

Characterization of Non-Deterministic Quantum Query and Quantum Communication Complexity

Ronald de Wolf*

Centrum voor Wiskunde en Informatica (CWI)

Kruislaan 413

1098 SJ Amsterdam, the Netherlands

rdewolf@cwi.nl

Abstract

It is known that the classical and quantum query complexities of a total Boolean function f are polynomially related to the degree of its representing polynomial, but the optimal exponents in these relations are unknown. We show that the non-deterministic quantum query complexity of f is linearly related to the degree of a “non-deterministic” polynomial for f . We also prove a quantum-classical gap of 1 vs. n for non-deterministic query complexity for a total f . In the case of quantum communication complexity there is a (partly undetermined) relation between the complexity of f and the logarithm of the rank of its communication matrix. We show that the non-deterministic quantum communication complexity of f is linearly related to the logarithm of the rank of a non-deterministic version of the communication matrix, and that it can be exponentially smaller than its classical counterpart.

1 Introduction and statement of results

There are two ways to view a classical non-deterministic algorithm for some Boolean function (or language) f . First, we may think of it as a deterministic algorithm A which receives the input x and a “certificate” y . For all inputs x , if $f(x) = 1$ then there is a certificate y such that $A(x, y) = 1$; if $f(x) = 0$ then $A(x, y) = 0$ for all y . Secondly, we may view A as a *randomized* algorithm whose acceptance probability $P(x)$ is positive if $f(x) = 1$ and $P(x) = 0$ if $f(x) = 0$. It is easy to see that these two views are equivalent in the case of classical computation: there is a view 1 algorithm for f iff there is a view 2 algorithm for f of roughly the same complexity.

Both views may be generalized to the quantum case, yielding three possibly non-equivalent definitions of non-deterministic quantum algorithms. The quantum algorithm may be required to output the right answer $f(x)$ when given an appropriate certificate (which may be quantum or classical); or the quantum algorithm may be required to have positive acceptance probability iff $f(x) = 1$. An example is given by two alternative definitions of “quantum NP”. Kitaev [28] (see also [26]) defines this class as the set of languages which are accepted by polynomial-time quantum algorithms that are given a polynomial-size quantum certificate. On the other hand, Adleman et.al. [1] and Fenner et.al. [21] define quantum NP as the set of languages L for which there is a polynomial-time quantum algorithm whose acceptance probability is positive iff $x \in L$. This quantum class was shown equal to the classical counting class co-C=P in [21], using tools from [22].

We will here adopt the latter view: a non-deterministic quantum algorithm for f is a quantum algorithm which outputs 1 with positive probability if $f(x) = 1$ and which always outputs 0 if $f(x) = 0$. (In the appendix we will show that for non-uniform settings, this definition is at least as strong as the other possible definitions.) We will study the complexity of such non-deterministic quantum algorithms in two different settings: query complexity and communication complexity. Our main results are characterizations of these complexities in algebraic terms and large gaps between quantum and classical non-deterministic complexity in both settings.

First consider the model of *query complexity*, also known as decision tree complexity or black-box complexity. Most existing quantum algorithms can naturally be expressed in this model and achieve provable speed-ups there over the best classical algorithms (e.g. [19, 39, 23, 7, 8, 9]) and also the order-finding problem on which Shor’s factoring algorithm is based [38, 15]). Let $D_q(f)$ and $Q_q(f)$ denote the query complexities of optimal deterministic and quan-

*Partially supported by the EU fifth framework project QAIP, IST-1999-11234. Also affiliated with the University of Amsterdam (ILLC).

tum algorithms that compute some $f : \{0,1\}^n \rightarrow \{0,1\}$ exactly.¹ Let $\deg(f)$ denote the degree of the multilinear polynomial that represents f . The following relations are known (see [3]; the last inequality is due to Nisan and Smolensky—unpublished, but see [13]):

$$\frac{\deg(f)}{2} \leq Q_q(f) \leq D_q(f) \leq O(\deg(f)^4).$$

Thus $\deg(f)$, $Q_q(f)$ and $D_q(f)$ are all polynomially related for all total f (the situation is very different for partial f [19, 39]). A function is known with a near-quadratic gap between $D_q(f)$ and $\deg(f)$ [33], but no function is known where $Q_q(f)$ is significantly larger than $\deg(f)$, and it may in fact be true that $Q_q(f)$ and $\deg(f)$ are linearly related. In Section 3 we show that such a linear relation holds between the *non-deterministic* versions of $Q_q(f)$ and $\deg(f)$:

$$\frac{n\deg(f)}{2} \leq NQ_q(f) \leq n\deg(f) \leq N_q(f).$$

Here $N_q(f)$ and $NQ_q(f)$ denote the query complexities of optimal non-deterministic classical and quantum algorithms for f , respectively, and $n\deg(f)$ is the minimal degree of a polynomial p which is non-zero iff $f(x) = 1$. Thus we have an algebraic characterization of the non-deterministic quantum query complexity $NQ_q(f)$, up to a factor of 2. We also show that $NQ_q(f)$ may be much smaller than $N_q(f)$: we exhibit an f where $NQ_q(f) = 1$ and $N_q(f) = n$, which is the biggest possible gap allowed by this model. Accordingly, while the case of exact computation allows at most polynomial quantum-classical gaps, the non-deterministic case allows *unbounded* gaps.

