THEORY OF COMPUTING, Volume 1 (2005), pp. 1–28
http://theoryofcomputing.org

Limitations of Quantum Advice and One-Way Communication

Scott Aaronson*

Received: June 29, 2004; published: February 9, 2005.

Abstract: Although a quantum state requires exponentially many classical bits to describe, the laws of quantum mechanics impose severe restrictions on how that state can be accessed. This paper shows in three settings that quantum messages have only limited advantages over classical ones.

First, we show that $\mathsf{BQP/qpoly} \subseteq \mathsf{PP/poly}$, where $\mathsf{BQP/qpoly}$ is the class of problems solvable in quantum polynomial time, given a polynomial-size "quantum advice state" that depends only on the input length. This resolves a question of Buhrman, and means that we should not hope for an unrelativized separation between quantum and classical advice. Underlying our complexity result is a general new relation between deterministic and quantum one-way communication complexities, which applies to partial as well as total functions.

Second, we construct an oracle relative to which $\mathsf{NP} \not\subset \mathsf{BQP/qpoly}$. To do so, we use the polynomial method to give the first correct proof of a *direct product theorem* for quantum search. This theorem has other applications; for example, it can be used to fix a result of Klauck about quantum time-space tradeoffs for sorting.

Third, we introduce a new *trace distance method* for proving lower bounds on quantum one-way communication complexity. Using this method, we obtain optimal quantum

*Work supported by an NSF Graduate Fellowship.

ACM Classification: F.1.2, F.1.3

AMS Classification: 81P68, 81P05, 68Q10, 68Q15, 42A05

Key words and phrases: quantum computation, advice, relativized complexity, direct product theorem, nonuniform, BQP, PP, communication, polynomial method, oracle

DOI: 10.4086/toc.2005.v001a001

lower bounds for two problems of Ambainis, for which no nontrivial lower bounds were previously known even for classical randomized protocols.

A preliminary version of this paper appeared in the 2004 Conference on Computational Complexity (CCC).

1 Introduction

How many classical bits can "really" be encoded into n qubits? Is it n, because of Holevo's Theorem [19]; $2n$, because of dense quantum coding [13] and quantum teleportation [9]; exponentially many, because of quantum fingerprinting [12]; or infinitely many, because amplitudes are continuous? The best general answer to this question is probably *mu*, the Zen word that "unasks" a question.[1]

To a computer scientist, however, it is natural to formalize the question in terms of *quantum one-way communication complexity* [6, 12, 20, 40]. The setting is as follows: Alice has an n-bit string x, Bob has an m-bit string y, and together they wish to evaluate $f(x,y)$ where $f : \{0,1\}^n \times \{0,1\}^m \rightarrow \{0,1\}$ is a Boolean function. After examining her input $x = x_1 \ldots x_n$, Alice can send a single quantum message ρ_x to Bob, whereupon Bob, after examining his input $y = y_1 \ldots y_m$, can choose some basis in which to measure ρ_x. He must then output a claimed value for $f(x,y)$. We are interested in how long Alice's message needs to be, for Bob to succeed with high probability on any x,y pair. Ideally the length will be much smaller than if Alice had to send a classical message.

Communication complexity questions have been intensively studied in theoretical computer science (see the book of Kushilevitz and Nisan [23] for example). In both the classical and quantum cases, though, most attention has focused on *two-way* communication, meaning that Alice and Bob get to send messages back and forth. We believe that the study of one-way quantum communication presents two main advantages. First, many open problems about two-way communication look gruesomely difficult—for example, are the randomized and quantum communication complexities of every total Boolean function polynomially related? We might gain insight into these problems by tackling their one-way analogues first. And second, because of its greater simplicity, the one-way model more directly addresses our opening question: how much "useful stuff" can be packed into a quantum state? Thus, results on one-way communication fall into the quantum information theory tradition initiated by Holevo [19] and others, as much as the communication complexity tradition initiated by Yao [38].

Related to quantum one-way communication is the notion of *quantum advice*. As pointed out by Nielsen and Chuang [28, p.203], there is no compelling physical reason to assume that the starting state of a quantum computer is a computational basis state:[2]

> [W]e know that many systems in Nature 'prefer' to sit in highly entangled states of many systems; might it be possible to exploit this preference to obtain extra computational power?

[1] Another *mu*-worthy question is, "Where does the power of quantum computing come from? Superposition? Interference? The large size of Hilbert space?"

[2] One might object that the starting state is itself the outcome of some computational process, which began no earlier than the Big Bang. However, (1) for all we know highly entangled states were created in the Big Bang, and (2) 14 billion years is a long time.

> It might be that having access to certain states allows particular computations to be done much more easily than if we are constrained to start in the computational basis.

One way to interpret Nielsen and Chuang's provocative question is as follows. Suppose we could request the *best possible* starting state for a quantum computer, knowing the language to be decided and the input length n but not knowing the input itself.[3] Denote the class of languages that we could then decide by BQP/qpoly—meaning quantum polynomial time, given an arbitrarily-entangled but polynomial-size quantum advice state.[4] How powerful is this class? If BQP/qpoly contained (for example) the NP-complete problems, then we would need to rethink our most basic assumptions about the power of quantum computing. We will see later that quantum advice is closely related to quantum one-way communication, since we can think of an advice state as a one-way message sent to an algorithm by a benevolent "advisor."

This paper is about the *limitations* of quantum advice and one-way communication. It presents three contributions which are basically independent of one another.

First, Section 3 shows that $D^1(f) = O\left(mQ_2^1(f)\log Q_2^1(f)\right)$ for any Boolean function f, partial or total. Here $D^1(f)$ is deterministic one-way communication complexity, $Q_2^1(f)$ is bounded-error one-way quantum communication complexity, and m is the length of Bob's input. Intuitively, whenever the set of Bob's possible inputs is not too large, Alice can send him a short classical message that lets him learn the outcome of any measurement he would have wanted to make on the quantum message ρ_x. It is interesting that a slightly tighter bound for total functions—$D^1(f) = O\left(mQ_2^1(f)\right)$—follows easily from a result of Klauck [20] together with a lemma of Sauer [34] about VC-dimension. However, the proof of the latter bound is highly nonconstructive, and seems to fail for partial f.

Using our communication complexity result, in Section 3.1 we show that BQP/qpoly $\subseteq$ PP/poly—in other words, BQP with polynomial-size quantum advice can be simulated in PP with polynomial-size classical advice.[5] This resolves a question of Harry Buhrman (personal communication), who asked whether quantum advice can be simulated in *any* classical complexity class with short classical advice. A corollary of our containment is that we cannot hope to show an unrelativized separation between quantum and classical advice (that is, that BQP/poly $\neq$ BQP/qpoly), without also showing that PP does not have polynomial-size circuits.

What makes this result surprising is that, in the minds of many computer scientists, a quantum state is basically an exponentially long vector. Indeed, this belief seems to fuel skepticism of quantum computing (see Goldreich [17] for example). But given an exponentially long advice string, even a classical computer could decide any language whatsoever. So one might imagine naïvely that quantum advice would let us solve problems that are not even recursively enumerable given classical advice of

[3] If we knew the input, we would simply request a starting state that contains the right answer!

[4] BQP/qpoly might remind readers of a better-studied class called QMA (Quantum Merlin-Arthur). But there are two key differences: first, advice can be trusted while proofs cannot; second, proofs can be tailored to a particular input while advice cannot.

[5] Here PP is Probabilistic Polynomial-Time, or the class of languages for which there exists a polynomial-time classical randomized algorithm that accepts with probability greater than $1/2$ if and only if an input x is in the language. Also, given a complexity class C, the class C/poly consists of all languages decidable by a C machine, given a polynomial-size classical advice string that depends only on the input length. See www.complexityzoo.com for more information about standard complexity classes mentioned in this paper.

a similar size! The failure of this naïve intuition supports the view that a quantum superposition over n-bit strings is "more similar" to a probability distribution over n-bit strings than to a 2^n-bit string.

Our second contribution, in Section 4, is an oracle relative to which NP is not contained in BQP/qpoly. Underlying this oracle separation is the first correct proof of a *direct product theorem* for quantum search. Given an N-item database with K marked items, the direct product theorem says that if a quantum algorithm makes $o\left(\sqrt{N}\right)$ queries, then the probability that the algorithm finds all K of the marked items decreases exponentially in K. Notice that such a result does not follow from any existing quantum lower bound. Earlier Klauck [21] claimed a weaker direct product theorem, based on the hybrid method of Bennett et al. [8], in a paper on quantum time-space tradeoffs for sorting. Unfortunately, Klauck's proof contained a bug. Our proof uses the polynomial method of Beals et al. [7], with the novel twist that we examine all *higher* derivatives of a polynomial (not just the first derivative). Our proof has already been improved by Klauck, Špalek, and de Wolf [22], who were able to recover and even extend Klauck's original results about quantum sorting.

Our final contribution, in Section 5, is a new *trace distance method* for proving lower bounds on quantum one-way communication complexity. Previously there was only one basic lower bound technique: the VC-dimension method of Klauck [20], which relied on lower bounds for quantum random access codes due to Ambainis et al. [5] and Nayak [27]. Using VC-dimension one can show, for example, that $Q_2^1(\text{DISJ}) = \Omega(n)$, where the *disjointness function* $\text{DISJ} : \{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}$ is defined by $\text{DISJ}(x,y) = 1$ if and only if $x_i y_i = 0$ for all $i \in \{1,\ldots,n\}$.

For some problems, however, the VC-dimension method yields no nontrivial quantum lower bound. Seeking to make this point vividly, Ambainis posed the following problem. Alice is given two elements x,y of a finite field $\mathbb{F}_p$ (where p is prime); Bob is given another two elements $a,b \in \mathbb{F}_p$. Bob's goal is to output 1 if $y \equiv ax + b \pmod{p}$ and 0 otherwise. For this problem, the VC-dimension method yields no randomized *or* quantum lower bound better than constant. On the other hand, the well-known fingerprinting protocol for the equality function [31] seems to fail for Ambainis' problem, because of the interplay between addition and multiplication. So it is natural to conjecture that the randomized and even quantum one-way complexities are $\Theta(\log p)$—that is, that no nontrivial protocol exists for this problem.

Ambainis posed a second problem in the same spirit. Here Alice is given $x \in \{1,\ldots,N\}$, Bob is given $y \in \{1,\ldots,N\}$, and both players know a subset $S \subset \{1,\ldots,N\}$. Bob's goal is to decide whether $x - y \in S$ where subtraction is modulo N. The conjecture is that if S is chosen uniformly at random with $|S|$ about $\sqrt{N}$, then with high probability the randomized and quantum one-way complexities are both $\Theta(\log N)$.

Using our trace distance method, we are able to show optimal quantum lower bounds for both of Ambainis' problems. Previously, no nontrivial lower bounds were known even for randomized protocols. The key idea is to consider two probability distributions over Alice's quantum message ρ_x. The first distribution corresponds to x chosen uniformly at random; the second corresponds to x chosen uniformly conditioned on $f(x,y) = 1$. These distributions give rise to two mixed states ρ and ρ_y, which Bob must be able to distinguish with non-negligible bias assuming he can evaluate $f(x,y)$. We then show an upper bound on the trace distance $\|\rho - \rho_y\|_{\text{tr}}$, which implies that Bob cannot distinguish the distributions.

Theorem 5.1 gives a very general condition under which our trace distance method works; Corollar-

ies 5.2 and 5.3 then show that the condition is satisfied for Ambainis' two problems. Besides showing a significant limitation of the VC-dimension method, we hope our new method is a non-negligible step towards proving that $R_2^1(f) = O\left(Q_2^1(f)\right)$ for all total Boolean functions f, where $R_2^1(f)$ is randomized one-way complexity. We conclude in Section 6 with some open problems.

