All Quantum Adversary Methods Are Equivalent

Robert Špalek[1,⋆] and Mario Szegedy[2,⋆⋆]

[1] CWI, Amsterdam
sr@cwi.nl
[2] Rutgers University
szegedy@cs.rutgers.edu

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 equal: spectral adversary [1], weighted adversary [2], strong weighted adversary [3], and the Kolmogorov complexity adversary [4]. 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 all versions of the quantum adversary method easily follow.

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 was the hybrid method of Bennett, Bernstein, Brassard, and Vazirani [5] to show an $\Omega(\sqrt{n})$ lower bound on the quantum database search. Their proof 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 [6]. It is based on the observation that the measure-

⋆ Supported in part by the EU fifth framework project RESQ, IST-2001-37559.
⋆⋆ 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.

L. Caires et al. (Eds.): ICALP 2005, LNCS 3580, pp. 1299–1311, 2005.

ment 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 some degree lower bounds for polynomials and thus obtain good lower bounds for quantum query complexity.

The third lower-bound method is the quantum adversary method of Ambainis [7]. 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 [1]. 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 quantum adversary method, let us call it unweighted, was invented by Ambainis [7]. It was successfully used to obtain the following tight lower bounds: $\Omega(\sqrt{n})$ for Grover search [8], $\Omega(\sqrt{n})$ for two-level And-Or trees (see [9] for a matching upper bound), and $\Omega(\sqrt{n})$ for inverting a permutation.

Some functions, such as sorting or ordered search, could not be well lower-bounded by the unweighted method. Høyer, Neerbek, and Shi used a weighting argument [10] to obtain tight bounds for these problems. Barnum, Saks, and Szegedy proposed a general method [1] that gives necessary and sufficient conditions for the existence of a quantum query algorithm. They also described a special case, 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* [2]. He applied it to get a lower bound for several iterated functions. Zhang observed that Ambainis had generalized his oldest method [7] in two independent ways, so he unified them, and published a *strong weighted adversary method* [3]. Finally, Laplante and Magniez used Kolmogorov complexity in an unusual way and described a *Kolmogorov complexity method* [4].

A few relations between the methods are known. The strong weighted adversary is clearly at least as good as the weighted adversary. Laplante and Magniez showed [4] 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. [10] is a special case of the weighted adversary method.

In addition it was known that there were some limitations for lower bounds obtained by the adversary method. Szegedy observed [11] 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 [4], 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 [3] and independently by us.

1.3 Our Results

In this paper, we clean up the forest of adversary methods. We show that there is essentially only one quantum adversary method and that all the former methods [1, 2, 3, 4, 10] are just different formulations of the same method. This means that the quantum adversary method is a very robust concept. Furthermore, we also present a new simple proof of the $\min(\sqrt{C_0 n}, \sqrt{C_1 n})$ limitation of the quantum adversary method for partial f, resp. $\sqrt{C_0 C_1}$ for total f.

This paper is an extended abstract with some proofs omitted. The full version can be downloaded from `http://arxiv.org/abs/quant-ph/0409116`.

1.4 Separation Between the Polynomial and Adversary Method

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

The polynomial method was successfully applied to obtain tight lower bound $\Omega\big(n^{1/3}\big)$ for the collision problem and $\Omega\big(n^{2/3}\big)$ for the element distinctness problem [12] (see [13] for a matching upper bound). The 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 [6]. The adversary method can only provide bounded-error lower bounds.

On the other hand, Ambainis exhibited some iterated functions [2], for which the adversary method gives better lower bounds than the polynomial method. The biggest proved gap between the two methods is $n^{1.321}$. Furthermore, the polynomial method did not yet succeed in proving 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 $\Omega(\sqrt{n})$ [7], whereas the best bound obtained by the polynomial method is $\Omega\big(n^{1/3}\big)$ and it follows from the element distinctness lower bound [12].

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 just implements the classical zero-error algorithm by Saks and Wigderson [14] running in expected time $\mathrm{O}\big(n^{0.753}\big)$. Adversary lower bounds are limited by $\sqrt{C_0 C_1} = \sqrt{n}$. Recently, Laplante, Lee, and Szegedy showed [15] 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. The polynomial method might prove a stronger lower bound. Another example is triangle finding. The best upper bound is $\mathrm{O}\big(n^{1.3}\big)$ [16] and the best lower bound is $\Omega(n)$. Again, the adversary method cannot give better than a linear bound, but the polynomial method might.