In the case of *communication complexity*, the goal is for two distributed parties, Alice and Bob, to compute some function $f : \{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}$. Alice receives an $x \in \{0,1\}^n$ and Bob receives a $y \in \{0,1\}^n$, and they want to compute $f(x,y)$, exchanging as few bits of communication as possible. This setting was introduced by Yao [41] and is fairly well understood for the case where Alice and Bob are classical players exchanging classical bits [30]. Much less is known about *quantum* communication complexity, where Alice and Bob have a quantum computer and can exchange qubits. This was first studied by Yao [42] and it was shown later that quantum communication complexity can be significantly smaller than classical communication complexity [16, 10, 2, 35].

Let $D_c(f)$ and $Q_c(f)$ denote the communication required for optimal deterministic and quantum protocols for computing f , respectively (we assume Alice and Bob do not share any prior entanglement). Let $\text{rank}(f)$ be the

rank of the $2^n \times 2^n$ communication matrix M_f defined by $M_f(x,y) = f(x,y)$. The following relations are known:

$$\frac{\log \text{rank}(f)}{2} \leq Q_c(f) \leq D_c(f).$$

The first inequality follows from work of Kremer [29] and Yao [42], as first noted in [10] (in [12] it is shown that this lower bound also holds if the quantum protocol can make use of unlimited prior entanglement between Alice and Bob). It is an open question whether $D_c(f)$ can in turn be upper bounded by some polynomial in $\log \text{rank}(f)$. This is known as the *log-rank conjecture*. If this conjecture holds, then $D_c(f)$ and $Q_c(f)$ are polynomially related for all total f (which may well be true). It is known that $\log \text{rank}(f)$ and $D_c(f)$ are not linearly related [34]. In Section 4 we show that the *non-deterministic* versions of $\log \text{rank}(f)$ and $Q_c(f)$ are in fact linearly related:

$$\frac{\log n\text{rank}(f)}{2} \leq NQ_c(f) \leq \log n\text{rank}(f) \leq N_c(f).$$

Here $n\text{rank}(f)$ denotes the minimal rank of a matrix whose (x,y) -entry is non-zero iff $f(x,y) = 1$. Thus we can characterize the non-deterministic quantum communication complexity as the logarithm of the rank of its non-deterministic matrix. Two other log-rank-style characterizations of certain variants of communication complexity are known: the communication complexity of quantum sampling [2] and modular communication complexity [31].

We also show an exponential gap between quantum and classical non-deterministic communication complexity: we exhibit an f where $NQ_c(f) \leq \log(n+1)$ and $N_c(f) \in \Omega(n)$. Cleve and Massar [18] earlier found another gap: $NQ_c(\text{NE}) = 1$ versus $N_c(\text{NE}) = \log n + 1$, where NE is the non-equality function.

2 Preliminaries

2.1 Functions and polynomials

For $x \in \{0,1\}^n$ we use $|x|$ for the Hamming weight (number of 1s) of x , and x_i for its i th bit, $i \in \{1, \dots, n\}$. We use $\vec{0}$ for a string of n zeroes. If $x, y \in \{0,1\}^n$ then $x \wedge y$ denotes the n -bit string obtained by bitwise ANDing x and y . Let $f : \{0,1\}^n \rightarrow \{0,1\}$ be a total Boolean function. For example, $\text{OR}(x) = 1$ iff $|x| \geq 1$, $\text{AND}(x) = 1$ iff $|x| = n$, $\text{PARITY}(x) = 1$ iff $|x|$ is odd. We use \bar{f} for the function $1 - f$.

For $b \in \{0,1\}$, a b -certificate for f is an assignment $C : S \rightarrow \{0,1\}$ to some set S of variables, such that $f(x) = b$ whenever x is consistent with C . The *size* of C is $|S|$. The *certificate complexity* $C_x(f)$ of f on input x is the minimal size of an $f(x)$ -certificate that is consistent with x . We define the 1-certificate complexity of f

¹Unfortunately, the notation $D(f)$ is used for deterministic complexity in both the field of decision tree complexity and in communication complexity. To avoid confusion, we will consistently add subscripts ‘ q ’ for query complexity and ‘ c ’ for communication complexity.

as $C^{(1)}(f) = \max_{x:f(x)=1} C_x(f)$. Similarly we define $C^{(0)}(f)$. For example, $C^{(1)}(\text{OR}) = 1$ and $C^{(0)}(\text{OR}) = n$.

An n -variate *multilinear polynomial* is a function $p : \mathbf{R}^n \rightarrow \mathbf{R}$ which can be written as

$$p(x) = \sum_{S \subseteq \{1, \dots, n\}} a_S X_S.$$

Here S ranges over all sets of indices of variables, a_S is a real number, and the monomial X_S is the product $\prod_{i \in S} x_i$ of all variables in S . The *degree* $\deg(p)$ of p is the degree of a largest monomial with non-zero coefficient. It is well known that every total Boolean f has a unique polynomial p such that $p(x) = f(x)$ for all $x \in \{0, 1\}^n$. Let $\deg(f)$ be the degree of this polynomial, which is at most n . For example, $\text{OR}(x_1, x_2) = x_1 + x_2 - x_1 x_2$, which has degree 2. Every multilinear polynomial $p = \sum_S a_S X_S$ can also be written out uniquely in the so-called *Fourier basis*:

$$p(x) = \sum_S c_S (-1)^{x \cdot S}.$$

Again S ranges over all sets of indices of variables (we often identify a set S with its characteristic n -bit vector), c_S is a real number, and $x \cdot S$ denotes the inner product of the n -bit strings x and S , equivalently $x \cdot S = |x \wedge S| = \sum_{i \in S} x_i$. It is easy to see that $\deg(p) = \max\{|S| \mid c_S \neq 0\}$. For example, $\text{OR}(x_1, x_2) = \frac{3}{4} - \frac{1}{4}(-1)^{x_1} - \frac{1}{4}(-1)^{x_2} - \frac{1}{4}(-1)^{x_1+x_2}$ in the Fourier basis. We refer to [4, 33, 13] for more details about polynomial representations of Boolean functions.