This paper is a moderately revised version of an extended abstract that appeared in CCC 2004 [2]. The proofs of Theorems 3.4, 3.5, 5.1, 5.2, and 5.4 have been written out in more detail; and discussions have been added to Sections 2.2 and 3.1, about the group membership problem and the class $\mathsf{PQP/qpoly}$ respectively. Also, an error has been fixed in Section 4: the direct product theorem in [2] based on Bernstein's inequality is incorrect. Fortunately, the easier version based on V. A. Markov's inequality is still perfectly sufficient to show $\mathsf{NP} \not\subset \mathsf{BQP/qpoly}$ relative to an oracle; and in any case, the Bernstein's version has been superseded by the results of Klauck, Špalek, and de Wolf [22].

2 Preliminaries

This section reviews basic definitions and results about quantum one-way communication (in Section 2.1) and quantum advice (in Section 2.2); then Section 2.3 proves a quantum information lemma that will be used throughout the paper.

2.1 Quantum One-Way Communication

Following standard conventions, we denote by $D^1(f)$ the deterministic one-way complexity of f, or the minimum number of bits that Alice must send if her message is a function of x. Also, $R_2^1(f)$, the bounded-error randomized one-way complexity, is the minimum k such that for every x,y, if Alice sends Bob a k-bit message drawn from some distribution $\mathcal{D}_x$, then Bob can output a bit a such that $a = f(x,y)$ with probability at least $2/3$. (The subscript 2 means that the error is two-sided.) The zero-error randomized complexity $R_0^1(f)$ is similar, except that Bob's answer can never be wrong: he must output $f(x,y)$ with probability at least $1/2$ and otherwise declare failure.

The bounded-error quantum one-way complexity $Q_2^1(f)$ is the minimum k such that, if Alice sends Bob a mixed state ρ_x of k qubits, there exists a joint measurement of ρ_x and y enabling Bob to output an a such that $a = f(x,y)$ with probability at least $2/3$. The zero-error and exact complexities $Q_0^1(f)$ and $Q_E^1(f)$ are defined analogously. Requiring Alice's message to be a pure state would increase these complexities by at most a factor of 2, since by Kraus' Theorem, every k-qubit mixed state can be realized as half of a $2k$-qubit pure state. (Winter [37] has shown that this factor of 2 is tight.) See Klauck [20] for more detailed definitions of quantum and classical one-way communication complexity measures.

It is immediate that $D^1(f) \geq R_0^1(f) \geq R_2^1(f) \geq Q_2^1(f)$, that $R_0^1(f) \geq Q_0^1(f) \geq Q_2^1(f)$, and that $D^1(f) \geq Q_E^1(f)$. Also, for total f, Duriš et al. [14] showed that $R_0^1(f) = \Theta\left(D^1(f)\right)$, while Klauck [20] showed that $Q_E^1(f) = D^1(f)$ and that $Q_0^1(f) = \Theta\left(D^1(f)\right)$. In other words, randomized and quantum messages yield no improvement for total functions if we are unwilling to tolerate a bounded probability of error. This remains true even if Alice and Bob share arbitrarily many EPR pairs [20]. As is often the case, the situation is dramatically different for partial functions: there it is easy to see that $R_0^1(f)$ can be constant even though $D^1(f) = \Omega(n)$: let $f(x,y) = 1$ if $x_1y_1 + \cdots + x_{n/2}y_{n/2} \geq n/4$ and

$x_{n/2+1}y_{n/2+1}+\cdots+x_ny_n=0$ and $f(x,y)=0$ if $x_1y_1+\cdots+x_{n/2}y_{n/2}=0$ and $x_{n/2+1}y_{n/2+1}+\cdots+x_ny_n\geq n/4$, promised that one of these is the case.

Moreover, Bar-Yossef, Jayram, and Kerenidis [6] have *almost* shown that $Q_E^1(f)$ can be exponentially smaller than $R_2^1(f)$. In particular, they proved that separation for a *relation*, meaning a problem for which Bob has many possible valid outputs. For a partial function f based on their relation, they also showed that $Q_E^1(f)=\Theta(\log n)$ whereas $R_0^1(f)=\Theta(\sqrt{n})$; and they conjectured (but did not prove) that $R_2^1(f)=\Theta(\sqrt{n})$.

2.2 Quantum Advice

Informally, BQP/qpoly is the class of languages decidable in polynomial time on a quantum computer, given a polynomial-size quantum advice state that depends only on the input length. We now make the definition more formal.

Definition 2.1. A language L is in BQP/qpoly if there exists a polynomial-size quantum circuit family $\{C_n\}_{n\geq 1}$, and a polynomial-size family of quantum states $\{|\psi_n\rangle\}_{n\geq 1}$, such that for all $x\in\{0,1\}^n$,

(i) If $x\in L$ then $q(x)\geq 2/3$, where $q(x)$ is the probability that the first qubit is measured to be $|1\rangle$, after C_n is applied to the starting state $|x\rangle\otimes|0\cdots0\rangle\otimes|\psi_n\rangle$.

(ii) If $x\notin L$ then $q(x)\leq 1/3$.[6]

The central open question about BQP/qpoly is whether it equals BQP/poly, or BQP with polynomial-size *classical* advice. We do have a candidate for an oracle problem separating the two classes: the *group membership problem* of Watrous [36], which we describe for completeness. Let G_n be a black box group[7] whose elements are uniquely labeled by n-bit strings, and let H_n be a subgroup of G_n. Both G_n and H_n depend only on the input length n, so we can assume that a nonuniform algorithm knows generating sets for both of them. Given an element $x\in G_n$ as input, the problem is to decide whether $x\in H_n$.

If G_n is "sufficiently nonabelian" and H_n is exponentially large, we do not know how to solve this problem in BQP or even BQP/poly. On the other hand, we can solve it in BQP/qpoly as follows. Let our quantum advice state be an equal superposition over all elements of H_n:

$$|H_n\rangle=\frac{1}{\sqrt{|H_n|}}\sum_{y\in H_n}|y\rangle$$

We can transform $|H_n\rangle$ into

$$|xH_n\rangle=\frac{1}{\sqrt{|H_n|}}\sum_{y\in H_n}|xy\rangle$$

[6] If the starting state is $|x\rangle\otimes|0\cdots0\rangle\otimes|\varphi\rangle$ for some $|\varphi\rangle\neq|\psi_n\rangle$, then we do not require the acceptance probability to lie in $[0,1/3]\cup[2/3,1]$. Therefore, what we call BQP/qpoly corresponds to what Nishimura and Yamakami [30] call BQP/*Qpoly. Also, it does not matter whether the circuit family $\{C_n\}_{n\geq 1}$ is uniform, since we are giving it advice anyway.

[7] In other words, we have a quantum oracle available that given $x,y\in G_n$ outputs xy (i.e. exclusive-OR's xy into an answer register), and that given $x\in G_n$ outputs x^{-1}.

by mapping $|y\rangle|0\rangle$ to $|y\rangle|xy\rangle$ to $|y\oplus x^{-1}xy\rangle|xy\rangle = |0\rangle|xy\rangle$ for each $y \in H_n$. Our algorithm will first prepare the state $(|0\rangle|H_n\rangle + |1\rangle|xH_n\rangle)/\sqrt{2}$, then apply a Hadamard gate to the first qubit, and finally measure the first qubit in the standard basis, in order to distinguish the cases $|H_n\rangle = |xH_n\rangle$ and $\langle H_n|xH_n\rangle = 0$ with constant bias. The first case occurs whenever $x \in H_n$, and the second occurs whenever $x \notin H_n$.

Although the group membership problem provides intriguing evidence for the power of quantum advice, we have no idea how to show that it is not also solvable using classical advice. Indeed, apart from a result of Nishimura and Yamakami [30] that $\mathsf{EESPACE} \not\subset \mathsf{BQP/qpoly}$, essentially nothing was known about the class $\mathsf{BQP/qpoly}$ before the present work.

2.3 The Almost As Good As New Lemma

The following simple lemma, which was implicit in [5], is used three times in this paper—in Theorems 3.4, 3.5, and 4.7. It says that, if the outcome of measuring a quantum state ρ could be predicted with near-certainty given knowledge of ρ, then measuring ρ will damage it only slightly. Recall that the trace distance $\|\rho-\sigma\|_{\text{tr}}$ between two mixed states ρ and σ equals $\frac{1}{2}\sum_i |\lambda_i|$, where $\lambda_1,\ldots,\lambda_N$ are the eigenvalues of $\rho-\sigma$.

Lemma 2.2. *Suppose a 2-outcome measurement of a mixed state ρ yields outcome 0 with probability $1-\varepsilon$. Then after the measurement, we can recover a state $\widetilde{\rho}$ such that $\|\widetilde{\rho}-\rho\|_{\text{tr}} \leq \sqrt{\varepsilon}$. This is true even if the measurement is a POVM (that is, involves arbitrarily many ancilla qubits).*

Proof. Let $|\psi\rangle$ be a purification of the entire system (ρ plus ancilla). We can represent any measurement as a unitary U applied to $|\psi\rangle$, followed by a 1-qubit measurement. Let $|\varphi_0\rangle$ and $|\varphi_1\rangle$ be the two possible pure states after the measurement; then $\langle\varphi_0|\varphi_1\rangle = 0$ and $U|\psi\rangle = \alpha|\varphi_0\rangle + \beta|\varphi_1\rangle$ for some α,β such that $|\alpha|^2 = 1-\varepsilon$ and $|\beta|^2 = \varepsilon$. Writing the measurement result as $\sigma = (1-\varepsilon)|\varphi_0\rangle\langle\varphi_0| + \varepsilon|\varphi_1\rangle\langle\varphi_1|$, it is easy to show that

$$\left\|\sigma - U|\psi\rangle\langle\psi|U^{-1}\right\|_{\text{tr}} = \sqrt{\varepsilon(1-\varepsilon)}.$$

So applying U^{-1} to σ,

$$\left\|U^{-1}\sigma U - |\psi\rangle\langle\psi|\right\|_{\text{tr}} = \sqrt{\varepsilon(1-\varepsilon)}.$$

Let $\widetilde{\rho}$ be the restriction of $U^{-1}\sigma U$ to the original qubits of ρ. Theorem 9.2 of Nielsen and Chuang [28] shows that tracing out a subsystem never increases trace distance, so $\|\widetilde{\rho}-\rho\|_{\text{tr}} \leq \sqrt{\varepsilon(1-\varepsilon)} \leq \sqrt{\varepsilon}$. □

3 Simulating Quantum Messages

Let $f : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$ be a Boolean function. In this section we first combine existing results to obtain the relation $D^1(f) = O\left(mQ_2^1(f)\right)$ for total f, and then prove using a new method that $D^1(f) = O\left(mQ_2^1(f)\log Q_2^1(f)\right)$ for all f (partial or total).

Define the *communication matrix* M_f to be a $2^n \times 2^m$ matrix with $f(x,y)$ in the x^{th} row and y^{th} column. Then letting $\text{rows}(f)$ be the number of distinct rows in M_f, the following is immediate.

Proposition 3.1. *For total f,*

$$D^1(f) = \lceil \log_2 \text{rows}(f) \rceil,$$
$$Q_2^1(f) = \Omega(\log \log \text{rows}(f)).$$

Also, let the VC-dimension $\text{VC}(f)$ equal the maximum k for which there exists a $2^n \times k$ submatrix M_g of M_f with $\text{rows}(g) = 2^k$. Then Klauck [20] observed the following, based on a lower bound for quantum random access codes due to Nayak [27].

Proposition 3.2 (Klauck). $Q_2^1(f) = \Omega(\text{VC}(f))$ *for total* f.

