communicate qubits. We use $Q^*(f)$ to denote the communication complexity in this third model. By teleportation, 1 EPR-pair and 2 classical bits can replace 1 qubit of communication, so we have $Q^*(f) \leq C^*(f) \leq 2Q^*(f)$. Similarly we define $Q_2(f)$, $Q_2^*(f)$, and $C_2^*(f)$ for bounded-error quantum protocols. Note that a shared EPR-pair can simulate a public coin toss: if Alice and Bob each measure their half of the pair, they get the same random bit.

Before continuing to study this model, we first have to face an important question: *is there anything to be gained here?* At first sight, the following argument seems to rule out any significant gain. By definition, in the classical world $D(f)$ bits have to be communicated in order to compute f. Since Holevo's theorem says that k qubits cannot contain more information than k classical bits, it seems that the quantum communication complexity should be roughly $D(f)$ qubits as well (maybe $D(f)/2$ to account for superdense coding, but not less). Fortunately and surprisingly, this argument is false, and quantum communication can sometimes be much less than classical communication complexity. The information-theoretic argument via Holevo's theorem fails, because Alice and Bob do not need to communicate the information in the $D(f)$ bits of the classical protocol; they are only interested in the value $f(x, y)$, which is just 1 bit. Below we will survey the main examples that have so far been found of gaps between quantum and classical communication complexity.

5 Quantum Communication Complexity: Upper bounds

5.1 Initial steps

Quantum communication complexity was introduced by Yao [65] and studied by Kremer [44], but neither showed any advantages of quantum over classical communication. Cleve and Buhrman [25] introduced the variant with classical communication and prior entanglement, and exhibited the first quantum protocol provably better than any classical protocol. It uses quantum entanglement to save 1 bit of classical communication. This gap was extended by Buhrman, Cleve, and van Dam [19] and, for arbitrary k parties, by Buhrman, van Dam, Høyer, and Tapp [23].

The first impressively large gaps between quantum and classical communication complexity were exhibited by Buhrman, Cleve, and Wigderson [21]. Their protocols are distributed versions of known quantum query algorithms, like the Deutsch-Jozsa and Grover algorithms. The following lemma shows how a query algorithm induces a communication protocol:

Lemma 1 (BCW [21]) *Let $g : \{0,1\}^n \to \{0,1\}$ and $f(x,y) = g(x \star y)$, where $\star$ is any binary connective (for instance $\oplus$ or $\wedge$). If there is a T-query quantum algorithm for g, then there is a protocol for f that communicates $T(2\log n + 4)$ qubits (and uses no prior entanglement) and that has the same error probability as the query algorithm.*

Proof. The quantum protocol consists of Alice's simulating the quantum query algorithm A on input $x \star y$. Every query in A will correspond to 2 rounds of communication. Namely, suppose Alice at some point wants to apply a query to the state $|\phi\rangle = \sum_{i,b} \alpha_{ib} |i,b\rangle$ (for simplicity we omit Alice's workspace). Then she adds a $|0\rangle$-qubit to the state, applies the unitary mapping $|i,b,0\rangle \to |i,b,x_i\rangle$, and sends the resulting state to Bob. Bob now applies the unitary mapping $|i,b,x_i\rangle \to |i,b \oplus (x_i \star y_i), x_i\rangle$ and sends the result back to Alice. Alice applies $|i,b,x_i\rangle \to |i,b,0\rangle$, takes off the last qubit, and ends up with the state $\sum_{i,b} \alpha_{ib} |i, b \oplus (x_i \star y_i)\rangle$, which is exactly the result of applying an $x \star y$-query to $|\phi\rangle$. Thus every query to $x \star y$ can be simulated using $2\log n + 4$ qubits of communication. The final quantum protocol will have $T(2\log n + 4)$ qubits of communication and computes $f(x,y)$ with the same error probability as A has on input $x \star y$. □