We introduce the notion of a *non-deterministic polynomial* for f . This is a polynomial p such that $p(x) \neq 0$ iff $f(x) = 1$. Let the *non-deterministic degree* of f , denoted $\text{ndeg}(f)$, be the minimum degree among all non-deterministic polynomials p for f . Without loss of generality we can assume $p(x) \in [-1, 1]$ for all $x \in \{0, 1\}^n$ (if not, just divide by $\max_x |p(x)|$).

We mention some upper and lower bounds for $\text{ndeg}(f)$. For example, $p(x) = \sum_i x_i/n$ is a non-deterministic polynomial for OR, hence $\text{ndeg}(\text{OR}) = 1$. More generally, let f be a non-constant symmetric function (i.e. $f(x)$ only depends on $|x|$). Suppose f achieves value 0 on z Hamming weights, k_1, \dots, k_z . Since $|x| = \sum_i x_i$, it is easy to see that $(|x| - k_1)(|x| - k_2) \cdots (|x| - k_z)$ is a non-deterministic polynomial for f , hence $\text{ndeg}(f) \leq z$. This upper bound is tight for AND (see below) but not for PARITY. For example, $p(x_1, x_2) = x_1 - x_2$ is a degree-1 non-deterministic polynomial for PARITY on 2 variables: it assumes value 0 on x -weights 0 and 2, and ± 1 on weight 1. Using standard symmetrization techniques (as used for instance in [32, 33, 3]) we can also show the general lower bound $\text{ndeg}(f) \geq z/2$ for symmetric f . Furthermore, it is easy to show that $\text{ndeg}(f) \leq C^{(1)}(f)$ for every f (take a polynomial which is the ‘‘sum’’ over all 1-certificates for f).

Finally, we mention a general lower bound on $\text{ndeg}(f)$. Let $\Pr[p \neq 0] = |\{x \in \{0, 1\}^n \mid p(x) \neq 0\}|/2^n$ denote the probability that a random Boolean input x makes a function p non-zero. A lemma of Schwartz [37] (see also [33, Section 2.2]) states that if p is a non-constant multilinear polynomial of degree d , then $\Pr[p \neq 0] \geq 2^{-d}$, hence $d \geq \log(1/\Pr[p \neq 0])$. Since a non-deterministic polynomial p for f is non-zero iff $f(x) = 1$, it follows that

$$\text{ndeg}(f) \geq \log(1/\Pr[f \neq 0]) = \log(1/\Pr[f = 1]).$$

Accordingly, functions with a very small fraction of 1-inputs will have high non-deterministic degree. For instance, $\Pr[\text{AND} = 1] = 2^{-n}$, so $\text{ndeg}(\text{AND}) = n$.

2.2 Query complexity

We assume familiarity with classical computation and briefly sketch the setting of quantum computation (see e.g. [5, 27, 14] for more details). An m -qubit state is a linear combination of all classical m -bit states

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

where $|i\rangle$ denotes the basis state i (a classical m -bit string), and α_i is a complex number which is called the *amplitude* of $|i\rangle$. We require $\sum_i |\alpha_i|^2 = 1$. Viewing $|\phi\rangle$ as a 2^m -dimensional column vector, we use $\langle\phi|$ for the row vector which is the conjugate transpose of $|\phi\rangle$. Note that the inner product $\langle i|j\rangle$ is 1 if $i = j$ and is 0 otherwise. When we observe $|\phi\rangle$ we will see $|i\rangle$ with probability $|\langle i|\phi\rangle|^2 = |\alpha_i|^2$, and the state will collapse to the observed $|i\rangle$. A quantum operation which is not an observation, corresponds to a unitary (=norm-preserving) transformation U on the 2^m -dimensional vector of amplitudes.

For some input $x \in \{0, 1\}^n$, a *query* corresponds to the unitary transformation O which maps $|i, b, z\rangle \rightarrow |i, b \oplus x_i, z\rangle$. Here $b \in \{0, 1\}$; the z -part corresponds to the workspace, which is not affected by the query. We assume that the input can only be accessed via such queries. A T -query quantum algorithm has the form $A = U_T O U_{T-1} \dots O U_1 O U_0$, where the U_k are fixed unitary transformations, independent of the input x . This A depends on x via the T applications of O . We sometimes write A_x to emphasize this. The algorithm starts in initial state $|\vec{0}\rangle$ and its *output* is the bit obtained from observing the leftmost qubit of the final superposition $A|\vec{0}\rangle$. The *acceptance probability* of A (on input x) is its probability of outputting 1 (on x).

We will consider classical and quantum algorithms, and will only count the number of queries these algorithms make on the worst-case input (see [3, 13] for more details). Let $D_q(f)$ and $Q_q(f)$ be the query complexities of optimal

deterministic classical and quantum algorithms for computing f , respectively. $D_q(f)$ is also known as the decision tree complexity of f . A *non-deterministic algorithm* for f is an algorithm that has positive acceptance probability on input x iff $f(x) = 1$. Let $N_q(f)$ and $NQ_q(f)$ be the query complexities of optimal non-deterministic classical and quantum algorithms for f , respectively (in the appendix we show that this definition of $NQ_q(f)$ is at least as powerful as the other possible definitions).

The 1-certificate complexity characterizes the classical non-deterministic complexity of f :

Proposition 1 $N_q(f) = C^{(1)}(f)$.

Proof

$N_q(\mathbf{f}) \leq C^{(1)}(\mathbf{f})$: a classical algorithm that guesses a 1-certificate, queries its variables, and outputs 1 iff the certificate holds, is a non-deterministic algorithm for f .

