

Negative weights make adversaries stronger

Peter Høyer*

hoyer@cpsc.ucalgary.ca

Troy Lee†

troyjlee@gmail.com

Robert Špalek‡

spalek@eecs.berkeley.edu

Abstract

The quantum adversary method is one of the most successful techniques for proving lower bounds on quantum query complexity. It gives optimal lower bounds for many problems, has application to classical complexity in formula size lower bounds, and is versatile with equivalent formulations in terms of weight schemes, eigenvalues, and Kolmogorov complexity. All these formulations rely on the principle that if an algorithm successfully computes a function then, in particular, it is able to distinguish between inputs which map to different values.

We present a stronger version of the adversary method which goes beyond this principle to make explicit use of the stronger condition that the algorithm actually computes the function. This new method, which we call ADV^\pm , has all the advantages of the old: it is a lower bound on bounded-error quantum query complexity, its square is a lower bound on formula size, and it behaves well with respect to function composition. Moreover ADV^\pm is always at least as large as the adversary method ADV , and we show an example of a monotone function for which $\text{ADV}^\pm(f) = \Omega(\text{ADV}(f)^{1.098})$. We also give examples showing that ADV^\pm does not face limitations of ADV like the certificate complexity barrier and the property testing barrier.

1 Introduction

Quantum query complexity is a popular model for study as it seems to capture much of the power of quantum computing—in particular, the search algorithm of Grover [Gro96] and the period finding routine of Shor’s factoring algorithm [Sho97] can be formulated in this model—yet is still

*Department of Computer Science, University of Calgary. Supported by Canada’s Natural Sciences and Engineering Research Council (NSERC), the Canadian Institute for Advanced Research (CIAR), and The Mathematics of Information Technology and Complex Systems (MITACS).

†LRI, Université Paris-Sud. Supported by a Rubicon grant from the Netherlands Organisation for Scientific Research (NWO) and by the European Commission under the Integrated Project Qubit Applications (QAP) funded by the IST directorate as Contract Number 015848. Part of this work conducted while at CWI, Amsterdam, and while visiting the University of Calgary.

‡University of California, Berkeley. Supported by NSF Grant CCF-0524837 and ARO Grant DAAD 19-03-1-0082. Work conducted in part while at CWI and the University of Amsterdam, supported by the European Commission under project QAP, IST-015848, and while visiting the University of Calgary.

simple enough that we can often hope to prove tight lower bounds. In this model, complexity is measured by the number of queries made to the input, and other operations are for free. For most known quantum algorithms, the time complexity is bigger than the query complexity by only a polylogarithmic factor.

The two most successful techniques for proving lower bounds on quantum query complexity are the polynomial method [BBC⁺01] and the quantum adversary method [Amb02]. The adversary method gives tight lower bounds for many problems and is quite versatile with formulations in terms of weight schemes [Amb06, Zha05], eigenvalues [BSS03], and Kolmogorov complexity [LM04]. Špalek and Szegedy [ŠS06] show that in fact all these formulations are equivalent. All these versions of the adversary method rest on the principle that, if an algorithm is able to *compute* a function f , then in particular it is able to *distinguish* inputs which map to different values. The method actually bounds the difficulty of this distinguishing task.

We present a stronger version of the adversary method which goes beyond this principle to essentially make use of the stronger condition that the algorithm actually computes the function—namely, we make use of the existence of a measurement which gives the correct answer with high probability from the final state of the algorithm. This new method, which we call ADV^\pm , is always at least as large as the adversary bound ADV , and we show an example of a monotone function f for which $\text{ADV}^\pm(f) = \Omega(\text{ADV}(f)^{1.098})$. Moreover, ADV^\pm possesses all the nice properties of the old adversary method: it is a lower bound on bounded-error quantum query complexity, its square is a lower bound on formula size, and it behaves well with respect to function composition. Using this last property, and the fact that our bound is larger than the adversary bound for the base function of Ambainis, we improve the best known separation between quantum query complexity and polynomial degree giving an f such that $Q_\epsilon(f) = \Omega(\text{deg}(f)^{1.329})$.

The limitations of the adversary method are fairly well understood. One limitation is the “certificate complexity barrier.” This says that $\text{ADV}(f) \leq \sqrt{C_0(f)C_1(f)}$ for a total function f with Boolean output [Zha05, ŠS06], where $C_0(f)$ is the certificate complexity of the inputs x which evaluate to zero on f , and $C_1(f)$ is the certificate complexity of inputs which evaluate to one. This means that for problems like determining if a graph contains a triangle, or element distinctness, where one of the certificate complexities is constant, the best bound which can be proven by the adversary method is $\Omega(\sqrt{N})$. For triangle finding, the best known upper bound is $O(N^{13/20})$ [MSS05], and for element distinctness the polynomial method is able to prove a tight lower bound of $\Omega(N^{2/3})$ [AS04]. We show that our new method can break the certificate complexity barrier—we give an example where $\text{ADV}^\pm(f) = \Omega((C_0(f)C_1(f))^{0.549})$.

Another limitation of the adversary method is the “property testing barrier.” For a partial Boolean function f where all zero-inputs have relative Hamming distance at least ϵ from all one-inputs, it holds that $\text{ADV}(f) \leq 1/\epsilon$. A prime example where this limitation applies is the collision problem of determining if a function is 2-to-1 or 1-to-1. Here all zero-inputs have relative Hamming distance at least $1/2$ from all one inputs and so the best bound provable by the adversary method is 2, while the polynomial method is able to prove a tight lower bound of $\Omega(n^{1/3})$ [AS04]. We show the property testing barrier does not apply in this strict sense to ADV^\pm , although we do not know of an asymptotic separation for constant ϵ .

Breaking these barriers opens the possibility that ADV^\pm can prove tight lower bounds for

problems like element distinctness and the collision problem, and improve the best known $\Omega(\sqrt{N})$ lower bound for triangle finding.

1.1 Comparison with previous methods

We now take a closer look at our new method and how it compares with previous adversary methods. We use the setting of the spectral formulation of the adversary method [BSS03].

Let $f : S \rightarrow \Sigma_O$ be a function, with $S \subseteq \Sigma_I^n$ the set of inputs. We assume $\Sigma_I = \{0, 1, \dots, |\Sigma_I| - 1\}$, and call this the input alphabet and Σ_O the output alphabet. Let Γ be a Hermitian matrix with rows and columns labeled by elements of S . We say that Γ is an *adversary matrix for f* if $\Gamma[x, y] = 0$ whenever $f(x) = f(y)$. We let $\|M\|$ denote the spectral norm of the matrix M , and for a real matrix M use $M \geq 0$ to say the entries of M are nonnegative. We now give the spectral formulation of the adversary method:

Definition 1

$$\text{ADV}(f) = \max_{\substack{\Gamma \geq 0 \\ \Gamma \neq 0}} \frac{\|\Gamma\|}{\max_i \|\Gamma \circ D_i\|}.$$

Here the maximum is taken over nonnegative symmetric adversary matrices Γ , and D_i is a zero-one matrix where $D_i[x, y] = 1$ if $x_i \neq y_i$ and $D_i[x, y] = 0$ otherwise. $\Gamma \circ D_i$ denotes the entry-wise (Hadamard) product of Γ and D_i .

Let $Q_\epsilon(f)$ be the two-sided ϵ -bounded error quantum query complexity of f . Barnum, Saks, and Szegedy show that the spectral version of the adversary method is a lower bound on $Q_\epsilon(f)$:

Theorem 1 ([BSS03]) For any function f , $Q_\epsilon(f) \geq \frac{1-2\sqrt{\epsilon(1-\epsilon)}}{2} \text{ADV}(f)$.

Note that the definition of $\text{ADV}(f)$ restricts the maximization to adversary matrices whose entries are all nonnegative and real. Our new bound removes these restrictions:

Definition 2

$$\text{ADV}^\pm(f) = \max_{\Gamma \neq 0} \frac{\|\Gamma\|}{\max_i \|\Gamma \circ D_i\|}.$$

It is clear that $\text{ADV}^\pm(f) \geq \text{ADV}(f)$ for any function f as the maximization is taken over a larger set. Our main theorem, presented in [Section 3](#), states that $\text{ADV}^\pm(f)$ is a lower bound on $Q_\epsilon(f)$.

Theorem 2 For any function f , $Q_\epsilon(f) \geq \frac{1-2\sqrt{\epsilon(1-\epsilon)}-2\epsilon}{2} \text{ADV}^\pm(f)$. If f has Boolean output, i.e. if $|\Sigma_O| = 2$, then $Q_\epsilon(f) \geq \frac{1-2\sqrt{\epsilon(1-\epsilon)}}{2} \text{ADV}^\pm(f)$.

While it is clear that ADV^\pm is always least as large as ADV , it might at first seem surprising that ADV^\pm can achieve bounds super-linear in ADV . An intuition for why negative weights help is that it is good to give negative weight to entries with large Hamming distance, entries which are easier to distinguish by queries. Consider an entry (x, y) where x and y have large Hamming distance. This entry appears in several $\Gamma \circ D_i$ matrices but only appears in the Γ matrix once. Thus

by giving this entry negative weight we can simultaneously decrease $\|\Gamma \circ D_i\|$ for several i 's, while doing relatively little damage to the large Γ matrix.

While in form the ADV^\pm bound is very similar to the ADV bound, our proof of [Theorem 2](#) departs from the standard adversary principle. The standard adversary principle is based on the fact that an algorithm A which is able to *compute* a function f is, in particular, able to *distinguish* inputs x, y such that $f(x) \neq f(y)$. Distinguishing quantum states is closely related to the inner product of the states as given by the following quantitative principle:

Theorem 3 *Suppose we are given one of two known states $|\psi_x\rangle, |\psi_y\rangle$. Let $0 \leq \epsilon \leq 1/2$. There is a measurement which correctly identifies which of the two states we are given with error probability ϵ if and only if $\langle \psi_x | \psi_y \rangle \leq 2\sqrt{\epsilon(1-\epsilon)}$.*

Let $|\psi_x^t\rangle$ be the state of an algorithm on input x after t queries. The adversary method works by defining a “progress function” based on the inner product $\langle \psi_x^t | \psi_y^t \rangle$. Initially, before the algorithm has made any queries, all inputs look the same and thus $\langle \psi_x^0 | \psi_y^0 \rangle = 1$ for all x, y , and thus the progress function is large. On the other hand, if a T -query algorithm computes a function f within error ϵ , then by [Theorem 3](#) for x, y with $f(x) \neq f(y)$ we must have $\langle \psi_x^T | \psi_y^T \rangle \leq 2\sqrt{\epsilon(1-\epsilon)}$, and thus the final progress function is small. In [\[BSS03\]](#) this is termed the Ambainis output condition. The adversary method then works by showing an upper bound on how much the progress function can change by a single query.