Now let $\text{cols}(f)$ be the number of distinct columns in M_f. Then Theorem 3.2 yields the following general lower bound:

Corollary 3.3. $D^1(f) = O\left(mQ_2^1(f)\right)$ *for total* f*, where* m *is the size of Bob's input.*

Proof. It follows from a lemma of Sauer [34] that

$$\text{rows}(f) \leq \sum_{i=0}^{\text{VC}(f)} \binom{\text{cols}(f)}{i} \leq \text{cols}(f)^{\text{VC}(f)+1}.$$

Hence $\text{VC}(f) \geq \log_{\text{cols}(f)} \text{rows}(f) - 1$, so

$$\begin{aligned} Q_2^1(f) = \Omega(\text{VC}(f)) &= \Omega\left(\frac{\log \text{rows}(f)}{\log \text{cols}(f)}\right) \\ &= \Omega\left(\frac{D^1(f)}{m}\right). \end{aligned}$$

□

In particular, $D^1(f)$ and $Q_2^1(f)$ are polynomially related for total f, whenever Bob's input is polynomially smaller than Alice's, and Alice's input is not "padded." More formally, $D^1(f) = O\left(Q_2^1(f)^{1/(1-c)}\right)$ whenever $m = O(n^c)$ for some $c < 1$ and $\text{rows}(f) = 2^n$ (i.e. all rows of M_f are distinct). For then $D^1(f) = n$ by Theorem 3.1, and $Q_2^1(f) = \Omega\left(D^1(f)/n^c\right) = \Omega\left(n^{1-c}\right)$ by Corollary 3.3.

We now give a new method for replacing quantum messages by classical ones when Bob's input is small. Although the best bound we know how to obtain with this method—$D^1(f) = O\left(mQ_2^1(f) \log Q_2^1(f)\right)$—is slightly weaker than the $D^1(f) = O\left(mQ_2^1(f)\right)$ of Corollary 3.3, our method works for *partial* Boolean functions as well as total ones. It also yields a (relatively) efficient procedure by which Bob can reconstruct Alice's quantum message, a fact we will exploit in Section 3.1 to show $\text{BQP/qpoly} \subseteq \text{PP/poly}$. By contrast, the method based on Sauer's Lemma seems to be nonconstructive.

Theorem 3.4. $D^1(f) = O\left(mQ_2^1(f) \log Q_2^1(f)\right)$ *for all* f *(partial or total).*

Proof. Let $f : \mathcal{D} \to \{0,1\}$ be a partial Boolean function with $\mathcal{D} \subseteq \{0,1\}^n \times \{0,1\}^m$, and for all $x \in \{0,1\}^n$, let $\mathcal{D}_x = \{y \in \{0,1\}^m : (x,y) \in \mathcal{D}\}$. Suppose Alice can send Bob a quantum state with $Q_2^1(f)$ qubits, that enables him to compute $f(x,y)$ for any $y \in \mathcal{D}_x$ with error probability at most $1/3$. Then she can also send him a boosted state ρ with $K = O\left(Q_2^1(f) \log Q_2^1(f)\right)$ qubits, such that for all $y \in \mathcal{D}_x$,

$$|P_y(\rho) - f(x,y)| \leq \frac{1}{Q_2^1(f)^{10}},$$

where $P_y(\rho)$ is the probability that some measurement $\Lambda[y]$ yields a '1' outcome when applied to ρ. We can assume for simplicity that ρ is a pure state $|\psi\rangle\langle\psi|$; as discussed in Section 2.1, this increases the message length by at most a factor of 2.

Let $\mathcal{Y}$ be any subset of $\mathcal{D}_x$ satisfying $|\mathcal{Y}| \leq Q_2^1(f)^2$. Then starting with ρ, Bob can measure $\Lambda[y]$ for each $y \in \mathcal{Y}$ in lexicographic order, reusing the same message state again and again but uncomputing whatever garbage he generates while measuring. Let ρ_t be the state after the t^{th} measurement; thus $\rho_0 = \rho = |\psi\rangle\langle\psi|$. Since the probability that Bob outputs the wrong value of $f(x,y)$ on any given y is at most $1/Q_2^1(f)^{10}$, Lemma 2.2 implies that

$$\|\rho_t - \rho_{t-1}\|_{\text{tr}} \leq \sqrt{\frac{1}{Q_2^1(f)^{10}}} = \frac{1}{Q_2^1(f)^5}.$$

Since trace distance satisfies the triangle inequality, this in turn implies that

$$\|\rho_t - \rho\|_{\text{tr}} \leq \frac{t}{Q_2^1(f)^5} \leq \frac{1}{Q_2^1(f)^3}.$$

Now imagine an "ideal scenario" in which $\rho_t = \rho$ for every t; that is, the measurements do not damage ρ at all. Then the maximum bias with which Bob could distinguish the actual from the ideal scenario is

$$\|\rho_0 - \rho\|_{\text{tr}} + \cdots + \left\|\rho_{|\mathcal{Y}|-1} - \rho\right\|_{\text{tr}} \leq \frac{|\mathcal{Y}|}{Q_2^1(f)^3} \leq \frac{1}{Q_2^1(f)}.$$

So by the union bound, Bob will output $f(x,y)$ for every $y \in \mathcal{Y}$ simultaneously with probability at least

$$1 - \frac{|\mathcal{Y}|}{Q_2^1(f)^{10}} - \frac{1}{Q_2^1(f)} \geq 0.9$$

for sufficiently large $Q_2^1(f)$.

Now imagine that the communication channel is blocked, so Bob has to guess what message Alice wants to send him. He does this by using the K-qubit maximally mixed state I in place of ρ. We can write I as

$$I = \frac{1}{2^K} \sum_{j=1}^{2^K} |\psi_j\rangle\langle\psi_j|,$$

where $|\psi_1\rangle, \ldots, |\psi_{2^K}\rangle$ are orthonormal vectors such that $|\psi_1\rangle = |\psi\rangle$. So if Bob uses the same procedure as above except with I instead of ρ, then for any $\mathcal{Y} \subseteq \mathcal{D}_x$ with $|\mathcal{Y}| \leq Q_2^1(f)^2$, he will output $f(x,y)$ for every $y \in \mathcal{Y}$ simultaneously with probability at least $0.9/2^K$.

We now give the classical simulation of the quantum protocol. Alice's message to Bob consists of $T \leq K$ inputs $y_1,\ldots,y_T \in \mathcal{D}_x$, together with $f(x,y_1),\ldots,f(x,y_T)$.[8] Thus the message length is $mT+T = O\left(mQ_2^1(f)\log Q_2^1(f)\right)$. Here are the semantics of Alice's message: *"Bob, suppose you looped over all $y \in \mathcal{D}_x$ in lexicographic order; and for each one, guessed that $f(x,y) = \text{round}(P_y(I))$, where* round$(p)$ *is* 1 *if* $p \geq 1/2$ *and* 0 *if* $p < 1/2$. *Then y_1 is the first y for which you would guess the wrong value of $f(x,y)$. In general, let I_t be the state obtained by starting from I and then measuring $\Lambda[y_1],\ldots,\Lambda[y_t]$ in that order, given that the outcomes of the measurements are $f(x,y_1),\ldots,f(x,y_t)$ respectively. (Note that I_t is not changed by measurements of every $y \in \mathcal{D}_x$ up to y_t, only by measurements of $y_1,\ldots,y_t$.) If you looped over all $y \in \mathcal{D}_x$ in lexicographic order beginning from y_t, then y_{t+1} is the first y you would encounter for which* round$(P_y(I_t)) \neq f(x,y)$."

Given the sequence of y_t's as defined above, it is obvious that Bob can compute $f(x,y)$ for any $y \in \mathcal{D}_x$. First, if $y = y_t$ for some t, then he simply outputs $f(x,y_t)$. Otherwise, let t^* be the largest t for which $y_t < y$ lexicographically. Then Bob prepares a classical description of the state I_{t^*}—which he can do since he knows $y_1,\ldots,y_{t^*}$ and $f(x,y_1),\ldots,f(x,y_{t^*})$—and then outputs round$(P_y(I_{t^*}))$ as his claimed value of $f(x,y)$. Notice that, although Alice uses her knowledge of $\mathcal{D}_x$ to prepare her message, Bob does not need to know $\mathcal{D}_x$ in order to interpret the message. That is why the simulation works for partial as well as total functions.

But why can we assume that the sequence of y_t's stops at y_T for some $T \leq K$? Suppose $T > K$; we will derive a contradiction. Let $\mathcal{Y} = \{y_1,\ldots,y_{K+1}\}$. Then $|\mathcal{Y}| = K+1 \leq Q_2^1(f)^2$, so we know from previous reasoning that if Bob starts with I and then measures $\Lambda[y_1],\ldots,\Lambda[y_{K+1}]$ in that order, he will observe $f(x,y_1),\ldots,f(x,y_{K+1})$ simultaneously with probability at least $0.9/2^K$. But by the definition of y_t, the probability that $\Lambda[y_t]$ yields the correct outcome is at most $1/2$, conditioned on $\Lambda[y_1],\ldots,\Lambda[y_{t-1}]$ having yielded the correct outcomes. Therefore $f(x,y_1),\ldots,f(x,y_{K+1})$ are observed simultaneously with probability at most $1/2^{K+1} < 0.9/2^K$, contradiction. □

3.1 Simulating Quantum Advice

We now apply our new simulation method to upper-bound the power of quantum advice.

Theorem 3.5. BQP/qpoly $\subseteq$ PP/poly.

Proof. For notational convenience, let $L_n(x) = 1$ if input $x \in \{0,1\}^n$ is in language L, and $L_n(x) = 0$ otherwise. Suppose L_n is computed by a BQP machine using quantum advice of length $p(n)$. We will give a PP machine that computes L_n using classical advice of length $O(np(n)\log p(n))$. Because of the close connection between advice and one-way communication, the simulation method will be essentially identical to that of Theorem 3.4.

By using a boosted advice state on $K = O(p(n)\log p(n))$ qubits, a polynomial-time quantum algorithm A can compute $L_n(x)$ with error probability at most $1/p(n)^{10}$. Now the classical advice to the PP machine consists of $T \leq K$ inputs $x_1,\ldots,x_T \in \{0,1\}^n$, together with $L_n(x_1),\ldots,L_n(x_T)$. Let I be the maximally mixed state on K qubits. Also, let $P_x(\rho)$ be the probability that A outputs '1' on input x, given ρ as its advice state. Then x_1 is the lexicographically first input x for which

[8] Strictly speaking, Bob will be able to compute $f(x,y_1),\ldots,f(x,y_T)$ for himself given $y_1,\ldots,y_T$; he does not need Alice to tell him the f values.

$\text{round}(P_x(I)) \neq L_n(x)$. In general, let I_t be the state obtained by starting with I as the advice and then running A on $x_1,\ldots,x_t$ in that order (uncomputing garbage along the way), if we postselect on A correctly outputting $L_n(x_1),\ldots,L_n(x_t)$. Then x_{t+1} is the lexicographically first $x > x_t$ for which $\text{round}(P_x(I_t)) \neq L_n(x)$.

Given the classical advice, we can compute $L_n(x)$ as follows: if $x \in \{x_1,\ldots,x_T\}$ then output $L_n(x_t)$. Otherwise let t^* be the largest t for which $x_t < x$ lexicographically, and output $\text{round}(P_x(I_{t^*}))$. The proof that this algorithm works is the same as in Theorem 3.4, and so is omitted for brevity. All we need to show is that the algorithm can be implemented in PP.

Adleman, DeMarrais, and Huang [4] (see also Fortnow and Rogers [16]) showed that $\mathsf{BQP} \subseteq \mathsf{PP}$, by using what physicists would call a "Feynman sum-over-histories." Specifically, let C be a polynomial-size quantum circuit that starts in the all-0 state, and that consists solely of Toffoli and Hadamard gates (Shi [35] has shown that this gate set is universal). Also, let α_z be the amplitude of basis state $|z\rangle$ after all gates in C have been applied. We can write α_z as a sum of exponentially many contributions, $a_1 + \cdots + a_N$, where each a_i is a rational real number computable in classical polynomial time. So by evaluating the sum

$$|\alpha_z|^2 = \sum_{i,j=1}^{N} a_i a_j,$$

putting positive and negative terms on "opposite sides of the ledger," a PP machine can check whether $|\alpha_z|^2 > \beta$ for any rational constant β. It follows that a PP machine can also check whether

$$\sum_{z\,:\,S_1(z)} |\alpha_z|^2 > \sum_{z\,:\,S_0(z)} |\alpha_z|^2$$

(or equivalently, whether $\Pr[S_1] > \Pr[S_0]$) for any classical polynomial-time predicates S_1 and S_0.