Now consider the disjointness function: $\mathrm{DISJ}(x,y) = \mathrm{NOR}(x \wedge y)$. Since Grover's algorithm can compute the NOR of n variables with $O(\sqrt{n})$ queries with bounded-error, the previous lemma implies a bounded-error protocol for disjointness with $O(\sqrt{n}\log n)$ qubits. On the other hand, the linear lower bound for disjointness is a well-known result of classical communication complexity [39,56]. Thus we obtain the following near-quadratic separation:

Theorem 2 (BCW [21]) *$Q_2(\mathrm{DISJ}) \in O(\sqrt{n}\log n)$ and $R_2(\mathrm{DISJ}) \in \Omega(n)$.*

Høyer and de Wolf [37] slightly improved the upper bound on $Q_2(\mathrm{DISJ})$ to $O(\sqrt{n}c^{\log^* n})$ for some constant $c > 1$, thus showing that the $\log n$ in the upper bound can be replaced by a function that grows slower than any iterated logarithm.

Another separation is given by a distributed version of the Deutsch-Jozsa problem of Section 2.3: define $\mathrm{EQ}'(x,y) = \mathrm{DeJo}(x \oplus y)$. This is a promise version of equality, where the promise is that x and y are either equal or are at Hamming distance $n/2$. Since there is an exact 1-query quantum algorithm for DeJo, Lemma 1 implies $Q(\mathrm{EQ}') \in O(\log n)$. In contrast, Buhrman, Cleve, and Wigderson use a combinatorial result of Frankl and Rödl [33] to prove the classical lower bound $D(\mathrm{EQ}') \in \Omega(n)$. Thus we have the following exponential separation for exact protocols:

Theorem 3 (BCW [21]) *$Q(\mathrm{EQ}') \in O(\log n)$ and $D(\mathrm{EQ}') \in \Omega(n)$.*

5.3 Raz

Notice the contrast between the two separations of the previous section. For the Deutsch-Jozsa problem we get an *exponential* quantum-classical separation, but the separation only holds if we force the classical protocol to be exact; it is easy to see that $O(\log n)$ bits are sufficient if we allow some error (the classical protocol can just try a few random positions i and check if $x_i = y_i$ or not). On the other hand, the gap for the disjointness function is only *quadratic*, but it holds even if we allow classical protocols to have some error probability. Ran Raz [55] has exhibited a function where the quantum-classical separation has both features: the quantum protocol is exponentially better than the classical protocol, even if the latter is allowed some error probability. Consider the following promise problem **P**:

> Alice receives a unit vector $v \in R^m$ and a decomposition of the corresponding space in two orthogonal subspaces $H^{(0)}$ and $H^{(1)}$. Bob receives an $m \times m$ unitary transformation U. Promise: Uv is either "close" to $H^{(0)}$ or to $H^{(1)}$. Question: which of the two?

As stated, this is a problem with continuous input, but it can be discretized in a natural way by approximating each real number by $O(\log m)$ bits. Alice and Bob's input is now $n = O(m^2 \log m)$ bits long. There is a simple yet efficient 2-round quantum protocol for this problem: Alice views v as a $\log m$-qubit vector and sends this to Bob. Bob applies U and sends back the result. Alice then measures in which subspace $H^{(i)}$ the vector Uv lies and outputs the resulting i. This takes only $2 \log m = O(\log n)$ qubits of communication.

The efficiency of this protocol comes from the fact that an m-dimensional vector can be "compressed" or "represented" as a $\log m$-qubit state. Similar compression is not possible with classical bits, which suggests that any classical protocol for **P** will have to send the vector v more or less literally and hence will require a lot of communication. This turns out to be true but the proof (given in [55]) is surprisingly hard. The result is the first exponential gap

between Q_2 and R_2:

Theorem 4 (Raz [55]) *$Q_2(\mathbf{P}) \in O(\log n)$ and $R_2(\mathbf{P}) \in \Omega(n^{1/4}/\log n)$.*

6 Quantum Communication Complexity: Lower Bounds

In the previous section we exhibited some of the power of quantum communication complexity. Here we will look at its limitations, first for exact protocols and then for the bounded-error case.