The semidefinite programming method [1] 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 [17] and sketch the model of quantum query complexity, referring to [18] for more details, also on the relation between query complexity and certificate complexity. Suppose we want to compute some function f. For input $x \in \{0,1\}^N$, a *query* gives us access to the input bits. It corresponds to the unitary transformation, which depends on input x in the following way: $O_x : |i,b,z\rangle \mapsto |i, b \oplus x_i, z\rangle$. Here $i \in [N] = \{1,\dots,N\}$ and $b \in \{0,1\}$; 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 $|0\rangle$. For a Boolean function f, the output of A is obtained by observing the leftmost qubit of the final superposition $A|0\rangle$, and its acceptance probability on input x is its probability of outputting 1.

2.2 Kolmogorov Complexity

An excellent book about Kolmogorov complexity is the book [19] by Li and Vitányi. A 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 [4], 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, $K(x|y) = \min\{|P| : P \in S, M(P,y) = x\}$.

2.3 Notation

Let $[n] = \{1,2,\dots,n\}$. 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. For a column vector x, let $|x|$ denote the ℓ_2-norm of x. Let $\lambda(A)$ denote the *spectral norm* of A, formally $\lambda(A) = \max_{x:|x|\neq 0} |Ax|/|x|$. We say that a matrix is *Boolean*, if it contains only zeroes and ones. Let AB denote the usual *matrix product* and let $A \circ B$ denote the *Hadamard (point-wise) product* [20]. Let $A \geq B$ denote the *point-wise comparison* and let $C \succeq D$ denote that $C - D$ is *positive semidefinite*. 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. Let $r(M) = \max_x r_x(M)$ and $c(M) = \max_y c_y(M)$.

We call a function $f : S \to \{0,1\}$ *total*, if $S = \{0,1\}^n$, otherwise it 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 bits $i \in I$ to x_i determines the function value. Formally, $\forall y \in S : y|_I = x|_I \Rightarrow f(y) = f(x)$, where $x|_I$ denotes the substring of x indexed by I. A certificate I for x is called minimal, if $|I| \le |J|$ for every certificate J for x. Let $\mathcal{C}_f(x)$ denote the lexicographically smallest *minimal certificate* for x. Let $C_0(f) = \max_{x: f(x)=0} |\mathcal{C}_f(x)|$ be the *zero-certificate complexity* of f and let $C_1(f) = \max_{x: f(x)=1} |\mathcal{C}_f(x)|$ be the *one-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 [1]: 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.

Theorem 1. *Let $n \ge 1$ be an integer, let $S \subseteq \{0,1\}^n$, and let $f : S \to \{0,1\}$ be a partial Boolean function. Let $Q_\varepsilon(f)$ be the ε-error quantum query complexity of f. Then* $\frac{Q_\varepsilon(f)}{1-2\sqrt{\varepsilon(1-\varepsilon)}} \ge \mathrm{SA}(f) = \mathrm{WA}(f) = \mathrm{SWA}(f) = \mathrm{MM}(f) = \mathrm{SMM}(f) = \mathrm{GSA}(f) = \Theta(\mathrm{KA}(f))$, *where* SA, WA, SWA, MM, SMM, GSA, *and* KA *are lower bounds given by the following methods:*

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

$$\mathrm{SA}(f) = \max_{\Gamma} \frac{\lambda(\Gamma)}{\max_i \lambda(\Gamma \circ D_i)}. \tag{1}$$

- **Weighted adversary [2].**[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 $w'(x,y,i)w'(y,x,i) \ge w^2(x,y)$ for all x,y,i such that $x_i \ne y_i$.*

[1] We use a different formulation [4] than in the original Ambainis papers [7, 2]. 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.