Now suppose the circuit C does the following, in the case $x \notin \{x_1,\ldots,x_T\}$. It first prepares the K-qubit maximally mixed state I (as half of a $2K$-qubit pure state), and then runs A on $x_1,\ldots,x_{t^*},x$ in that order, using I as its advice state. The claimed values of $L_n(x_1),\ldots,L_n(x_{t^*}),L_n(x)$ are written to output registers but not measured. For $i \in \{0,1\}$, let the predicate $S_i(z)$ hold if and only if basis state $|z\rangle$ contains the output sequence $L_n(x_1),\ldots,L_n(x_{t^*}),i$. Then it is not hard to see that

$$P_x(I_{t^*}) = \frac{\Pr[S_1]}{\Pr[S_1] + \Pr[S_0]},$$

so $P_x(I_{t^*}) > 1/2$ and hence $L_n(x) = 1$ if and only if $\Pr[S_1] > \Pr[S_0]$. Since the case $x \in \{x_1,\ldots,x_T\}$ is trivial, this shows that $L_n(x)$ is computable in $\mathsf{PP/poly}$. □

We make five remarks about Theorem 3.5. First, for the same reason that Theorem 3.4 works for partial as well as total functions, we actually obtain the stronger result that $\mathsf{PromiseBQP/qpoly} \subseteq \mathsf{PromisePP/poly}$, where PromiseBQP and PromisePP are the promise-problem versions of BQP and PP respectively.

Second, as pointed out to us by Lance Fortnow, a corollary of Theorem 3.5 is that we cannot hope to show an unrelativized separation between $\mathsf{BQP/poly}$ and $\mathsf{BQP/qpoly}$, without also showing that PP does not have polynomial-size circuits. For $\mathsf{BQP/poly} \neq \mathsf{BQP/qpoly}$ clearly implies that $\mathsf{P/poly} \neq \mathsf{PP/poly}$. But the latter then implies that $\mathsf{PP} \not\subset \mathsf{P/poly}$, since assuming $\mathsf{PP} \subset \mathsf{P/poly}$ we

could also obtain polynomial-size circuits for a language $L \in \mathsf{PP/poly}$ by defining a new language $L' \in \mathsf{PP}$, consisting of all (x,a) pairs such that the PP machine would accept x given advice string a. The reason this works is that PP is a syntactically defined class.

Third, an earlier version of this paper showed that $\mathsf{BQP/qpoly} \subseteq \mathsf{EXP/poly}$, by using a simulation in which an EXP machine keeps track of a subspace H of the advice Hilbert space to which the 'true' advice state must be close. In that simulation, the classical advice specifies inputs $x_1,\ldots,x_T$ for which $\dim(H)$ is at least halved; the observation that $\dim(H)$ must be at least 1 by the end then implies that $T \leq K = O(p(n)\log p(n))$, meaning that the advice is of polynomial size. The huge improvement from EXP to PP came solely from working with *measurement outcomes* and their *probabilities* instead of with *subspaces* and their *dimensions*. We can compute the former using the same "Feynman sum-over-histories" that Adleman et al. [4] used to show $\mathsf{BQP} \subseteq \mathsf{PP}$, but could not see any way to compute the latter without explicitly storing and diagonalizing exponentially large matrices.

Fourth, assuming $\mathsf{BQP/poly} \neq \mathsf{BQP/qpoly}$, Theorem 3.5 is *almost* the best result of its kind that one could hope for, since the only classes known to lie between BQP and PP and not known to equal either are obscure ones such as AWPP [16]. Initially the theorem seemed to us to prove something stronger, namely that $\mathsf{BQP/qpoly} \subseteq \mathsf{PostBQP/poly}$. Here PostBQP is the class of languages decidable by polynomial-size quantum circuits with *postselection*—meaning the ability to measure a qubit that has a nonzero probability of being $|1\rangle$, and then *assume* that the measurement outcome will be $|1\rangle$. Clearly PostBQP lies somewhere between BQP and PP; one can think of it as a quantum analogue of the classical complexity class $\mathsf{BPP}_{\mathsf{path}}$ [18]. We have since shown, however, that $\mathsf{PostBQP} = \mathsf{PP}$ [3].

Fifth, it is clear that Adleman et al.'s $\mathsf{BQP} \subseteq \mathsf{PP}$ result [4] can be extended to show that $\mathsf{PQP} = \mathsf{PP}$. Here PQP is the quantum analogue of PP—that is, quantum polynomial time but where the probability of a correct answer need only be bounded above $1/2$, rather than above $2/3$. A reviewer asked whether Theorem 3.5 could similarly be extended to show that $\mathsf{PQP/qpoly} = \mathsf{PP/poly}$. The answer is no—for indeed, $\mathsf{PQP/qpoly}$ contains every language whatsoever! To see this, given any function $L_n : \{0,1\}^n \rightarrow \{0,1\}$, let our quantum advice state be

$$|\psi_n\rangle = \frac{1}{2^{n/2}} \sum_{x \in \{0,1\}^n} |x\rangle\,|L_n(x)\rangle\,.$$

Then a PQP algorithm to compute L_n is as follows: given an input $x \in \{0,1\}^n$, first measure $|\psi_n\rangle$ in the standard basis. If $|x\rangle\,|L_n(x)\rangle$ is observed, output $L_n(x)$; otherwise output a uniform random bit.

4 Oracle Limitations

Can quantum computers solve NP-complete problems in polynomial time? In the early days of quantum computing, Bennett et al. [8] gave an oracle relative to which $\mathsf{NP} \not\subset \mathsf{BQP}$, providing what is still the best evidence we have that the answer is no. It is easy to extend Bennett et al.'s result to give an oracle relative to which $\mathsf{NP} \not\subset \mathsf{BQP/poly}$; that is, NP is hard even for nonuniform quantum algorithms. But when we try to show $\mathsf{NP} \not\subset \mathsf{BQP/qpoly}$ relative to an oracle, a new difficulty arises: even if the oracle encodes 2^n exponentially hard search problems for each input length n, the quantum advice, being an "exponentially large object" itself, might somehow encode information about all 2^n problems. We need

to argue that even if so, only a miniscule fraction of that information can be extracted by measuring the advice.

How does one prove such a statement? As it turns out, the task can be reduced to proving a *direct product theorem* for quantum search. This is a theorem that in its weakest form says the following: given N items, K of which are marked, if we lack enough time to find even *one* marked item, then the probability of finding all K items decreases exponentially in K. For intuitively, suppose there were a quantum advice state that let us efficiently find any one of K marked items. Then by "guessing" the advice (i.e. replacing it by a maximally mixed state), and then using the guessed advice multiple times, we could efficiently find all K of the items with a success probability that our direct product theorem shows is impossible. This reduction is formalized in Theorem 4.7.

But what about the direct product theorem itself? It seems like it should be trivial to prove—for surely there are no devious correlations by which success in finding one marked item leads to success in finding all the others! So it is surprising that even a weak direct product theorem eluded proof for years. In 2001, Klauck [21] gave an attempted proof using the hybrid method of Bennett et al. [8]. His motivation was to show a limitation of space-bounded quantum sorting algorithms. Unfortunately, Klauck's proof contained a bug.[9]

In this section we give the first correct proof of a direct product theorem, based on the polynomial method of Beals et al. [7]. Besides showing that $\mathsf{NP} \not\subset \mathsf{BQP/qpoly}$ relative to an oracle, our result can be used to recover the claims made in [21] about the hardness of quantum sorting (see Klauck, Špalek, and de Wolf [22] for details). We expect the result to have other applications as well.

We will need the following lemma of Beals et al. [7], which builds on ideas due to Minsky and Papert [26] and Nisan and Szegedy [29].

Lemma 4.1 (Beals et al.). *Suppose a quantum algorithm makes T queries to an oracle string $X \in \{0,1\}^N$, and accepts with probability $A(X)$. Then there exists a real polynomial p, of degree at most $2T$, such that*

$$p(i) = \mathop{\mathrm{EX}}_{|X|=i}[A(X)]$$

for all integers $i \in \{0,\ldots,N\}$, where $|X|$ denotes the Hamming weight of X.

Lemma 4.1 implies that, to lower-bound the number of queries T made by a quantum algorithm, it suffices to lower-bound $\deg(p)$, where p is a real polynomial representing the algorithm's expected acceptance probability. As an example, any quantum algorithm that computes the OR function on N bits, with success probability at least $2/3$, yields a polynomial p such that $p(0) \in [0,1/3]$ and $p(i) \in [2/3,1]$ for all integers $i \in \{1,\ldots,N\}$. To lower-bound the degree of such a polynomial, one can use an inequality proved by A. A. Markov in 1890 ([24]; see also [32]):

Theorem 4.2 (A. A. Markov). *Given a real polynomial p and constant $N > 0$, let $r^{(0)} = \max_{x\in[0,N]} |p(x)|$ and $r^{(1)} = \max_{x\in[0,N]} |p'(x)|$. Then*

$$\deg(p) \geq \sqrt{\frac{Nr^{(1)}}{2r^{(0)}}}.$$

[9] Specifically, the last sentence in the proof of Lemma 5 in [21] ("Clearly this probability is at least $q_x(p_x - \alpha)$") is not justified by what precedes it.

Theorem 4.2 deals with the entire range $[0,N]$, whereas in our setting $p(x)$ is constrained only at the integer points $x \in \{0,\ldots,N\}$. But as shown in [15, 29, 33], this is not a problem. For by elementary calculus, $p(0) \leq 1/3$ and $p(1) \geq 2/3$ imply that $p'(x) \geq 1/3$ for some real $x \in [0,1]$, and therefore $r^{(1)} \geq 1/3$. Furthermore, let x^* be a point in $[0,N]$ where $|p(x^*)| = r^{(0)}$. Then $p(\lfloor x^* \rfloor) \in [0,1]$ and $p(\lceil x^* \rceil) \in [0,1]$ imply that $r^{(1)} \geq 2\left(r^{(0)}-1\right)$. Thus

$$\deg(p) \geq \sqrt{\frac{Nr^{(1)}}{2r^{(0)}}} \geq \sqrt{\frac{N \max\left\{1/3, 2\left(r^{(0)}-1\right)\right\}}{2r^{(0)}}} = \Omega\left(\sqrt{N}\right).$$

This is the proof of Beals et al. [7] that quantum search requires $\Omega\left(\sqrt{N}\right)$ queries.

When proving a direct product theorem, we can no longer apply Theorem 4.2 so straightforwardly. The reason is that the success probabilities in question are extremely small, and therefore the maximum derivative $r^{(1)}$ could also be extremely small. Fortunately, though, we can still prove a good lower bound on the degree of the relevant polynomial p. The key is to look not just at the first derivative of p, but at higher derivatives.

To start, we need a lemma about the behavior of functions under repeated differentiation.

Lemma 4.3. *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be an infinitely differentiable function such that for some positive integer K, we have $f(i) = 0$ for all $i \in \{0,\ldots,K-1\}$ and $f(K) = \delta > 0$. Also, let $r^{(m)} = \max_{x \in [0,N]} \left|f^{(m)}(x)\right|$, where $f^{(m)}(x)$ is the m^{th} derivative of f evaluated at x (thus $f^{(0)} = f$). Then $r^{(m)} \geq \delta/m!$ for all $m \in \{0,\ldots,K\}$.*

Proof. We claim, by induction on m, that there exist $K-m+1$ points $0 \leq x_0^{(m)} < \cdots < x_{K-m}^{(m)} \leq K$ such that $f^{(m)}\left(x_i^{(m)}\right) = 0$ for all $i \leq K-m-1$ and $f^{(m)}\left(x_{K-m}^{(m)}\right) \geq \delta/m!$. If we define $x_i^{(0)} = i$, then the base case $m = 0$ is immediate from the conditions of the lemma. Suppose the claim is true for m; then by elementary calculus, for all $i \leq K-m-2$ there exists a point $x_i^{(m+1)} \in \left(x_i^{(m)}, x_{i+1}^{(m)}\right)$ such that $f^{(m+1)}\left(x_i^{(m+1)}\right) = 0$. Notice that $x_i^{(m+1)} \geq x_i^{(m)} \geq \cdots \geq x_i^{(0)} = i$. So there is also a point $x_{K-m-1}^{(m+1)} \in \left(x_{K-m-1}^{(m)}, x_{K-m}^{(m)}\right)$ such that

$$\begin{aligned} f^{(m+1)}\left(x_{K-m-1}^{(m+1)}\right) &\geq \frac{f^{(m)}\left(x_{K-m}^{(m)}\right) - f^{(m)}\left(x_{K-m-1}^{(m)}\right)}{x_{K-m}^{(m)} - x_{K-m-1}^{(m)}} \\ &\geq \frac{\delta/m! - 0}{K-(K-m-1)} \\ &= \frac{\delta}{(m+1)!}. \end{aligned}$$

□

With the help of Lemma 4.3, we can sometimes lower-bound the degree of a real polynomial even its first derivative is small throughout the region of interest. To do so, we use the following generalization of A. A. Markov's inequality (Theorem 4.2), which was proved by A. A. Markov's younger brother V. A. Markov in 1892 ([25]; see also [32]).

