Downloaded from UvA-DARE, the institutional repository of the University of Amsterdam (UvA)
http://hdl.handle.net/11245/2.41078

File ID uvapub:41078
Filename 191182.pdf
Version unknown

SOURCE (OR PART OF THE FOLLOWING SOURCE):
Type conference contribution
Title All Quantum Adversary Methods are Equivalent
Author(s) R. Spalek, M. Szegedy
Faculty FNWI: Institute for Logic, Language and Computation (ILLC)
Year 2005

FULL BIBLIOGRAPHIC DETAILS:
http://hdl.handle.net/11245/1.243149

Copyright

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content licence (like Creative Commons).

UvA-DARE is a service provided by the library of the University of Amsterdam (http://dare.uva.nl)
(pagedate: 2014-11-17)

THEORY OF COMPUTING, Volume 2 (2006), pp. 1–18
http://theoryofcomputing.org

All Quantum Adversary Methods are Equivalent

Robert Špalek Mario Szegedy

Received: August 4, 2005; published: January 31, 2006.

Abstract: The quantum adversary method is one of the most versatile lower-bound methods for quantum algorithms. We show that all known variants of this method are equivalent: spectral adversary (Barnum, Saks, and Szegedy, 2003), weighted adversary (Ambainis, 2003), strong weighted adversary (Zhang, 2005), and the Kolmogorov complexity adversary (Laplante and Magniez, 2004). We also present a few new equivalent formulations of the method. This shows that there is essentially *one* quantum adversary method. From our approach, all known limitations of these versions of the quantum adversary method easily follow.

ACM Classification: F.1.2, F.1.3

AMS Classification: 81P68, 68Q17

Key words and phrases: Quantum computing, query complexity, adversary lower bounds

1 Introduction

1.1 Lower-bound methods for quantum query complexity

In the query complexity model, the input is accessed using oracle queries and the query complexity of the algorithm is the number of calls to the oracle. The query complexity model is helpful in obtaining time complexity lower bounds, and often this is the only way to obtain time bounds in the random access model.

The first lower-bound method on quantum computation was the hybrid method of Bennett, Bernstein, Brassard, and Vazirani [9] to show an $\Omega(\sqrt{n})$ lower bound on the quantum database search. Their proof

Authors retain copyright to their papers and grant "Theory of Computing" unlimited rights to publish the paper electronically and in hard copy. Use of the article is permitted as long as the author(s) and the journal are properly acknowledged. For the detailed copyright statement, see http://theoryofcomputing.org/copyright.html.

© 2006 Robert Špalek, Mario Szegedy

is based on the following simple observation: If the value of function f differs on two inputs x, y, then the output quantum states of any bounded-error algorithm for f on x and y must be almost orthogonal. On the other hand, the inner product is 1 at the beginning, because the computation starts in a fixed state. By upper-bounding the change of the inner product after one query, we lower bound the number of queries that need to be made.

The second lower-bound method is the polynomial method of Beals, Buhrman, Cleve, Mosca, and de Wolf [8]. It is based on the observation that the measurement probabilities can be described by low-degree polynomials in the input bits. If t queries have been made, then the degree is at most $2t$. Since the measurement probabilities are always inside $[0,1]$, one can apply degree lower bounds for polynomials to obtain good lower bounds for quantum query complexity.

The third lower-bound method is the quantum adversary method of Ambainis [2]. It extends the hybrid method. Instead of examining a fixed input pair, Ambainis takes an average over many pairs of inputs. In this paper, we study different variants of the quantum adversary method.

The fourth lower-bound method is the semidefinite programming method of Barnum, Saks, and Szegedy [7]. It exactly characterizes quantum query complexity by a semidefinite program. The dual of this program gives a lower bound that encompasses the quantum adversary bound.

1.2 The variants of the quantum adversary method

The original version of the quantum adversary method, let us call it *unweighted*, was invented by Ambainis [2]. It was successfully used to obtain the following tight lower bounds: $\Omega(\sqrt{n})$ for Grover search [12], $\Omega(\sqrt{n})$ for two-level And-Or trees (see [13] for a matching upper bound), and $\Omega(\sqrt{n})$ for inverting a permutation. The method starts with choosing a set of pairs of inputs on which f takes different values. Then the lower bound is determined by some combinatorial properties of the graph of all pairs chosen.

Some functions, such as sorting or ordered search, could not be satisfactorily lower-bounded by the unweighted adversary method. Høyer, Neerbek, and Shi used a novel argument [14] to obtain tight bounds for these problems. They weighted the input pairs and obtained the lower bound by evaluating the spectral norm of the Hilbert matrix. Barnum, Saks, and Szegedy proposed a general method [7] that gives necessary and sufficient conditions for the existence of a quantum query algorithm. They also described a special case, the so-called *spectral method*, which gives a lower bound in terms of spectral norms of an adversary matrix. Ambainis also published a *weighted* version of his adversary method [3]. He showed that it is stronger than the unweighted method and successfully applied it to get a lower bound for several iterated functions. This method is slightly harder to apply, because it requires one to design a so-called *weight scheme*, which can be seen as a quantum counterpart of the classical *hard distribution* on the inputs. Zhang observed that Ambainis had generalized his oldest method [2] in two independent ways, so he unified them, and published a *strong weighted adversary method* [27]. Finally, Laplante and Magniez used Kolmogorov complexity in an unusual way and described a *Kolmogorov complexity method* [17].

All adversary lower-bound methods above except the Kolmogorov complexity method were defined and proved only for Boolean functions, that is for functions with Boolean input bits and a Boolean output.

A few relations between the methods are known. It is a trivial fact that the strong weighted adversary is at least as good as the weighted adversary. Laplante and Magniez showed [17] that the Kolmogorov complexity method is at least as strong as all the following methods: the Ambainis unweighted and weighted method, the strong weighted method, and the spectral method. The method of Høyer et al. [14] is a special case of the weighted adversary method. It seemed that there were several incompatible variants of the quantum adversary method of different strength.

In addition it was known that there were some limitations for lower bounds obtained by the adversary method. Let f be Boolean. Szegedy observed [26] that the weighted adversary method is limited by $\min(\sqrt{C_0 n}, \sqrt{C_1 n})$, where C_0 is the zero-certificate complexity of f and C_1 is the one-certificate complexity of f. Laplante and Magniez proved the same limitation for the Kolmogorov complexity method [17], which implies that all other methods are also bounded. Finally, this bound was improved to $\sqrt{C_0 C_1}$ for total f by Zhang [27] and independently by us.

1.3 Our results

In this paper, we clean up the forest of adversary methods. First, we extend all adversary lower bound methods to general non-Boolean functions. Second, we show that there is essentially only one quantum adversary method and that all the former methods [7, 3, 27, 17] are just different formulations of the same method. Since one method can be defined in several seemingly unrelated ways and yet one always obtains the same bound, it implies that the quantum adversary method is a very robust concept.

Third, we present a new simple proof of the limitation of the quantum adversary method. If we order the letters in the output alphabet by their certificate complexities such that $C_0 \geq C_1 \geq \ldots$, then all adversary lower bounds are at most $2\sqrt{C_1 n}$ for partial f and $\sqrt{C_0 C_1}$ for total f.

1.4 Separation between the polynomial and adversary method

The polynomial method and the adversary method are generally incomparable. There are examples when the polynomial method gives better bounds and vice versa.