Our proof follows the same basic reasoning, but the Ambainis output condition no longer seems to suffice to show that the final progress function is small. We use in an essential way the stronger output condition that if a T -query algorithm A computes a function f , then there exists orthogonal projectors $\{\Pi_b\}_{b \in \Sigma_O}$ which sum to the identity such that $\|\Pi_b |\psi_x^T\rangle\|^2 \geq 1 - \epsilon$ when $f(x) = b$.

2 Preliminaries

We assume standard background from quantum computing and Boolean function complexity, see [\[NC00\]](#) and [\[BW02\]](#) for nice references. In this section, we restrict ourselves to more specific background.

2.1 Linear algebra

The background we need about matrices can be found in, for example, [\[Bha97\]](#). We use standard notations such as $|\cdot|$ for absolute value, \overline{A} for the entrywise complex conjugate of a matrix A , A^* for the conjugate transpose of A , and $\|x\| = \sqrt{x^*x}$ for the ℓ_2 -norm of a vector x . For two matrices A, B of the same size, the Hadamard product or entrywise product is the matrix $(A \circ B)[x, y] = A[x, y]B[x, y]$.

For an indexed set of vectors $\{|\psi_x\rangle : x \in S\}$, we associate an $|S|$ -by- $|S|$ Gram matrix $M = \text{Gram}(|\psi_x\rangle : s \in S)$ where

$$M[x, y] = \langle \psi_x | \psi_y \rangle.$$

It is easy to see that M is Hermitian and positive semidefinite.

We will make use of several matrix norms. For a matrix A let $\|A\|$ be the spectral norm of A

$$\|A\| = \max_{x,y} \frac{|x^*Ay|}{\|x\|\|y\|}.$$

For two matrices A, B let $\langle A, B \rangle$ be the Hilbert-Schmidt inner product. This is the inner product of A, B viewed as long vectors,

$$\langle A, B \rangle = \text{Tr}(A^*B) = \sum_{i,j} \overline{A[i,j]}B[i,j].$$

The Frobenius norm, denoted $\|A\|_F$, is the norm associated with this inner product,

$$\|A\|_F = \sqrt{\langle A, A \rangle} = \sqrt{\sum_{i,j} |A[i,j]|^2}.$$

Finally, we will use the trace norm, denoted $\|A\|_{tr}$, where

$$\|A\|_{tr} = \max_B \frac{|\langle A, B \rangle|}{\|B\|},$$

and B runs over all complex matrices of the same size as A . The following theorem is an easy consequence of this definition.

Theorem 4 *Let A, B be n -by- n matrices. Then $|\langle A, B \rangle| \leq \|A\| \cdot \|B\|_{tr}$.*

In our proof that ADV^\pm is a lower bound on quantum query complexity we will need one more tool for bounding norms:

Theorem 5 (HÖLDER'S INEQUALITY, [BHA97] COROLLARY IV.2.6) *Let A, B be matrices such that AB^* is defined. Then*

$$\|AB^*\|_{tr} \leq \|A\|_F \|B\|_F.$$

A *partial trace* is a linear mapping $\text{Tr}_A : \mathcal{L}(A \otimes B) \rightarrow \mathcal{L}(A)$ mapping linear operators (e.g., density matrices) over the joint system AB to linear operators over A . This mapping is uniquely determined by the requirement

$$\text{Tr}_A(\rho_A \otimes \rho_B) = \rho_A \cdot \text{Tr}(\rho_B).$$

and linearity. For example, since $\text{Tr}_A(|i\rangle\langle j|_A \otimes |\psi_i\rangle\langle\psi_j|_B) = |i\rangle\langle j|_A \cdot \langle\psi_j|\psi_i\rangle$, by linearity

$$\text{Tr}_A\left(\sum_{i,j} |i\rangle\langle j|_A \otimes |\psi_i\rangle\langle\psi_j|_B\right) = \sum_{i,j} |i\rangle\langle j|_A \cdot \langle\psi_j|\psi_i\rangle.$$

2.2 Quantum query complexity

As with the classical model of decision trees, in the quantum query model we wish to compute some function f and we access the input through queries. The complexity of f is the number of queries needed to compute f on a worst-case input x . Unlike the classical case, however, we can now make queries in superposition.

The memory of a quantum query algorithm is described by three registers: the input register, \mathcal{H}_I , which holds the input $x \in S \subseteq \Sigma_I^n$, the query register, \mathcal{H}_Q , which holds two integers $1 \leq i \leq n$ and $0 \leq p < |\Sigma_I|$, and the working memory, \mathcal{H}_W , which holds an arbitrary value. The query register and working memory together form the accessible memory, denoted \mathcal{H}_A .

The accessible memory of a quantum query algorithm A is initialized to a fixed state. For convenience, on input x we assume the state of the algorithm is $|x\rangle_I |1, 0\rangle_Q |0\rangle_W$ where all qubits in the working memory are initialized to 0. The state of the algorithm then evolves through queries, which depend on the input register, and accessible memory operators which do not. We now describe these operations.

We will model a query by a unitary operator where the oracle answer is given in the phase. This operator O is defined by its action on the basis state $|x\rangle|i, p\rangle$ as

$$O|x\rangle|i, p\rangle = e^{\frac{2\pi i}{|\Sigma_I|} p x_i} |x\rangle|i, p\rangle,$$

where $1 \leq i \leq n$ is the index of the queried input variable and $0 \leq p < |\Sigma_I|$ is the phase multiplier. This operation can be extended to act on the whole space by interpreting it as $O \times I_W$, where I_W is the identity operation on the workspace \mathcal{H}_W . In the sequel, we will refer to the action of O both on $\mathcal{H}_I \otimes \mathcal{H}_Q$ and the full space $\mathcal{H}_I \otimes \mathcal{H}_Q \otimes \mathcal{H}_W$, and let context dictate which we mean.

For a function with $\Sigma_I = \{0, 1\}$, the query operator simply becomes

$$O|x\rangle|i, p\rangle = (-1)^{p x_i} |x\rangle|i, p\rangle,$$

An alternative, perhaps more common, way to model a quantum query is through an operator $O' : |x\rangle|i, p\rangle \rightarrow |x\rangle|i, (x_i + p) \bmod |\Sigma_I|\rangle$ that encodes the result in a register. These two query models are equivalent, as can be seen by conjugating with the quantum Fourier transform on $|p\rangle$. For our results, it is more convenient to work with the phase oracle.

An accessible memory operator is an arbitrary unitary operation U on the accessible memory \mathcal{H}_A . This operation is extended to act on the whole space by interpreting it as $I_I \otimes U$, where I_I is the identity operation on the input space \mathcal{H}_I . Thus the state of the algorithm on input x after t queries can be written

$$|\phi_x^t\rangle = U_t O U_{t-1} \cdots U_1 O U_0 |x\rangle |1, 0\rangle |0\rangle.$$

As the input register is left unchanged by the algorithm, we can decompose $|\phi_x^t\rangle$ as $|\phi_x^t\rangle = |x\rangle |\psi_x^t\rangle$, where $|\psi_x^t\rangle$ is the state of the accessible memory after t queries.

The output of a T -query algorithm A on input x is chosen according to a probability distribution which depends on the final state of the accessible memory $|\psi_x^T\rangle$. Namely, the probability that the algorithm outputs some $b \in \Sigma_O$ on input x is $\|\Pi_b |\psi_x^T\rangle\|^2$, for a fixed set of projectors $\{\Pi_b\}$ which are orthogonal and complete, that is, sum to the identity. The ϵ -error quantum query complexity of a function f , denoted $Q_\epsilon(f)$, is the minimum number of queries made by an algorithm which outputs $f(x)$ with probability at least $1 - \epsilon$ for every x .

3 Bounded-error quantum query complexity

We now show that $\text{ADV}^\pm(f)$ is a lower bound on the bounded-error quantum query complexity of f .

Proof of Theorem 2 Let $f : S \rightarrow \Sigma_O$, where $S \subseteq \Sigma_I^n$, be a function and let Γ be a $|S|$ -by- $|S|$ Hermitian matrix such that $\Gamma[x, y] = 0$ if $f(x) \neq f(y)$. As Γ is Hermitian, its eigenvalues will be equal to its singular values, up to sign. Notice that Γ and $-\Gamma$ have the same adversary value, thus without loss of generality we will assume Γ has largest eigenvalue equal to its spectral norm. Therefore, let δ be an eigenvector of Γ corresponding to the eigenvalue $\|\Gamma\|$.

We imagine that we initially prepare the state $|\Psi^0\rangle = \sum_x \delta_x |x\rangle |1, 0\rangle |0\rangle$ and run the algorithm on this superposition. Thus after t queries we have the state

$$|\Psi^t\rangle = U_t O U_{t-1} \dots U_1 O U_0 \sum_x \delta_x |x\rangle |1, 0\rangle |0\rangle = \sum_x \delta_x |x\rangle |\psi_x^t\rangle,$$

where $|\psi_x^t\rangle$ is the state of the accessible memory of the algorithm on input x after t queries. We define $\rho^{(t)} = \text{Tr}_I |\Psi^t\rangle \langle \Psi^t|$ to be the reduced density matrix of the state $|\Psi^t\rangle$ on the input register, that is we trace out the accessible memory. In other words, $\rho^{(t)} = \text{Gram}(\delta_x |\psi_x^t\rangle : x \in S)$.