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

$$\mathrm{WA}(f) = \max_{w,w'} \min_{\substack{x,i \\ f(x)=0 \\ v(x,i)>0}} \sqrt{\frac{wt(x)}{v(x,i)}} \cdot \min_{\substack{y,j \\ f(y)=1 \\ v(y,j)>0}} \sqrt{\frac{wt(y)}{v(y,j)}}. \tag{2}$$

– **Strong weighted adversary [3].** *Let* w, w' *denote a weight scheme as above. Then*

$$\mathrm{SWA}(f) = \max_{w,w'} \min_{\substack{x,y,i \\ w(x,y)>0 \\ x_i \neq y_i}} \sqrt{\frac{wt(x)wt(y)}{v(x,i)v(y,i)}}. \tag{3}$$

– **Kolmogorov complexity [4].**[2] *Let* $\sigma \in \{0,1\}^*$ *denote a finite string. Then*

$$\mathrm{KA}(f) = \min_{\sigma} \max_{\substack{x,y \\ f(x)\neq f(y)}} \frac{1}{\sum_{i:x_i \neq y_i} \sqrt{2^{-K(i|x,\sigma)-K(i|y,\sigma)}}}. \tag{4}$$

– **Minimax over probability distributions [4].** *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*

$$\mathrm{MM}(f) = \min_{p} \max_{\substack{x,y \\ f(x)\neq f(y)}} \frac{1}{\sum_{i:x_i \neq y_i} \sqrt{p_x(i)\, p_y(i)}} \tag{5}$$

$$= 1 \Big/ \max_{p} \min_{\substack{x,y \\ f(x)\neq f(y)}} \sum_{i:x_i \neq y_i} \sqrt{p_x(i)\, p_y(i)}. \tag{6}$$

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

$$\begin{array}{ll} \text{maximize} & \mu \\ \text{subject to} & \forall i : \quad R_i \succeq 0, \\ & \sum_i R_i \circ I = I, \\ & \sum_i R_i \circ D_i \geq \mu F. \end{array} \tag{7}$$

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

$$\begin{array}{ll} \text{minimize} & \mu = \operatorname{tr} \Delta \\ \text{subject to} & \Delta \text{ is diagonal} \\ & Z \geq 0 \\ & Z \cdot F = 1 \\ & \forall i : \ \Delta - Z \circ D_i \succeq 0. \end{array} \tag{8}$$

[2] We use a different formulation than Laplante and Magniez [4]. They minimize over all algorithms A computing f and substitute σ = source code of A, whereas we minimize over all finite strings σ.

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 2. *Let $S \subseteq \{0,1\}^n$ and let $f : S \to \{0,1\}$ be a partial Boolean function. The max-min bound (6) is upper-bounded by* $\mathrm{MM}(f) \leq \min(\sqrt{C_0(f)n}, \sqrt{C_1(f)n})$. *If f is total, then* $\mathrm{MM}(f) \leq \sqrt{C_0(f)C_1(f)}$.

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

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

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

Let f be partial. For every $x \in f^{-1}(0)$, distribute the probability uniformly over any minimal certificate $\mathcal{C}_f(x)$, and for every $y \in f^{-1}(1)$, distribute the probability uniformly over all input bits. Formally, $p_x(i) = 1/|\mathcal{C}_f(x)|$ iff $i \in \mathcal{C}_f(x)$, $p_x(i) = 0$ for $i \notin \mathcal{C}_f(x)$, and $p_y(i) = 1/n$. Take any x, y such that $f(x) = 0$ and $f(y) = 1$, and the zero-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 (6):

$$\sum_{i:x_i \neq y_i} \sqrt{p_x(i)\, p_y(i)} \geq \sqrt{p_x(j)\, p_y(j)} = \sqrt{\frac{1}{|\mathcal{C}_f(x)|} \cdot \frac{1}{n}} \geq \frac{1}{\sqrt{C_0 n}}.$$

If f is total, then we can do even better. For every $x \in \{0,1\}^n$, 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 0 and 1.) After we have found a j, we lower-bound the sum of (6) in the same way as above. □

Some parts of the following statement have already been observed for individual methods by Szegedy [11], Laplante and Magniez [4], and Zhang [3]. This corollary rules out all adversary attempts to prove good lower bounds for problems with small certificate complexity, such as element distinctness [12], binary And-Or trees [14, 21, 9], or triangle finding [16].

Corollary 1. *All quantum adversary lower-bounds are at most* $\min(\sqrt{C_0(f)n}, \sqrt{C_1(f)n})$ *for partial functions and* $\sqrt{C_0(f)C_1(f)}$ *for total functions.*

4 Equivalence of Spectral and Strong Weighted Adversary

In this section, we give a linear-algebraic proof that the spectral bound [1] and the strong weighted bound [3] are equal. The proof has three steps. First, we show that the weighted bound [2] 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 third step is trivial, since the strong weighted bound is always at least as good as the weighted bound. The generalization of the weighted adversary method thus does not make the bound stronger, however its formulation is easier to use.