6.1 Lower bounds on exact protocols

Quite good lower bounds are known for exact quantum protocols for total functions. For a total function $f : X \times Y \to \{0,1\}$ let $M_f[x,y] = f(x,y)$ be the *communication matrix* of f. This is an $|X| \times |Y|$ Boolean matrix that completely describes f. Let $rank(f)$ denote the rank of M_f over the reals. Mehlhorn and Schmidt [47] proved that $D(f) \geq \log rank(f)$, which is the main source of lower bounds on $D(f)$. For $Q(f)$ a similar lower bound follows from techniques of Yao and Kremer [65,44], as first observed in [21]. This bound was later extended to the case where Alice and Bob share unlimited prior entanglement by Buhrman and de Wolf [24]. Their result turned out to be equivalent to a result in Nielsen's thesis [52, Section 6.4.2]. The result is:

Theorem 5 *$Q^*(f) \geq \frac{1}{2}\log rank(f)$ and $C^*(f) \geq \log rank(f)$.*

Hence quantum communication complexity in the exact model with prior entanglement is maximal whenever M_f has full rank, which happens for almost all functions, including equality, (the complement of) inner product, and disjointness. For $Q(f)$, the model without prior entanglement, the same bounds apply and it is open whether the factor of $\frac{1}{2}$ can be removed in this case. For the equality and disjointness functions, the optimal bounds $Q(\mathrm{EQ}) = Q(\mathrm{DISJ}) = n+1$ were shown recently by Høyer and de Wolf [37].

How tight is the $\log rank(f)$ lower bound? It has been conjectured that $D(f) \leq (\log rank(f))^{O(1)}$ for all total functions, in which case $\log rank(f)$ would characterize $D(f)$ up to polynomial factors. If this *log-rank conjecture* is true, then Theorem 5 implies that $Q^*(f)$ and $D(f)$ are polynomially close for all total f, since then $Q^*(f) \leq D(f) \leq (\log rank(f))^{O(1)} \leq (2Q^*(f))^{O(1)}$. Some small classes of functions where this provably holds are identified in [24]. It should be noted that, in fact, no total f is known where $Q^*(f)$ is more than a factor of 2 smaller than $D(f)$ (the factor of 2 can be achieved by superdense coding).

The previous section showed some strong lower bounds for exact quantum protocols. The situation is worse in the case of bounded-error protocols, for which only a few good lower bounds are known. One of the few general lower bound techniques known to hold for bounded-error quantum complexity (without prior entanglement), is the so-called "discrepancy method". This was shown by Kremer [44], who used it to derive an $\Omega(n)$ lower bound for $Q_2(\mathrm{IP})$. Cleve, van Dam, Nielsen, and Tapp [26] later independently proved such a lower bound for $Q_2^*(\mathrm{IP})$.

We will sketch the very elegant proof of [26] here for the case of exact protocols. The proof uses the IP-protocol to communicate Alice's n-bit input to Bob, and then invokes Holevo's theorem to conclude that many qubits must have been communicated in order to achieve this. Suppose Alice and Bob have some protocol for IP. They can use this to compute the following mapping:

$$|x\rangle|y\rangle \to |x\rangle(-1)^{x\cdot y}|y\rangle.$$

Now suppose Alice starts with an arbitrary n-bit state $|x\rangle$ and Bob starts with the uniform superposition $\frac{1}{\sqrt{2^n}}\sum_{y\in\{0,1\}^n}|y\rangle$. If they apply the above mapping, the final state becomes

$$|x\rangle\frac{1}{\sqrt{2^n}}\sum_{y\in\{0,1\}^n}(-1)^{x\cdot y}|y\rangle.$$

If Bob applies a Hadamard transform to each of his n qubits, then he obtains the basis state $|x\rangle$, so Alice's n classical bits have been communicated to Bob. Theorem 1 now implies that the IP-protocol must communicate $n/2$ qubits, even if Alice and Bob share unlimited prior entanglement. With some more technical complication, the same idea gives an $\frac{1}{2}(1-2\varepsilon)^2 n$ lower bound for ε-error protocols [26]. The constant factor in this bound was recently improved to the optimal $\frac{1}{2}$ by Nayak and Salzman [49].