The polynomial method has been successfully applied to obtain tight lower bounds for the following problems: $\Omega\left(n^{1/3}\right)$ for the collision problem and $\Omega\left(n^{2/3}\right)$ for the element distinctness problem [1] (see [4] for a matching upper bound). The quantum adversary method is incapable of proving such lower bounds due to the small certificate complexity of the function. Furthermore, the polynomial method often gives tight lower bounds for the exact and zero-error quantum complexity, such as n for the Or function [8]. The adversary method completely fails in this setting and the only lower bound it can offer is the bounded-error lower bound.

On the other hand, Ambainis exhibited some iterated functions [3] for which the adversary method gives better lower bounds than the polynomial method. The largest established gap between the two methods is $n^{1.321}$. Furthermore, it is unknown how to apply the polynomial method to obtain several lower bounds that are very simple to prove by the adversary method. A famous example is the two-level And-Or tree. The adversary method gives a tight lower bound $\Omega(\sqrt{n})$ [2], whereas the best bound obtained by the polynomial method is $\Omega\left(n^{1/3}\right)$ and it follows [5] from the element distinctness lower bound [1].

There are functions for which none of the methods is known to give a tight bound. A long-standing open problem is the binary And-Or tree. The best known quantum algorithm is just an implementation of the classical zero-error algorithm by Snir [25] running in expected time $\mathrm{O}\big(n^{0.753}\big)$, which is optimal for both zero-error [23] and bounded-error [24] algorithms. The adversary lower bounds are limited by $\sqrt{C_0 C_1} = \sqrt{n}$. In a recent development, Laplante, Lee, and Szegedy showed [16] that this limitation $\sqrt{n}$ holds for every read-once $\{\wedge, \vee\}$ formula. The best known lower bound obtained by the polynomial method is also $\Omega(\sqrt{n})$ and it follows from embedding the parity function. It could be that the polynomial method can prove a stronger lower bound. Two other examples are triangle finding and verification of matrix products. For triangle finding, the best upper bound is $\mathrm{O}\big(n^{1.3}\big)$ [20] and the best lower bound is $\Omega(n)$. For verification of matrix products, the best upper bound is $\mathrm{O}\big(n^{5/3}\big)$ [10] and the best lower bound is $\Omega\big(n^{3/2}\big)$. Again, the adversary method cannot give better bounds, but the polynomial method might.

The semidefinite programming method [7] gives an exact characterization of quantum query complexity. However, it is too general to be applied directly. It is an interesting open problem to find a lower bound that cannot be proved by the adversary or polynomial method.

2 Preliminaries

2.1 Quantum query algorithms

We assume familiarity with quantum computing [22] and sketch the model of quantum query complexity, referring to [11] for more details, also on the relation between query complexity and certificate complexity. Suppose we want to compute some function $f : S \to H$, where $S \subseteq G^N$ and G, H are some finite alphabets. For input $x \in S$, a *query* gives us access to the input variables. It corresponds to the unitary transformation, which depends on input x in the following way:

$$O_x : |i, b, z\rangle \mapsto |i, (b + x_i) \bmod |G|, z\rangle \quad .$$

Here $i \in [N] = \{1, \ldots, N\}$ and $b \in G$; the z-part corresponds to the workspace, which is not affected by the query. We assume the input can be accessed only via such queries. A T-query quantum algorithm has the form $A = U_T O_x U_{T-1} \cdots O_x U_1 O_x U_0$, where the U_k are fixed unitary transformations, independent of x. This A depends on x via the T applications of O_x. The algorithm starts in initial S-qubit state $|0\rangle$. The output of A is obtained by observing the first few qubits of the final superposition $A|0\rangle$, and its success probability on input x is the probability of outputting $f(x)$.

2.2 Kolmogorov complexity

An excellent book about Kolmogorov complexity is the book [18] by Li and Vitányi. Deep knowledge of Kolmogorov complexity is not necessary to understand this paper. Some results on the relation between various classical forms of the quantum adversary method and the Kolmogorov complexity method are taken from Laplante and Magniez [17], and the others just use basic techniques.

A set is called *prefix-free* if none of its members is a prefix of another member. Fix a universal Turing machine M and a prefix-free set S. The *prefix-free Kolmogorov complexity* of x given y, denoted

by $K(x|y)$, is the length of the shortest program from S that prints x if it gets y on the input. Formally,

$$K(x|y) = \min\{|P| : P \in S, M(P,y) = x\} \quad .$$

2.3 Semidefinite programming

In this paper, we use the duality theory of semidefinite programming [19]. There are various forms of the duality principle in the literature. We use a semidefinite extension of Farkas's lemma [19, Theorem 3.4].

2.4 Notation

Let $[n] = \{1,2,\ldots,n\}$. Let Σ^* denote the set of all finite strings over alphabet Σ. All logarithms are binary. Let I denote the *identity* matrix. Let A^T denote the *transpose* of A. Let $\operatorname{diag}(A)$ denote the column vector containing the *main diagonal* of A. Let $\operatorname{tr}(A)$ be the *trace* of A and let $A \cdot B$ be the scalar product of A and B, formally $A \cdot B = \sum_{x,y} A[x,y]B[x,y]$. For a column vector x, let $|x|$ denote the ℓ_2-norm of x, formally $|x| = \sqrt{x^T x}$. Let $\lambda(A)$ denote the *spectral norm* of A, formally $\lambda(A) = \max_{x:|x|\neq 0} |Ax|/|x|$. Let AB denote the usual *matrix product* and let $A \circ B$ denote the *Hadamard (point-wise) product* [21]. Formally, $(AB)[x,y] = \sum_i A[x,i]B[i,y]$ and $(A \circ B)[x,y] = A[x,y]B[x,y]$. Let $A \geq B$ denote the *point-wise comparison* and let $C \succeq D$ denote that $C - D$ is *positive semidefinite*. Formally, $\forall x,y : A[x,y] \geq B[x,y]$ and $\forall v : v^T(C-D)v \geq 0$. Let $r_x(M)$ denote the ℓ_2-*norm* of the x-th row of M and let $c_y(M)$ denote the ℓ_2-*norm* of the y-th column of M. Formally,

$$r_x(M) = \sqrt{\sum_y M[x,y]^2} \qquad \text{and} \qquad c_y(M) = \sqrt{\sum_x M[x,y]^2} \quad .$$

Let $r(M) = \max_x r_x(M)$ and $c(M) = \max_y c_y(M)$.

Let $S \subseteq G^n$ be a set of inputs. We say that a function $f : S \to H$ is *total* if $S = G^n$. A general function is called *partial*. Let f be a partial function. A *certificate* for an input $x \in S$ is a subset $I \subseteq [n]$ such that fixing the input variables $i \in I$ to x_i determines the function value. Formally,

$$\forall y \in S : y|_I = x|_I \Rightarrow f(y) = f(x) \quad ,$$

where $x|_I$ denotes the substring of x indexed by I. A certificate I for x is called *minimal* if $|I| \leq |J|$ for every certificate J for x. Let $\mathcal{C}_f(x)$ denote the lexicographically smallest *minimal certificate* for x. For an $h \in H$, let $C_h(f) = \max_{x:f(x)=h} |\mathcal{C}_f(x)|$ be the *h-certificate complexity* of f.

