

The Quantum Query Complexity of Certification

Andris Ambainis* Andrew M. Childs† François Le Gall‡ Seiichiro Tani§

Abstract

We study the quantum query complexity of finding a certificate for a d -regular, k -level balanced NAND formula. We show that the query complexity is $\tilde{\Theta}(d^{(k+1)/2})$ for 0-certificates, and $\tilde{\Theta}(d^{k/2})$ for 1-certificates. In particular, this shows that the zero-error quantum query complexity of evaluating such formulas is $\tilde{O}(d^{(k+1)/2})$. Our lower bound relies on the fact that the quantum adversary method obeys a direct sum theorem.

1 Introduction

Recently, there has been considerable progress in understanding the quantum query complexity of evaluating Boolean formulas. Here, we consider the closely related problem of *certifying* the value of a formula.

An n -bit Boolean formula ϕ corresponds to a rooted tree with n leaves, where each leaf represents a binary variable x_i for $i \in \{1, \dots, n\}$, and each internal vertex represents a logic gate acting on its children. We can evaluate ϕ on an input $x \in \{0, 1\}^n$ by applying the logic gates from the leaves toward the root, giving some binary value $\phi(x)$ at the root. Suppose we are given a black box that, on input i , produces the bit x_i . The minimum number of queries to such a black box required to learn $\phi(x)$ for an arbitrary $x \in \{0, 1\}^n$ is called the query complexity of evaluating ϕ .

In this article, we focus on the case of NAND formulas, in which each internal vertex corresponds

to a NAND gate. (Note that formulas consisting of AND, OR, and NOT gates are equivalent to NAND formulas.) Furthermore, we suppose that the tree is regular, with internal vertices of degree d , and balanced, with a total of k levels. Thus, there are $n = d^k$ input variables.

To evaluate such a formula on a black-box input, a deterministic classical computer must make all n queries. However, the expected number of queries for a randomized classical computer to evaluate the formula with zero error is $O(\lambda_d^k)$ [15, 17], where $\lambda_d := (d-1 + \sqrt{d^2 + 14d + 1})/4$, using a simple recursive strategy. Indeed, this is the optimal query complexity not just for zero-error algorithms [15], but also for bounded-error ones [16]. If we fix d and let k grow, this represents a polynomial speedup over the deterministic strategy of evaluating all leaves, since only $O(n^{\log_d \lambda_d})$ queries are required; for example, with $d = 2$, the classical query complexity is $\Theta(n^{0.753\dots})$. On the other hand, if we fix k and let d grow, then the query complexity is essentially n .

For quantum computers, the situation is quite different. Grover's algorithm shows that if $k = 1$, so that $d = n$, a quantum computer can evaluate the formula with bounded error in $O(\sqrt{n})$ queries [11], which is optimal [6]. Similar results can be obtained for constant k : a recursive strategy with straightforward error reduction uses $O(\sqrt{n} \log^{k-1} n)$ queries [7], and a more delicate approach to error reduction uses $O(\sqrt{n} c^k)$ queries for some constant c [13]. In particular, this shows that the quantum query complexity is $O(\sqrt{n})$ provided k is constant. However, for k growing with n (say, with constant d), no better quantum algorithm than the classical one was known for nearly ten years. A breakthrough occurred with the revolutionary quantum walk approach of [10], which—combined with a simple observation on the simulation of Hamiltonian dynamics by quantum circuits [9]—showed that $\sqrt{n^{1+o(1)}}$ queries suffice for the

*Department of Computing, University of Latvia

†Department of Combinatorics & Optimization and Institute for Quantum Computing, University of Waterloo

‡ERATO-SORST Quantum Computation and Information Project / Graduate School of Informatics, Kyoto University

§ERATO-SORST Quantum Computation and Information Project / NTT Communication Science Laboratories, NTT Corporation

case where $d = 2$. Subsequent work showed that in fact the quantum query complexity is $O(\sqrt{n})$ for any d, k [3]. Since this is optimal [4], the quantum query complexity of evaluating balanced NAND formulas is now a closed problem.