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 [20]. The second one is our generalization and its proof is omitted.

Lemma 1. [20] *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 $\lambda(S) \leq r(M)c(N) = \max_{x,y} r_x(M) c_y(N)$. Moreover, for every S there exist M, N such that $S = M \circ N$ and $\lambda(S) = r(M)c(N)$.*

Lemma 2. *Let S be a non-negative symmetric matrix and let M and N be non-negative matrices such that $S \leq M \circ N$. Let $B(M,N) = \max_{\substack{x,y \\ S[x,y]>0}} r_x(M) c_y(N)$. Then $\lambda(S) \leq B(M,N)$.*

Now we reduce the spectral adversary to the weighted adversary.

Theorem 3. $\mathrm{SA}(f) \leq \mathrm{WA}(f)$.

Proof. Let Γ be any non-negative symmetric matrix with $\Gamma \circ F = \Gamma$ as in equation (1). 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: $w(x,y) = w(y,x) = \Gamma[x,y] \cdot \delta[x]\delta[y]$. Furthermore, for every i, decompose every $\Gamma_i = \Gamma \circ D_i$ into a Hadamard product of two non-negative matrices $\Gamma_i = M_i \circ N_i$ such that $\lambda(\Gamma_i) = r(M_i)c(N_i)$. This is always possible by Lemma 1. We ensure that $r(M_i) = c(N_i) = \sqrt{\lambda(\Gamma_i)}$ by multiplying M_i and dividing N_i by the same constant. Define w':

$$w'(x,y,i) = \begin{cases} (M_i[x,y]\,\delta[x])^2 & \text{iff } f(x)=0,\ f(y)=1, \text{ and } x_i \neq y_i, \\ (N_i[y,x]\,\delta[y])^2 & \text{iff } f(x)=1,\ f(y)=0, \text{ and } x_i \neq y_i, \\ 0 & \text{otherwise.} \end{cases}$$

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 $f(x) = 0$, $f(y) = 1$, and $x_i \neq y_i$, then $w'(x,y,i)w'(y,x,i) = (M_i[x,y]\,\delta[x])^2 (N_i[x,y]\,\delta[y])^2 = (\Gamma_i[x,y]\,\delta[x]\delta[y])^2 = w(x,y)^2$. Finally, let us compute the bound (2) given by the weight scheme. Let $v_b = \max_{\substack{x,i \\ f(x)=b}} \frac{v(x,i)}{wt(x)}$. Then

$$wt(x) = \sum_y w(x,y) = \delta[x] \sum_y \Gamma[x,y]\delta[y] = \delta[x]\,(\Gamma\delta)[x] = \delta[x]^2,$$

$$v_0 = \max_{\substack{x,i \\ f(x)=0}} \frac{\sum_y w'(x,y,i)}{wt(x)} \leq \max_{\substack{x,i \\ f(x)=0}} \frac{\sum_y (M_i[x,y])^2\,\delta[x]^2}{\delta[x]^2} \leq \max_i (r(M_i))^2,$$

and, analogously, $v_1 \leq \max_i (c(N_i))^2$. Since $r(M_i) = c(N_i) = \sqrt{\lambda(\Gamma_i)}$, both $v_0, v_1 \leq \max_i \lambda(\Gamma_i)$. Hence $1/\sqrt{v_0 v_1} \geq 1/\max_i \lambda(\Gamma_i)$, and the weight scheme w, w' gives at least as good bound as the matrix Γ. □

Now we reduce the strong weighted adversary to the spectral adversary.

Theorem 4. $\mathrm{SWA}(f) \leq \mathrm{SA}(f)$.

Proof. Let w, w' be any weight scheme as in equation (2). 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 $\lambda(\Gamma) \geq \delta^T \Gamma \delta / |\delta|^2 = W/W = 1$. Next, we show that, for every i, we have $\lambda(\Gamma_i) \leq \sqrt{u_i}$ for $u_i = \max_{\substack{x,y \\ w(x,y)>0, x_i \neq y_i}} \frac{v(x,i)v(y,i)}{wt(x)wt(y)}$. Once we prove this, we are done, since the strong weighted bound (3) is $1/\max_i \sqrt{u_i}$. Let $A = f^{-1}(0)$ and $B = f^{-1}(1)$. Fix i and define the following rectangular matrices on the index set $A \times B$:

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

Every weight scheme satisfies $w'(x,y,i)w'(y,x,i) \geq w^2(x,y)$ for all x, y, i such that $x_i \neq y_i$. It follows that if we reorder Γ_i to put A first and B last, then

$$\Gamma_i = \Gamma \circ D_i \leq \begin{pmatrix} 0 & M_i \circ N_i \\ M_i^T \circ N_i^T & 0 \end{pmatrix} = \begin{pmatrix} 0 & \overline{M}_i \\ \overline{N}_i^{\,T} & 0 \end{pmatrix} \circ \begin{pmatrix} 0 & \overline{N}_i \\ \overline{M}_i^{\,T} & 0 \end{pmatrix},$$

where $\overline{M}_i = \sqrt{\frac{c(N_i)}{r(M_i)}}\, M_i$ and $\overline{N}_i = \sqrt{\frac{r(M_i)}{c(N_i)}}\, N_i$. This is done for balancing the row norm of M_i and column norm of N_i: $r(\overline{M}_i) = c(\overline{N}_i) = \sqrt{r(M_i)c(N_i)}$. Evaluate

$$B(M_i, N_i^T) = \max_{\substack{x,y \\ \Gamma_i[x,y]>0}} \sqrt{\sum_k \frac{w'(x,k,i)}{wt(x)} \sum_\ell \frac{w'(y,\ell,i)}{wt(y)}} = \max_{\substack{x,y \\ w(x,y)>0 \\ x_i \neq y_i}} \sqrt{\frac{v(x,i)v(y,i)}{wt(x)wt(y)}}.$$

By Lemma 2, $\lambda(\Gamma_i) \leq B(\overline{M}_i, \overline{N}_i^{\,T}) = B(M_i, N_i^T) = \sqrt{u_i}$, as claimed. □

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 get rid of 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. $\mathrm{MM}(f) = \mathrm{SMM}(f)$.

Proof. Let p be a set of probability distributions as in equation (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 $\min_{\substack{x,y \\ f(x) \neq f(y)}} \sum_{i: x_i \neq y_i} R_i[x,y] \geq \mu$ is rewritten into $\forall x,y : f(x) \neq f(y) \implies \sum_{i:x_i \neq y_i} R_i[x,y] \geq \mu$, which is $\sum_i R_i \circ D_i \geq \mu F$. However, the matrices R_i are rank-1 and they have non-negative entries. We have replaced that condition by $R_i \succeq 0$ to get semidefinite program (7). Hence the program (7) is a relaxation of the condition of (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,

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

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

Theorem 6. $\mathrm{SMM}(f) = \mathrm{GSA}(f)$.

Proof. Omitted; it only uses the duality theory of semidefinite programming.

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 a positive diagonal matrix Δ as a new variable.

Theorem 7. $\mathrm{GSA}(f) = \mathrm{SA}(f)$.

Proof. Let Z, Δ be a solution of (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, $\Delta - Z \circ D_i \succeq 0$ implies that $Z \circ D_i$ is symmetric for every i. It follows that Z must be also symmetric.

Take a column vector $a = \mathrm{diag}\,(\Delta^{-1/2})$ 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

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

Γ is clearly symmetric, $\Gamma \geq 0$, and $\Gamma \circ F = \Gamma$. The bound (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 (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 [4] 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 8. [4–Theorem 2] $\mathrm{KA}(f) = \Omega(\mathrm{WA}(f))$.

Theorem 9. $\mathrm{KA}(f) = \mathrm{O}(\mathrm{MM}(f))$.

Proof. Take a set of probability distributions p as in equation (5). The query information lemma [4–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 prefix-free. 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 (4) includes among others also the descriptions of all probability distributions p. □

Proof (of Theorem 1). We have to prove that $\frac{Q_\varepsilon(f)}{1-2\sqrt{\varepsilon(1-\varepsilon)}} \geq \mathrm{SA}(f) = \mathrm{WA}(f) = \mathrm{SWA}(f) = \mathrm{MM}(f) = \mathrm{SMM}(f) = \mathrm{GSA}(f) = \Theta(\mathrm{KA}(f))$, Put together all known equalities and inequalities:

– $\mathrm{SA}(f) = \mathrm{WA}(f) = \mathrm{SWA}(f)$ by Theorem 3 and Theorem 4,
– $\mathrm{MM}(f) = \mathrm{SMM}(f)$ by Theorem 5,
– $\mathrm{SMM}(f) = \mathrm{GSA}(f)$ by Theorem 6,
– $\mathrm{GSA}(f) = \mathrm{SA}(f)$ by Theorem 7,
– $\mathrm{KA}(f) = \Theta(\mathrm{WA}(f))$ by Theorem 8 and Theorem 9.