3 Main result

In this section, we present several equivalent quantum adversary methods and a new simple proof of the limitations of these methods. We can categorize these methods into two groups. Some of them solve conditions on the primal of the quantum system [7]: these are the spectral, weighted, strong weighted, and generalized spectral adversary; and some of them solve conditions on the dual: these are the Kolmogorov complexity bound, minimax, and the semidefinite version of minimax. Primal methods

are mostly used to give lower bounds on the query complexity, while we can use the duals to give limitations of the method.

The primal methods, that is the spectral, weighted, and strong weighted adversary, have been stated only for Boolean functions. The generalization to the more general non-Boolean case is straightforward and hence we state them here in the generalized form.

Theorem 3.1. *Let $S \subseteq G^n$ and let $f : S \to H$ be a partial function. Let $Q_\varepsilon(f)$ denote the ε-error quantum query complexity of f. Then*

$$\frac{Q_\varepsilon(f)}{1-2\sqrt{\varepsilon(1-\varepsilon)}} \geq \mathrm{SA}(f) = \mathrm{WA}(f) = \mathrm{SWA}(f) = \mathrm{MM}(f) = \mathrm{SMM}(f) = \mathrm{GSA}(f) = \Theta(\mathrm{KA}(f)) \quad ,$$

where SA, WA, SWA, MM, SMM, GSA, *and* KA *are lower bounds given by the following methods.*

- **Spectral adversary [7].** *Let D_i, F be $|S| \times |S|$ zero-one valued matrices that satisfy $D_i[x,y] = 1$ iff $x_i \neq y_i$ for $i \in [n]$, and $F[x,y] = 1$ iff $f(x) \neq f(y)$. Let Γ denote an $|S| \times |S|$ non-negative symmetric matrix such that $\Gamma \circ F = \Gamma$. Then*

$$\mathrm{SA}(f) = \max_{\Gamma} \frac{\lambda(\Gamma)}{\max_i \lambda(\Gamma \circ D_i)} \quad . \tag{3.1}$$

- **Weighted adversary [3].**[1] *Let w, w' denote a weight scheme as follows:*
  - *Every pair $(x,y) \in S^2$ is assigned a non-negative weight $w(x,y) = w(y,x)$ that satisfies $w(x,y) = 0$ whenever $f(x) = f(y)$.*
  - *Every triple $(x,y,i) \in S^2 \times [n]$ is assigned a non-negative weight $w'(x,y,i)$ that satisfies $w'(x,y,i) = 0$ whenever $x_i = y_i$ or $f(x) = f(y)$, and $w'(x,y,i)w'(y,x,i) \geq w^2(x,y)$ for all x,y,i such that $x_i \neq y_i$ and $f(x) \neq f(y)$.*

  For all x,i, let $wt(x) = \sum_y w(x,y)$ and $v(x,i) = \sum_y w'(x,y,i)$. Then

$$\mathrm{WA}(f) = \max_{w,w'} \min_{\substack{x,y,\ i,j \\ f(x)\neq f(y) \\ v(x,i)v(y,j)>0}} \sqrt{\frac{wt(x)wt(y)}{v(x,i)v(y,j)}} \quad . \tag{3.2}$$

- **Strong weighted adversary [27].** *Let w, w' denote a weight scheme as above. Then*

$$\mathrm{SWA}(f) = \max_{w,w'} \min_{\substack{x,y,i \\ w(x,y)>0 \\ x_i\neq y_i}} \sqrt{\frac{wt(x)wt(y)}{v(x,i)v(y,i)}} \quad . \tag{3.3}$$

[1] We use a different formulation [17] than in the original Ambainis papers [2, 3]. In particular, we omit the relation $R \subseteq A \times B$ on which the weights are required to be nonzero, and instead allow zero weights. It is simple to prove that both formulations are equivalent.

- **Kolmogorov complexity [17].**[2] *Let $\sigma \in \{0,1\}^*$ denote a finite string. Then*

$$\mathrm{KA}(f) = \min_{\sigma} \max_{\substack{x,y \\ f(x)\neq f(y)}} \frac{1}{\sum_{i:x_i\neq y_i} \sqrt{2^{-K(i|x,\sigma)-K(i|y,\sigma)}}} \quad . \tag{3.4}$$

- **Minimax over probability distributions [17].** *Let $p : S \times [n] \to \mathbb{R}$ denote a set of probability distributions, that is $p_x(i) \geq 0$ and $\sum_i p_x(i) = 1$ for every x. Then*

$$\mathrm{MM}(f) = \min_{p} \max_{\substack{x,y \\ f(x)\neq f(y)}} \frac{1}{\sum_{i:x_i\neq y_i} \sqrt{p_x(i)\, p_y(i)}} \tag{3.5}$$

$$= 1 \Big/ \max_{p} \min_{\substack{x,y \\ f(x)\neq f(y)}} \sum_{i:x_i\neq y_i} \sqrt{p_x(i)\, p_y(i)} \quad . \tag{3.6}$$

- **Semidefinite version of minimax.** *Let D_i, F be matrices as above. Then $\mathrm{SMM}(f) = 1/\mu_{\max}$, where $\mu_{\max}$ is the maximal solution of the following semidefinite program:*

$$\begin{aligned} &\text{maximize } \mu \\ &\text{subject to} \quad \forall i: \quad R_i \succeq 0 \\ &\qquad\qquad\qquad \textstyle\sum_i R_i \circ I = I \\ &\qquad\qquad\qquad \textstyle\sum_i R_i \circ D_i \geq \mu F \quad . \end{aligned} \tag{3.7}$$

- **Generalized spectral adversary.** *Let D_i, F be matrices as above. Then $\mathrm{GSA}(f) = 1/\mu_{\min}$, where $\mu_{\min}$ is the minimal solution of the following semidefinite program:*

$$\begin{aligned} &\text{minimize } \mu = \mathrm{tr}\,\Delta \\ &\text{subject to} \qquad \Delta \text{ is diagonal} \\ &\qquad\qquad\qquad Z \geq 0 \\ &\qquad\qquad\qquad Z \cdot F = 1 \\ &\qquad\qquad \forall i: \ \Delta - Z \circ D_i \succeq 0 \quad . \end{aligned} \tag{3.8}$$

Before we prove the main theorem in the next sections, let us draw some consequences. We show that there are limits that none of these quantum adversary methods can go beyond.

Theorem 3.2. *Let $S \subseteq G^n$ and let $f : S \to H$ be a partial function. Let the output alphabet be $H = \{0,1,\ldots,|H|-1\}$ and order the letters $h \in H$ by their h-certificate complexities such that $C_0 \geq C_1 \geq \cdots \geq C_{|H|-1}$. Then the max-min bound (3.6) is upper-bounded by $\mathrm{MM}(f) \leq 2\sqrt{C_1(f)\cdot n}$. If f is total, that is if $S = G^n$, then $\mathrm{MM}(f) \leq \sqrt{C_0(f)\cdot C_1(f)}$.*

[2] We use a different formulation than Laplante and Magniez [17]. They minimize over all algorithms A computing f and substitute $\sigma =$ source code of A, whereas we minimize over all finite strings σ. Our way is equivalent. One can easily argue that any finite string σ can be "embedded" into any algorithm B: Let C be the source code of B with appended comment σ that is never executed. Now, the programs B and C are equivalent, and $K(x|\sigma) \leq K(x|C) + \mathrm{O}(1)$ for every x.