Here, we consider the problem of determining a minimal certificate for the value of a NAND formula. A *certificate* for ϕ on input $x \in \{0, 1\}^n$ is a subset $c \subseteq \{1, \dots, n\}$ such that the values of x_i for $i \in c$ determine $\phi(x)$. We say that a certificate c is *minimal* if no proper subset of c is also a certificate.

We can view the size of the minimal certificate as a kind of nondeterministic query complexity, since this is the number of queries required to verify the value of the formula if we can guess the certificate. The certificate size of a regular, balanced NAND formula can be understood simply, as follows. Let $C_b(k)$ denote the size of a minimal certificate for a k -level balanced NAND formula evaluating to b (a *b-certificate*). We have $C_0(0) = C_1(0) = 1$, and

$$C_0(k) = dC_1(k-1) \quad C_1(k) = C_0(k-1).$$

Solving this recursion, we obtain

$$C_0(k) = \begin{cases} d^{k/2}k & \text{even} \\ d^{(k+1)/2}k & \text{odd} \end{cases}$$

$$C_1(k) = \begin{cases} d^{k/2}k & \text{even} \\ d^{(k-1)/2}k & \text{odd.} \end{cases}$$

Thus, the certificate size is precisely \sqrt{n} for a formula with an even number of levels, and is close to that value when the number of levels is large and odd. Note that a minimal 0-certificate is a set of 0 inputs for k even and a set of 1 inputs for k odd; a minimal 1-certificate is a set of 1 inputs for k even and a set of 0 inputs for k odd.

Given a regular, balanced NAND formula and a black box for the input, we are interested in the number of queries required to output a minimal certificate. This question is of interest for several reasons. Whereas the optimal classical algorithm for evaluating a NAND formula produces a certificate in the course of its evaluation, the same does not hold for the known optimal quantum algorithms, so it is natural to ask whether quantum computers have a comparable advantage for certification as they do for evaluation. (The algorithm of [3] produces a quantum state that is related to certificates,

which suggests a potential approach to certificate finding, but it turns out that a simpler strategy is nearly optimal.) In addition, the query complexity of certificate finding gives an upper bound on the zero-error query complexity of the problem (as emphasized in [8]), although it is not known if this bound is optimal in general.

For a constant number of levels k , the previous best quantum algorithm for certifying a NAND formula uses $O(d^{(k+2)/2})$ queries in the case of 0-certificates, and $O(d^{(k+1)/2})$ in the case of 1-certificates [8, Lemma 1]. We improve upon this, and also give a result for non-constant k , as follows:

Theorem 1. *Consider a d -regular, k -level balanced NAND formula ϕ with black-box input x . The bounded-error and zero-error quantum query complexities of certifying (ϕ, x) are*

$$O(k^2 d^{(k+1)/2}) \quad \text{and} \quad \Omega(d^{(k+1)/2}) \quad \text{if } \phi(x) = 0$$

$$O(k^2 d^{k/2}) \quad \text{and} \quad \Omega(d^{k/2}) \quad \text{if } \phi(x) = 1.$$

In particular, the zero-error quantum query complexity of evaluating $\phi(x)$ is $O(k^2 d^{(k+1)/2})$.

The upper bound comes from a strategy that applies the formula evaluation algorithm of [3] recursively, as described in Section 2. The lower bound appeals to direct sum theorems. We discuss a direct sum theorem for adversary lower bounds in general in Section 3, and apply it to certification in Section 4. We conclude in Section 5 by mentioning a few open questions.

2 Algorithm

To find a certificate for a NAND formula, consider the following simple recursive strategy. First evaluate the root using the $O(\sqrt{n})$ -query algorithm of [3]. If the result is 0, each child of the root must evaluate to 1, and we can find a 1-certificate for each in turn. If the result is 1, then we search for a child evaluating to 0 using Grover's algorithm combined with formula evaluation, and then find a 0-certificate for it.

This strategy gives a nearly optimal algorithm for certificate finding. However, we must be careful when applying the formula evaluation algorithm recursively, since it only succeeds with bounded error, and we must ensure that this error can be kept under control.

To find a b -certificate for a k -level tree, we can apply the procedure $A_b(k)$, defined as follows.