$N_q(\mathbf{f}) \geq C^{(1)}(\mathbf{f})$: a non-deterministic algorithm for f can only output 1 if the outcomes of the queries that it has made force the function to 1. Hence if x is an input where all 1-certificates have size at least $C^{(1)}(f)$, then the algorithm will have to query at least $C^{(1)}(f)$ variables before it can output 1 (which it must do on some runs). \square

In Section 3 we will characterize $NQ_q(f)$ in terms of $ndeg(f)$, using the following result from [3].

Lemma 1 (BBCMW) *The amplitudes of the basis states in the final superposition of a T -query quantum algorithm can be written as multilinear complex-valued polynomials of degree $\leq T$ in the n x_i -variables. Therefore the acceptance probability of the algorithm (which is the sum of squares of some of those amplitudes) can be written as an n -variate multilinear polynomial $P(x)$ of degree $\leq 2T$.*

2.3 Communication complexity

Below we sketch the setting of communication complexity. For more details and results we refer to the book of Kushilevitz and Nisan [30].

Let $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$. For example, $EQ(x, y) = 1$ iff $x = y$, $NE(x, y) = 1$ iff $x \neq y$, $DISJ(x, y) = 1$ iff $|x \wedge y| = 0$. A *rectangle* is a subset $R = S \times T$ of the domain of f . R is a *1-rectangle* (for f) if $f(x, y) = 1$ for all $(x, y) \in R$. A *1-cover* for f is a set of 1-rectangles which covers all 1-inputs of f . $C^1(f)$ denotes the minimal size (i.e. minimal number of rectangles) of a 1-cover for f . Similarly we define 0-rectangles, 0-covers, and $C^0(f)$. (These $C^1(f)$ and $C^0(f)$ should not be confused with the certificate complexities $C^{(1)}(f)$ and $C^{(0)}(f)$.)

The *communication matrix* M_f of f is the $2^n \times 2^n$ Boolean matrix whose x, y entry is $f(x, y)$, and $rank(f)$ denotes the rank of M_f over the reals. An $2^n \times 2^n$ matrix

M is called a *non-deterministic communication matrix* for f if it has the property that $M(x, y) \neq 0$ iff $f(x, y) = 1$. Thus M is any matrix obtainable by replacing 1-entries in M_f by non-zero reals. Let the *non-deterministic rank* of f , denoted $nrank(f)$, be the minimum rank over all non-deterministic matrices M for f . Without loss of generality we can assume all M -entries are in $[-1, 1]$.

We consider classical and quantum communication protocols, and only count the amount of communication (bits or qubits) these protocols make on the worst-case input. For classical randomized protocols we assume Alice and Bob each have their own private coin flips. Let $D_c(f)$ and $Q_c(f)$ be the communication complexities of optimal deterministic classical and quantum protocols for computing f , respectively. A *non-deterministic protocol* for f is a protocol that has positive acceptance probability iff $f(x, y) = 1$. Let $N_c(f)$ and $NQ_c(f)$ be the communication complexities of optimal non-deterministic classical and quantum protocols for f , respectively. $N_c(f)$ is called $N^1(f)$ in [30].

It is not hard to show that $N_c(f) = \lceil \log C^1(f) \rceil$. In Section 4 we will characterize $NQ_c(f)$ in terms of $nrank(f)$. As noticed in [10], the following very useful lemma is implied by results in [42, 29]:

Lemma 2 (Kremer/Yao) *The acceptance probabilities of an ℓ -qubit quantum communication protocol can be written as a $2^n \times 2^n$ matrix $P(x, y)$ of rank $\leq 2^{2\ell}$.*

3 Non-deterministic quantum query complexity

Here we show a tight relation between non-deterministic quantum query complexity $NQ_q(f)$ and non-deterministic degree $ndeg(f)$. The upper bound uses a trick similar to the one used in [21] to show $\text{co-C=P} \subseteq \text{quantum-NP}$.

Theorem 1 $\frac{ndeg(f)}{2} \leq NQ_q(f) \leq ndeg(f)$.

Proof Suppose we have an $NQ_q(f)$ -query non-deterministic quantum algorithm A for f . By Lemma 1, its acceptance probability can be written as a polynomial $P(x)$ of degree $\leq 2NQ_q(f)$. Because A is a non-deterministic algorithm for f , $P(x)$ is a non-deterministic polynomial for f . Hence $ndeg(f) \leq 2NQ_q(f)$.

For the upper bound: let $p(x)$ be a non-deterministic polynomial for f of degree $d = ndeg(f)$. Recall that $x \cdot S$ denotes $|x \wedge S|$, identifying $S \subseteq \{1, \dots, n\}$ with its characteristic n -bit vector. We write p in the Fourier basis:

$$p(x) = \sum_S c_S (-1)^{x \cdot S}.$$

Since $deg(p) = \max\{|S| \mid c_S \neq 0\}$, we have that $c_S \neq 0$ only if $|S| \leq d$.

We can make a unitary transformation F which uses d queries and maps $|S\rangle \rightarrow (-1)^{x \cdot S} |S\rangle$ whenever $|S| \leq d$. Informally, this transformation does a controlled parity-computation: it computes $|x \cdot S| \pmod{2}$ using $|S|/2$ queries [3, 20] and then reverses the computation to clean up the workspace (at the cost of another $|S|/2$ queries). By a standard trick, the answer $|x \cdot S| \pmod{2}$ can then be turned into a phase factor $(-1)^{|x \cdot S| \pmod{2}} = (-1)^{x \cdot S}$.