The above proof works for IP, but does not easily yield good bounds in general. Neither does the discrepancy method, or an approximate version of the rank lower bound that was noted in [21]. New lower bound techniques for quantum communication are required. Of particular interest is whether the upper bound of roughly $\sqrt{n}$ on $Q_2(\mathrm{DISJ})$ is tight. Because disjointness can be reduced to many other problems (it is in fact "coNP-complete" [4]), a good lower bound for disjointness would imply many other lower bounds as well. Høyer and de Wolf proved an $\Omega(\sqrt{n})$ lower bound for a restricted class of protocols, namely those whose acceptance probability is a function of $x \wedge y$, rather than x and y separately. All known quantum protocols for disjointness fall in this class, but it is still rather limited. Klauck [42] extended the classical Fourier analysis-

based lower bound technique of Raz [54] to the quantum case. The technique gives good lower bounds on $Q_2(f)$ for threshold functions (where $f(x,y)=1$ iff $|x \wedge y| \geq t$) with sufficiently large threshold t. Unfortunately, it fails to give good bounds for the $t=1$ case, which is exactly the complement of disjointness.

Then, in a recent breakthrough, Razborov [57] established the expected lower bound, using a clever multidimensional version of the discrepancy method:

Theorem 6 (Razborov [57]) $Q_2^*(\mathrm{DISJ}) = \Omega(\sqrt{n})$.

Moreover, his lower bound method gives nearly optimal lower bounds on $Q_2^*(f)$ for *all* functions that are of the form $f(x,y) = g(x \wedge y)$ for some g depending only on the Hamming weight of its input.

7 Quantum Communication Complexity: Applications

The main applications of classical communication complexity have been in proving lower bounds for various models like VLSI, Boolean circuits, formula size, Turing machine complexity, data structures, automata size etc. We refer to [45,38] for many examples. Typically one proceeds by showing that a communication complexity problem f is "embedded" in the computational problem P of interest, and then uses communication complexity lower bounds on f to establish lower bounds on P. Similarly, quantum communication complexity has been used to establish lower bounds in various models of quantum computation, though such applications have received relatively little attention so far. We will briefly mention some.

Yao [65] initially introduced quantum communication complexity as a tool for proving a superlinear lower bound on the quantum *formula size* of the majority function (a "formula" is a circuit of restricted form). More recently, Klauck [40] used one-round quantum communication complexity lower bounds to prove lower bounds on the size of quantum formulae.

Since upper bounds on query complexity give upper bounds on communication complexity (Lemma 1), lower bounds on communication complexity give *lower bounds on query complexity*. For instance, $\mathrm{IP}(x,y) = \mathrm{PARITY}(x \wedge y)$, so the $\Omega(n)$ bound for IP (Section 6.2) implies an $\Omega(n/\log n)$ lower bound for the quantum query complexity of the parity function, as observed by Buhrman, Cleve, and Wigderson [21]. This query lower bound was later strengthened to $n/2$ in [7,32].

Furthermore, as in the classical case, lower bounds on one-round communica-

tion complexity imply lower bounds on the size of quantum *finite automata*. This was used by Klauck [40] to show that Las Vegas (zero-error) quantum finite automata cannot be much smaller than classical deterministic finite automata.

Again, as in the classical case, lower bounds on quantum communication complexity give rise to lower bounds on the size of certain quantum *data structures*. For example, tradeoffs between size and access time for the "static predecessor" and "static membership" problems were obtained recently by Sen and Venkatesh [58].

Finally, Ben-Or [9] has applied the lower bounds for IP in a new proof of the security of *quantum key distribution*.