Procedure $A_0(k)$.

- If $k = 1$, query every leaf. If they are all 1, then output every index; otherwise, go into an infinite loop.
- If $k > 1$, run $A_1(k-1)$ on each subtree.

Procedure $A_1(k)$.

- If $k = 1$, repeat Grover's algorithm until a 0-leaf is found, verifying the result to ensure one-sided error. When a 0-leaf is found, output its index.
- If $k > 1$, repeat the following until passing the verification:
 - Perform robust quantum search [13] for a 0-subtree, using the NAND tree evaluation algorithm [3] as a subroutine. (Note that robust quantum search has two-sided error.) Repeat the search until a 0-subtree T is found.
 - Verify that T is a 0-subtree with error probability at most $1/n^2$ by running the NAND tree evaluation algorithm $O(\log n)$ times and taking a majority vote.
- Once T passes the verification, run $A_0(k-1)$ on T .

A simple inductive argument shows that $A_b(k)$ will never terminate when given a formula evaluating to \bar{b} . Furthermore, when $A_b(k)$ does halt, it is guaranteed to return a correct b -certificate.

It remains to understand how long it takes for $A_b(k)$ to produce an b -certificate.

Lemma 2. *Suppose that $A_b(k)$ is given a d -regular, k -level balanced NAND formula evaluating to b . Assume that the verification step in procedure A_1 never makes an error. Then the expected number of queries before $A_0(k), A_1(k)$ terminates is*

$$O\left(k\left(1 + k\frac{\log d}{\sqrt{d}}\right)\sqrt{d^{k+1}}\right), \quad O\left(k\left(1 + k\frac{\log d}{\sqrt{d}}\right)\sqrt{d^k}\right),$$

respectively.

Proof. Let $\mathbb{E}[A_b(k)]$ denote the expected number of queries made by $A_b(k)$ when given a formula evaluating to b , assuming the procedure A_1 never

makes an erroneous verification. For $k = 1$, we have $\mathbb{E}[A_0(1)] = d$ and $\mathbb{E}[A_1(1)] = O(\sqrt{d})$. For $k \geq 2$,

$$\begin{aligned} \mathbb{E}[A_1(k)] &= O(d^{k/2}) + O(d^{(k-1)/2} \log n) + \mathbb{E}[A_0(k-1)] \\ &= O\left(\left(1 + k\frac{\log d}{\sqrt{d}}\right)d^{k/2}\right) + \mathbb{E}[A_0(k-1)] \end{aligned}$$

and

$$\mathbb{E}[A_0(k)] = d\mathbb{E}[A_1(k-1)].$$

Solving this recurrence gives

$$\mathbb{E}[A_1(k)] = O\left(k\left(1 + k\frac{\log d}{\sqrt{d}}\right)\sqrt{d^k}\right)$$

and

$$\mathbb{E}[A_0(k)] = O\left(k\left(1 + k\frac{\log d}{\sqrt{d}}\right)\sqrt{d^{k+1}}\right)$$

as claimed. \square

Of course, we cannot actually assume that A_1 never makes a verification error, but we can nevertheless obtain a zero-error algorithm for b -certification as follows. Fix some sufficiently large constant c . Run A_b on the input; if it does not halt after $ck(1 + k\log d/\sqrt{d})d^{(k+1)/2}$ queries (for $b = 0$) or $ck(1 + k\log d/\sqrt{d})d^{k/2}$ queries (for $b = 1$), then restart. The probability that one or more verifications fail during each trial is $O((\log^2 n)/n)$, since the error probability of each verification is at most $1/n^2$, and the total number of verifications per trial is at most the total number of queries, which is upper bounded by $O(n\log^2 n)$. If a failure occurs during a given trial, then the algorithm will either output a correct certificate or will not terminate during that trial. Thus we only expect $O(1)$ repetitions to be required before the process terminates with a correct certificate, and we obtain an unconditional zero-error quantum algorithm with an expected running time given by Lemma 2.