Proof. The following simple argument is due to Ronald de Wolf. We exhibit two sets of probability distributions p such that

$$m(p) = \min_{\substack{x,y \\ f(x)\neq f(y)}} \sum_{i:x_i\neq y_i} \sqrt{p_x(i)\,p_y(i)} \geq \frac{1}{2\sqrt{C_1 n}} \ , \text{ resp. } \frac{1}{\sqrt{C_0 C_1}} \ .$$

The max-min bound (3.6) is $\mathrm{MM}(f) = 1/\max_p m(p)$ and the statement follows.

Let f be partial. For every $x \in S$, distribute one half of the probability uniformly over any minimal certificate $\mathcal{C}_f(x)$, and one half of the probability uniformly over all input variables. Formally,

$$p_x(i) = \frac{1}{2n} + \frac{1}{2|\mathcal{C}_f(x)|} \text{ iff } i \in \mathcal{C}_f(x) \ , \qquad \text{and} \qquad p_x(i) = \frac{1}{2n} \text{ for } i \notin \mathcal{C}_f(x) \ .$$

Take any x,y such that $f(x) \neq f(y)$. Assume that $C_x \leq C_y$, and take the $f(x)$-certificate $I = \mathcal{C}_f(x)$. Since $y|_I \neq x|_I$, there is a $j \in I$ such that $x_j \neq y_j$. Now we lower-bound the sum of (3.6).

$$\sum_{i:x_i\neq y_i} \sqrt{p_x(i)\,p_y(i)} \geq \sqrt{p_x(j)\,p_y(j)} \geq \sqrt{\frac{1}{|2\mathcal{C}_f(x)|}\cdot\frac{1}{2n}} \geq \frac{1}{2\sqrt{C_{f(x)}n}} \geq \frac{1}{2\sqrt{C_1 n}} \ .$$

Since this inequality holds for any x,y such that $f(x) \neq f(y)$, also $m(p) \geq 1/2\sqrt{C_1 n}$. Take the reciprocal and conclude that $\mathrm{MM}(f) \leq 2\sqrt{C_1 n}$.

For Boolean output alphabet $H = \{0,1\}$, we can prove a slightly stronger bound $\mathrm{MM}(f) \leq \sqrt{C_1 n}$ as follows. Define p as a uniform distribution over some minimal certificate for all one-inputs, and a uniform distribution over all input bits for all zero-inputs. The same computation as above gives the bound.

If f is total, then we can do even better. For every $x \in G^n$, distribute the probability uniformly over any minimal certificate $\mathcal{C}_f(x)$. Formally, $p_x(i) = 1/|\mathcal{C}_f(x)|$ iff $i \in \mathcal{C}_f(x)$, and $p_x(i) = 0$ otherwise. Take any x,y such that $f(x) \neq f(y)$, and let $I = \mathcal{C}_f(x) \cap \mathcal{C}_f(y)$. There must exist a $j \in I$ such that $x_j \neq y_j$, otherwise we could find an input z that is consistent with both certificates. (That would be a contradiction, because f is total and hence $f(z)$ has to be defined and be equal to both $f(x)$ and $f(y)$.) After we have found a j, we lower-bound the sum of (3.6) by $1/\sqrt{C_{f(x)}C_{f(y)}}$ in the same way as above. Since $\sqrt{C_{f(x)}C_{f(y)}} \leq \sqrt{C_0 C_1}$, the bound follows. □

Some parts of the following statement have been observed for individual methods by Szegedy [26], Laplante and Magniez [17], and Zhang [27]. This corollary rules out all adversary attempts to prove good lower bounds for problems with small certificate complexity, such as element distinctness [1], binary And-Or trees [6, 13], triangle finding [20], or verification of matrix products [10].

Corllary 3.3. *All quantum adversary lower bounds are at most* $\min(\sqrt{C_0(f)n}, \sqrt{C_1(f)n})$ *for partial Boolean functions and* $\sqrt{C_0(f)C_1(f)}$ *for total Boolean functions.*

4 Equivalence of spectral and strong weighted adversary

In this section, we give a linear-algebraic proof that the spectral bound [7] and the strong weighted bound [27] are equal. The proof has three steps. First, we show that the weighted bound [3] is at least as

good as the spectral bound. Second, using a small combinatorial lemma, we show that the spectral bound is at least as good as the strong weighted bound. The strong weighted bound is always at least as good as the weighted bound, because every term in the minimization of (3.3) is included in the minimization of (3.2): if $w(x,y) > 0$ and $x_i \neq y_i$, then $f(x) \neq f(y)$ and both $w'(x,y,i) > 0$ and $w'(y,x,i) > 0$. The generalization of the weighted adversary method thus does not make the bound stronger, however its formulation is easier to use.

4.1 Reducing spectral adversary to weighted adversary

First, let us state two useful statements upper-bounding the spectral norm of a Hadamard product of two non-negative matrices. The first one is due to Mathias [21]. The second one is our generalization and its proof is postponed to Appendix A.

Lemma 4.1. [21] *Let S be a non-negative symmetric matrix and let M and N be non-negative matrices such that $S \leq M \circ N$. Then*

$$\lambda(S) \leq r(M)c(N) = \max_{x,y} r_x(M)c_y(N) \quad . \tag{4.1}$$

Moreover, for every symmetric $S \geq 0$ there exists an $M \geq 0$ such that $S = M \circ M^T$ and $r(M) = c(M^T) = \sqrt{\lambda(S)}$. This optimal matrix can be written as $M[x,y] = \sqrt{S[x,y] \cdot d[y]/d[x]}$, where d is the principal eigenvector of S.

Lemma 4.2. *Let S be a non-negative symmetric matrix and let M and N be non-negative matrices such that $S \leq M \circ N$. Then*

$$\lambda(S) \leq \max_{\substack{x,y \\ S[x,y]>0}} r_x(M)c_y(N) \quad . \tag{4.2}$$

Now we use the first bound to reduce the spectral adversary to the weighted adversary.

Theorem 4.3. $\mathrm{SA}(f) \leq \mathrm{WA}(f)$.

Proof. Let Γ be a non-negative symmetric matrix with $\Gamma \circ F = \Gamma$ as in equation (3.1) that gives the optimal spectral bound. Assume without loss of generality that $\lambda(\Gamma) = 1$. Let δ be the principal eigenvector of Γ, that is $\Gamma\delta = \delta$. Define the following weight scheme:

$$w(x,y) = w(y,x) = \Gamma[x,y] \cdot \delta[x]\delta[y] \quad .$$

Furthermore, for every i, using Lemma 4.1, decompose $\Gamma_i = \Gamma \circ D_i$ into a Hadamard product of two non-negative matrices $\Gamma_i = M_i \circ M_i^T$ such that $r(M_i) = \sqrt{\lambda(\Gamma_i)}$. Define w' as follows:

$$w'(x,y,i) = M_i[x,y]^2 \delta[x]^2 \quad .$$

Let us verify that w, w' is a weight scheme. From the definition, $w(x,y) = w'(x,y,i) = 0$ if $f(x) = f(y)$, and also $w'(x,y,i) = 0$ if $x_i = y_i$. Furthermore, if $f(x) \neq f(y)$ and $x_i \neq y_i$, then

$$w'(x,y,i)w'(y,x,i) = (M_i[x,y]\,\delta[x])^2 (M_i[y,x]\,\delta[y])^2 = (\Gamma_i[x,y]\,\delta[x]\delta[y])^2 = w(x,y)^2 \quad .$$