8 Other Developments and Open Problems

Here we mention some other results in quantum communication complexity or related models:

- **Quantum sampling.** For the *sampling* problem, Alice and Bob do not want to compute some $f(x, y)$, but instead want to sample an (x, y)-pair according to some known joint probability distribution, using as little communication as possible. Ambainis et al. [3] give a tight algebraic characterization of quantum sampling complexity, and exhibit an exponential gap between the quantum and classical communication required for a sampling problem related to disjointness.
- **Spooky communication.** Referring to Einstein's description of certain quantum effects as "spooky action at a distance" ("spukhafte Fernwirkungen"), Brassard, Cleve, and Tapp [17] exhibit tasks that can be achieved in the quantum world with entanglement and *no communication whatsoever*, but that would require communication in the classical world. They also give upper and lower bounds on the amount of classical communication needed to "simulate" EPR-pairs. Their results—and subsequent ones [46]—may be viewed as quantitative extensions of the famous Bell inequalities [8].
- **Las Vegas protocols.** In this survey we just considered two modes of computation: exact and bounded-error. An intermediate type of protocols are *zero-error* or *Las Vegas* protocols. These never output an incorrect answer, but may claim ignorance with probability at most 1/2. Some quantum-classical separations for zero-error protocols may be found in [22,40].
- **One-round communication.** Suppose the communication consists of just one round: Alice sends a message (depending on x) to Bob, who should then compute $f(x, y)$. Klauck [40] showed for all total functions that quantum one-round communication is not significantly better than classical one-round

communication in the case of exact or zero-error protocols. For the case of bounded-error protocols this is still open.

- **Quantum fingerprinting**. The model of *simultaneous message passing* is even more restricted than the one-round setting. Here there are three parties: Alice has x, Bob has y, they don't share entanglement or randomness but they can each send one message to a referee, who wants to determine $f(x, y)$. Buhrman, Cleve, Watrous, and de Wolf [20] gave an efficient quantum protocol for the equality function in this model: Alice and Bob send $O(\log n)$-qubit "quantum fingerprints" of their respective inputs to the referee, who can then decide with high success probability whether $x = y$. In contrast, classically the equality function requires $\Theta(\sqrt{n})$ bits of communication in this model [1,51,5].
- **Rounds.** In classical communication complexity it is well known that allowing Alice and Bob $k + 1$ rounds of communication instead of k reduces the required communication exponentially for some functions. An analogous result has been shown for quantum communication by Klauck, Nayak, Ta-Shma, and Zuckerman [43], using a quantum version of the classical "round elimination" technique. This technique has been further strengthened by Sen and Venkatesh [58], giving for instance tight communication-rounds tradeoffs for the "greater than" function.
- **Non-deterministic communication complexity.** A *non-deterministic* protocol has positive acceptance probability on input (x, y) iff $f(x, y) = 1$. Classically, the non-deterministic communication complexity is characterized by the logarithm of the cover number of the communication matrix M_f. The *quantum* non-deterministic communication complexity has been shown equal to the logarithm of the rank of a non-deterministic version of M_f [61,37]. There exist total functions where the quantum non-deterministic complexity is exponentially smaller than the classical one [61].

Finally, here's a list of interesting open problems in quantum communication complexity:

- Very few interesting quantum protocols are known. For what other problems can quantum mechanics save communication?
- Raz's exponential gap only holds for a promise problem. Are $R_2(f)$ and $Q_2^*(f)$ polynomially related for all *total* f? A similar question can be posed for the relation between $D(f)$ and $Q^*(f)$. As we showed in Section 6.1, a positive answer to this last question would be implied by the classical log-rank conjecture.
- Does entanglement add much power to qubit communication? That is, what are the biggest gaps between $Q(f)$ and $Q^*(f)$, and between $Q_2(f)$ and $Q_2^*(f)$? (We only know $Q_2(\mathrm{EQ}) \in \Theta(\log n)$ and $Q_2^*(\mathrm{EQ}) \in O(1)$.)
- Classically, Yao [63] used the minimax theorem from game theory to show an equivalence between deterministic protocols with a probability distribution on the inputs, and bounded-error protocols. Is some relation like this true

in the quantum case as well? Some preliminary results on quantum versions of Yao's result may be found in [34].

Acknowledgment

Thanks to Andris Ambainis for pointing out a mistake in an earlier version of this paper.

References

[1] A. Ambainis. Communication complexity in a 3-computer model. *Algorithmica*, 16(3):298–301, 1996.

[2] A. Ambainis, A. Nayak, A. Ta-Shma, and U. Vazirani. Quantum dense coding and a lower bound for 1-way quantum finite automata. In *Proceedings of 31st ACM STOC*, pages 376–383, 1999. quant-ph/9804043.