3 A direct sum theorem for the adversary method

We now turn to the question of lower bounds. Our approach to the lower bound for certificate finding is based on the concept of a *direct sum theorem*. A direct sum theorem says that if solving one instance

of a problem requires q queries, then solving t instances requires $\Omega(tq)$ queries. (Several papers on quantum lower bounds—e.g., [2, 14]—have studied direct *product* theorems, which are stronger statements. In particular, a *strong direct product theorem* says that if an algorithm attempts to solve t instances with $o(tq)$ queries, then it can only succeed with an exponentially small probability. In this paper, we only need direct sum theorems.)

One of the main methods for proving lower bounds on quantum query complexity is the quantum adversary method [1], which can be reformulated as follows [5]:

Theorem 3. *Let $S \subseteq \{0, 1\}^n$, and let Σ be a finite set. For any function $f : S \rightarrow \Sigma$, define*

$$\text{ADV}(f) := \max_{\Gamma} \frac{\|\Gamma\|}{\max_{i=1}^n \|\Gamma \circ D_i\|},$$

where $\|\cdot\|$ denotes the spectral norm; $\Gamma \neq 0$ is a symmetric, entrywise non-negative $|S| \times |S|$ matrix satisfying $\Gamma_{xy} = 0$ if $f(x) \neq f(y)$; $(D_i)_{xy} = 1$ if $x_i \neq y_i$ and $(D_i)_{xy} = 0$ otherwise; and \circ denotes the Hadamard (i.e., entrywise) product of matrices. Finally, let $Q_\epsilon(f)$ denote the minimum number of quantum queries to x needed to compute $f(x)$ with error at most ϵ . Then

$$Q_\epsilon(f) \geq \frac{1 - 2\sqrt{\epsilon(1-\epsilon)}}{2} \text{ADV}(f).$$

We show that the quantity $\text{ADV}(f)$ obeys a direct sum theorem, as follows.

Theorem 4. *Let $S \subseteq \{0, 1\}^n$, and let Σ be a finite set. Given a function $f^{(1)} : S \rightarrow \Sigma$, let $f^{(t)} : S^t \rightarrow \Sigma^t$ be the function defined by $f^{(t)}(x) := (f^{(1)}(x_1 \dots x_n), \dots, f^{(1)}(x_{(t-1)n+1} \dots x_{tn}))$. Then $\text{ADV}(f^{(t)}) = t \text{ADV}(f^{(1)})$.*

Proof. Let $\Gamma^{(1)}$ be an optimal adversary matrix for $f^{(1)}$, and let $\Gamma^{(t)} := \Gamma^{(1)} \otimes I^{\otimes t-1} + I \otimes \Gamma^{(1)}$, where I denotes the $|S| \times |S|$ identity matrix. Choosing $\Gamma = \Gamma^{(t)}$, the numerator of $\text{ADV}(f^{(t)})$ is $\|\Gamma^{(t)}\| = t\|\Gamma^{(1)}\|$.

Now let $D_i^{(t)}$ denote the $|S|^t \times |S|^t$ matrix with $(D_i^{(t)})_{xy} = 1$ if $x_i \neq y_i$ and $(D_i^{(t)})_{xy} = 0$ otherwise. For $i \in \{1, \dots, n\}$, we have $D_i^{(t)} = D_i^{(1)} \otimes J^{\otimes t-1}$, where J denotes the $|S| \times |S|$ matrix with every entry equal to 1. Matrices $D_i^{(t)}$ for $i \in \{n+1, \dots, tn\}$

can be expressed similarly by suitable permutation of the tensor factors. Then for $i \in \{1, \dots, n\}$,

$$\begin{aligned} \|\Gamma^{(t)} \circ D_i^{(t)}\| &= \|(\Gamma^{(1)} \otimes I^{\otimes t-1} + I \otimes \Gamma^{(1)}) \\ &\quad \circ (D_i^{(1)} \otimes J^{\otimes t-1})\| \\ &= \|(\Gamma^{(1)} \circ D_i^{(1)}) \otimes I^{\otimes t-1}\| \\ &= \|\Gamma^{(1)} \circ D_i^{(1)}\|. \end{aligned}$$

By symmetry, $\|\Gamma^{(t)} \circ D_i^{(t)}\|$ for general $i \in \{1, \dots, tn\}$ only depends on the congruence class of i modulo n . It follows that $\text{ADV}(f^{(t)}) \geq t \text{ADV}(f^{(1)})$.