Finally, let us compute the bound (3.2) given by the weight scheme.

$$wt(x) = \sum_y w(x,y) = \delta[x] \sum_y \Gamma[x,y]\delta[y] = \delta[x]\,(\Gamma\delta)[x] = \delta[x]^2 \ ,$$
$$\frac{v(x,i)}{wt(x)} = \frac{\sum_y w'(x,y,i)}{wt(x)} = \frac{\sum_y M_i[x,y]^2\delta[x]^2}{\delta[x]^2} = r_x(M_i)^2 \le r(M_i)^2 = \lambda(\Gamma_i) \ .$$

The weighted adversary lower bound (3.2) is thus at least

$$\mathrm{WA}(f) \ge \min_{\substack{x,y,\ i,j \\ f(x)\neq f(y) \\ v(x,i)v(y,j)>0}} \sqrt{\frac{wt(x)wt(y)}{v(x,i)v(y,j)}} \ge \min_{i,j} \frac{1}{\sqrt{\lambda(\Gamma_i)\cdot\lambda(\Gamma_j)}} = \frac{\lambda(\Gamma)}{\max_i \lambda(\Gamma_i)} = \mathrm{SA}(f) \ .$$

Hence the weighted adversary is at least as strong as the spectral adversary (3.1). □

4.2 Reducing strong weighted adversary to spectral adversary

Theorem 4.4. $\mathrm{SWA}(f) \le \mathrm{SA}(f)$.

Proof. Let w, w' be a weight scheme as in Equation (3.2) that gives the optimal weighted bound. Define the following symmetric matrix Γ on $S \times S$:

$$\Gamma[x,y] = \frac{w(x,y)}{\sqrt{wt(x)wt(y)}} \ .$$

We also define column vector δ on S such that $\delta[x] = \sqrt{wt(x)}$. Let $W = \sum_x wt(x)$. Then

$$\lambda(\Gamma) \ge \delta^T\Gamma\delta/|\delta|^2 = W/W = 1 \ .$$

Define the following matrix on the index set $S \times S$:

$$M_i[x,y] = \sqrt{\frac{w'(x,y,i)}{wt(x)}} \ .$$

Every weight scheme satisfies $w'(x,y,i)w'(y,x,i) \ge w^2(x,y)$ for all x,y,i such that $x_i \neq y_i$. Hence

$$M_i[x,y]\cdot M_i[y,x] = \frac{\sqrt{w'(x,y,i)w'(y,x,i)}}{\sqrt{wt(x)wt(y)}} \ge \frac{w(x,y)\cdot D_i[x,y]}{\sqrt{wt(x)wt(y)}} = \Gamma_i[x,y] \ .$$

This means that $\Gamma \le M \circ M^T$. By Lemma 4.2 and using $c_y(M^T) = r_y(M)$,

$$\lambda(\Gamma_i) \le \max_{\substack{x,y \\ \Gamma_i[x,y]>0}} r_x(M)r_y(M) = \max_{\substack{x,y \\ w(x,y)>0 \\ x_i\neq y_i}} \sqrt{\sum_k \frac{w'(x,k,i)}{wt(x)} \sum_\ell \frac{w'(y,\ell,i)}{wt(y)}} = \max_{\substack{x,y \\ w(x,y)>0 \\ x_i\neq y_i}} \sqrt{\frac{v(x,i)v(y,i)}{wt(x)wt(y)}} \ .$$

The spectral adversary lower bound (3.1) is thus at least

$$\mathrm{SA}(f) \geq \frac{\lambda(\Gamma)}{\max_i \lambda(\Gamma_i)} \geq \min_i \min_{\substack{x,y \\ w(x,y)>0 \\ x_i \neq y_i}} \sqrt{\frac{wt(x)wt(y)}{v(x,i)v(y,i)}} = \mathrm{SWA}(f) \ .$$

Hence the spectral adversary is at least as strong as the weighted adversary (3.3). □

Remark 4.5. The strength of the obtained reduction depends on which statement is used for upper-bounding the spectral norm of Γ_i.

- Lemma 4.2 has just given us $\mathrm{SWA}(f) \leq \mathrm{SA}(f)$.
- Lemma 4.1 would give a weaker bound $\mathrm{WA}(f) \leq \mathrm{SA}(f)$.
- Høyer, Neerbek, and Shi used an explicit expression for the norm of the Hilbert matrix to get a lower bound for ordered search [14]. Their method is thus also a special case of the spectral method.
- Both versions of the original unweighted adversary method [2] are obtained by using a spectral matrix Γ corresponding to a zero-one valued weight scheme w, the lower bound $\lambda(\Gamma) \geq d^T \Gamma d / |d|^2$, and Lemma 4.1, resp. Lemma 4.2.

5 Equivalence of minimax and generalized spectral adversary

In this section, we prove that the minimax bound is equal to the generalized spectral bound. We first remove the reciprocal by taking the max-min bound. Second, we write this bound as a semidefinite program. An application of duality theory of semidefinite programming finishes the proof.

Theorem 5.1. $\mathrm{MM}(f) = \mathrm{SMM}(f)$.

Proof. Let p be a set of probability distributions as in Equation (3.6). Define $R_i[x,y] = \sqrt{p_x(i)\, p_y(i)}$. Since p_x is a probability distribution, we get that $\sum_i R_i$ must have all ones on the diagonal. The condition $\min_{\substack{x,y \\ f(x)\neq f(y)}} \sum_{i:x_i\neq y_i} R_i[x,y] \geq \mu$ may be rewritten

$$\forall x,y : f(x) \neq f(y) \Longrightarrow \sum_{i:x_i \neq y_i} R_i[x,y] \geq \mu \ ,$$

which is to say $\sum_i R_i \circ D_i \geq \mu F$. Each matrix R_i should be an outer product of a non-negative vector with itself: $R_i = r_i r_i^T$ for a column vector $r_i[x] = \sqrt{p_x(i)}$. We have, however, replaced that condition by $R_i \succeq 0$ to get semidefinite program (3.7). Since $r_i r_i^T \succeq 0$, the program (3.7) is a relaxation of the condition of (3.6) and $\mathrm{SMM}(f) \leq \mathrm{MM}(f)$.

Let us show that every solution R_i of the semidefinite program can be changed to an at least as good rank-1 solution R_i'. Take a Cholesky decomposition $R_i = X_i X_i^T$. Define a column-vector $q_i[x] = \sqrt{\sum_j X_i[x,j]^2}$ and a rank-1 matrix $R_i' = q_i q_i^T$. It is not hard to show that all R_i' satisfy the same constraints

as R_i. First, R'_i is positive semidefinite. Second, $R'_i[x,x] = \sum_j X_i[x,j]^2 = R_i[x,x]$, hence $\sum_i R'_i \circ I = I$. Third, by a Cauchy-Schwarz inequality,

$$R_i[x,y] = \sum_j X_i[x,j]X_i[y,j] \leq \sqrt{\sum_k X_i[x,k]^2}\sqrt{\sum_\ell X_i[y,\ell]^2} = q_i[x]q_i[y] = R'_i[x,y] \quad ,$$

hence $\sum_i R'_i \circ D_i \geq \sum_i R_i \circ D_i \geq \mu F$. We conclude that $\mathrm{MM}(f) \leq \mathrm{SMM}(f)$. □

The equivalence of the semidefinite version of minimax and the generalized spectral adversary is proved using the duality theory of semidefinite programming. We use a semidefinite version of Farkas's lemma [19, Theorem 3.4].