We define a progress function W^t based on $\rho^{(t)}$ as $W^t = \langle \Gamma, \rho^{(t)} \rangle$. Although phrased differently, this is in fact the same progress function used by Høyer and Špalek [HŠ05] in their proof that the regular adversary method is a lower bound on bounded-error quantum query complexity. Note that W^t is real, as both Γ and $\rho^{(t)}$ are Hermitian. Our proof rests on three claims:

1. At the beginning of the algorithm $W^0 = \|\Gamma\|$.
2. With any one query, the progress measure changes by at most $W^t - W^{t+1} \leq 2 \max_i \|\Gamma \circ D_i\|$.
3. At the end of the algorithm $W^T \leq (2\sqrt{\epsilon(1-\epsilon)} + 2\epsilon)\|\Gamma\|$.

The theorem clearly follows from these three claims. The main novelty of the proof lies in the third step. This is where we depart from the standard adversary principle in using a stronger output condition implied by a successful algorithm.

Item 1 As the state of the accessible memory $|\psi_x^0\rangle$ is independent of the oracle, $\langle \psi_x^0 | \psi_y^0 \rangle = 1$ for every x, y , and so $\rho^{(0)} = \delta\delta^*$. Thus $W^0 = \langle \Gamma, \delta\delta^* \rangle = \text{Tr}(\delta^* \Gamma^* \delta) = \|\Gamma\|$.

Item 2 After the $t + 1^{\text{st}}$ query, the quantum state is $|\Psi^{t+1}\rangle = U_{t+1} O |\Psi^t\rangle$ and thus

$$\rho^{(t+1)} = \text{Tr}_I (U_{t+1} O |\Psi^t\rangle \langle \Psi^t| O^* U_{t+1}^*) = \text{Tr}_I (O |\Psi^t\rangle \langle \Psi^t| O^*),$$

because the unitary operator U_{t+1} acts as identity on the input register. The oracle operator O only acts on the input register and the query register, hence we can trace out the working memory.

Denote $\rho = \text{Tr}_{I,Q} |\Psi^t\rangle\langle\Psi^t|$ and $\rho' = \text{O}\rho\text{O}^*$. Then $\rho^{(t)} = \text{Tr}_I(\rho)$ and $\rho^{(t+1)} = \text{Tr}_I(\rho')$. We express the progress function in terms of ρ, ρ' . Define two block-diagonal matrices on $\mathcal{H}_I \otimes \mathcal{H}_Q$:

$$G = \Gamma \otimes \mathbb{I}_n \otimes \mathbb{I}_{|\Sigma_I|} = \bigoplus_{i=1}^n \bigoplus_{p=0}^{|\Sigma_I|-1} \Gamma$$

$$D = \bigoplus_{i=1}^n \bigoplus_{p=0}^{|\Sigma_I|-1} D_i,$$

where D_i is the zero-one symmetric matrix from [Definition 1](#). Then $W^t = \langle \Gamma, \rho^{(t)} \rangle = \langle G, \rho \rangle$ and $W^{t+1} = \langle \Gamma, \rho^{(t+1)} \rangle = \langle G, \rho' \rangle$. We upper-bound the change of the progress function as follows:

$$\begin{aligned} W^t - W^{t+1} &= \langle G, \rho \rangle - \langle G, \rho' \rangle \\ &= \langle G, \rho - \text{O}\rho\text{O}^* \rangle && \rho' = \text{O}\rho\text{O}^* \\ &= \langle G, (\rho - \text{O}\rho\text{O}^*) \circ D \rangle && \rho - \text{O}\rho\text{O}^* = (\rho - \text{O}\rho\text{O}^*) \circ D \\ &= \langle G \circ D, \rho - \text{O}\rho\text{O}^* \rangle && D \text{ is real} \\ &\leq \|G \circ D\| \cdot \|\rho - \text{O}\rho\text{O}^*\|_{tr} && \text{Theorem 4} \\ &\leq \|G \circ D\| \cdot (\|\rho\|_{tr} + \|\text{O}\rho\text{O}^*\|_{tr}) && \text{triangle inequality} \\ &= 2\|G \circ D\| \cdot \|\rho\|_{tr} && \text{O is unitary} \\ &= 2\|G \circ D\| && \|\rho\|_{tr} = 1 \\ &= 2 \max_i \|\Gamma \circ D_i\|. \end{aligned}$$

The third equality holds because O is diagonal in the computational basis and thus, in the block corresponding to the value $|i, p\rangle$ of the query register, $(\rho - \text{O}\rho\text{O}^*)[x, y] = (1 - e^{\frac{2\pi i}{|\Sigma_I|} p(x_i - y_i)})\rho[x, y]$, which is 0 if $x_i = y_i$. For the remaining equalities note that conjugating with a unitary operator does not change the trace norm and that density matrices have trace norm 1.

Item 3 Now consider the algorithm at the final time T . We want to upper bound $\langle \Gamma, \rho^{(T)} \rangle$. The first thing to notice is that as $\Gamma[x, y] = 0$ when $f(x) = f(y)$, we have $\Gamma = \Gamma \circ F$, where F is a zero-one matrix such that $F[x, y] = 1$ if $f(x) \neq f(y)$ and $F[x, y] = 0$ otherwise.

As F is a real matrix, it is clear from the definition of the Hilbert-Schmidt inner product that $\langle \Gamma \circ F, \rho^{(T)} \rangle = \langle \Gamma, F \circ \rho^{(T)} \rangle$. Now applying [Theorem 4](#) we have $\langle \Gamma, F \circ \rho^{(T)} \rangle \leq \|\Gamma\| \cdot \|\rho^{(T)} \circ F\|_{tr}$. It remains to upper bound $\|\rho^{(T)} \circ F\|_{tr}$, which we do using [Theorem 5](#). To be able to apply this theorem, we would like to find matrices X, Y such that $XY^* = \rho^{(T)} \circ F$, and the product $\|X\|_F \|Y\|_F$ is small.

Let $\{\Pi_z\}_{z \in \Sigma_O}$ be a complete set of orthogonal projectors that determine the output probabilities, that is the probability that the algorithm outputs z on input x is $\|\Pi_z |\psi_x^T\rangle\|^2$. We will use these projectors to help decompose $\rho^{(T)}$ in such a way as to apply [Theorem 5](#). The correctness of the algorithm tells us that $\|\Pi_{f(x)} |\psi_x^T\rangle\|^2 \geq 1 - \varepsilon$. For an $i \in \{0, 1, \dots, |\Sigma_O| - 1\}$, let X_i denote the matrix with $|S|$ rows $\{\Pi_{f(x)+i} \delta_x |\psi_x^T\rangle\}_{x \in S}$, where $f(x) + i$ is computed modulo $|\Sigma_O|$. Intuitively, X_0 is the matrix where we project onto the correct answers, and X_i for $i \geq 1$ where we project onto some incorrect answer. The matrices X_i for $i \geq 1$ will therefore have small Frobenius norm.

We claim $\rho^{(T)} \circ F = \sum_{i \neq j} X_i X_j^*$:

$$\begin{aligned} \left(\sum_{i \neq j} X_i X_j^* \right) [x, y] &= \sum_{i \neq j} \delta_x \delta_y^* \cdot \langle \psi_y^T | \Pi_{f(y)+j} \Pi_{f(x)+i} | \psi_x^T \rangle \\ &= \begin{cases} 0 & f(x) = f(y) \\ \delta_x \delta_y^* \langle \psi_y^T | \psi_x^T \rangle & f(x) \neq f(y) \end{cases} \\ &= (\rho^{(T)} \circ F)[x, y], \end{aligned}$$

because $\Pi_{z_1} \Pi_{z_2} = 0$ for $z_1 \neq z_2$, $\Pi_z^2 = \Pi_z$, and $\sum_z \Pi_z = I$. Then

$$\begin{aligned} \|\rho^{(T)} \circ F\|_{tr} &= \left\| \sum_{i \neq j} X_i X_j^* \right\|_{tr} \\ &\leq \left\| \sum_{i \geq 1} (X_0 X_i^* + X_i X_0^*) \right\|_{tr} + \left\| \sum_{\substack{i \neq j \\ i, j \geq 1}} X_i X_j^* \right\|_{tr} && \text{triangle inequality} \\ &= \|X_0 (X_0^\perp)^* + X_0^\perp X_0^*\|_{tr} + \underbrace{\left\| \sum_{i, j \geq 1} X_i X_j^* - \sum_{i \geq 1} X_i X_i^* \right\|_{tr}}_{\text{only present for } |\Sigma_O| > 2} && \text{define } X_0^\perp = \sum_{i \geq 1} X_i \\ &\leq 2\|X_0 (X_0^\perp)^*\|_{tr} + \|X_0^\perp (X_0^\perp)^*\|_{tr} + \left\| \sum_{i \geq 1} X_i X_i^* \right\|_{tr} && \text{triangle inequality} \\ &\leq 2\|X_0\|_F \|X_0^\perp\|_F + \|X_0^\perp\|_F^2 + \left\| \sum_{i \geq 1} X_i X_i^* \right\|_{tr}. && \text{Theorem 5} \end{aligned}$$

We now bound the term $\left\| \sum_{i \geq 1} X_i X_i^* \right\|_{tr}$. As each $X_i X_i^*$ is positive semidefinite, the trace norm of this sum is equal to its trace. Notice that $\text{Tr}(X_i X_j^*) = 0$ if $i \neq j$, because $(X_i X_j^*)[x, x] = |\delta_x|^2 \langle \psi_x^T | \Pi_{f(x)+i} \Pi_{f(x)+j} | \psi_x^T \rangle = 0$. Thus

$$\left\| \sum_{i \geq 1} X_i X_i^* \right\|_{tr} = \text{Tr} \left(\sum_{i \geq 1} X_i X_i^* \right) = \text{Tr}(X_0^\perp (X_0^\perp)^*) = \|X_0^\perp\|_F^2.$$

Therefore $\|\rho^{(T)} \circ F\|_{tr} \leq 2ab + 2b^2$, where $a = \|X_0\|_F$ and $b = \|X_0^\perp\|_F$. We know that a, b satisfy the following constraints:

$$a^2 + b^2 = \|X_0\|_F^2 + \|X_0^\perp\|_F^2 = \sum_{x \in S} |\delta_x|^2 (\|\Pi_{f(x)} | \psi_x^T \rangle\|^2 + \|(I - \Pi_{f(x)}) | \psi_x^T \rangle\|^2) = \|\delta\|^2 = 1$$

and

$$a^2 = \|X_0\|_F^2 = \sum_{x \in S} |\delta_x|^2 \|\Pi_{f(x)} | \psi_x^T \rangle\|^2 \geq (1 - \epsilon) \sum_{x \in S} |\delta_x|^2 = 1 - \epsilon.$$

Assuming $\epsilon \leq \frac{1}{2}$, the maximum of the expression $2ab + 2b^2$ under these constraints is the boundary case $a = \sqrt{1 - \epsilon}$. Hence $\|\rho^{(T)} \circ F\|_{tr} \leq 2\sqrt{\epsilon(1 - \epsilon)} + 2\epsilon$. As we note above, the bound can be strengthened to $2\sqrt{\epsilon(1 - \epsilon)}$ if the function has Boolean output, that is $|\Sigma_O| = 2$. \square

4 Formula size

Laplante, Lee, and Szegedy [LLS06] show that the adversary method can also be used to prove classical lower bounds—they show that $\text{ADV}(f)^2$ is a lower bound on the formula size of f . A formula is circuit with AND, OR, and NOT gates with the restriction that every gate has out-degree exactly one. The size of a formula is the number of leaves and the size of a smallest formula computing f is denoted $L(f)$. We show that $\text{ADV}^\pm(f)^2$ remains a lower bound on the formula size of f .