Although we do not need the converse to establish a direct sum theorem, it can be proven as follows. Using semidefinite programming duality, we can express $\text{ADV}(f)$ as a minimization problem; in particular [19],

$$\text{ADV}(f) = \min_p \max_{\substack{x, y \in S \\ f(x) \neq f(y)}} \frac{1}{\sum_{i: x_i \neq y_i} \sqrt{p_x(i) p_y(i)}},$$

where the minimization is over sets of probability distributions $p_x(i)$, i.e., $\sum_{i=1}^n p_x(i) = 1$ for all $x \in S$. Let $p^{(1)}$ achieve the minimum in (3) for $f = f^{(1)}$. Define a set of probability distributions $p^{(t)}$ by $p_x^{(t)}(i) := \frac{1}{t} p_{x_{n \lfloor i/n \rfloor + 1} \dots x_{n \lfloor i/n \rfloor + n}}^{(1)}(i \bmod n)$ for $i \in \{1, \dots, tn\}$; this is clearly a probability distribution for any $x \in S^t$. To ensure $f^{(t)}(x) \neq f^{(t)}(y)$, we must have $f^{(1)}(x_{jn+1} \dots x_{jn+n}) \neq f^{(1)}(y_{jn+1} \dots y_{jn+n})$ for some $j \in \{0, \dots, t-1\}$. Without loss of generality, suppose $j = 1$. The maximum over $x, y \in S^t$ is achieved by taking $x_i = y_i$ for all $i \in \{n+1, \dots, tn\}$ so that there is no contribution to the denominator from these indices. Thus, we find $\text{ADV}(f^{(t)}) \leq t \text{ADV}(f^{(1)})$, which completes the proof. \square

As a consequence of this fact, whenever a function f has an optimal adversary lower bound, it obeys a direct sum theorem. This can be compared with the multiplicative adversary method [18], in which a multiplicative adversary lower bound implies a strong direct product theorem.

Note that the lower bound of Theorem 4 applies to the negative adversary method as well [12], so negative adversary lower bounds also imply direct sum theorems.

4 Lower bound

We now apply the result of the previous section to give a nearly optimal lower bound for certificate finding.

The query complexity of certification is clearly $\Omega(d^{k/2})$, as knowing a certificate allows one to evaluate the formula. (Recall that a minimal certificate for a balanced formula is a subset of inputs all taking the same value; hence we can learn this value by evaluating one input.) This immediately shows that the algorithm of Section 2 is close to optimal for 1-certificate finding.

It remains to prove the lower bound for 0-certificate finding. First, suppose that k is even. In this case, a 0-certificate consists of one 0-leaf from each of $d^{k/2}$ bottom subtrees, where each subtree has size d . Suppose we are told which of the bottom subtrees we need to consider; this can only make the problem easier. Then we must find a 0-leaf in each of these subtrees. Intuitively, the best strategy for finding these leaves is to run Grover's algorithm for an input of size d independently on each of the $d^{k/2}$ subtrees. To make this intuition rigorous, we can apply the direct sum theorem for search, which says we need at least $\Omega(t\sqrt{d})$ queries to solve t independent instances of the search problem on inputs of size d . While we could prove this using Theorem 4, it is also implied by the strong direct product theorem for search from [2] (which applies even to the problem of finding t marked items with the promise that they exist). Thus, with $t = d^{k/2}$, we find a lower bound of $\Omega(d^{(k+1)/2})$ queries. (A similar argument works for 1-certificate finding with k odd, but we do not need it.)

We can apply a similar approach when k is odd. Suppose we could prove a direct sum theorem for 1-certifying two-level NAND formulas—in other words, that finding 1-certificates to t independent two-level formulas requires t times the number of queries required to certify one such formula. A 0-certificate with k odd consists of all d 1-leaves of each of $d^{(k-1)/2}$ bottom subtrees. If we are told which subtrees we must consider one level above, then we must evaluate $d^{(k-1)/2}$ two-level trees. Thus, a direct sum theorem would say that we need $\Omega(d^{(k-1)/2} \cdot d) = \Omega(d^{(k+1)/2})$ queries.