Now consider the following quantum algorithm:

1. Start with $c \sum_S c_S |S\rangle$ (an n -qubit state, where $c = 1/\sqrt{\sum_S c_S^2}$ is a normalizing constant)
2. Apply F to the state
3. Apply a Hadamard transform H to each qubit
4. Measure the final state and output 1 if the outcome is the all-zero state $|\vec{0}\rangle$ and output 0 otherwise.

The acceptance probability (i.e. the probability of observing $|\vec{0}\rangle$ at the end) is

$$\begin{aligned} P(x) &= |\langle \vec{0} | H^{\otimes n} F c \sum_S c_S |S\rangle|^2 \\ &= \frac{c^2}{2^n} \left| \sum_{S'} \langle S' | \sum_S c_S (-1)^{x \cdot S} |S\rangle \right|^2 \\ &= \frac{c^2}{2^n} \left| \sum_S c_S (-1)^{x \cdot S} \right|^2 = \frac{c^2 p(x)^2}{2^n}. \end{aligned}$$

Since $p(x)$ is non-zero iff $f(x) = 1$, $P(x)$ will be positive iff $f(x) = 1$. Hence we have a non-deterministic quantum algorithm for f with $d = ndeg(f)$ queries. \square

The upper bound in this theorem is tight: by a proof similar to [3, Proposition 6.1] we can show $NQ_q(\text{AND}) = ndeg(\text{AND}) = n$. We do not know if the factor of 2 in the lower bound can be dispensed with. If we were to change the output requirement of the quantum algorithm a little bit, by saying that the algorithm accepts iff measuring the final superposition gives basis state $|\vec{0}\rangle$, then the required number of queries is exactly $ndeg(f)$. The upper bound of $ndeg(f)$ queries in this changed model is the same as above. The lower bound of $ndeg(f)$ queries follows since the amplitude of the basis state $|\vec{0}\rangle$ in the final superposition must now be non-zero iff $f(x) = 1$, and this polynomial has degree at most the number of queries (Lemma 1).

What is the biggest possible gap between quantum and classical non-deterministic query complexity? Consider the Boolean function f defined by

$$f(x) = 1 \text{ iff } |x| \neq 1.$$

It is easy to see that $N_q(f) = C^{(1)}(f) = C^{(0)}(f) = n$. On the other hand, the following is a degree-1 non-deterministic polynomial for f :

$$p(x) = \frac{\sum_i x_i - 1}{n - 1}. \quad (1)$$

Thus $ndeg(f) = 1$ and by Theorem 1 we have $NQ_q(f) = 1$. For the complement of f , we can easily show $NQ_q(\bar{f}) \geq n/2$ using Lemma 1, since the acceptance probability of a non-deterministic algorithm for \bar{f} must be 0 on n Hamming weights and hence have degree at least n (this $NQ_q(\bar{f}) \geq n/2$ is tight for $n = 2$, witness $p(x) = x_1 - x_2$). In sum:

Theorem 2 For the above f we have $NQ_q(f) = 1$, $NQ_q(\bar{f}) \geq n/2$ and $N_q(f) = N_q(\bar{f}) = n$.

A slightly smaller gap holds for the function defined by $\text{DeJo}(x) = 1$ iff $|x| \neq n/2$. This is a total version of the well known Deutsch-Jozsa promise problem [19]. The algorithm of [19] (in its 1-query version [17]) turns out to be a non-deterministic algorithm for DeJo, so $NQ_q(\text{DeJo}) = 1$. In contrast, $N_q(\text{DeJo}) = C^{(1)}(\text{DeJo}) = n/2 + 1$.

4 Non-deterministic quantum communication complexity

Here we characterize the non-deterministic quantum communication complexity $NQ_c(f)$ in terms of the non-deterministic rank $nrank(f)$:

Theorem 3 $\frac{\log nrank(f)}{2} \leq NQ_c(f) \leq \lceil \log nrank(f) \rceil$.

Proof Consider an $NQ_c(f)$ -qubit non-deterministic quantum protocol for f . By Lemma 2, its acceptance probability $P(x, y)$ determines a matrix of rank $\leq 2^{2NQ_c(f)}$. It is easy to see that this is a non-deterministic matrix for f , hence $nrank(f) \leq 2^{2NQ_c(f)}$ and the first inequality follows.

For the upper bound, let $r = nrank(f)$ and M be a rank- r non-deterministic matrix for f . Let $M^T = U\Sigma V$ be the singular value decomposition of M^T (see [25, Chapter 3]), 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 1-round non-deterministic protocol for f , using $\lceil \log r \rceil$ qubits. First Alice prepares the vector $|\phi_x\rangle = c_x \Sigma V |x\rangle$, 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 $\lceil \log r \rceil$ 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 outputs 1, otherwise he outputs 0. The acceptance probability of this protocol is

$$\begin{aligned} P(x, y) &= |\langle y | U | \phi_x \rangle|^2 = c_x^2 |\langle y | U \Sigma V | x \rangle|^2 \\ &= c_x^2 |M^T(y, x)|^2 = c_x^2 |M(x, y)|^2. \end{aligned}$$

Since $M(x, y)$ is non-zero iff $f(x, y) = 1$, $P(x, y)$ will be positive iff $f(x, y) = 1$. Thus we have a non-deterministic protocol for f with $\lceil \log r \rceil$ qubits. \square

Thus classically we have $N_c(f) = \lceil \log C^1(f) \rceil$ and quantumly we have $NQ_c(f) \approx \log nrank(f)$. We now give an f with an exponential gap between $N_c(f)$ and $NQ_c(f)$. For $n > 1$, define f by

$$f(x, y) = 1 \text{ iff } |x \wedge y| \neq 1.$$

We first show that the quantum complexity $NQ_c(f)$ is low:

Theorem 4 *For the above f we have $NQ_c(f) \leq \lceil \log(n+1) \rceil$.*

Proof By Theorem 3, it suffices to prove $nrank(f) \leq n+1$. We will derive a low-rank non-deterministic matrix from the polynomial p of equation 1, using a technique from [34]. Let M_i be the matrix defined by $M_i(x, y) = 1$ if $x_i = y_i = 1$, and $M_i(x, y) = 0$ otherwise. Notice that M_i has rank 1. Now define a $2^n \times 2^n$ matrix M by

$$M(x, y) = \frac{\sum_i M_i(x, y) - 1}{n-1}.$$

Note that $M(x, y) = p(x \wedge y)$. Since p is a non-deterministic polynomial for the function which is 1 iff its input does not have weight 1, it can be seen that M is a non-deterministic matrix for f . Because M is the sum of $n+1$ rank-1 matrices, M itself has rank at most $n+1$. \square

Now we show that the classical $N_c(f)$ is high (both for f and its complement):

Theorem 5 *For the above f we have $N_c(f) \in \Omega(n)$ and $N_c(\bar{f}) \geq n-1$.*

Proof Let R_1, \dots, R_k be a minimal 1-cover for f . We use the following result from [30, Example 3.22 and Section 4.6], which is essentially due to Razborov [36].

There exist sets $A, B \subseteq \{0, 1\}^n \times \{0, 1\}^n$ and a probability distribution $\mu : \{0, 1\}^n \times \{0, 1\}^n \rightarrow [0, 1]$ such that all $(x, y) \in A$ have $|x \wedge y| = 0$, all $(x, y) \in B$ have $|x \wedge y| = 1$, $\mu(A) = 3/4$, and there are $\alpha, \delta > 0$ such that for all rectangles R , $\mu(R \cap B) \geq \alpha \cdot \mu(R \cap A) - 2^{-\delta n}$.

Since the R_i are 1-rectangles, they cannot contain elements from B . Hence $\mu(R_i \cap B) = 0$ and $\mu(R_i \cap A) \leq 2^{-\delta n}/\alpha$. But since all elements of A are covered by the R_i we have

$$\frac{3}{4} = \mu(A) = \mu\left(\bigcup_{i=1}^k (R_i \cap A)\right) \leq \sum_{i=1}^k \mu(R_i \cap A) \leq k \cdot \frac{2^{-\delta n}}{\alpha}.$$

Therefore $N_c(f) = \lceil \log k \rceil \geq \delta n + \log(3\alpha/4)$.

For the lower bound on $N_c(\bar{f})$, consider the set $S = \{(x, y) \mid x_1 = y_1 = 1, x_i = \bar{y}_i \text{ for } i > 1\}$. This S contains 2^{n-1} elements, all of which are 1-inputs for \bar{f} . Note that if (x, y) and (x', y') are two elements from S then $|x \wedge y'| > 1$ or $|x' \wedge y| > 1$, so a 1-rectangle for \bar{f} can contain at most one element of S . This shows that a minimal 1-cover for \bar{f} requires at least 2^{n-1} rectangles and $N_c(\bar{f}) \geq n-1$. \square

Another quantum-classical separation was obtained earlier by Richard Cleve and Serge Massar [18]:

Theorem 6 (Cleve & Massar) *For the non-equality problem on n bits, we have $NQ_c(\text{NE}) = 1$ versus $N_c(\text{NE}) = \log n + 1$.*

Proof $N_c(\text{NE}) = \log n + 1$ is well known (see [30, Example 2.5]). Below we give Cleve and Massar's 1-qubit non-deterministic protocol for NE.

Viewing her input x as a number $\in [0, 2^n - 1]$, Alice rotates a $|0\rangle$ -qubit over an angle $x\pi/2^n$, obtaining a qubit $\cos(x\pi/2^n)|0\rangle + \sin(x\pi/2^n)|1\rangle$ which she sends to Bob. Bob rotates the qubit back over an angle $y\pi/2^n$, obtaining $\cos((x-y)\pi/2^n)|0\rangle + \sin((x-y)\pi/2^n)|1\rangle$. Bob now measures the qubit and outputs the observed bit. If $x = y$ then $\sin((x-y)\pi/2^n) = 0$, so Bob will always output 0. If $x \neq y$ then $\sin((x-y)\pi/2^n) \neq 0$, so Bob will output 1 with positive probability. \square

Note that $nrank(\text{EQ}) = 2^n$, since any non-deterministic matrix for equality will be a diagonal $2^n \times 2^n$ matrix with non-zero diagonal entries. Thus $NQ_c(\text{EQ}) \geq (\log nrank(\text{EQ}))/2 = n/2$, which contrasts sharply with the non-deterministic quantum complexity $NQ_c(\text{NE}) = 1$ of its complement.

5 Future work

One of the main reasons for the usefulness of non-deterministic query and communication complexities in the classical case, is the tight relation of these complexities with deterministic complexity. In the query complexity (decision tree) setting we have

$$\max\{N_q(f), N_q(\bar{f})\} \leq D_q(f) \leq N_q(f)N_q(\bar{f}).$$

This was independently shown in [6, 24, 40]. We conjecture that something similar holds in the quantum case:

$$\max\left\{\frac{ndeg(f)}{2}, \frac{ndeg(\bar{f})}{2}\right\} \leq \frac{deg(f)}{2} \leq Q_q(f) \leq D_q(f) \stackrel{?}{\leq} O(NQ_q(f)NQ_q(\bar{f})) = O(ndeg(f)ndeg(\bar{f})).$$

Here the $?$ -part is open. This conjecture would imply $D_q(f) \in O(Q_0(f)^2)$ ($Q_0(f)$ is zero-error quantum query complexity; the quadratic relation would be close to optimal [11]). It would also imply $D_q(f) \in O(\deg(f)^2)$, which is again close to optimal [33]. The currently best known relations have a fourth power instead of a square.