Finally, one has to prove one of the lower bounds. For example, Ambainis proved [2] that $Q_2(f) \geq (1 - 2\sqrt{\varepsilon(1-\varepsilon)})\,\mathrm{WA}(f)$. □

Acknowledgments

We thank Ronald de Wolf for many fruitful discussions, for his suggestions concerning Theorem 2, and for proofreading, and Troy Lee for discussions.

References

1. Barnum, H., Saks, M., Szegedy, M.: Quantum decision trees and semidefinite programming. In: Proc. of the 18th IEEE Conference on Computational Complexity. (2003) 179–193
2. Ambainis, A.: Polynomial degree vs. quantum query complexity. In: Proc. of the 44th IEEE Symp. on Foundations of Computer Science. (2003) 230–239
3. Zhang, S.: On the power of Ambainis's lower bounds. In: Proceedings of 31st ICALP. (2004) 1238–1250 quant-ph/0311060.
4. Laplante, S., Magniez, F.: Lower bounds for randomized and quantum query complexity using Kolmogorov arguments. In: Proceedings of 19th IEEE Conference on Computational Complexity. (2004) 294–304 quant-ph/0311189.
5. Bennett, H., Bernstein, E., Brassard, G., Vazirani, U.: Strengths and weaknesses of quantum computing. SIAM Journal on Computing **26** (1997) 1510–1523 quant-ph/9701001.
6. Beals, R., Buhrman, H., Cleve, R., Mosca, M., de Wolf, R.: Quantum lower bounds by polynomials. Journal of the ACM **48** (2001) 778–797 Earlier version in FOCS'98. quant-ph/9802049.
7. Ambainis, A.: Quantum lower bounds by quantum arguments. Journal of Computer and System Sciences **64** (2002) 750–767 Earlier version in STOC'2000. quant-ph/0002066.
8. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Proceedings of 28th ACM STOC. (1996) 212–219 quant-ph/9605043.
9. Høyer, P., Mosca, M., de Wolf, R.: Quantum search on bounded-error inputs. In: Proceedings of 30th ICALP. (2003) 291–299 LNCS 2719, quant-ph/0304052.
10. Høyer, P., Neerbek, J., Shi, Y.: Quantum complexities of ordered searching, sorting, and element distinctness. Algorithmica **34** (2002) 429–448 Special issue on Quantum Computation and Cryptography. quant-ph/0102078.
11. Szegedy, M.: On the quantum query complexity of detecting triangles in graphs. quant-ph/0310107 (2003)
12. Aaronson, S., Shi, Y.: Quantum lower bounds for the collision problem. Journal of the ACM **51** (2004) 595–605 quant-ph/0111102.

13. Ambainis, A.: Quantum walk algorithm for element distinctness. In: Proceedings of 45th IEEE FOCS. (2004) 22–31 quant-ph/0311001.
14. Saks, M., Wigderson, A.: Probabilistic Boolean decision trees and the complexity of evaluating games trees. In: Proc. of the 27th Annual Symp. on FOCS. (1986) 29–38
15. Laplante, S., Lee, T., Szegedy, M.: The quantum adversary method and formula size lower bounds. In: Proceedings of 20th IEEE Conference on Computational Complexity. (2005) To appear. quant-ph/0501057.
16. Magniez, F., Santha, M., Szegedy, M.: Quantum algorithms for the triangle problem. In: Proceedings of 16th SODA. (2005) 1109–1117 quant-ph/0310134.
17. Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information. Cambridge University Press (2000)
18. Buhrman, H., de Wolf, R.: Complexity measures and decision tree complexity: A survey. Theoretical Computer Science **288** (2002) 21–43
19. Li, M., Vitányi, P.M.B.: An Introduction to Kolmogorov Complexity and its Applications. Second edn. Springer, Berlin (1997)
20. Mathias, R.: The spectral norm of a nonnegative matrix. Linear Algebra and its Applications **139** (1990) 269–284
21. Barnum, H., Saks, M.: A lower bound on the quantum query complexity of read-once functions. Journal of Computer and Systems Sciences **69** (2004) 244–258 quant-ph/0201007.