[3] 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.

[4] L. Babai, P. Frankl, and J. Simon. Complexity classes in communication complexity theory. In *Proceedings of 27th IEEE FOCS*, pages 337–347, 1986.

[5] L. Babai and P. G. Kimmel. Randomized simultaneous messages: Solution of a problem of Yao in communication complexity. In *Proceedings of the 12th IEEE Conference on Computational Complexity*, pages 239–246, 1997.

[6] A. Barenco, C.H. Bennett, R. Cleve, D.P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. Smolin, and H. Weinfurter. Elementary gates for quantum computation. *Physical Review A*, 52:3457–3467, 1995. quant-ph/9503016.

[7] 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.

[8] J. S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1965.

[9] M. Ben-Or. Security of quantum key distribution. Unpublished manuscript, 1999.

[10] C. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. Wootters. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters*, 70:1895–1899, 1993.

[11] C. Bennett and S. Wiesner. Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states. *Physical Review Letters*, 69:2881–2884, 1992.

[12] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26(5):1510–1523, 1997. quant-ph/9701001.

[13] C. H. Bennett and P. W. Shor. Quantum information theory. *IEEE Transactions on Information Theory*, 44(6):2724–2742, 1998.

[14] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997. Earlier version in STOC'93.

[15] M. Boyer, G. Brassard, P. Høyer, and A. Tapp. Tight bounds on quantum searching. *Fortschritte der Physik*, 46(4–5):493–505, 1998. Earlier version in Physcomp'96. quant-ph/9605034.

[16] G. Brassard. Quantum communication complexity (a survey). quant-ph/0101005, 1 Jan 2001.

[17] G. Brassard, R. Cleve, and A. Tapp. The cost of exactly simulating quantum entanglement with classical communication. *Physical Review Letters*, 83(9):1874–1877, 1999. quant-ph/9901035.

[18] H. Buhrman. Quantum computing and communication complexity. *EATCS Bulletin*, 70:131–141, February 2000.

[19] H. Buhrman, R. Cleve, and W. van Dam. Quantum entanglement and communication complexity. *SIAM Journal on Computing*, 30(8):1829–1841, 2001. quant-ph/9705033.

[20] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16), September 26, 2001. quant-ph/0102001.

[21] 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.

[22] H. Buhrman, R. Cleve, R. de Wolf, and Ch. Zalka. Bounds for small-error and zero-error quantum algorithms. In *Proceedings of 40th IEEE FOCS*, pages 358–368, 1999. cs.CC/9904019.

[23] 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.

[24] 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.

[25] R. Cleve and H. Buhrman. Substituting quantum entanglement for communication. *Physical Review A*, 56(2):1201–1204, 1997. quant-ph/9704026.

[26] 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.

[27] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca. Quantum algorithms revisited. In *Proceedings of the Royal Society of London*, volume A454, pages 339–354, 1998. quant-ph/9708016.

[28] D. Deutsch. Quantum theory, the Church-Turing principle, and the universal quantum Turing machine. In *Proceedings of the Royal Society of London*, volume A400, pages 97–117, 1985.

[29] D. Deutsch. Quantum computational networks. In *Proceedings of the Royal Society of London*, volume A425, pages 73–90, 1989.

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

[31] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, 47:777–780, 1935.

[32] 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.

[33] P. Frankl and V. Rödl. Forbidden intersections. *Transactions of the American Mathematical Society*, 300(1):259–286, 1987.

[34] M. de Graaf and R. de Wolf. On quantum versions of the Yao principle. In *Proceedings of 19th Annual Symposium on Theoretical Aspects of Computer Science (STACS'2002)*, volume 2285 of *Lecture Notes in Computer Science*, pages 347–358. Springer, 2002. quant-ph/0109070.

[35] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of 28th ACM STOC*, pages 212–219, 1996. quant-ph/9605043.

[36] A. S. Holevo. Bounds for the quantity of information transmitted by a quantum communication channel. *Problemy Peredachi Informatsii*, 9(3):3–11, 1973. English translation in *Problems of Information Transmission*, 9:177–183, 1973.