Before we prove this statement, note that it implies a limitation of $\text{ADV}^\pm(f)$ —it is upper bounded by the square root of the formula size of f . Thus for the binary AND-OR tree—or read-once formulae in general—the largest lower bounds provable by ADV^\pm are \sqrt{n} . Laplante, Lee, and Szegedy conjecture that this is not a limitation at all—that is, they conjecture that bounded-error quantum query complexity squared is in general a lower bound on quantum query complexity. A major step has recently been taken toward proving this conjecture by [FGG07, CRŠZ07], who show that $Q_2(f) \leq L(f)^{1/2+\epsilon}$ for any $\epsilon > 0$.

We will work in the setting of Karchmer and Wigderson, who characterize formula size in terms of a communication complexity game [KW88]. Since this seminal work, nearly all formula size lower bounds have been formulated in the language of communication complexity.

Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function. Following Karchmer and Wigderson, we associate with f a relation $R_f \subseteq \{0, 1\}^n \times \{0, 1\}^n \times [n]$ where

$$R_f = \{(x, y, z) : f(x) = 0, f(y) = 1, x_z \neq y_z\}.$$

For a relation R , let $C^P(R)$ denote the number of leaves in a smallest communication protocol for R , and let $L(f)$ be the number of leaves in a smallest formula for f . Karchmer and Wigderson show the following:

Theorem 6 $L(f) = C^P(R)$.

We say that a set $S \subseteq X \times Y$ is *monochromatic* with respect to R if there exists $z \in Z$ such that $(x, y, z) \in R$ for all $(x, y) \in S$. It is well known, see for example [KN97], that a successful communication protocol for a relation $R \subseteq X \times Y \times Z$ partitions $X \times Y$ into disjoint combinatorial rectangles which are monochromatic with respect to R . Let $C^D(R)$ be the size of a smallest decomposition of $X \times Y$ into disjoint rectangles monochromatic with respect to R . Clearly, $C^D(R) \leq C^P(R)$. We are actually able to show the stronger statement that the square of $\text{ADV}^\pm(f)$ is a lower bound on the size of a smallest rectangle decomposition of R_f .

Theorem 7 $L(f) \geq C^D(R_f) \geq (\text{ADV}^\pm(f))^2$.

Proof. Laplante, Lee, and Szegedy [LLS06] show that two conditions are sufficient for a measure to lower bound formula size. The first is rectangle subadditivity—they show that the spectral norm squared is subadditive over rectangles, and this result holds for an arbitrary, possibly negative, matrix.

Lemma 8 (Laplante, Lee, Szegedy) *Let A be an arbitrary $|X|$ -by- $|Y|$ matrix and \mathcal{R} a rectangle partition of $|X| \times |Y|$. Then $\|A\|^2 \leq \sum_{R \in \mathcal{R}} \|A_R\|^2$.*

The second property is monotonicity, and here we need to modify their argument to handle negative entries. They use the property that if A, B are nonnegative matrices, and if $A \leq B$ entrywise, then $\|A\| \leq \|B\|$. In our application, however, we actually know more: if R is a rectangle monochromatic with respect to a color i , then A_R is a *submatrix* of A_i . And, for arbitrary matrices A, B , if A is a submatrix of B then $\|A\| \leq \|B\|$.

This allows us to complete the proof: let \mathcal{R} be a monochromatic partition of R_f with $|\mathcal{R}| = C^D(R_f)$. Then for any matrix A

$$\begin{aligned} \|A\|^2 &\leq \sum_{R \in \mathcal{R}} \|A_R\|^2 \leq C^D(R_f) \cdot \max_R \|A_R\|^2 \\ &\leq C^D(R_f) \cdot \max_i \|A_i\|^2. \end{aligned}$$

And so we conclude

$$L(f) \geq C^D(R_f) \geq \max_{A \neq 0} \frac{\|A\|^2}{\max_i \|A_i\|^2}.$$

□

5 Automorphism Principle

In practice, many of the functions we are interested in possess a high degree of symmetry. We now show how to take advantage of this symmetry to simplify the computation of the adversary bound. We will state this principle in a general way for possibly non-Boolean functions. Thus let Σ, T be two finite sets, and let $f : \Sigma^n \rightarrow T$ be a function. We will define a group action of $S_n \times S_\Sigma^n$ on our set of inputs Σ^n . A natural action of an element $\tau \in S_n$ on input $x \in \Sigma^n$ is to permute the indices of x . Namely, we define $\tau \cdot x = y$ to be the string where $x_i = y_{\tau(i)}$. An element $\sigma \in S_\Sigma^n = (\sigma_1, \dots, \sigma_n)$ similarly has a natural action on an input $x \in \Sigma^n$. Namely, we define $\sigma \cdot x = y$ to be the string where $y_i = \sigma_i(x_i)$. In general we can combine these actions and for $(\tau, \sigma) \in S_n \times S_\Sigma^n$ define $(\tau, \sigma) \cdot x = \tau \cdot (\sigma \cdot x)$. From now on, we will use the more convenient functional notation $(\tau, \sigma)(x) = \tau \cdot (\sigma \cdot x)$.

Definition 3 *Let $f : \Sigma^n \rightarrow T$ be a function, and $\pi \in S_n \times S_\Sigma^n$. We say that π is an automorphism of f if $f(x) \neq f(y) \Rightarrow f(\pi(x)) \neq f(\pi(y))$ for all $x, y \in \Sigma^n$.*

Note that the automorphisms of a function form a group.

Intuitively, when choosing a weight matrix Γ , it seems that pairs (x, y) and $(\pi(x), \pi(y))$ “look the same” when π is an automorphism of f and therefore should be given the same weight. The automorphism principle makes this intuition rigorous. This principle can vastly simplify the computation of the adversary bound, helping both in choosing good weight matrices, and in showing upper bounds on the adversary value.

Definition 4 Let G be a group of automorphisms for a function f . We say that G is f -transitive if for every x, y such that $f(x) = f(y)$, there is $\pi \in G$ such that $\pi(x) = y$.

Theorem 9 (Automorphism Principle) Let G be a group of automorphisms of f . There is an optimal adversary matrix Γ for which $\Gamma[x, y] = \Gamma[\pi(x), \pi(y)]$ for all $\pi \in G$ and x, y . Furthermore, if G is f -transitive then Γ has a principal eigenvector β for which $\beta[x] = \beta[y]$ whenever $f(x) = f(y)$.

Proof. Let Γ be an optimal adversary matrix for f . By normalizing as necessary, we assume that $\text{ADV}^\pm(f) = \|\Gamma\|$ —that is, $\max_i \|\Gamma \circ D_i\| = 1$. For an automorphism π , we let Γ_π be the matrix obtained from Γ by permuting the rows and columns by π , that is $\Gamma_\pi[x, y] = \Gamma[\pi(x), \pi(y)]$. Letting P be the permutation matrix representing π , where $P[x, y] = 1$ if $x = \pi(y)$ and 0 otherwise, we see that $P^T \Gamma P = \Gamma_\pi$. As P is unitary, this means that $\|\Gamma\| = \|\Gamma_\pi\|$. Notice that if π sends the index i to j , then if $x_i \neq y_i$ it follows $\pi(x)_j \neq \pi(y)_j$. Thus $P^T(\Gamma \circ D_i)P = \Gamma_\pi \circ D_j$, and these matrices also have the same spectral norm. It follows that Γ_π achieves the same adversary bound as Γ , and so is an optimal adversary matrix.

Let δ be a principal eigenvector of Γ . We may assume without loss of generality that all entries of δ are nonzero, as the rows and columns of Γ corresponding to zero entries of δ can be removed without affecting the adversary bound. We now see how δ relates to a principal eigenvector of Γ_π : note that $P^T \Gamma P (P^T \delta) = P^T \Gamma P P^T \delta = \|\Gamma\| P^T \delta$, thus $P^T \delta$ will be a principal eigenvector of Γ_π . The vector $P^T \delta$ has entries $P^T \delta[x] = \delta[\pi(x)]$. For convenience we set $\delta_\pi = P^T \delta$.