Theorem 5.2. $\mathrm{SMM}(f) = \mathrm{GSA}(f)$.

Proof. Let us compute the dual of a semidefinite program without converting it to/from the standard form, but using Lagrange multipliers. Take the objective function μ of the semidefinite version of minimax (3.7) and add negative penalty terms for violating the constraints.

$$\begin{aligned} \mu \quad + \quad & \sum_i Y_i \cdot R_i + D \cdot \Big(\sum_i R_i \circ I - I\Big) + Z \cdot \Big(\sum_i R_i \circ D_i - \mu F\Big) = \\ & \text{for } Y_i \succeq 0, \text{ unconstrained } D, \text{ and } Z \geq 0 \\ = \quad & \sum_i R_i \cdot \Big(Y_i + D \circ I + Z \circ D_i\Big) + \mu\Big(1 - Z \cdot F\Big) - D \cdot I \quad . \end{aligned}$$

Its dual system is formed by the constraints on Y_i, D, and Z plus the requirements that both expression in the parentheses are zero. The duality principle [19, Theorem 3.4] says that any primal solution is smaller than or equal to any dual solution. Moreover, if any of the two systems has a strictly feasible solution, then the maximal primal solution equals to the minimal dual solution.

Since $Y_i \succeq 0$ only appears once, we get rid of it by requiring that $D \circ I + Z \circ D_i \preceq 0$. We substitute $\Delta = -D \circ I$ and obtain $\Delta - Z \circ D_i \succeq 0$. The objective function is $-D \cdot I = \mathrm{tr}\,\Delta$. We have obtained the generalized spectral adversary (3.8). Let us prove its strong feasibility. Assume that the function f is not constant, hence $F \neq 0$. Take Z a uniform probability distribution over nonzero entries of F and a large enough constant Δ. This is a strictly feasible solution. We conclude that $\mu_{\max} = \mu_{\min}$. □

6 Equivalence of generalized spectral and spectral adversary

In this section, we prove that the generalized spectral adversary bound is equal to the spectral adversary bound. The main difference between them is that the generalized method uses an arbitrary positive diagonal matrix Δ as a new variable instead of the identity matrix I.

Theorem 6.1. $\mathrm{GSA}(f) = \mathrm{SA}(f)$.

Proof. Let Z, Δ be a solution of (3.8). First, let us prove that $\Delta \succ 0$. Since both $Z \geq 0$ and $D_i \geq 0$, it holds that $\mathrm{diag}\,(-Z \circ D_i) \leq 0$. We know that $\Delta - Z \circ D_i \succeq 0$, hence $\mathrm{diag}\,(\Delta - Z \circ D_i) \geq 0$, and $\mathrm{diag}\,(\Delta) \geq 0$ follows. Moreover, $\mathrm{diag}\,(\Delta) > 0$ unless Z contains an empty row, in which case we delete it (together

with the corresponding column) and continue. Second, since positive semidefinite real matrices are symmetric, $\Delta - Z \circ D_i \succeq 0$ implies that $Z \circ D_i$ is symmetric (for every i). For every $x \neq y$ there is a bit i such that $x_i \neq y_i$, hence Z must be also symmetric.

Take a column vector $a = \operatorname{diag}(\Delta^{-1/2})$ and a rank-1 matrix $A = aa^T \succ 0$. It is simple to prove that $A \circ X \succeq 0$ for every matrix $X \succeq 0$. Since $\Delta - Z \circ D_i \succeq 0$, also $A \circ (\Delta - Z \circ D_i) = I - Z \circ D_i \circ A \succeq 0$ and hence $\lambda(Z \circ D_i \circ A) \leq 1$. Now, define the spectral adversary matrix $\Gamma = Z \circ F \circ A$. Since $0 \leq Z \circ F \leq Z$, it follows that

$$\lambda(\Gamma \circ D_i) = \lambda(Z \circ F \circ A \circ D_i) \leq \lambda(Z \circ D_i \circ A) \leq 1 \ .$$

It remains to show that $\lambda(\Gamma) \geq 1/\operatorname{tr}\Delta$. Let $b = \operatorname{diag}(\sqrt{\Delta})$ and $B = bb^T$. Then

$$1 = Z \cdot F = \Gamma \cdot B = b^T \Gamma b \leq \lambda(\Gamma) \cdot |b|^2 = \lambda(\Gamma) \cdot \operatorname{tr}\Delta \ .$$

It is obvious that Γ is symmetric, $\Gamma \geq 0$, and $\Gamma \circ F = \Gamma$. The bound (3.1) given by Γ is bigger than or equal to $1/\operatorname{tr}\Delta$, hence $\mathrm{SA}(f) \geq \mathrm{GSA}(f)$.

For the other direction, let Γ be a non-negative symmetric matrix satisfying $\Gamma \circ F = \Gamma$. Let δ be its principal eigenvector with $|\delta| = 1$. Assume without loss of generality that $\lambda(\Gamma) = 1$ and let $\mu = \max_i \lambda(\Gamma_i)$. Take $A = \delta\delta^T$, $Z = \Gamma \circ A$, and $\Delta = \mu I \circ A$. Then $Z \cdot F = \Gamma \cdot A = \delta^T \Gamma \delta = 1$ and $\operatorname{tr}\Delta = \mu$. For every i, $\lambda(\Gamma_i) \leq \mu$, hence $\mu I - \Gamma \circ D_i \succeq 0$. It follows that $0 \preceq A \circ (\mu I - \Gamma \circ D_i) = \Delta - Z \circ D_i$. The semidefinite program (3.8) is satisfied and hence its optimum is $\mu_{\min} \leq \mu$. We conclude that $\mathrm{GSA}(f) \geq \mathrm{SA}(f)$. □

7 Proof of the main theorem

In this section, we close the circle of reductions. We use the results of Laplante and Magniez, who recently proved [17] that the Kolmogorov complexity bound is asymptotically lower-bounded by the weighted adversary bound and upper-bounded by the minimax bound. The upper bound is implicit in their paper, because they did not state the minimax bound as a separate theorem.

Theorem 7.1. [17, Theorem 2] $\mathrm{KA}(f) = \Omega(\mathrm{WA}(f))$.

Theorem 7.2. $\mathrm{KA}(f) = \mathrm{O}(\mathrm{MM}(f))$.

Proof. Take a set of probability distributions p as in Equation (3.5). The query information lemma [17, Lemma 3] says that $K(i|x,p) \leq \log \frac{1}{p_x(i)} + \mathrm{O}(1)$ for every x, i such that $p_x(i) > 0$. This is true, because any i of nonzero probability can be encoded in $\lceil \log \frac{1}{p_x(i)} \rceil$ bits using the Shannon-Fano code of distribution p_x, and the Shannon-Fano code is a prefix-free code. Rewrite the inequality as $p_x(i) = \mathrm{O}\big(2^{-K(i|x,p)}\big)$. The statement follows, because the set of all strings σ in (3.4) includes among others also the descriptions of all probability distributions p. □

Remark 7.3. The constant in the equality $\mathrm{KA}(f) = \Theta(\mathrm{WA}(f))$ depends on the choice of the universal Turing machine and the prefix-free set.