To prove this direct sum theorem, by Theorem 4, it suffices to give an adversary lower bound for the problem of 1-certifying a two-level NAND formula.

Since we are promised that the formula evaluates to 1, at least one of its d subtrees evaluates to 0, i.e., all d inputs to that subtree are 1. We further promise that precisely one subtree evaluates to 0, and that in all other subtrees, $d-1$ of the d inputs evaluate to 1, with the remaining input evaluating to 0. These promises only make the problem easier. Let $S \subseteq \{0, 1\}^{d^2}$ denote the set of such inputs, and let $f : S \rightarrow \{1, \dots, d\}$ be defined by $f(x) = i$, where i is the unique index of the subtree that evaluates to 0.

We claim that $\text{ADV}(f) = \Omega(d)$. To see this, let Γ be the $|S| \times |S|$ matrix in which $\Gamma_{xy} = 0$ unless x and y differ in exactly two bits and $f(x) \neq f(y)$, in which case $\Gamma_{xy} = 1$. In the latter case,

- the subtree which evaluates to 0 in x evaluates to 1 in y , meaning that one of its d inputs switches from 1 in x to 0 in y , and
- exactly one of $d-1$ subtrees which evaluates to 1 in x evaluates to 0 in y , meaning that its unique 0-input in x switches to 1 in y .

Thus, for any fixed x , there are $d(d-1)$ possible y with $f(x) \neq f(y)$. In other words, Γ is the adjacency matrix of a $d(d-1)$ -regular graph, which means that $\|\Gamma\| = d(d-1)$.

Now, fixing x as well as an index i , the maximum number of y s with $x_i \neq y_i$ is d : if i is an input to the subtree evaluating to 0 in x , then we only have the freedom to choose one of $d-1$ other subtrees in which we switch a fixed input, whereas if i is an input to a subtree evaluating to 1 in x , then we only have the freedom to choose one of d inputs to the 0-subtree to switch. This means that $D_i \circ \Gamma$ is the adjacency matrix of a graph of maximum degree d , which implies $\|D_i \circ \Gamma\| \leq d$. Overall, we have $\text{ADV}(f) \geq d-1 = \Omega(d)$.

Note that, more generally, a similar argument can be used to establish a direct sum theorem for the problem of certifying k -level NAND formulas. In particular, the bounded-error quantum query complexity of certifying t independent d -regular, k -level balanced NAND formulas, each of which evaluates to $b \in \{0, 1\}$, is $\Omega(td^{(k+1)/2})$ for $b = 0$ and $\Omega(td^{k/2})$ for $b = 1$.

5 Open questions

We have described a nearly optimal quantum algorithm for certifying the value of a formula. This leaves several open questions:

- For constant-degree formulas, there is a logarithmic gap between our lower and upper bounds. Is it possible to improve the lower bound or the algorithm (or both)?
- Since a certificate can be used to verify the value of a formula with zero error in only $O(\sqrt{n})$ queries, our result shows in particular that the zero-error quantum query complexity of evaluating a d -regular, k -level balanced NAND formula is $O(k^2 d^{(k+1)/2})$. However, the best known lower bound is $\Omega(d^{k/2})$, from the bounded-error case. Is there a faster quantum algorithm for evaluating a formula with zero error that does not work by producing a certificate? (More generally, the relationship between bounded- and zero-error quantum query complexity is poorly understood.)
- We have restricted our attention to regular, balanced formulas. What is the quantum query complexity of certifying an arbitrary NAND formula? Of course, we can apply a similar recursive strategy to unbalanced formulas, but in general, it is not clear how Procedure A_1 should select a 0-subtree, since some subtrees may require more queries to certify than others.

Acknowledgments

We thank Harry Buhrman and Rahul Jain for helpful discussions. AA was supported by University of Latvia research grant ZP01-100 and a Marie Curie International Reintegration Grant (IRG). AMC received support from MITACS, NSERC, and the US ARO/DTO.