Theorem 4.4 (V. A. Markov). *Given a real polynomial p of degree d and positive real number N, let $r^{(m)} = \max_{x\in[0,N]} \left|p^{(m)}(x)\right|$. Then for all $m \in \{1,\ldots,d\}$,*

$$\begin{aligned} r^{(m)} &\leq \left(\frac{2r^{(0)}}{N}\right)^m T_d^{(m)}(1) \\ &\leq \left(\frac{2r^{(0)}}{N}\right)^m \frac{d^2\left(d^2-1^2\right)\left(d^2-2^2\right)\cdot\cdots\cdot\left(d^2-(m-1)^2\right)}{1\cdot 3\cdot 5\cdot\cdots\cdot(2m-1)}. \end{aligned}$$

Here $T_d(x) = \cos(d\arccos x)$ is the d^{th} Chebyshev polynomial of the first kind.

As we demonstrate below, combining Theorem 4.4 with Lemma 4.3 yields a lower bound on $\deg(p)$.

Lemma 4.5. *Let p be a real polynomial such that*

(i) $p(x) \in [0,1]$ at all integer points $x \in \{0,\ldots,N\}$, and

(ii) for some positive integer $K \leq N$ and real $\delta > 0$, we have $p(K) = \delta$ and $p(i) = 0$ for all $i \in \{0,\ldots,K-1\}$.

Then $\deg(p) = \Omega\left(\sqrt{N\delta^{1/K}}\right)$.

Proof. Let $p^{(m)}$ and $r^{(m)}$ be as in Theorem 4.4. Then for all $m \in \{1,\ldots,\deg(p)\}$, Theorem 4.4 yields

$$r^{(m)} \leq \left(\frac{2r^{(0)}}{N}\right)^m \frac{\deg(p)^{2m}}{1\cdot 3\cdot 5\cdot\cdots\cdot(2m-1)}$$

Rearranging,

$$\deg(p) \geq \sqrt{\frac{N}{2r^{(0)}}\left(1\cdot 3\cdot 5\cdot\cdots\cdot(2m-1)\cdot r^{(m)}\right)^{1/m}}$$

for all $m \geq 1$ (if $m > \deg(p)$ then $r^{(m)} = 0$ so the bound is trivial).

There are now two cases. First suppose $r^{(0)} \geq 2$. Then as discussed previously, condition (i) implies that $r^{(1)} \geq 2\left(r^{(0)}-1\right)$, and hence that

$$\deg(p) \geq \sqrt{\frac{Nr^{(1)}}{2r^{(0)}}} \geq \sqrt{\frac{N\left(r^{(0)}-1\right)}{r^{(0)}}} = \Omega\left(\sqrt{N}\right)$$

by Theorem 4.2. Next suppose $r^{(0)} < 2$. Then $r^{(m)} \geq \delta/m!$ for all $m \leq K$ by Lemma 4.3. So setting $m = K$ yields

$$\deg(p) \geq \sqrt{\frac{N}{4}\left(1\cdot 3\cdot 5\cdot\cdots\cdot(2K-1)\cdot\frac{\delta}{K!}\right)^{1/K}} = \Omega\left(\sqrt{N\delta^{1/K}}\right).$$

Either way we are done. □

Strictly speaking, we do not need the full strength of Theorem 4.4 to prove a lower bound on $\deg(p)$ that suffices for an oracle separation between NP and BQP/qpoly. For we can show a "rough-and-ready" version of V. A. Markov's inequality by applying A. A. Markov's inequality (Theorem 4.2) repeatedly, to $p, p^{(1)}, p^{(2)}$, and so on. This yields

$$r^{(m)} \leq \frac{2}{N} \deg(p)^2 r^{(m-1)} \leq \left(\frac{2}{N} \deg(p)^2\right)^m r^{(0)}$$

for all m. If $\deg(p)$ is small, then this upper bound on $r^{(m)}$ contradicts the lower bound of Lemma 4.3. However, the lower bound on $\deg(p)$ that one gets from A. A. Markov's inequality is only $\Omega\left(\sqrt{N\delta^{1/K}/K}\right)$, as opposed to $\Omega\left(\sqrt{N\delta^{1/K}}\right)$ from Lemma 4.5.[10]

Shortly after seeing our proof of a weak direct product theorem, Klauck, Špalek, and de Wolf [22] managed to improve the lower bound on $\deg(p)$ to the essentially tight $\Omega\left(\sqrt{NK\delta^{1/K}}\right)$. In particular, their bound implies that δ decreases exponentially in K whenever $\deg(p) = o\left(\sqrt{NK}\right)$. They obtained this improvement by *factoring* p instead of differentiating it as in Lemma 4.3.

In any case, a direct product theorem follows trivially from what has already been said.

Theorem 4.6 (Direct Product Theorem). *Suppose a quantum algorithm makes T queries to an oracle string $X \in \{0,1\}^N$. Let δ be the minimum probability, over all X with Hamming weight $|X| = K$, that the algorithm finds all K of the '1' bits. Then $\delta \leq \left(cT^2/N\right)^K$ for some constant c.*

Proof. Have the algorithm accept if it finds K or more '1' bits and reject otherwise. Let $p(i)$ be the expected probability of acceptance if X is drawn uniformly at random subject to $|X| = i$. Then we know the following about p:

(i) $p(i) \in [0,1]$ at all integer points $i \in \{0,\ldots,N\}$, since $p(i)$ is a probability.

(ii) $p(i) = 0$ for all $i \in \{0,\ldots,K-1\}$, since there are not K marked items to be found.

(iii) $p(K) \geq \delta$.

Furthermore, Lemma 4.1 implies that p is a polynomial in i satisfying $\deg(p) \leq 2T$. It follows from Lemma 4.5 that $T = \Omega\left(\sqrt{N\delta^{1/K}}\right)$, or rearranging, that $\delta \leq \left(cT^2/N\right)^K$ for some constant c. □

We can now prove the desired oracle separation using standard complexity theory tricks.

Theorem 4.7. *There exists an oracle relative to which* NP $\not\subset$ BQP/qpoly.

[10] An earlier version of this paper claimed to prove $\deg(p) = \Omega\left(\sqrt{NK}/\log^{3/2}(1/\delta)\right)$, by applying *Bernstein's inequality* [11] rather than A. A. Markov's to all derivatives $p^{(m)}$. We have since discovered a flaw in that argument. In any case, the Bernstein lower bound is both unnecessary for an oracle separation, and superseded by the later results of Klauck et al. [22].

Proof. Given an oracle $A : \{0,1\}^* \to \{0,1\}$, define the language L_A by $(y,z) \in L_A$ if and only if $y \leq z$ lexicographically and there exists an x such that $y \leq x \leq z$ and $A(x) = 1$. Clearly $L_A \in \mathsf{NP}^A$ for all A. We argue that for some A, no $\mathsf{BQP/qpoly}$ machine M with oracle access to A can decide L_A. Without loss of generality we assume M is fixed, so that only the advice states $\{|\psi_n\rangle\}_{n\geq 1}$ depend on A. We also assume the advice is boosted, so that M's error probability on any input (y,z) is $2^{-\Omega\left(n^2\right)}$.

Choose a set $S \subset \{0,1\}^n$ subject to $|S| = 2^{n/10}$; then for all $x \in \{0,1\}^n$, set $A(x) = 1$ if and only if $x \in S$. We claim that by using M, an algorithm could find all $2^{n/10}$ elements of S with high probability after only $2^{n/10}\,\mathrm{poly}(n)$ queries to A. Here is how: first use binary search (repeatedly halving the distance between y and z) to find the lexicographically first element of S. By Lemma 2.2, the boosted advice state $|\psi_n\rangle$ is good for $2^{\Omega\left(n^2\right)}$ uses, so this takes only $\mathrm{poly}(n)$ queries. Then use binary search to find the lexicographically second element, and so on until all elements have been found.

Now replace $|\psi_n\rangle$ by the maximally mixed state as in Theorem 3.4. This yields an algorithm that uses no advice, makes $2^{n/10}\,\mathrm{poly}(n)$ queries, and finds all $2^{n/10}$ elements of S with probability $2^{-O(\mathrm{poly}(n))}$. But taking $\delta = 2^{-O(\mathrm{poly}(n))}$, $T = 2^{n/10}\,\mathrm{poly}(n)$, $N = 2^n$, and $K = 2^{n/10}$, such an algorithm would satisfy $\delta \gg \left(cT^2/N\right)^K$, which violates the bound of Theorem 4.6. □

Indeed one can show that $\mathsf{NP} \not\subset \mathsf{BQP/qpoly}$ relative a random oracle with probability 1.[11]

5 The Trace Distance Method

This section introduces a new method for proving lower bounds on quantum one-way communication complexity. Unlike in Section 3, here we do not try to simulate quantum protocols using classical ones. Instead we prove lower bounds for quantum protocols directly, by reasoning about the trace distance between two possible distributions over Alice's quantum message (that is, between two mixed states). The result is a method that works even if Alice's and Bob's inputs are the same size.

We first state our method as a general theorem; then, in Section 5.1, we apply the theorem to prove lower bounds for two problems of Ambainis. Let $\|\mathcal{D} - \mathcal{E}\|$ denote the variation distance between probability distributions $\mathcal{D}$ and $\mathcal{E}$.

Theorem 5.1. *Let $f : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$ be a total Boolean function. For each $y \in \{0,1\}^m$, let $\mathcal{A}_y$ be a distribution over $x \in \{0,1\}^n$ such that $f(x,y) = 1$. Let $\mathcal{B}$ be a distribution over $y \in \{0,1\}^m$, and let $\mathcal{D}_k$ be the distribution over $(\{0,1\}^n)^k$ formed by first choosing $y \in \mathcal{B}$ and then choosing k samples independently from $\mathcal{A}_y$. Suppose that $\Pr_{x\in\mathcal{D}_1, y\in\mathcal{B}}[f(x,y) = 0] = \Omega(1)$ and that $\left\|\mathcal{D}_2 - \mathcal{D}_1^2\right\| \leq \delta$. Then $Q_2^1(f) = \Omega(\log 1/\delta)$.*

Proof. Suppose that if Alice's input is x, then she sends Bob the l-qubit mixed state ρ_x. Suppose also that for every $x \in \{0,1\}^n$ and $y \in \{0,1\}^m$, Bob outputs $f(x,y)$ with probability at least $2/3$. Then by amplifying a constant number of times, Bob's success probability can be made $1-\varepsilon$ for any constant $\varepsilon > 0$. So with $L = O(l)$ qubits of communication, Bob can distinguish the following two cases with constant bias:

[11] First group the oracle bits into polynomial-size blocks as Bennett and Gill [10] do, then use the techniques of Aaronson [1] to show that the acceptance probability is a low-degree univariate polynomial in the number of all-0 blocks. The rest of the proof follows Theorem 4.7.

Case I. y was drawn from $\mathcal{B}$ and x from $\mathcal{D}_1$.

Case II. y was drawn from $\mathcal{B}$ and x from $\mathcal{A}_y$.