Proof of Theorem 3.1. We have to prove that

$$\frac{Q_\varepsilon(f)}{1-2\sqrt{\varepsilon(1-\varepsilon)}} \geq \mathrm{SA}(f) = \mathrm{WA}(f) = \mathrm{SWA}(f) = \mathrm{MM}(f) = \mathrm{SMM}(f) = \mathrm{GSA}(f) = \Theta(\mathrm{KA}(f)) \quad .$$

Put together all known equalities and inequalities.

- $\mathrm{SA}(f) = \mathrm{WA}(f) = \mathrm{SWA}(f)$ by Theorem 4.3 and Theorem 4.4,
- $\mathrm{MM}(f) = \mathrm{SMM}(f)$ by Theorem 5.1,
- $\mathrm{SMM}(f) = \mathrm{GSA}(f)$ by Theorem 5.2,
- $\mathrm{GSA}(f) = \mathrm{SA}(f)$ by Theorem 6.1,
- $\mathrm{KA}(f) = \Theta(\mathrm{WA}(f))$ by Theorem 7.1 and Theorem 7.2.

Finally, one has to prove one of the lower bounds. For example, Ambainis proved [3] that $Q_2(f) \geq (1-2\sqrt{\varepsilon(1-\varepsilon)})\,\mathrm{WA}(f)$ for every Boolean f. Laplante and Magniez proved [17] that $Q_2(f) = \Omega(\mathrm{KA}(f))$ for general f. Høyer and Špalek present in their survey [15] an alternative proof of the spectral adversary bound that can easily be adapted to the non-Boolean case. □

A Proof of the upper bound on the spectral norm

Proof of Lemma 4.2. Let $S = M \circ N$. Define a shortcut

$$B(M,N) = \max_{\substack{x,y \\ S[x,y]>0}} r_x(M)c_y(N) \quad .$$

Without loss of generality, we assume that $M[x,y] = 0 \Leftrightarrow N[x,y] = 0 \Leftrightarrow S[x,y] = 0$. Let us prove the existence of matrices M', N' with $B(M',N') = r(M')c(N')$ such that

$$M \circ N = M' \circ N', \text{ and } B(M,N) = B(M',N') \quad . \tag{A.1}$$

We then apply Lemma 4.1 and obtain

$$\lambda(S) \leq \lambda(M \circ N) = \lambda(M' \circ N') \leq r(M')c(N') = B(M',N') = B(M,N) \quad .$$

Take as M', N' any pair of matrices that satisfies (A.1) and the following constraints:

- $b = r(M')c(N')$ is minimal, that is there is no pair M'', N'' giving a smaller b,
- and, among those, the set R of maximum-norm rows of M' and the set C of maximum-norm columns of N' are both minimal (in the same sense).

Let (r,c) be any "maximal" entry, that is $S[r,c] > 0$ and $r_r(M')c_c(N') = B(M',N')$. Let $\overline{R}$ denote the *complement* of R and let $S[R,C]$ denote the *sub-matrix* of S indexed by $R \times C$. Then one of the following cases happens:

1. $(r,c) \in R \times C$: Then $B(M',N') = r(M')c(N')$ and we are done. If this is not the case, then we know that $S[R,C] = 0$.

2. $(r,c) \in R \times \overline{C}$: Then $S[\overline{R},C] = 0$, otherwise we get a contradiction with one of the minimality assumptions. If $S[x,y] \neq 0$ for some $(x,y) \in \overline{R} \times C$, multiply $M'[x,y]$ by $1+\varepsilon$ and divide $N'[x,y]$ by $1+\varepsilon$ for some small $\varepsilon > 0$ such that the norm of the x-th row of M' is still smaller than $r(M')$. Now, we have either deleted the y-th column from C or, if $|C| = 1$, decreased $c(N')$. Both cases are a contradiction. Finally, if $S[\overline{R},C] = 0$, then $c(N') = 0$ due to $S[R,C] = 0$ and the fact that C are the maximum-norm columns. Hence S is a zero matrix, and we are done.

3. $(r,c) \in \overline{R} \times C$: This case is similar to the previous case.

4. $(r,c) \in \overline{R} \times \overline{C}$: First, note that $S[R,c] = 0$, otherwise (r,c) would not be "maximal". Now we divide all entries in $M'[R,\overline{C}]$ by $1+\varepsilon$ and multiply all entries in $N'[R,\overline{C}]$ by $1+\varepsilon$ for some small $\varepsilon > 0$ such that the "maximal" entries are unchanged. Since $S[R,C] = 0$, it follows that either $S[R,\overline{C}] = 0$ and S is a zero matrix, or there is a nonzero number in every row of $M'[R,\overline{C}]$. Therefore, unless S is a zero matrix, we have preserved $B(M',N')$ and $c(N')$, and decreased $r(M')$, which is a contradiction.

We conclude that $(r,c) \in R \times C$, $B(M',N') = r(M')c(N')$, and hence $\lambda(S) \leq B(M,N)$. □

Acknowledgments

We thank Ronald de Wolf for many fruitful discussions, for his suggestions concerning Theorem 3.2, and for proofreading, and Troy Lee for discussions. We thank anonymous referees for their helpful comments.

Robert Špalek is supported in part by the EU fifth framework project RESQ, IST-2001-37559. Mario Szegedy is supported by NSF grant 0105692, and in part by the National Security Agency (NSA) and Advanced Research and Development Activity (ARDA) under Army Research Office (ARO), contract number DAAD19-01-1-0506.

References

[1] * S. AARONSON AND Y. SHI: Quantum lower bounds for the collision and the element distinctness problem. *Journal of the ACM*, 51(4):595–605, 2004. [JACM:1008731.1008735, arXiv:quant-ph/0111102]. 1.4, 3

[2] * A. AMBAINIS: Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64(4):750–767, 2002. Earlier version in STOC '00. [JCSS:10.1006/jcss.2002.1826, STOC:335305.335394, arXiv:quant-ph/0002066]. 1.1, 1.2, 1.4, 1, 4.5

[3] * A. AMBAINIS: Polynomial degree vs. quantum query complexity. In *Proc. of 44th IEEE FOCS*, pp. 230–239, 2003. [FOCS:10.1109/SFCS.2003.1238197, arXiv:quant-ph/0305028]. 1.2, 1.3, 1.4, 3.1, 1, 4, 7

[4] * A. AMBAINIS: Quantum walk algorithm for element distinctness. In *Proc. of 45th IEEE FOCS*, pp. 22–31, 2004. [arXiv:quant-ph/0311001]. 1.4

[5] * A. AMBAINIS: Polynomial degree and lower bounds in quantum complexity: Collision and element distinctness with small range. *Theory of Computing*, 1:37–46, 2005. [ToC:v001/a003, arXiv:quant-ph/0305179]. 1.4

[6] * H. BARNUM AND M. SAKS: A lower bound on the quantum query complexity of read-once functions. *Journal of Computer and System Sciences*, 69(2):244–258, 2004. [JCSS:10.1016/j.jcss.2004.02.002, arXiv:quant-ph/0201007]. 3