To prove the automorphism principle, we will now “average” the matrices Γ_π over $\pi \in G$ in the following way. We define a vector β as

$$\beta[x] = \sqrt{\sum_{\pi \in G} \delta[\pi(x)]^2}$$

Notice that β has norm $\sqrt{|G|}$. Now form the matrix Γ' , where

$$\Gamma'[x, y] = \frac{\sum_{\pi \in G} \Gamma[\pi(x), \pi(y)] \delta[\pi(x)]^* \delta[\pi(y)]}{\beta[x]^* \beta[y]} \quad (1)$$

We claim that Γ' is an optimal adversary matrix. The spectral norm of Γ' is at least

$$\|\Gamma'\| \geq \frac{1}{|G|} \beta^* \Gamma' \beta = \frac{1}{|G|} \sum_{\pi \in G} \|\Gamma_\pi\| = \|\Gamma\|.$$

We will now show that $\|\Gamma' \circ D_i\| \leq 1$ for all $i \in [n]$. This is equivalent to showing $I \pm \Gamma' \circ D_i \succeq 0$. As argued above, we have $I \pm \Gamma_\pi \circ D_i \succeq 0$, for all $\pi \in G$. It follows that also $\delta_\pi \delta_\pi^* \circ (I \pm \Gamma_\pi \circ D_i) \succeq 0$. Now adding these equations over $\pi \in G$ we obtain

$$\left(\sum_{\pi \in G} \delta_\pi \delta_\pi^* \right) \circ I \pm \left(\sum_{\pi \in G} \Gamma_\pi \circ \delta_\pi \delta_\pi^* \right) \circ D_i \succeq 0.$$

We can further take the Hadamard product of this matrix with the rank one matrix A , where $A[x, y] = [1/\beta[x]^* \beta[y]]$, with the result remaining positive semidefinite. This gives

$$I \pm \left(\sum_{\pi \in G} \Gamma_{\pi} \circ \delta_{\pi} \delta_{\pi}^* \right) \circ D_i \circ A = I \pm \Gamma' \circ D_i \succeq 0,$$

which concludes the proof that Γ' is an optimal adversary matrix.

Note that this argument has actually shown that β is a principal eigenvector of Γ' . Let σ be an arbitrary element of G . We have

$$\beta[\sigma(x)] = \sqrt{\sum_{\pi \in G} \delta[\pi \sigma(x)]^2} = \sqrt{\sum_{\pi \sigma^{-1} : \pi \in G} \delta[\pi(x)]^2} = \beta[x].$$

If G is f -transitive, then for every x, y with $f(x) = f(y)$, there is some σ such that $\sigma(x) = y$, and thus $\beta(x) = \beta(y)$. This proves the “furthermore” of the theorem.

Now we will show that $\Gamma'[x, y] = \Gamma'[\sigma(x), \sigma(y)]$. We have just argued that $\beta[x] = \beta[\sigma(x)]$, which gives that the denominators of these terms, defined by [Equation \(1\)](#), are equal. That the numerators are equal follows similarly since summing over $\pi \in G$ is the same as summing over $\pi \sigma^{-1} : \pi \in G$ as G is a group. \square

We single out another class of functions which arise frequently in practice and where the automorphism principle can give the adversary bound a particularly simple form.

Definition 5 *Let $f : \Sigma^n \rightarrow T$ be a function, and let $G \subseteq S_n \times S_{\Sigma}^n$ be its group of automorphisms. We say that G is index transitive if for every $i, j \in [n]$ there is an automorphism $\pi = (\sigma, \tau) \in G$ with $\sigma(i) = j$.*

Corollary 10 *If f has a group of automorphisms which is index transitive, then there is an optimal adversary matrix Γ such that $\|\Gamma \circ D_i\| = \|\Gamma \circ D_j\|$ for all $i, j \in [n]$.*

This corollary means that if a function has an automorphism group which is index transitive we can do away with the maximization in the denominator of the adversary bound—all $\Gamma \circ D_i$ will have the same spectral norm.

Proof. Let $i, j \in [n]$, and let G be the group of automorphisms of f . By assumption, there is a $\pi = (\sigma, \tau) \in G$ such that $\sigma(i) = j$. Applying [Theorem 9](#), there is an optimal adversary matrix Γ such that $\Gamma[\pi(x), \pi(y)] = \Gamma[x, y]$, for all x, y . Notice that $x_i \neq y_i$ if and only if $\pi(x)_j \neq \pi(y)_j$. Letting P be the permutation matrix representing π , it then follows that $P^T(\Gamma \circ D_i)P = \Gamma \circ D_j$, and so $\Gamma \circ D_i$ and $\Gamma \circ D_j$ have the same spectral norm. \square

6 Composition theorem

One nice property of the adversary method is that it behaves very well with respect to iterated functions. In this section we will exclusively deal with Boolean functions. For a function $f :$

$\{0, 1\}^n \rightarrow \{0, 1\}$ we define the d^{th} iteration of f , $f^d : \{0, 1\}^{n^d} \rightarrow \{0, 1\}$ recursively as $f^1 = f$ and $f^d = f \circ (f^{d-1}, \dots, f^{d-1})$ for $d > 1$. Ambainis [Amb06] shows that $\text{ADV}(f^d) \geq \text{ADV}(f)^d$. Thus by proving a good adversary bound on the base function f , one can easily obtain good lower bounds on the iterates of f . In this way, Ambainis shows a super-linear gap between the bound given by the polynomial degree of a function and the adversary method, thus separating polynomial degree and quantum query complexity.

Laplante, Lee, and Szegedy [LLS06] show a matching upper bound for iterated functions, namely that if $\text{ADV}(f) \leq a$ then $\text{ADV}(f^d) \leq a^d$. Thus we conclude that the adversary method possesses the following composition property.

Theorem 11 ([Amb06, LLS06]) *For any function $f : S \rightarrow \{0, 1\}$, with $S \subseteq \{0, 1\}^n$ and natural number $d > 0$,*

$$\text{ADV}(f^d) = \text{ADV}(f)^d.$$

Høyer, Lee, and Špalek [HLŠ05] generalize this composition theorem to functions that can be written in the form

$$h = f \circ (g_1, \dots, g_k). \quad (2)$$

They give an exact expression for the adversary bound of h in terms of the adversary bounds of f and g_i for $1 \leq i \leq k$. We will also look at the composition of the ADV^\pm bound in this general setting.

One may think of h as a two-level decision tree with the top node being labeled by a function $f : \{0, 1\}^k \rightarrow \{0, 1\}$, and each of the k internal nodes at the bottom level being labeled by a function $g_i : \{0, 1\}^{n_i} \rightarrow \{0, 1\}$. We do not require that the inputs to the inner functions g_i have the same length. An input $x \in \{0, 1\}^n$ to h is a bit string of length $n = \sum_i n_i$, which we think of as being comprised of k parts, $x = (x^1, x^2, \dots, x^k)$, where $x^i \in \{0, 1\}^{n_i}$. We may evaluate h on input x by first computing the k bits $\tilde{x}_i = g_i(x^i)$, and then evaluating f on input $\tilde{x} = (\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_k)$.

Adversary bound with costs To show their composition theorem, [HLŠ05] consider as an intermediate step a generalization of the adversary method allowing input bits to be given an arbitrary positive cost. For any function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, and any vector $\alpha \in \mathbb{R}_+^n$ of length n of positive reals, they define a quantity $\text{ADV}_\alpha(f)$ as follows:

$$\text{ADV}_\alpha(f) = \max_{\substack{\Gamma \geq 0 \\ \Gamma \neq 0}} \min_i \left\{ \alpha_i \frac{\|\Gamma\|}{\|\Gamma \circ D_i\|} \right\}.$$

We define the analogous quantity $\text{ADV}_\alpha^\pm(f)$ by enlarging the maximization over all nonzero adversary matrices. We will use the notation $\text{ADV}^{(\pm)}$ to simultaneously refer to both ADV and ADV^\pm . One may think of α_i as expressing the cost of querying the i^{th} input bit x_i . For example, x_i could be equal to the parity of α_i new input bits, or, alternatively, each query to x_i could reveal only a fraction of $1/\alpha_i$ bits of information about x_i . When $\alpha = (a, \dots, a)$ and all costs are equal to a , the new adversary bound $\text{ADV}_\alpha^{(\pm)}(f)$ reduces to $a \cdot \text{ADV}^{(\pm)}(f)$, the product of a and the adversary bound $\text{ADV}^{(\pm)}(f)$. In particular, when all costs $a = 1$ we have $Q_\epsilon(f) = \Omega(\text{ADV}_1^{(\pm)}(f))$. When α is not the all-one vector, then $\text{ADV}_\alpha^{(\pm)}(f)$ will not necessarily be a lower bound on the quantum

query complexity of f , but this quantity can still be very useful in computing the adversary bound of composed functions. We will show the following composition theorem for the nonnegative ADV bound:

Theorem 12 (Exact expression for adversary bound of composed functions) *For any function $h : S \rightarrow \{0, 1\}$ of the form $h = f \circ (g_1, \dots, g_k)$ with domain $S \subseteq \{0, 1\}^n$, and any cost function $\alpha \in \mathbb{R}_+^n$,*

$$\text{ADV}_\alpha(h) = \text{ADV}_\beta(f),$$

where $\beta_i = \text{ADV}_{\alpha^i}(g_i)$, $\alpha = (\alpha^1, \alpha^2, \dots, \alpha^k)$, and $\beta = (\beta_1, \dots, \beta_k)$.

We show that one direction of this theorem, the lower bound, also holds for the ADV^\pm bound. This is the direction which is useful for proving separations.

Theorem 13 *Let h, f, g_i be as in the previous theorem. Then*

$$\text{ADV}_\alpha^\pm(h) \geq \text{ADV}_\beta^\pm(f),$$

where $\beta_i = \text{ADV}_{\alpha^i}^\pm(g_i)$, $\alpha = (\alpha^1, \alpha^2, \dots, \alpha^k)$, and $\beta = (\beta_1, \dots, \beta_k)$.

As with the proof that ADV^\pm is a lower bound on quantum query complexity, the presence of negative entries again causes new difficulties in the proof of the composition theorem. In particular, previous proofs of composition theorems do not seem to work for ADV^\pm and we prove **Theorem 13** in a quite different manner. Also, as the dual of the ADV^\pm bound is more complicated than that of the ADV bound, we have not yet been able to show the upper bound in this theorem.

The usefulness of such a theorem is that it allows one to divide and conquer—it reduces the computation of the adversary bound for h into the disjoint subproblems of first computing the adversary bound for each g_i , and then, having determined $\beta_i = \text{ADV}_{\alpha^i}^\pm(g_i)$, computing $\text{ADV}_\beta^\pm(f)$, the adversary bound for f with costs β .

One need not compute exactly the adversary bound for each g_i to apply the theorem. Indeed, a bound of the form $a \leq \text{ADV}(g_i) \leq b$ for all i already gives some information about h .