For in Case I, we assumed that $f(x,y)=0$ with constant probability, whereas in Case II, $f(x,y)=1$ always. An equivalent way to say this is that with constant probability over y, Bob can distinguish the mixed states $\rho = \mathrm{EX}_{x\in\mathcal{D}_1}[\rho_x]$ and $\rho_y = \mathrm{EX}_{x\in\mathcal{A}_y}[\rho_x]$ with constant bias. Therefore

$$\mathop{\mathrm{EX}}_{y\in\mathcal{B}}\left[\left\|\rho-\rho_y\right\|_{\mathrm{tr}}\right] = \Omega(1).$$

We need an upper bound on the trace distance $\left\|\rho-\rho_y\right\|_{\mathrm{tr}}$ that is more amenable to analysis. Let $\lambda_1,\ldots,\lambda_{2^L}$ be the eigenvalues of $\rho-\rho_y$. Then

$$\begin{aligned}
\left\|\rho-\rho_y\right\|_{\mathrm{tr}} &= \frac{1}{2}\sum_{i=1}^{2^L}|\lambda_i| \\
&\leq \frac{1}{2}\sqrt{2^L\sum_{i=1}^{2^L}\lambda_i^2} \\
&= 2^{L/2-1}\sqrt{\sum_{i,j=1}^{2^L}\left|(\rho)_{ij}-(\rho_y)_{ij}\right|^2}
\end{aligned}$$

where $(\rho)_{ij}$ is the (i,j) entry of ρ. Here the second line uses the Cauchy-Schwarz inequality, and the third line uses the unitary invariance of the Frobenius norm.

We claim that

$$\mathop{\mathrm{EX}}_{y\in\mathcal{B}}\left[\sum_{i,j=1}^{2^L}\left|(\rho)_{ij}-(\rho_y)_{ij}\right|^2\right] \leq 2\delta.$$

From this claim it follows that

$$\begin{aligned}
\mathop{\mathrm{EX}}_{y\in\mathcal{B}}\left[\left\|\rho-\rho_y\right\|_{\mathrm{tr}}\right] &\leq 2^{L/2-1}\mathop{\mathrm{EX}}_{y\in\mathcal{B}}\left[\sqrt{\sum_{i,j=1}^{2^L}\left|(\rho)_{ij}-(\rho_y)_{ij}\right|^2}\right] \\
&\leq 2^{L/2-1}\sqrt{\mathop{\mathrm{EX}}_{y\in\mathcal{B}}\left[\sum_{i,j=1}^{2^L}\left|(\rho)_{ij}-(\rho_y)_{ij}\right|^2\right]} \\
&\leq \sqrt{2^{L-1}\delta}.
\end{aligned}$$

Therefore the message length L must be $\Omega(\log 1/\delta)$ to ensure that $\mathrm{EX}_{y\in\mathcal{B}}\left[\left\|\rho-\rho_y\right\|_{\mathrm{tr}}\right]=\Omega(1)$.

Let us now prove the claim. We have

$$\begin{aligned}
\mathop{\mathrm{EX}}_{y\in\mathcal{B}}\left[\sum_{i,j=1}^{2^L}\left|(\rho)_{ij}-(\rho_y)_{ij}\right|^2\right] &= \sum_{i,j=1}^{2^L}\left(\left|(\rho)_{ij}\right|^2 - 2\,\mathrm{Re}\left((\rho)_{ij}^*\mathop{\mathrm{EX}}_{y\in\mathcal{B}}\left[(\rho_y)_{ij}\right]\right)+\mathop{\mathrm{EX}}_{y\in\mathcal{B}}\left[\left|(\rho_y)_{ij}\right|^2\right]\right) \\
&= \sum_{i,j=1}^{2^L}\left(\mathop{\mathrm{EX}}_{y\in\mathcal{B}}\left[\left|(\rho_y)_{ij}\right|^2\right]-\left|(\rho)_{ij}\right|^2\right),
\end{aligned}$$

since $\mathrm{EX}_{y\in\mathcal{B}}\left[(\rho_y)_{ij}\right]=(\rho)_{ij}$. For a given (i,j) pair,

$$\begin{aligned}\mathop{\mathrm{EX}}_{y\in\mathcal{B}}\left[\left|(\rho_y)_{ij}\right|^2\right]-\left|(\rho)_{ij}\right|^2 &= \mathop{\mathrm{EX}}_{y\in\mathcal{B}}\left[\left|\mathop{\mathrm{EX}}_{x\in\mathcal{A}_y}\left[(\rho_x)_{ij}\right]\right|^2\right]-\left|\mathop{\mathrm{EX}}_{x\in\mathcal{D}_1}\left[(\rho_x)_{ij}\right]\right|^2\\ &= \mathop{\mathrm{EX}}_{y\in\mathcal{B},x,z\in\mathcal{A}_y}\left[(\rho_x)^*_{ij}(\rho_z)_{ij}\right]-\mathop{\mathrm{EX}}_{x,z\in\mathcal{D}_1}\left[(\rho_x)^*_{ij}(\rho_z)_{ij}\right]\\ &= \sum_{x,z}\left(\Pr_{\mathcal{D}_2}[x,z]-\Pr_{\mathcal{D}_1^2}[x,z]\right)(\rho_x)^*_{ij}(\rho_z)_{ij}.\end{aligned}$$

Now for all x,z,

$$\left|\sum_{i,j=1}^{2^L}(\rho_x)^*_{ij}(\rho_z)_{ij}\right|\leq\sum_{i,j=1}^{2^L}\left|(\rho_x)_{ij}\right|^2\leq 1.$$

Hence

$$\begin{aligned}\sum_{x,z}\left(\Pr_{\mathcal{D}_2}[x,z]-\Pr_{\mathcal{D}_1^2}[x,z]\right)\sum_{i,j=1}^{2^L}(\rho_x)^*_{ij}(\rho_z)_{ij} &\leq \sum_{x,z}\left(\Pr_{\mathcal{D}_2}[x,z]-\Pr_{\mathcal{D}_1^2}[x,z]\right)\\ &= 2\left\|\mathcal{D}_2-\mathcal{D}_1^2\right\|\\ &\leq 2\delta,\end{aligned}$$

and we are done. □

The difficulty in extending Theorem 5.1 to partial functions is that the distribution $\mathcal{D}_1$ might not make sense, since it might assign a nonzero probability to some x for which $f(x,y)$ is undefined.

5.1 Applications

In this subsection we apply Theorem 5.1 to prove lower bounds for two problems of Ambainis. To facilitate further research and to investigate the scope of our method, we state the problems in a more general way than Ambainis did. Given a group G, the *coset problem* $\mathrm{Coset}(G)$ is defined as follows. Alice is given a left coset C of a subgroup in G, and Bob is given an element $y\in G$. Bob must output 1 if $y\in C$ and 0 otherwise. By restricting the group G, we obtain many interesting and natural problems. For example, if p is prime then $\mathrm{Coset}(\mathbb{Z}_p)$ is just the equality problem, so the protocol of Rabin and Yao [31] yields $Q_2^1(\mathrm{Coset}(\mathbb{Z}_p))=\Theta(\log\log p)$.

Theorem 5.2. $Q_2^1\left(\mathrm{Coset}\left(\mathbb{Z}_p^2\right)\right)=\Theta(\log p)$.

Proof. The upper bound is obvious. For the lower bound, it suffices to consider a function f_p defined as follows. Alice is given $\langle x,y\rangle\in\mathbb{F}_p^2$ and Bob is given $\langle a,b\rangle\in\mathbb{F}_p^2$; then

$$f_p(x,y,a,b)=\begin{cases}1 & \text{if } y\equiv ax+b\,(\mathrm{mod}\,p)\\ 0 & \text{otherwise.}\end{cases}$$

Let $\mathcal{B}$ be the uniform distribution over $\langle a,b\rangle \in \mathbb{F}_p^2$, and let $\mathcal{A}_{a,b}$ be the uniform distribution over $\langle x,y\rangle$ such that $y \equiv ax+b \pmod p$. Thus $\mathcal{D}_1$ is the uniform distribution over $\langle x,y\rangle \in \mathbb{F}_p^2$; note that

$$\Pr_{\langle x,y\rangle\in\mathcal{D}_1,\langle a,b\rangle\in\mathcal{B}}[f_p(x,y,a,b)=0] = 1-\frac{1}{p}.$$

But what about the distribution $\mathcal{D}_2$, which is formed by first drawing $\langle a,b\rangle \in \mathcal{B}$, and then drawing $\langle x,y\rangle$ and $\langle z,w\rangle$ independently from $\mathcal{A}_{a,b}$? Given a pair $\langle x,y\rangle, \langle z,w\rangle \in \mathbb{F}_p^2$, there are three cases regarding the probability of its being drawn from $\mathcal{D}_2$:

(1) $\langle x,y\rangle = \langle z,w\rangle$ (p^2 pairs). In this case

$$\begin{aligned}\Pr_{\mathcal{D}_2}[\langle x,y\rangle,\langle z,w\rangle] &= \sum_{\langle a,b\rangle\in\mathbb{F}_p^2} \Pr[\langle a,b\rangle]\Pr[\langle x,y\rangle,\langle z,w\rangle \mid \langle a,b\rangle] \\ &= p\left(\frac{1}{p^2}\cdot\frac{1}{p^2}\right) = \frac{1}{p^3}.\end{aligned}$$

(2) $x \neq z$ (p^4-p^3 pairs). In this case there exists a unique $\langle a^*,b^*\rangle$ such that $y \equiv a^*x+b^* \pmod p$ and $w \equiv a^*z+b^* \pmod p$, so

$$\begin{aligned}\Pr_{\mathcal{D}_2}[\langle x,y\rangle,\langle z,w\rangle] &= \Pr[\langle a^*,b^*\rangle]\Pr[\langle x,y\rangle,\langle z,w\rangle \mid \langle a^*,b^*\rangle] \\ &= \frac{1}{p^2}\cdot\frac{1}{p^2} = \frac{1}{p^4}.\end{aligned}$$

(3) $x = z$ but $y \neq w$ (p^3-p^2 pairs). In this case $\Pr_{\mathcal{D}_2}[\langle x,y\rangle,\langle z,w\rangle] = 0$.

Putting it all together,

$$\begin{aligned}\left\|\mathcal{D}_2-\mathcal{D}_1^2\right\| &= \frac{1}{2}\left(p^2\left|\frac{1}{p^3}-\frac{1}{p^4}\right| + \left(p^4-p^3\right)\left|\frac{1}{p^4}-\frac{1}{p^4}\right| + \left(p^3-p^2\right)\left|0-\frac{1}{p^4}\right|\right) \\ &= \frac{1}{p}-\frac{1}{p^2}.\end{aligned}$$

So taking $\delta = 1/p-1/p^2$, we have $Q_2^1\left(\text{Coset}\left(\mathbb{Z}_p^2\right)\right) = \Omega(\log(1/\delta)) = \Omega(\log p)$ by Theorem 5.1. □

We now consider Ambainis' second problem. Given a group G and nonempty set $S \subset G$ with $|S| \leq |G|/2$, the *subset problem* $\text{Subset}(G,S)$ is defined as follows. Alice is given $x \in G$ and Bob is given $y \in G$; then Bob must output 1 if $xy \in S$ and 0 otherwise.

Let $\mathcal{M}$ be the distribution over $st^{-1} \in G$ formed by drawing s and t uniformly and independently from S. Then let $\Delta = \|\mathcal{M}-\mathcal{D}_1\|$, where $\mathcal{D}_1$ is the uniform distribution over G.