Similarly, for communication complexity the following is known [30, Section 2.11]:

$$\max\{N_c(f), N_c(\bar{f})\} \leq D_c(f) \leq O(N_c(f)N_c(\bar{f})).$$

An analogous result might be true for quantum, but we have been unable to prove it.

Acknowledgments. I thank Harry Buhrman, Richard Cleve, Wim van Dam, and John Watrous for useful discussions and pointers to the literature.

References

- [1] L. M. Adleman, J. Demarrais, and M. A. Huang. Quantum computability. *SIAM Journal on Computing*, 26(5):1524–1540, 1997.
- [2] A. Ambainis, L. Schulman, A. Ta-Shma, U. Vazirani, and A. Wigderson. The quantum communication complexity of sampling. In *Proceedings of 39th FOCS*, pages 342–351, 1998.
- [3] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf. Quantum lower bounds by polynomials. In *Proceedings of 39th FOCS*, pages 352–361, 1998. quant-ph/9802049.
- [4] R. Beigel. The polynomial method in circuit complexity. In *Proceedings of the 8th IEEE Structure in Complexity Theory Conference*, pages 82–95, 1993.
- [5] A. Berthiaume. Quantum computation. In A. Selman and L. Hemaspaandra, editors, *Complexity Theory Retrospective II*, pages 23–51. Springer, 1997.
- [6] M. Blum and R. Impagliazzo. Generic oracles and oracle classes (extended abstract). In *Proceedings of 28th FOCS*, pages 118–126, 1987.
- [7] G. Brassard and P. Høyer. An exact quantum polynomial-time algorithm for Simon’s problem. In *Proceedings of the 5th Israeli Symposium on Theory of Computing and Systems (ISTCS’97)*, pages 12–23, 1997. quant-ph/9704027.
- [8] G. Brassard, P. Høyer, and A. Tapp. Quantum algorithm for the collision problem. *ACM SIGACT News (Cryptography Column)*, 28:14–19, 1997. quant-ph/9705002.
- [9] G. Brassard, P. Høyer, and A. Tapp. Quantum counting. In *Proceedings of 25th ICALP*, volume 1443 of *Lecture Notes in Computer Science*, pages 820–831. Springer, 1998. quant-ph/9805082.
- [10] H. Buhrman, R. Cleve, and A. Wigderson. Quantum vs. classical communication and computation (preliminary version). In *Proceedings of 30th STOC*, pages 63–68, 1998. quant-ph/9802040.
- [11] H. Buhrman, R. Cleve, R. de Wolf, and C. Zalka. Bounds for small-error and zero-error quantum algorithms. In *Proceedings of 40th FOCS*, pages 358–368, 1999. cs.CC/9904019.
- [12] H. Buhrman and R. de Wolf. Communication complexity lower bounds by polynomials. cs.CC/9910010, 1999.
- [13] H. Buhrman and R. de Wolf. Complexity measures and decision tree complexity: A survey. Submitted. Available at <http://www.cwi.nl/~rdewolf>, 1999.
- [14] R. Cleve. An introduction to quantum complexity theory. quant-ph/9906111, 28 Jun 1999.
- [15] R. Cleve. The query complexity of order-finding. quant-ph/9911124, 30 Nov 1999.
- [16] R. Cleve and H. Buhrman. Substituting quantum entanglement for communication. *Physical Review A*, 56(2):1201–1204, 1997. quant-ph/9704026.
- [17] 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.
- [18] R. Cleve and S. Massar. Paper in preparation, 1999.
- [19] 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.
- [20] 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.
- [21] S. Fenner, F. Green, S. Homer, and R. Pruim. Determining acceptance possibility for a quantum computation is hard for the polynomial hierarchy. In *Proceedings of the 6th Italian Conference on Theoretical Computer Science*, pages 241–252, 1998. quant-ph/9812056.
- [22] L. Fortnow and J. Rogers. Complexity limitations on quantum computation. *Journal of Computer and Systems Sciences*, 59(2):240–252, 1999. Earlier version in Complexity’98. Also cs.CC/9811023.
- [23] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of 28th STOC*, pages 212–219, 1996. quant-ph/9605043.
- [24] J. Hartmanis and L. Hemachandra. One-way functions, robustness and the non-isomorphism of NP-complete sets. In *Proceedings of the 2nd IEEE Structure in Complexity Theory Conference*, pages 160–174, 1987.
- [25] R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1991.
- [26] A. Kitaev and J. Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof systems. In *Proceedings of 32nd STOC*, 2000. To appear.
- [27] A. Y. Kitaev. Quantum computations: Algorithms and error correction. *Russian Mathematical Surveys*, 52(6):1191–1249, 1997.
- [28] A. Y. Kitaev. Quantum NP, January 1999. Talk given at AQIP’99, DePaul University, Chicago.
- [29] I. Kremer. Quantum communication. Master’s thesis, Hebrew University, Computer Science Department, 1995.
- [30] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [31] C. Meinel and S. Waack. The “log rank” conjecture for modular communication complexity. In *Proceedings of 13th Annual Symposium on Theoretical Aspects of Computer Science (STACS’96)*, volume 1046 of *Lecture Notes in Computer Science*, pages 619–630. Springer, 1996.