Corollary 14 *If $h = f \circ (g_1, \dots, g_k)$ and $a \leq \text{ADV}(g_i) \leq b$ for all i , then $a \cdot \text{ADV}(f) \leq \text{ADV}(h) \leq b \cdot \text{ADV}(f)$.*

One limitation of our theorem is that we require the sub-functions g_i to act on disjoint subsets of the input bits. Thus one cannot use this theorem to compute the adversary bound of any function by, say, proceeding inductively on the structure of a $\{\wedge, \vee, \neg\}$ -formula for the function. One general situation where the theorem can be applied, however, is to read-once functions, as by definition these functions are described by a formula over $\{\wedge, \vee, \neg\}$ where each variable appears only once.

To demonstrate how **Theorem 12** can be applied, we give a simple proof of the $\Omega(\sqrt{n})$ lower bound due to Barnum and Saks [BS04] on the bounded-error quantum query complexity of read-once functions.

Corollary 15 (Barnum-Saks) *Let h be a read-once Boolean function with n variables. Then $Q_\epsilon(h) = \Omega(\sqrt{n})$.*

Proof. We prove by induction on the number of variables n that $\text{ADV}(f) \geq \sqrt{n}$. If $n = 1$ then the formula is either x_i or $\neg x_i$ and taking $\Gamma = 1$ shows the adversary bound is at least 1.

Now assume the induction hypothesis holds for read-once formulas on n variables, and let h be given by a read-once formula with $n + 1$ variables. As usual, we can push any NOT gates down to the leaves, and assume that the root gate in the formula for h is labeled either by an AND gate or an OR gate. Assume it is AND—the other case follows similarly. In this case, h can be written as $h = g_1 \wedge g_2$ where g_1 is a read-once function on $n_1 \leq n$ bits and g_2 is a read-once function on $n_2 \leq n$ bits, where $n_1 + n_2 = n + 1$. We want to calculate $\text{ADV}_{\bar{\Gamma}}(h)$. Applying [Theorem 12](#), we proceed to first calculate $\beta_1 = \text{ADV}(g_1)$ and $\beta_2 = \text{ADV}(g_2)$. By the induction hypothesis, we know $\beta_1 \geq \sqrt{n_1}$ and $\beta_2 \geq \sqrt{n_2}$. We now proceed to calculate $\text{ADV}_{\bar{\Gamma}}(h) = \text{ADV}_{(\beta_1, \beta_2)}(\text{AND})$. We set up our AND adversary matrix as follows:

	00	01	10	11
00	0	0	0	0
01	0	0	0	β_1
10	0	0	0	β_2
11	0	β_1	β_2	0

This matrix has spectral norm $\sqrt{\beta_1^2 + \beta_2^2}$, and $\|\Gamma \circ D_1\| = \beta_1$, and $\|\Gamma \circ D_2\| = \beta_2$. Thus

$$\beta_1 \frac{\|\Gamma\|}{\|\Gamma \circ D_1\|} = \beta_2 \frac{\|\Gamma\|}{\|\Gamma \circ D_2\|} = \sqrt{\beta_1^2 + \beta_2^2} \geq \sqrt{n+1}.$$

□

6.1 Composition Lemma

We now turn to the proof of the composition theorem. Given an adversary matrix Γ_f realizing the adversary bound for f and adversary matrices Γ_{g_i} realizing the adversary bound for g_i where $i = 1, \dots, k$, we build an adversary matrix Γ_h for the function $h = f \circ (g_1, \dots, g_k)$. [Lemma 16](#) expresses the spectral norm of this Γ_h in terms of the spectral norms of Γ_f and Γ_{g_i} . Moreover, if Γ_f, Γ_{g_i} are nonnegative, then Γ_h will be nonnegative.

Let Γ_f be an adversary matrix for f , i.e. a Hermitian matrix satisfying $\Gamma_f[x, y] = 0$ if $f(x) \neq f(y)$, and let δ_f be a principal eigenvector of Γ_f with unit norm. Similarly, let Γ_{g_i} be a spectral matrix for g_i and let δ_{g_i} be a principal eigenvector of unit norm, for every $i = 1, \dots, k$.

It is helpful to visualize an adversary matrix in the following way. Let $X_f = f^{-1}(0)$ and $Y_f = f^{-1}(1)$. We order the rows first by elements from X_f and then by elements of Y_f . In this way, the matrix has the following form:

$$\Gamma_f = \begin{bmatrix} 0 & \Gamma_f^{(0,1)} \\ \Gamma_f^{(1,0)} & 0 \end{bmatrix}$$

where $\Gamma_f^{(0,1)}$ is the submatrix of Γ_f with rows labeled from X_f and columns labeled from Y_f and $\Gamma_f^{(1,0)}$ is the conjugate transpose of $\Gamma_f^{(0,1)}$.

Thus one can see that an adversary matrix for a Boolean function corresponds to a (weighted) bipartite graph where the two color classes are the domains where the function takes the values 0 and 1. For $b \in \{0, 1\}$ let $\delta_{g_i}^{[b]}[x] = \delta_{g_i}[x]$ if $g_i(x) = b$ and $\delta_{g_i}^{[b]}[x] = 0$ otherwise. In other words, $\delta_{g_i}^{[b]}$ is the vector δ_{g_i} restricted to the color class b .

Before we define our composition matrix, we need one more piece of notation. Let $\Gamma_f^{(0,0)} = \|\Gamma_f\| I_{|X_f|}$, where I is a $|X_f|$ -by- $|X_f|$ identity matrix and similarly $\Gamma_f^{(1,1)} = \|\Gamma_f\| I_{|Y_f|}$.

We are now ready to define the matrix Γ_h :

Definition 6 $\Gamma_h[x, y] = \Gamma_f[\tilde{x}, \tilde{y}] \cdot \left(\bigotimes_i \Gamma_{g_i}^{(\tilde{x}_i, \tilde{y}_i)} \right) [x, y]$

Lemma 16 *Let Γ_h be as in Definition 6. Then $\|\Gamma_h\| = \|\Gamma_f\| \cdot \prod_{i=1}^k \|\Gamma_{g_i}\|$ and a principal eigenvector of Γ_h is $\delta_h[x] = \delta_f[\tilde{x}] \cdot \prod_{i=1}^k \delta_{g_i}[x^i]$.*

Proof. The more difficult direction is to show $\|\Gamma_h\| \leq \|\Gamma_f\| \cdot \prod_{i=1}^k \|\Gamma_{g_i}\|$, and we do this first. The outline of this direction is as follows:

1. We first define 2^{k+n} many vectors $\delta_{\alpha,c} \in \mathbb{C}^{2^n}$.
2. We show that each $\delta_{\alpha,c}$ is an eigenvector of Γ_h .
3. We show that $\{\delta_{\alpha,c}\}_{\alpha,c}$ span a space of dimension 2^n . This implies that every eigenvalue of Γ_h is an eigenvalue associated to at least one of the $\delta_{\alpha,c}$ as eigenvectors corresponding to different eigenvalues of a symmetric matrix are orthogonal.
4. We upper bound the absolute value of the eigenvalues corresponding to the $\delta_{\alpha,c}$ by $\|\Gamma_f\| \cdot \prod_{i=1}^k \|\Gamma_{g_i}\|$.

Let $c = (c_1, \dots, c_k)$ where $c_i \in [2^{n_i}]$ for $i = 1, \dots, k$. Let δ_{c_i} be an eigenvector of unit norm corresponding to the c_i th largest eigenvalue of Γ_{g_i} —that is $\Gamma_{g_i} \delta_{c_i} = \lambda_{c_i}(\Gamma_{g_i}) \delta_{c_i}$.

It is helpful to look at the matrix Γ_h as composed of blocks labeled by $a, b \in \{0, 1\}^k$ where the (a, b) block of the matrix consists of all x, y pairs with $\tilde{x} = a$ and $\tilde{y} = b$. Notice that the (a, b) block of Γ_h is the matrix $\Gamma_f[a, b] \cdot \bigotimes_i \Gamma_{g_i}^{(a_i, b_i)}$.

Let $\lambda_{c_i}^0(A) = \|A\|$ and $\lambda_{c_i}^1(A) = \lambda_{c_i}(A)$. We claim that $\Gamma_{g_i}^{(a_i, b_i)} \delta_{c_i}^{[b_i]} = \lambda_{c_i}^{a_i \oplus b_i}(\Gamma_{g_i}) \delta_{c_i}^{[a_i]}$. This is because if $a_i \neq b_i$ then $\Gamma_{g_i}^{(a_i, b_i)}$ is one half of the bipartite matrix Γ_{g_i} and so $\Gamma_{g_i}^{(a_i, b_i)} \delta_{c_i}^{[b_i]} = \lambda_{c_i}(\Gamma_{g_i}) \delta_{c_i}^{[a_i]}$. On the other hand, if $a_i = b_i$ then $\Gamma_{g_i}^{(a_i, b_i)} = \|\Gamma_{g_i}\| I$ and so $\Gamma_{g_i}^{(a_i, b_i)} \delta_{c_i}^{[b_i]} = \|\Gamma_{g_i}\| \delta_{c_i}^{[b_i]} = \|\Gamma_{g_i}\| \delta_{c_i}^{[a_i]}$.