Proposition 5.3. *For all G,S such that $|S| \leq |G|/2$,*

$$Q_2^1(\text{Subset}(G,S)) = \Omega(\log 1/\Delta).$$

Proof. Let $\mathcal{B}$ be the uniform distribution over $y \in G$, and let $\mathcal{A}_y$ be the uniform distribution over x such that $xy \in S$. Thus $\mathcal{D}_1$ is the uniform distribution over $x \in G$; note that

$$\Pr_{x \in \mathcal{D}_1, y \in \mathcal{B}}[xy \notin S] = 1 - \frac{|S|}{|G|} \geq \frac{1}{2}.$$

We have

$$\begin{aligned}
\left\|\mathcal{D}_2 - \mathcal{D}_1^2\right\| &= \frac{1}{2}\sum_{x,z \in G}\left|\frac{\left|\{y \in G, s,t \in S : xy = s, zy = t\}\right|}{|G|\,|S|^2} - \frac{1}{|G|^2}\right| \\
&= \frac{1}{2}\sum_{x,z \in G}\left|\frac{\left|\left\{s,t \in S : xz^{-1} = st^{-1}\right\}\right|}{|S|^2} - \frac{1}{|G|^2}\right| \\
&= \frac{1}{2}\sum_{x \in G}\left|\frac{\left|\left\{s,t \in S : x = st^{-1}\right\}\right|}{|S|^2} - \frac{1}{|G|}\right| \\
&= \frac{1}{2}\sum_{x \in G}\left|\Pr_{\mathcal{M}}[x] - \frac{1}{|G|}\right| \\
&= \|\mathcal{M} - \mathcal{D}_1\| \\
&= \Delta.
\end{aligned}$$

Therefore $\log(1/\delta) = \Omega(\log 1/\Delta)$. □

Having lower-bounded $Q_2^1(\text{Subset}(G,S))$ in terms of $1/\Delta$, it remains only to upper-bound the variation distance Δ. The following proposition implies that for all constants $\varepsilon > 0$, if S is chosen uniformly at random subject to $|S| = |G|^{1/2+\varepsilon}$, then $Q_2^1(\text{Subset}(G,S)) = \Omega(\log(|G|))$ with constant probability over S.

Theorem 5.4. *For all groups G and integers $K \in \{1,\ldots,|G|\}$, if $S \subset G$ is chosen uniformly at random subject to $|S| = K$, then $\Delta = O\left(\sqrt{|G|}/K\right)$ with $\Omega(1)$ probability over S.*

Proof. We have

$$\Delta = \frac{1}{2}\sum_{x \in G}\left|\Pr_{\mathcal{M}}[x] - \frac{1}{|G|}\right| \leq \frac{\sqrt{|G|}}{2}\sqrt{\sum_{x \in G}\left(\Pr_{\mathcal{M}}[x] - \frac{1}{|G|}\right)^2}$$

by the Cauchy-Schwarz inequality. We claim that

$$\mathop{\mathrm{E}}_{S}\left[\sum_{x \in G}\left(\Pr_{\mathcal{M}}[x] - \frac{1}{|G|}\right)^2\right] \leq \frac{c}{K^2}$$

for some constant c. From this it follows by Markov's inequality that

$$\Pr_{S}\left[\sum_{x \in G}\left(\Pr_{\mathcal{M}}[x] - \frac{1}{|G|}\right)^2 \geq \frac{2c}{K^2}\right] \leq \frac{1}{2}$$

and hence

$$\Delta \leq \frac{\sqrt{|G|}}{2}\sqrt{\frac{2c}{K^2}} = O\left(\frac{\sqrt{|G|}}{K}\right)$$

with probability at least $1/2$.

Let us now prove the claim. We have

$$\Pr_{\mathcal{M}}[x] = \Pr_{i,j}\left[s_i s_j^{-1} = x\right] = \Pr_{i,j}[s_i = xs_j],$$

where $S = \{s_1, \ldots, s_K\}$ and i, j are drawn uniformly and independently from $\{1, \ldots, K\}$. So by linearity of expectation,

$$\begin{aligned}\mathop{\mathrm{EX}}_{S}\left[\sum_{x\in G}\left(\Pr_{\mathcal{M}}[x] - \frac{1}{|G|}\right)^2\right] &= \mathop{\mathrm{EX}}_{S}\left[\sum_{x\in G}\left(\left(\Pr_{i,j}[s_i = xs_j]\right)^2 - \frac{2}{|G|}\Pr_{i,j}[s_i = xs_j] + \frac{1}{|G|^2}\right)\right] \\ &= \sum_{x\in G}\left(\frac{1}{K^4}\sum_{i,j,k,l=1}^{K} p_{x,ijkl}\right) - \frac{2}{|G|}\sum_{x\in G}\left(\frac{1}{K^2}\sum_{i,j=1}^{K} p_{x,ij}\right) + \frac{1}{|G|}\end{aligned}$$

where

$$\begin{aligned} p_{x,ij} &= \Pr_{S}[s_i = xs_j], \\ p_{x,ijkl} &= \Pr_{S}[s_i = xs_j \wedge s_k = xs_l]. \end{aligned}$$

First we analyze $p_{x,ij}$. Let $\mathrm{ord}(x)$ be the order of x in G. Of the K^2 possible ordered pairs (i, j), there are K pairs with the "pattern" ii (meaning that $i = j$), and $K(K-1)$ pairs with the pattern ij (meaning that $i \neq j$). If $\mathrm{ord}(x) = 1$ (that is, x is the identity), then we have $p_{x,ij} = \Pr_S[s_i = s_j]$, so $p_{x,ij} = 1$ under the pattern ii, and $p_{x,ij} = 0$ under the pattern ij. On the other hand, if $\mathrm{ord}(x) > 1$, then $p_{x,ij} = 0$ under the pattern ii, and $p_{x,ij} = \frac{1}{|G|-1}$ under the pattern ij. So

$$\frac{1}{K^2}\sum_{x\in G}\sum_{i,j=1}^{K} p_{x,ij} = \frac{1}{K^2}\left(K + (|G|-1)\frac{K(K-1)}{|G|-1}\right) = 1.$$

Though unnecessarily cumbersome, the above analysis was a warmup for the more complicated case of $p_{x,ijkl}$. The following table lists the expressions for $p_{x,ijkl}$, given $\mathrm{ord}(x)$ and the pattern of (i, j, k, l).

Pattern	Number of such 4-tuples	$\mathrm{ord}(x) = 1$	$\mathrm{ord}(x) = 2$	$\mathrm{ord}(x) > 2$
$iiii, iikk$	K^2	1	0	0
$ijij$	$K(K-1)$	0	$\frac{1}{\lvert G\rvert-1}$	$\frac{1}{\lvert G\rvert-1}$
$ijji$	$K(K-1)$	0	$\frac{1}{\lvert G\rvert-1}$	0
$iiil, iiki, ijii, ijjj$	$4K(K-1)$	0	0	0
$ijki, ijjk$	$2K(K-1)(K-2)$	0	0	$\frac{1}{(\lvert G\rvert-1)(\lvert G\rvert-2)}$
$iikl, ijkk, ijik, ijkj$	$4K(K-1)(K-2)$	0	0	0
$ijkl$	$K(K-1)(K-2)(K-3)$	0	$\frac{1}{(\lvert G\rvert-1)(\lvert G\rvert-3)}$	$\frac{1}{(\lvert G\rvert-1)(\lvert G\rvert-3)}$

Let r be the number of $x \in G$ such that $\mathrm{ord}(x) = 2$, and let $r' = |G| - r - 1$ be the number such that $\mathrm{ord}(x) > 2$. Then

$$\frac{1}{K^4} \sum_{x \in G} \sum_{i,j,k,l=1}^{K} p_{x,ijkl} = \frac{1}{K^4} \left(\begin{array}{c} K^2 + (2r + r') \frac{K(K-1)}{|G|-1} + 2r' \frac{K(K-1)(K-2)}{(|G|-1)(|G|-2)} \\ + (r + r') \frac{K(K-1)(K-2)(K-3)}{(|G|-1)(|G|-3)} \end{array} \right)$$
$$\leq \frac{1}{|G|-3} + O\left(\frac{1}{K^2}\right)$$

using the fact that $K \leq |G|$.

Putting it all together,

$$\mathop{\mathrm{EX}}_{S} \left[\sum_{x \in G} \left(\Pr_{\mathcal{M}}[x] - \frac{1}{|G|} \right)^2 \right] \leq \frac{1}{|G|-3} + O\left(\frac{1}{K^2}\right) - \frac{2}{|G|} + \frac{1}{|G|} = O\left(\frac{1}{K^2}\right)$$

and we are done. □

From fingerprinting we also have the following upper bound. Let q be the periodicity of S, defined as the number of distinct sets $gS = \{gs : s \in S\}$ where $g \in G$.

Proposition 5.5. $R_2^1(\mathrm{Subset}(G,S)) = O(\log|S| + \log\log q)$.

Proof. Assume for simplicity that $q = |G|$; otherwise we could reduce to a subgroup $H \leq G$ with $|H| = q$. The protocol is as follows: Alice draws a uniform random prime p from the range $\left[|S|^2 \log^2|G|, 2|S|^2 \log^2|G|\right]$; she then sends Bob the pair $(p, x \bmod p)$ where x is interpreted as an integer. This takes $O(\log|S| + \log\log|G|)$ bits. Bob outputs 1 if and only if there exists a $z \in G$ such that $zy \in S$ and $x \equiv z \pmod p$. To see the protocol's correctness, observe that if $x \neq z$, then there at most $\log|G|$ primes p such that $x - z \equiv 0 \pmod p$, whereas the relevant range contains $\Omega\left(\frac{|S|^2 \log^2|G|}{\log(|S|\log|G|)}\right)$ primes. Therefore, if $xy \notin S$, then by the union bound

$$\Pr_p[\exists z : zy \in S, x \equiv z \pmod p] = O\left(|S|\log|G| \frac{\log(|S|\log|G|)}{|S|^2 \log^2|G|}\right) = o(1).$$

□

6 Open Problems

- Are $R_2^1(f)$ and $Q_2^1(f)$ polynomially related for every total Boolean function f? Also, can we exhibit *any* asymptotic separation between these measures? The best separation we know of is a factor of 2: for the equality function we have $R_2^1(\mathrm{EQ}) \geq (1 - o(1))\log_2 n$, whereas Winter [37] has shown that $Q_2^1(\mathrm{EQ}) \leq (1/2 + o(1))\log_2 n$ using a protocol involving mixed states.[12] This factor-2 savings is tight for equality: a simple counting argument shows that $Q_2^1(\mathrm{EQ}) \geq$

[12] If we restrict ourselves to pure states, then $(1 - o(1))\log_2 n$ qubits are needed. Based on that fact, a previous version of this paper claimed incorrectly that $Q_2^1(\mathrm{EQ}) \geq (1 - o(1))\log_2 n$.

$(1/2-o(1))\log_2 n$; and although the usual randomized protocol for equality [31] uses $(2+o(1))$ $\log_2 n$ bits, there exist protocols based on error-correcting codes that use only $\log_2(cn) = \log_2 n + O(1)$ bits. All of this holds for any constant error probability $0 < \varepsilon < 1/2$.

- As a first step toward answering the above questions, can we lower-bound $Q_2^1(\text{Coset}(G))$ for groups other than $\mathbb{Z}_p^2$ (such as $\mathbb{Z}_2^n$, or nonabelian groups)? Also, can we characterize $Q_2^1(\text{Subset}(G,S))$ for all sets S, closing the gap between the upper and lower bounds?

- Is there an oracle relative to which $\mathsf{BQP/poly} \neq \mathsf{BQP/qpoly}$?

- Can we give oracles relative to which $\mathsf{NP} \cap \mathsf{coNP}$ and SZK are not contained in $\mathsf{BQP/qpoly}$? Bennett et al. [8] gave an oracle relative to which $\mathsf{NP} \cap \mathsf{coNP} \not\subset \mathsf{BQP}$, while Aaronson [1] gave an oracle relative to which $\mathsf{SZK} \not\subset \mathsf{BQP}$.

- Even more ambitiously, can we prove a direct product theorem for quantum query complexity that applies to any partial or total function (not just search)?

- For all f (partial or total), is $R_2^1(f) = O(\sqrt{n})$ whenever $Q_2^1(f) = O(\log n)$? In other words, is the separation of Bar-Yossef et al. [6] the best possible?

- Can the result $D^1(f) = O\left(mQ_2^1(f)\log Q_2^1(f)\right)$ for partial f be improved to $D^1(f) = O\left(mQ_2^1(f)\right)$? We do not even know how to rule out $D^1(f) = O\left(m + Q_2^1(f)\right)$.