References

- [1] A. Ambainis, *Quantum lower bounds by quantum arguments*, Journal of Computer and System Sciences **64** (2002), no. 4, 750–767, available at quant-ph/0002066. Preliminary version in STOC 2000.
- [2] ———, *A new quantum lower bound method, with an application to strong direct product theorem for quantum search*, quant-ph/0508200.
- [3] A. Ambainis, A. M. Childs, B. W. Reichardt, R. Špalek, and S. Zhang, *Any AND-OR formula of size N can be evaluated in time $N^{1/2+o(1)}$ on a quantum computer*, Proc. 48th IEEE Symposium on Foundations of Computer Science, 2007, pp. 363–372. Merged version of quant-ph/0703015 and arXiv:0704.3628.
- [4] H. Barnum and M. Saks, *A lower bound on the quantum query complexity of read-once functions*, Journal of Computer and System Sciences **69** (2004), no. 2, 244–258, available at quant-ph/0201007.
- [5] H. Barnum, M. Saks, and M. Szegedy, *Quantum query complexity and semidefinite programming*, Proc. 18th IEEE Conference on Computational Complexity, 2003, pp. 179–193.
- [6] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, *Strengths and weaknesses of quantum computing*, SIAM Journal on Computing **26** (1997), 1510–1523, available at quant-ph/9701001.
- [7] H. Buhrman, R. Cleve, and A. Wigderson, *Quantum vs. classical communication and computation*, Proc. 30th ACM Symposium on Theory of Computing, 1998, pp. 63–68, available at quant-ph/9802040.
- [8] H. Buhrman, R. Cleve, R. de Wolf, and C. Zalka, *Bounds for small-error and zero-error quantum algorithms*, Proc. 40th IEEE Symposium on Foundations of Computer Science, 1999, pp. 358–368, available at cs/9904019.
- [9] A. M. Childs, R. Cleve, S. P. Jordan, and D. Yeung, *Discrete-query quantum algorithm for NAND trees*, quant-ph/0702160.
- [10] E. Farhi, J. Goldstone, and S. Gutmann, *A quantum algorithm for the Hamiltonian NAND tree*, Theory of Computing **4** (2008), no. 1, 169–190, available at quant-ph/0702144.
- [11] L. K. Grover, *Quantum mechanics helps in searching for a needle in a haystack*, Physical Review Letters **79** (1997), 325–328, available at quant-ph/9706033. Preliminary version in STOC 1996.
- [12] P. Høyer, T. Lee, and R. Špalek, *Negative weights make adversaries stronger*, Proc. 39th ACM Symposium on Theory of Computing, 2007, pp. 526–535.
- [13] P. Høyer, M. Mosca, and R. de Wolf, *Quantum search on bounded-error inputs*, Proc. 30th International Colloquium on Automata, Languages, and Programming, 2003, pp. 291–299, available at quant-ph/0304052.
- [14] H. Klauck, R. Špalek, and R. de Wolf, *Quantum and classical strong direct product theorems and optimal time-space tradeoffs*, SIAM Journal on Computing **36** (2007), no. 5, 1472–1493, available at quant-ph/0402123. Preliminary version in FOCS 2004.
- [15] M. Saks and A. Wigderson, *Probabilistic boolean decision trees and the complexity of evaluating game trees*, Proc. 27th IEEE Symposium on Foundations of Computer Science, 1986, pp. 29–38.

- [16] M. Santha, *On the Monte Carlo Boolean decision tree complexity of read-once formulae*, Random Structures and Algorithms **6** (1995), no. 1, 75–87.
- [17] M. Snir, *Lower bounds on probabilistic linear decision trees*, Theoretical Computer Science **38** (1985), 69–82.
- [18] R. Špalek, *The multiplicative quantum adversary*, Proc. 23rd IEEE Conference on Computational Complexity, 2008, pp. 237–248, available at [quant-ph/0703237](http://arxiv.org/abs/quant-ph/0703237).
- [19] R. Špalek and M. Szegedy, *All quantum adversary methods are equivalent*, Theory of Computing **2** (2006), no. 1, 1–18, available at [quant-ph/0409116](http://arxiv.org/abs/quant-ph/0409116). Preliminary version in ICALP 2005.