Thus for the tensor product matrix $\bigotimes_i \Gamma_{g_i}^{(a_i, b_i)}$ we have that

$$\bigotimes_i \Gamma_{g_i}^{(a_i, b_i)} \otimes \delta_{c_i}^{[b_i]} = \prod_{i=1}^k \lambda_{c_i}^{a_i \oplus b_i}(\Gamma_{g_i}) \cdot \bigotimes_i \delta_{c_i}^{[a_i]}.$$

Expanding this equation gives that for every x such that $\tilde{x} = a$

$$\sum_{y: \tilde{y}=b} \bigotimes_i \Gamma_{g_i}^{(a_i, b_i)} [x, y] \cdot (\bigotimes_i \delta_{c_i}) [y] = \prod_{i=1}^k \lambda_{c_i}^{a_i \oplus b_i}(\Gamma_{g_i}) \cdot (\bigotimes_i \delta_{c_i}) [x]. \quad (3)$$

Now consider a 2^k -by- 2^k matrix A_c where

$$A_c[a, b] = \Gamma_f[a, b] \cdot \prod_{i=1}^k \lambda_{c_i}^{a_i \oplus b_i}(\Gamma_{g_i}).$$

Let α be a unit norm eigenvector of this matrix, say with eigenvalue $\mu_{\alpha,c}$. Explicitly writing out the eigenvalue equation means that for every a ,

$$\sum_b \Gamma_f[a, b] \cdot \prod_{i=1}^k \lambda_{c_i}^{a_i \oplus b_i}(\Gamma_{g_i}) \cdot \alpha[b] = \mu_{\alpha,c} \alpha[a]. \quad (4)$$

Item 1: We are ready to define our proposed eigenvectors of Γ_h . For any $c = (c_1, \dots, c_k)$ and α an eigenvector of A_c let

$$\delta_{\alpha,c}[x] = \alpha[\tilde{x}] \cdot \prod_{i=1}^k \delta_{c_i}[x^i] = \alpha[\tilde{x}] \cdot (\otimes \delta_{c_i})[x].$$

Item 2: We claim that $\delta_{\alpha,c}$ is an eigenvector of Γ_h with eigenvalue $\mu_{\alpha,c}$. This can be verified as follows: for any x ,

$$\begin{aligned} \sum_y \Gamma_h[x, y] \delta_{\alpha,c}[y] &= \sum_y \Gamma_f[\tilde{x}, \tilde{y}] \alpha[\tilde{y}] \cdot (\otimes \Gamma_{g_i}^{(\tilde{x}_i, \tilde{y}_i)})[x, y] \cdot (\otimes \delta_{c_i})[y] \\ &= \sum_b \Gamma_f[\tilde{x}, b] \alpha[b] \cdot \sum_{y: \tilde{y}=b} (\otimes \Gamma_{g_i}^{(\tilde{x}_i, \tilde{y}_i)})[x, y] \cdot (\otimes \delta_{c_i})[y] \end{aligned}$$

Applying [Equation \(3\)](#) gives

$$\begin{aligned} \sum_y \Gamma_h[x, y] \delta_{\alpha,c}[y] &= \sum_b \Gamma_f[\tilde{x}, b] \alpha[b] \cdot \prod_{i=1}^k \lambda_{c_i}^{\tilde{x}_i \oplus b_i}(\Gamma_{g_i}) \cdot (\otimes \delta_{c_i})[x] \\ &= (\otimes \delta_{c_i})[x] \cdot \sum_b \Gamma_f[\tilde{x}, b] \cdot \prod_{i=1}^k \lambda_{c_i}^{\tilde{x}_i \oplus b_i}(\Gamma_{g_i}) \alpha[b]. \end{aligned}$$

And now applying [Equation \(4\)](#) gives

$$\sum_y \Gamma_h[x, y] \delta_{\alpha,c}[y] = \mu_{\alpha,c} \alpha[\tilde{x}] \cdot (\otimes \delta_{c_i})[x] = \mu_{\alpha,c} \delta_{\alpha,c}[x].$$

Thus $\delta_{\alpha,c}$ is an eigenvector of Γ_h with eigenvalue $\mu_{\alpha,c}$. This completes the second step of the proof.

Item 3: We now claim that the vectors $\{\delta_{\alpha,c}\}_{\alpha,c}$ span \mathbb{C}^{2^n} . For a fixed c , the set of eigenvectors $\{\alpha_\ell\}_{\ell=1}^{2^k}$ of A_c forms an orthogonal basis for the space of vectors of dimension 2^k , hence there is a linear combination γ of α_ℓ 's such that $\sum_\ell \gamma_\ell \alpha_\ell = (1, 1, \dots, 1)$. Then $\sum_\ell \gamma_\ell \delta_{\alpha_\ell,c} = \otimes \delta_{c_i}$. Now, since $\{\delta_{c_i}\}_{c_i=1}^{2^{n_i}}$ form an orthogonal basis for every i , linear combinations of $\delta_{\alpha,c}$ span the whole space of dimension $2^{\sum_i n_i}$, which is the dimension of Γ_h . Hence every eigenvector of Γ_h can be expressed in this form. This completes step three of the proof.

Item 4: It now remains to show that $\mu_{\alpha,c} \leq \|\Gamma_f\| \cdot \prod_i \|\Gamma_{g_i}\|$ for every α, c . To do this, fix c and consider the matrix A_c .

$$\mu_{\alpha,c} = \alpha^* A_c \alpha = \sum_{a,b} \Gamma_f[a,b] \cdot \prod_{i=1}^k \lambda_{c_i}^{a_i \oplus b_i}(\Gamma_{g_i}) \cdot \alpha[a] \alpha[b]. \quad (5)$$

Notice that $-\|\Gamma_{g_i}\| \leq \lambda_{c_i}(\Gamma_{g_i}) \leq \|\Gamma_{g_i}\|$. Our first claim is that we can replace $\lambda_{c_i}(\Gamma_{g_i})$ by either $\|\Gamma_{g_i}\|$ or $-\|\Gamma_{g_i}\|$ in such a way that the sum in (5) does not decrease. To see this, we can first factor out $\lambda_{c_1}(\Gamma_{g_1})$ of the above sum and look at the term it multiplies. If this term is positive, then setting $\lambda_{c_1}(\Gamma_{g_1})$ to $\|\Gamma_{g_1}\|$ will not decrease the sum; on the other hand, if the term it multiplies is negative, then replacing $\lambda_{c_1}(\Gamma_{g_1})$ by $-\|\Gamma_{g_1}\|$ will not decrease the sum. We continue this process in turn with $i = 2, \dots, k$.

Let $d_i = 1$ if in this process we replaced $\lambda_{c_i}(\Gamma_{g_i})$ by $-\|\Gamma_{g_i}\|$ and $d_i = 0$ if $\lambda_{c_i}(\Gamma_{g_i})$ was replaced by $\|\Gamma_{g_i}\|$. Note that if $a_i = b_i$, then no replacement was made and the coefficient remains $\|\Gamma_{g_i}\|$. We thus now have

$$\mu_{\alpha,c} \leq \sum_{a,b} \Gamma_f[a,b] \alpha[a] \alpha[b] \cdot \prod_{i=1}^k (-1)^{d_i(a_i+b_i)} \|\Gamma_{g_i}\|, \quad (6)$$

A key fact here is that the sign of $\|\Gamma_{g_i}\|$ will be the same everywhere $a_i \neq b_i$ —the signs of entries cannot be flipped at will.

We now mimic the pattern of signs in Equation (6) by defining a new unit vector α' . Let $\alpha'[a] = \alpha[a] \prod_i (-1)^{d_i \cdot a_i}$. Then

$$\begin{aligned} \mu_{\alpha,c} &\leq \sum_{a,b} \Gamma_f[a,b] \alpha[a] \alpha[b] \cdot \prod_{i=1}^k (-1)^{d_i(a_i+b_i)} \|\Gamma_{g_i}\| \\ &= \prod_{i=1}^k \|\Gamma_{g_i}\| \sum_{a,b} \Gamma_f[a,b] \alpha'[a] \alpha'[b] \\ &\leq \|\Gamma_f\| \cdot \prod_i \|\Gamma_{g_i}\|, \end{aligned}$$

which we wished to show.

Other direction: We now show that $\|\Gamma_h\| \geq \|\Gamma_f\| \cdot \prod_{i=1}^k \|\Gamma_{g_i}\|$. Let δ_f be a principal eigenvector of Γ_f and δ_{g_i} a principal eigenvector for Γ_{g_i} for $i = 1, \dots, k$. We have already argued that $\delta_h =$

$\delta_f[\tilde{x}] \cdot \prod_{i=1}^k \delta_{g_i}[x^i]$ is an eigenvector of Γ_h whose eigenvalue is the eigenvalue of the matrix $A_{\tilde{\Gamma}}$ where

$$A_{\tilde{\Gamma}}[a, b] = \Gamma_f[a, b] \cdot \prod_{i=1}^k \|\Gamma_{g_i}\|.$$

Factoring out $\prod_{i=1}^k \|\Gamma_{g_i}\|$ from $A_{\tilde{\Gamma}}$ we are simply left with the matrix Γ_f , thus the largest eigenvalue of $A_{\tilde{\Gamma}}$ is $\|\Gamma_f\| \cdot \prod_{i=1}^k \|\Gamma_{g_i}\|$. \square

6.2 Composition lower bound

With [Lemma 16](#) in hand, it is a relatively easy matter to show a lower bound on the adversary value of the composed function h . Let $\text{ADV}^{(\pm)}$ denote either ADV or ADV^\pm .

Lemma 17 $\text{ADV}_\alpha^{(\pm)}(h) \geq \text{ADV}_\beta^{(\pm)}(f)$, where $\beta_i = \text{ADV}_{\alpha^i}^{(\pm)}(g_i)$,

Proof. Due to the maximization over all matrices Γ , the spectral bound of the composite function h is at least $\text{ADV}_\alpha^{(\pm)}(h) \geq \min_{\ell=1}^n (\alpha_\ell \|\Gamma_h\| / \|\Gamma_h \circ D_\ell\|)$, where Γ_h is defined as in [Lemma 16](#). Notice that in [Lemma 16](#), if the component matrices are nonnegative, then Gh will be as well, thus we can simultaneously treat both adversary bounds.