[37] P. Høyer and R. de Wolf. Improved quantum communication complexity bounds for disjointness and equality. In *Proceedings of 19th Annual Symposium on Theoretical Aspects of Computer Science (STACS'2002)*, volume 2285 of *Lecture Notes in Computer Science*, pages 299–310. Springer, 2002. quant-ph/0109068.

[38] J. Hromkovič. *Communication Complexity and Parallel Computing*. Springer, 1997.

[39] B. Kalyanasundaram and G. Schnitger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5(4):545–557, 1992. Earlier version in Structures'87.

[40] H. Klauck. On quantum and probabilistic communication: Las Vegas and one-way protocols. In *Proceedings of 32nd ACM STOC*, pages 644–651, 2000.

[41] H. Klauck. Quantum communication complexity. In *Proceedings of Workshop on Boolean Functions and Applications at 27th ICALP*, pages 241–252, 2000. quant-ph/0005032.

[42] H. Klauck. Lower bounds for quantum communication complexity. In *Proceedings of 42nd IEEE FOCS*, pages 288–297, 2001. quant-ph/0106160.

[43] 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.

[44] I. Kremer. Quantum communication. Master's thesis, Hebrew University, Computer Science Department, 1995.

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

[46] S. Massar, D. Bacon, N. Cerf, and R. Cleve. Classical simulation of quantum entanglement without local hidden variables. *Physical Review A*, 63(5), 2001. quant-ph/0009088.

[47] K. Mehlhorn and E. Schmidt. Las Vegas is better than determinism in VLSI and distributed computing. In *Proceedings of 14th ACM STOC*, pages 330–337, 1982.

[48] A. Nayak. Optimal lower bounds for quantum automata and random access codes. In *Proceedings of 40th IEEE FOCS*, pages 369–376, 1999. quant-ph/9904093.

[49] A. Nayak and J. Salzman. On communication over an entanglement-assisted quantum channel. In *Proceedings of 34th ACM STOC*, 2002.

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

[51] I. Newman and M. Szegedy. Public vs. private coin flips in one round communication games. In *Proceedings of 28th ACM STOC*, pages 561–570, 1996.

[52] M. A. Nielsen. *Quantum Information Theory.* PhD thesis, University of New Mexico, Albuquerque, 1998. quant-ph/0011036.

[53] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information.* Cambridge University Press, 2000.

[54] R. Raz. Fourier analysis for probabilistic communication complexity. *Computational Complexity*, 5(3/4):205–221, 1995.

[55] R. Raz. Exponential separation of quantum and classical communication complexity. In *Proceedings of 31st ACM STOC*, pages 358–367, 1999.

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

[57] A. Razborov. Quantum communication complexity of symmetric predicates. quant-ph/0204025, 4 Apr 2002.

[58] P. Sen and S. Venkatesh. Lower bounds in the quantum cell probe model. In *Proceedings of 28th ICALP*, volume 2076 of *Lecture Notes in Computer Science*, pages 358–369. Springer, 2001. More extensive version at quant-ph/0104100.

[59] A. Steane and W. van Dam. Physicists triumph at Guess my Number. *Physics Today*, 53(2):35–39, February 2000.

[60] A. Ta-Shma. Classical versus quantum communication complexity. *ACM SIGACT News (Complexity Column 23)*, 30:25–34, 1999.

[61] 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.

[62] W. K. Wootters and W. H. Zurek. A single quantum cannot be copied. *Nature*, 299:802–803, 1982.

[63] A. C-C. Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of 18th IEEE FOCS*, pages 222–227, 1977.

[64] A. C-C. Yao. Some complexity questions related to distributive computing. In *Proceedings of 11th ACM STOC*, pages 209–213, 1979.

[65] A. C-C. Yao. Quantum circuit complexity. In *Proceedings of 34th IEEE FOCS*, pages 352–360, 1993.

[66] Ch. Zalka. Grover's quantum searching algorithm is optimal. *Physical Review A*, 60:2746–2751, 1999. quant-ph/9711070.