- [32] M. Minsky and S. Papert. *Perceptrons*. MIT Press, Cambridge, MA, 1968. Second, expanded edition 1988.
- [33] N. Nisan and M. Szegedy. On the degree of Boolean functions as real polynomials. *Computational Complexity*, 4(4):301–313, 1994. Earlier version in STOC’92.
- [34] N. Nisan and A. Wigderson. On rank vs. communication complexity. *Combinatorica*, 15(4):557–565, 1995. Earlier version in FOCS’94.
- [35] R. Raz. Exponential separation of quantum and classical communication complexity. In *Proceedings of 31st STOC*, pages 358–367, 1999.
- [36] A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385–390, 1992.
- [37] J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27:701–717, 1980.
- [38] 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. Earlier version in FOCS’94. quant-ph/9508027.
- [39] D. Simon. On the power of quantum computation. *SIAM Journal on Computing*, 26(5):1474–1483, 1997. Earlier version in FOCS’94.
- [40] G. Tardos. Query complexity, or why is it difficult to separate $NP^A \cap coNP^A$ from P^A by random oracles A ? *Combinatorica*, 9(4):385–392, 1989.
- [41] A. C.-C. Yao. Some complexity questions related to distributive computing. In *Proceedings of 11th STOC*, pages 209–213, 1979.
- [42] A. C.-C. Yao. Quantum circuit complexity. In *Proceedings of 34th FOCS*, pages 352–360, 1993.

A Comparison with alternative definitions

As mentioned in the introduction, three different definitions of non-deterministic quantum complexity are possible. We may consider the complexity of quantum algorithms which either:

1. output 1 iff given an appropriate *classical* certificate (such certificates must exist iff $f(x) = 1$)
2. output 1 iff given an appropriate *quantum* certificate (such certificates must exist iff $f(x) = 1$)
3. output 1 with positive probability iff $f(x) = 1$

The third definition is the one we adopted for this paper. Clearly the second definition is at least as strong as the first. Here we will show that the third definition is at least as strong as the second. (We give the proof for the query complexity setting, but the same proof works for communication complexity and other non-uniform settings as well.) Thus our $NQ_q(f)$ is in fact the most powerful definition of non-deterministic quantum query complexity.

We formalize the second definition as follows: a T -query quantum verifier for f is a T -query quantum algorithm V

together with a set \mathcal{C} of m -qubit states, such that for all $x \in \{0, 1\}^n$ we have: (1) if $f(x) = 1$ then there is a $|\phi_x\rangle \in \mathcal{C}$ such that $V_x|\phi_x\rangle$ has acceptance probability 1, and (2) if $f(x) = 0$ then $V_x|\phi\rangle$ has acceptance probability 0 for every $|\phi\rangle \in \mathcal{C}$. Informally: the set \mathcal{C} contains all possible certificates, (1) for every 1-input there is a verifiable 1-certificate in \mathcal{C} , and (2) for 0-inputs there aren’t any. We do not put any constraints on \mathcal{C} . However, note that the definition implies that if $f(x) = 0$ for some x , then \mathcal{C} cannot contain *all* m -qubit states: otherwise $|\phi_x\rangle = V_x^{-1}|1\vec{0}\rangle$ would be a 1-certificate in \mathcal{C} even for x with $f(x) = 0$.

We now prove that a T -query quantum verifier can be turned into a T -query non-deterministic quantum algorithm according to our third definition. This shows that the third definition is at least as powerful as the second (in fact, this even holds if we replace the acceptance probability 1 in clause (1) of the definition of a quantum verifier by just positive acceptance probability — in this case both definitions are equivalent).

Theorem 7 *Suppose there exists a T -query quantum verifier V for f . Then $NQ_q(f) \leq T$.*

Proof The verifier V and the associated set \mathcal{C} satisfy:

1. if $f(x) = 1$ then there is a $|\phi_x\rangle \in \mathcal{C}$ such that $V_x|\phi_x\rangle$ has acceptance probability 1
2. if $f(x) = 0$ then $V_x|\phi\rangle$ has acceptance probability 0 for all $|\phi\rangle \in \mathcal{C}$

Let $X_1 = \{z \mid f(z) = 1\}$. For each $z \in X_1$ choose one specific 1-certificate $|\phi_z\rangle \in \mathcal{C}$. Now let us consider some input x and see what happens if we run $V_x \otimes I$ (where I is the $2^n \times 2^n$ identity operation) on the $m + n$ -qubit state

$$|\phi\rangle = \frac{1}{\sqrt{|X_1|}} \sum_{z \in X_1} |\phi_z\rangle |z\rangle.$$

V_x only acts on the first m qubits of $|\phi\rangle$, the $|z\rangle$ -part remains unaffected. Therefore running $V_x \otimes I$ on $|\phi\rangle$ gives the same acceptance probabilities as when we first randomly choose some $z \in X_1$ and then apply V_x to $|\phi_z\rangle$. In case $f(x) = 0$, this $V_x|\phi_z\rangle$ will have acceptance probability 0, so $(V_x \otimes I)|\phi\rangle$ will have acceptance probability 0 as well. In case the input x is such that $f(x) = 1$, the specific certificate $|\phi_x\rangle$ that we chose for this x will satisfy that $V_x|\phi_x\rangle$ has acceptance probability 1. But then $(V_x \otimes I)|\phi\rangle$ has acceptance probability at least $1/|X_1|$. Accordingly, $(V_x \otimes I)|\phi\rangle$ has positive acceptance probability iff $f(x) = 1$. By prefixing $V_x \otimes I$ with a unitary transformation which maps $|\vec{0}\rangle$ (of $m + n$ qubits) to $|\phi\rangle$, we have constructed a non-deterministic quantum algorithm for f with T queries. \square