- In the Simultaneous Messages (SM) model, there is no direct communication between Alice and Bob; instead, Alice and Bob both send messages to a third party called the *referee*, who then outputs the function value. The complexity measure is the sum of the two message lengths. Let $R_2^{||}(f)$ and $Q_2^{||}(f)$ be the randomized and quantum bounded-error SM complexities of f respectively, and let $R_2^{||,\text{pub}}(f)$ be the randomized SM complexity if Alice and Bob share an arbitrarily long random string. Building on work by Buhrman et al. [12], Yao [40] showed that $Q_2^{||}(f) = O(\log n)$ whenever $R_2^{||,\text{pub}}(f) = O(1)$. He then asked about the other direction: for some $\varepsilon > 0$, does $R_2^{||,\text{pub}}(f) = O\left(n^{1/2-\varepsilon}\right)$ whenever $Q_2^{||}(f) = O(\log n)$, and does $R_2^{||}(f) = O\left(n^{1-\varepsilon}\right)$ whenever $Q_2^{||}(f) = O(\log n)$? In an earlier version of this paper, we showed that $R_2^{||}(f) = O\left(\sqrt{n}\left(R_2^{||,\text{pub}}(f) + \log n\right)\right)$, which means that a positive answer to Yao's first question would imply a positive answer to the second. Later we learned that Yao independently proved the same result [39].

  Here we ask a related question: can $Q_2^{||}(f)$ ever be exponentially smaller than $R_2^{||,\text{pub}}(f)$? (Buhrman et al. [12] showed that $Q_2^{||}(f)$ can be exponentially smaller than $R_2^{||}(f)$.) Iordanis Kerenidis has pointed out to us that, based on the hidden matching problem of Bar-Yossef et al. [6] discussed in Section 2, one can define a *relation* for which $Q_2^{||}(f)$ is exponentially smaller than $R_2^{||,\text{pub}}(f)$. However, as in the case of $Q_2^1(f)$ versus $R_2^1(f)$, it remains to extend that result to functions.

7 Acknowledgments

I thank Oded Regev for pointing out the connection between advice and one-way communication; Andris Ambainis for posing the coset and subset problems; Harry Buhrman for asking about an upper bound on BQP/qpoly; Lance Fortnow for pointing out that P/poly $\neq$ PP/poly implies PP $\not\subset$ P/poly; Hartmut Klauck and Ronald de Wolf for discussions that led to improvements of Lemma 4.5; Iordanis Kerenidis and Andreas Winter for correcting me about the constant factors in $R_2^1(\mathrm{EQ})$ and $Q_2^1(\mathrm{EQ})$; and the anonymous reviewers and ToC editors for helpful comments and corrections.

References

[1] * S. AARONSON: Quantum lower bound for the collision problem. In *Proc. 34th ACM STOC*, pp. 635–642, 2002. [STOC:509907.509999, arXiv:quant-ph/0111102]. 11, 6

[2] * S. AARONSON: Limitations of quantum advice and one-way communication. In *19th IEEE Conf. Comput. Complexity (CCC)*, pp. 320–332, 2004. [CCC:10.1109/CCC.2004.1313854, arXiv:quant-ph/0402095]. 1

[3] * S. AARONSON: Quantum computing, postselection, and probabilistic polynomial-time. Submitted, 2004. [arXiv:quant-ph/0412187]. 3.1

[4] * L. ADLEMAN, J. DEMARRAIS, AND M.-D. HUANG: Quantum computability. *SIAM J. Computing*, 26(5):1524–1540, 1997. [SICOMP:10.1137/S0097539795293639]. 3.1

[5] * A. AMBAINIS, A. NAYAK, A. TA-SHMA, AND U. V. VAZIRANI: Quantum dense coding and quantum finite automata. *J. ACM*, 49:496–511, 2002. Earlier version in 31st ACM STOC, 1999, pp. 376-383. [JACM:581771.581773, arXiv:quant-ph/9804043]. 1, 2.3

[6] * Z. BAR-YOSSEF, T. S. JAYRAM, AND I. KERENIDIS: Exponential separation of quantum and classical one-way communication complexity. In *36th ACM STOC*, pp. 128–137, 2004. [STOC:1007352.1007379, ECCC:TR04-036]. 1, 2.1, 6

[7] * R. BEALS, H. BUHRMAN, R. CLEVE, M. MOSCA, AND R. DE WOLF: Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001. Earlier version in 39th IEEE FOCS, 1998, pp. 352-361. [JACM:502090.502097, arXiv:quant-ph/9802049]. 1, 4, 4

[8] * C. BENNETT, E. BERNSTEIN, G. BRASSARD, AND U. VAZIRANI: Strengths and weaknesses of quantum computing. *SIAM J. Computing*, 26(5):1510–1523, 1997. [SICOMP:30093, arXiv:quant-ph/9701001]. 1, 4, 6

[9] * C. H. BENNETT, G. BRASSARD, C. CRÉPEAU, R. JOZSA, A. PERES, AND W. WOOTTERS: Teleporting an unknown quantum state by dual classical and EPR channels. *Phys. Rev. Lett.*, 70:1895–1898, 1993. [PRL:10.1103/PhysRevLett.70.1895]. 1

[10] * C. H. BENNETT AND J. GILL: Relative to a random oracle A, $P^A \neq NP^A \neq coNP^A$ with probability 1. *SIAM J. Computing*, 10(1):96–113, 1981. [SICOMP:10/0210008]. 11

[11] * S. N. BERNSTEIN: Sur l'ordre de la meilleure approximation des fonctions continues par les polynômes de degré donné. *Mem. Cl. Sci. Acad. Roy. Belg.*, 4:1–103, 1912. French. 10

[12] * H. BUHRMAN, R. CLEVE, J. WATROUS, AND R. DE WOLF: Quantum fingerprinting. *Phys. Rev. Lett.*, 87(16), 2001. 4 pages. [PRL:10.1103/PhysRevLett.87.167902, arXiv:quant-ph/0102001]. 1, 6

[13] * H. BUHRMAN AND R. DE WOLF: Complexity measures and decision tree complexity: a survey. *Theoretical Computer Sci.*, 288:21–43, 2002. [TCS:10.1016/S0304-3975(01)00144-X]. 1

[14] * P. DURIŠ, J. HROMKOVIČ, J. D. P. ROLIM, AND G. SCHNITGER: Las Vegas versus determinism for one-way communication complexity, finite automata, and polynomial-time computations. In *14th Symp. Theoret. Aspects of Comp.Sci. (STACS)*, pp. 117–128, 1997. 2.1

[15] * H. EHLICH AND K. ZELLER: Schwankung von Polynomen zwischen Gitterpunkten. *Mathematische Zeitschrift*, 86:41–44, 1964. 4

[16] * L. FORTNOW AND J. ROGERS: Complexity limitations on quantum computation. *J. Comput. Sys. Sci.*, 59(2):240–252, 1999. [JCSS:10.1006/jcss.1999.1651, arXiv:cs.CC/9811023]. 3.1

[17] * O. GOLDREICH: On quantum computing. 2004. 1 page. 1

[18] * Y. HAN, L. HEMASPAANDRA, AND T. THIERAUF: Threshold computation and cryptographic security. *SIAM J. Computing*, 26(1):59–78, 1997. [SICOMP:24046]. 3.1

[19] * A. S. HOLEVO: Some estimates of the information transmitted by quantum communication channels. *Problems of Information Transmission*, 9:177–183, 1973. English translation. 1

[20] * H. KLAUCK: Quantum communication complexity. In *27th Int'l. Colloq. on Automata, Languages, and Programming (ICALP)*, pp. 241–252, 2000. [arXiv:quant-ph/0005032]. 1, 2.1, 3

[21] * H. KLAUCK: Quantum time-space tradeoffs for sorting. In *35th ACM STOC*, pp. 69–76, 2003. [STOC:780542.780553, arXiv:quant-ph/0211174]. 1, 4, 9

[22] * H. KLAUCK, R. ŠPALEK, AND R. DE WOLF: Quantum and classical strong direct product theorems and optimal time-space tradeoffs. In *45th IEEE FOCS*, pp. 12–21, 2004. [FOCS:10.1109/FOCS.2004.52, arXiv:quant-ph/0402123]. 1, 4, 4, 10

[23] * E. KUSHILEVITZ AND N. NISAN: *Communication Complexity*. Cambridge, 1997. 1

[24] * A. A. MARKOV: On a question by D. I. Mendeleev. *Zapiski Imperatorskoi Akademii Nauk*, SP6(62):1–24, 1890. Russian. English translation at `http://www.math.technion.ac.il/hat/fpapers/markov4.pdf`. 4

[25] * V. A. MARKOV: Über Polynome, die in einem gegebenen Intervalle möglichst wenig von Null abweichen. *Math. Ann.*, 77:213–258, 1916. German. Originally written in 1892. 4

[26] * M. MINSKY AND S. PAPERT: *Perceptrons (2nd edition)*. MIT Press, 1988. First appeared in 1968. 4

[27] * A. NAYAK: Optimal lower bounds for quantum automata and random access codes. In *40th IEEE FOCS*, pp. 369–377, 1999. [FOCS:10.1109/SFFCS.1999.814608, arXiv:quant-ph/9904093]. 1, 3

[28] * M. NIELSEN AND I. CHUANG: *Quantum Computation and Quantum Information*. Cambridge University Press, 2000. 1, 2.3

[29] * N. NISAN AND M. SZEGEDY: On the degree of Boolean functions as real polynomials. *Computational Complexity*, 4(4):301–313, 1994. 4, 4

[30] * H. NISHIMURA AND T. YAMAKAMI: Polynomial time quantum computation with advice. *Inform. Proc. Lett.*, 90:195–204, 2003. [IPL:10.1016/j.ipl.2004.02.005, ECCC:TR03-059, arXiv:quant-ph/0305100]. 6, 2.2

[31] * M. RABIN AND A. C-C. YAO: Manuscript, cited in [38], 1979. 1, 5.1, 6

[32] * T. J. RIVLIN: *Chebyshev Polynomials: From Approximation Theory to Algebra and Number Theory*. Wiley, 1990. 4, 4

[33] * T. J. RIVLIN AND E. W. CHENEY: A comparison of uniform approximations on an interval and a finite subset thereof. *SIAM J. Numerical Analysis*, 3(2):311–320, 1966. [SINUM:03/0703024]. 4

[34] * N. SAUER: On the density of families of sets. *J. Combinatorial Theory Series A*, 13:145–147, 1972. [JCombThA:10.1016/0097-3165(72)90019-2]. 1, 3

[35] * Y. SHI: Both Toffoli and controlled-NOT need little help to do universal quantum computation. *Quantum Information and Computation*, 3(1):84–92, 2002. [arXiv:quant-ph/0205115]. 3.1

[36] * J. WATROUS: Succinct quantum proofs for properties of finite groups. In *41st IEEE FOCS*, pp. 537–546, 2000. [FOCS:10.1109/SFCS.2000.892141, arXiv:cs.CC/0009002]. 2.2

[37] * A. WINTER: Quantum and classical message identification via quantum channels. In O. Hirota, editor, *Quantum Information, Statistics, Probability (A. S. Holevo festschrift)*, pp. 172–189. Rinton, 2004. [arXiv:quant-ph/0401060]. 2.1, 6

[38] * A. C-C. YAO: Some complexity questions related to distributive computing. In *11th ACM STOC*, pp. 209–213, 1979. [STOC:804414]. 1, 31

[39] * A. C-C. YAO: Princeton University course assignment, 2001. 6

[40] * A. C-C. YAO: On the power of quantum fingerprinting. In *35th ACM STOC*, pp. 77–81, 2003. [STOC:780542.780554, ECCC:TR02-074]. 1, 6

AUTHOR

Scott Aaronson[13]
postdoc
Institute for Advanced Study, Princeton
aaronson@ias.edu
http://www.cs.berkeley.edu/~aaronson

ABOUT THE AUTHOR

Scott Aaronson received his Ph.D. at the University of California, Berkeley, in 2004 under Umesh Vazirani. He is proud to be the first author in Theory of Computing.

[13]This work was performed while the author was a graduate student at UC Berkeley.