[7] * H. BARNUM, M. SAKS, AND M. SZEGEDY: Quantum decision trees and semidefinite programming. In *Proc. of 18th IEEE Complexity*, pp. 179–193, 2003. [CCC:10.1109/CCC.2003.1214419]. 1.1, 1.2, 1.3, 1.4, 3, 3.1, 4

[8] * R. BEALS, H. BUHRMAN, R. CLEVE, M. MOSCA, AND R. DE WOLF: Quantum lower bounds by polynomials. *Journal of the ACM*, 48(4):778–797, 2001. Earlier version in FOCS '98. [JACM:502090.502097, FOCS:10.1109/SFCS.1998.743485, arXiv:quant-ph/9802049]. 1.1, 1.4

[9] * H. BENNETT, E. BERNSTEIN, G. BRASSARD, AND U. VAZIRANI: Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26(5):1510–1523, 1997. [SICOMP:30093, arXiv:quant-ph/9701001]. 1.1

[10] * H. BUHRMAN AND R. ŠPALEK: Quantum verification of matrix products. In *Proc. of 17th ACM-SIAM SODA*, pp. 880–889, 2006. [SODA:1109557.1109654, arXiv:quant-ph/0409035]. 1.4, 3

[11] * H. BUHRMAN AND R. DE WOLF: Complexity measures and decision tree complexity: A survey. *Theoretical Computer Science*, 288(1):21–43, 2002. [TCS:10.1016/S0304-3975(01)00144-X]. 2.1

[12] * L. K. GROVER: A fast quantum mechanical algorithm for database search. In *Proc. of 28th ACM STOC*, pp. 212–219, 1996. [STOC:237814.237866, arXiv:quant-ph/9605043]. 1.2

[13] * P. HØYER, M. MOSCA, AND R. DE WOLF: Quantum search on bounded-error inputs. In *Proc. of 30th ICALP*, pp. 291–299, 2003. LNCS 2719. [ICALP:214dhep41d6vk3d2, arXiv:quant-ph/0304052]. 1.2, 3

[14] * P. HØYER, J. NEERBEK, AND Y. SHI: Quantum complexities of ordered searching, sorting, and element distinctness. *Algorithmica*, 34(4):429–448, 2002. Special issue on Quantum Computation and Cryptography. [Algorithmica:25gl9elr5rxr3q6a, arXiv:quant-ph/0102078]. 1.2, 4.5

[15] * P. HØYER AND R. ŠPALEK: Lower bounds on quantum query complexity. *EATCS Bulletin*, 87:78–103, October, 2005. [arXiv:quant-ph/0509153]. 7

[16] * S. LAPLANTE, T. LEE, AND M. SZEGEDY: The quantum adversary method and formula size lower bounds. In *Proc. of 20th IEEE Complexity*, pp. 76–90, 2005. [CCC:10.1109/CCC.2005.29, arXiv:quant-ph/0501057]. 1.4

[17] * S. LAPLANTE AND F. MAGNIEZ: Lower bounds for randomized and quantum query complexity using Kolmogorov arguments. In *Proc. of 19th IEEE Complexity*, pp. 294–304, 2004. [CCC:10.1109/CCC.2004.1313852, arXiv:quant-ph/0311189]. 1.2, 1.3, 2.2, 1, 3.1, 3.1, 2, 3, 7, 7.1, 7, 7

[18] * M. LI AND P. M. B. VITÁNYI: *An Introduction to Kolmogorov Complexity and its Applications*. Springer, Berlin, second edition, 1997. 2.2

[19] * L. LOVÁSZ: Semidefinite programs and combinatorial optimization. `http://research.microsoft.com/users/lovasz/semidef.ps`, 2000. 2.3, 5, 5

[20] * F. MAGNIEZ, M. SANTHA, AND M. SZEGEDY: Quantum algorithms for the triangle problem. In *Proc. of 16th ACM-SIAM SODA*, pp. 1109–1117, 2005. [SODA:1070432.1070591, arXiv:quant-ph/0310134]. 1.4, 3

[21] * R. MATHIAS: The spectral norm of a nonnegative matrix. *Linear Algebra and its Applications*, 139:269–284, 1990. [10.10160024-3795(90)90403-Y]. 2.4, 4.1, 4.1

[22] * M. A. NIELSEN AND I. L. CHUANG: *Quantum Computation and Quantum Information*. Cambridge University Press, 2000. 2.1

[23] * M. SAKS AND A. WIGDERSON: Probabilistic Boolean decision trees and the complexity of evaluating games trees. In *Proc. of 27th IEEE FOCS*, pp. 29–38, 1986. 1.4

[24] * M. SANTHA: On the Monte Carlo decision tree complexity of read-once formulae. *Random Structures and Algorithms*, 6(1):75–87, 1995. 1.4

[25] * M. SNIR: Lower bounds on probabilistic decision trees. *Theoretical Computer Science*, 38:69–82, 1985. [TCS:10.1016/0304-3975(85)90210-5]. 1.4

[26] * M. SZEGEDY: On the quantum query complexity of detecting triangles in graphs. quant-ph/0310107, 2003. [arXiv:quant-ph/0310107]. 1.2, 3

[27] * S. ZHANG: On the power of Ambainis's lower bounds. *Theoretical Computer Science*, 339(2–3):241–256, 2005. Earlier version in ICALP'04. [TCS:10.1016/j.tcs.2005.01.019, ICALP:gm2ff6wpc0q39v3x, arXiv:quant-ph/0311060]. 1.2, 1.3, 3.1, 3, 4

AUTHORS

Robert Špalek
graduate student
Centrum voor Wiskunde en Informatica
Amsterdam, The Netherlands
sr@cwi.nl
http://www.ucw.cz/~robert/

Mario Szegedy
professor
Rutgers, the State University of New Jersey
Piscataway, New Jersey, USA
szegedy@cs.rutgers.edu
http://athos.rutgers.edu/~szegedy/

ABOUT THE AUTHORS

ROBERT ŠPALEK received his Masters Degrees in Computer Science from Charles University, Prague and Vrije Universiteit, Amsterdam. He is currently a graduate student at CWI, advised by Harry Buhrman. His research interests include quantum computing, computational complexity, algorithms, data structures, and search engines. He loves dancing salsa, climbing, photography, travelling to distant countries, and playing guitar.

MARIO SZEGEDY received his Ph. D. in computer science at the University of Chicago under the supervision of Laci Babai and Janos Simon. He held a Lady Davis Postdoctoral Fellowship at the Hebrew University, Jerusalem (1989-90), a postdoc at the University of Chicago, 1991-92, and a postdoc at Bell Laboratories (1992). He was a permanent member of Bell Labs for 7 years and for two more years of AT&T Research. He left AT&T in September 1999 to conduct research at the Institute for Advanced Study in Princeton for a year. In 2000 he joined the faculty of Rutgers University.

He received the Gödel Prize twice, in 2001 for his part in the PCP Theorem and its connection to inapproximability and in 2005 for the analysis of data streams using limited memory.

His research interests include complexity theory, combinatorics, combinatorial geometry and quantum computing, but he also has an interest in algebra and in programming languages.

With a group of students he has founded QCteam, a quantum computing laboratory at Rutgers, which is his main project at the present time. The laboratory has received substantial funding from the university and from the National Science Foundation. It has a vigorous visitor program, and pursues collaboration with the local industry.