We compute $\|\Gamma_h \circ D_\ell\|$ for $\ell = 1, \dots, n$. Let the ℓ^{th} input bit be the q^{th} bit in the p^{th} block. Recall that

$$\Gamma_h[x, y] = \Gamma_f[\tilde{x}, \tilde{y}] \cdot \prod_{i=1}^k \Gamma_{g_i}^{(\tilde{x}_i, \tilde{y}_i)}[x^i, y^i].$$

We prove that

$$(\Gamma_h \circ D_\ell)[x, y] = (\Gamma_f \circ D_p)[\tilde{x}, \tilde{y}] \cdot (\Gamma_{g_p} \circ D_q)^{(\tilde{x}_p, \tilde{y}_p)}[x^p, y^p] \cdot \prod_{i \neq p} \Gamma_{g_i}^{(\tilde{x}_i, \tilde{y}_i)}[x^i, y^i].$$

If $x_\ell \neq y_\ell$ and $\tilde{x}_p \neq \tilde{y}_p$ then both sides are equal because all multiplications by D_p, D_q, D_ℓ are multiplications by 1. If this is not the case—that is, if $x_\ell = y_\ell$ or $\tilde{x}_p = \tilde{y}_p$ —then both sides are zero. We see this by means of two cases:

1. $x_\ell = y_\ell$: In this case the left hand side is zero due to $(\Gamma_h \circ D_\ell)[x, y] = 0$. The right hand side is also zero because
 - (a) if $\tilde{x}_p = \tilde{y}_p$ then the right hand side is zero as $(\Gamma_f \circ D_p)[\tilde{x}, \tilde{y}] = 0$.
 - (b) else if $\tilde{x}_p \neq \tilde{y}_p$ then the right hand side is zero as $(\Gamma_{g_p} \circ D_q)[x^p, y^p] = 0$.
2. $x_\ell \neq y_\ell, \tilde{x}_p = \tilde{y}_p$: The left side is zero because $\Gamma_{g_p}^{(\tilde{x}_p, \tilde{y}_p)}[x^p, y^p] = \|\Gamma_{g_p}\| I[x^p, y^p] = 0$ since $x^p \neq y^p$. The right side is also zero due to $(\Gamma_f \circ D_p)[\tilde{x}, \tilde{y}] = 0$.

Since $\Gamma_h \circ D_\ell$ has the same structure as Γ_h , by [Lemma 16](#), $\|\Gamma_h \circ D_\ell\| = \|\Gamma_f \circ D_p\| \cdot \|\Gamma_{g_p} \circ D_q\| \cdot \prod_{i \neq p} \|\Gamma_{g_i}\|$. By dividing the two spectral norms,

$$\frac{\|\Gamma_h\|}{\|\Gamma_h \circ D_\ell\|} = \frac{\|\Gamma_f\|}{\|\Gamma_f \circ D_p\|} \cdot \frac{\|\Gamma_{g_p}\|}{\|\Gamma_{g_p} \circ D_q\|}. \quad (7)$$

Since the spectral adversary maximizes over all adversary matrices, we conclude that

$$\begin{aligned} \text{ADV}_\alpha^{(\pm)}(h) &\geq \min_{\ell=1}^n \frac{\|\Gamma_h\|}{\|\Gamma_h \circ D_\ell\|} \cdot \alpha_\ell \\ &= \min_{i=1}^k \frac{\|\Gamma_f\|}{\|\Gamma_f \circ D_i\|} \cdot \min_{j=1}^{n_i} \frac{\|\Gamma_{g_i}\|}{\|\Gamma_{g_i} \circ D_j\|} \cdot \alpha_j^i \\ &= \min_{i=1}^k \frac{\|\Gamma_f\|}{\|\Gamma_f \circ D_i\|} \cdot \text{ADV}_{\alpha^i}^{(\pm)}(g_i) \\ &= \min_{i=1}^k \frac{\|\Gamma_f\|}{\|\Gamma_f \circ D_i\|} \cdot \beta_i \\ &= \text{ADV}_\beta^{(\pm)}(f), \end{aligned}$$

which we had to prove. □

6.3 Composition upper bound

The non-negative adversary bound ADV satisfies the matching upper bound $\text{ADV}_\alpha(h) \leq \text{ADV}_\beta(f)$. Interestingly, we do not know yet how to show this for the ADV^\pm bound.

We apply the duality theory of semidefinite programming to obtain an equivalent expression for ADV_α in terms of a minimization problem. We then upper bound $\text{ADV}_\alpha(h)$ by showing how to compose solutions to the minimization problems.

Definition 7 *Let $f : S \rightarrow \{0, 1\}$ be a partial boolean function, where $S \subseteq \{0, 1\}^n$, and let $\alpha \in \mathbb{R}_+^n$. The minimax bound of f with costs α is*

$$\text{MM}_\alpha(f) = \min_p \max_{\substack{x, y \\ f(x) \neq f(y)}} \frac{1}{\sum_{i: x_i \neq y_i} \sqrt{p_x(i)p_y(i)}/\alpha_i},$$

where $p : S \times [n] \rightarrow [0, 1]$ ranges over all sets of $|S|$ probability distributions over input bits, that is, $p_x(i) \geq 0$ and $\sum_i p_x(i) = 1$ for every $x \in S$.

This definition is a natural generalization of the minimax bound introduced in [\[LM04, ŠS06\]](#). As [\[ŠS06\]](#) show that the minimax bound is equal to the spectral norm formulation of the adversary method, one can similarly show that the versions of these methods with costs are equal.

Theorem 18 (Duality of adversary bounds) *For every $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $\alpha \in \mathbb{R}_+^n$,*

$$\text{ADV}_\alpha(f) = \text{MM}_\alpha(f).$$

Sketch of proof. We start with the minimax bound with costs, substitute $q_x(i)p_x(i)/\alpha_i$, and rewrite the condition $\sum_i p_x(i) = 1$ into $\sum_i \alpha_i q_x(i) = 1$. Using similar arguments as in [ŠS06], we rewrite the bound as a semidefinite program, compute its dual, and after a few simplifications, get the spectral bound with costs. \square

Lemma 19 $\text{ADV}_\alpha(h) \leq \text{ADV}_\beta(f)$.

Proof. Let p^f and p^{g_i} for $i = 1, \dots, k$ be optimal sets of probability distributions achieving the minimax bounds. Thus using **Theorem 18** we have

$$\begin{aligned} \text{ADV}_\beta(f) &= \max_{\substack{x,y \\ f(x) \neq f(y)}} \frac{1}{\sum_{i: x_i \neq y_i} \sqrt{p_x^f(i)p_y^f(i)/\beta_i}}, \\ \text{ADV}_{\alpha^i}(g_i) &= \max_{\substack{x,y \\ g_i(x) \neq g_i(y)}} \frac{1}{\sum_{j: x_j \neq y_j} \sqrt{p_x^{g_i}(j)p_y^{g_i}(j)/\alpha_j^i}}. \end{aligned}$$

Define the set of probability distributions p^h as $p_x^h(\ell) = p_{\tilde{x}}^f(i)p_{\tilde{y}}^{g_i}(j)$, where the ℓ^{th} input bit is the j^{th} bit in the i^{th} block. This construction was used by Laplante, Lee, and Szegedy [LLS06] to prove a similar bound in the stronger setting where the sub-functions g_i can act on the same input bits. We claim that p^h witnesses that $\text{ADV}_\alpha(h) \leq \text{ADV}_\beta(f)$:

$$\begin{aligned} \text{ADV}_\alpha(h) &\leq \max_{\substack{x,y \\ h(x) \neq h(y)}} \frac{1}{\sum_{\ell: x_\ell \neq y_\ell} \sqrt{p_x^h(\ell)p_y^h(\ell)/\alpha_\ell}} \\ &= 1 / \min_{\substack{x,y \\ h(x) \neq h(y)}} \sum_{\ell: x_\ell \neq y_\ell} \sqrt{p_{\tilde{x}}^f(i)p_{\tilde{y}}^f(i)} \sqrt{p_{x^i}^{g_i}(j)p_{y^i}^{g_i}(j)/\alpha_j^i} \\ &= 1 / \min_{\substack{\tilde{x}, \tilde{y} \\ f(\tilde{x}) \neq f(\tilde{y})}} \sum_i \sqrt{p_{\tilde{x}}^f(i)p_{\tilde{y}}^f(i)} \min_{\substack{x^i, y^i \\ g_i(x^i) = \tilde{x}_i \\ g_i(y^i) = \tilde{y}_i}} \sum_{j: x_j^i \neq y_j^i} \sqrt{p_{x^i}^{g_i}(j)p_{y^i}^{g_i}(j)/\alpha_j^i} \\ &\leq 1 / \min_{\substack{\tilde{x}, \tilde{y} \\ f(\tilde{x}) \neq f(\tilde{y})}} \sum_{i: \tilde{x}_i \neq \tilde{y}_i} \sqrt{p_{\tilde{x}}^f(i)p_{\tilde{y}}^f(i)} \min_{\substack{x^i, y^i \\ g_i(x^i) \neq g_i(y^i)}} \sum_{j: x_j^i \neq y_j^i} \sqrt{p_{x^i}^{g_i}(j)p_{y^i}^{g_i}(j)/\alpha_j^i} \\ &= 1 / \min_{\substack{\tilde{x}, \tilde{y} \\ f(\tilde{x}) \neq f(\tilde{y})}} \sum_{i: \tilde{x}_i \neq \tilde{y}_i} \sqrt{p_{\tilde{x}}^f(i)p_{\tilde{y}}^f(i)} / \text{ADV}_{\alpha^i}(g_i) \\ &= 1 / \min_{\substack{\tilde{x}, \tilde{y} \\ f(\tilde{x}) \neq f(\tilde{y})}} \sum_{i: \tilde{x}_i \neq \tilde{y}_i} \sqrt{p_{\tilde{x}}^f(i)p_{\tilde{y}}^f(i)} / \beta_i \\ &= \text{ADV}_\beta(f), \end{aligned}$$

where the second inequality follows from that fact that we have removed $i : \tilde{x}_i = \tilde{y}_i$ from the sum and the last equality follows from **Theorem 18**. \square

7 Examples

In this section, we look at some examples to see how negative weights can help to achieve larger lower bounds. We consider two examples in detail: a 4-bit function giving the largest known separation between the polynomial degree and quantum query complexity, and a 6-bit function breaking the certificate complexity and property testing barriers.

To help find good adversary matrices, we implemented both adversary bounds as semidefinite programs and used the convex optimization package SeDuMi [sed06] for Matlab. Using these programs, we tested both ADV and ADV^\pm bounds for all 222 functions on 4 or fewer variables which are not equivalent under negation of output and input variables and permutation of input variables (see sequence number A000370 in [Slo]). The ADV^\pm bound is strictly larger than the ADV bound for 128 of these functions. The source code of our semidefinite programs and more examples can be downloaded from [HLŠ06].

7.1 Ambainis function

In order to separate quantum query complexity and polynomial degree, Ambainis defines a Boolean function $f : \{0, 1\}^4 \rightarrow \{0, 1\}$ which is one if and only if the four input bits are sorted¹, that is they are either in a non-increasing or non-decreasing order. This function has polynomial degree 2, and an adversary bound of 2.5. Thus by the composition theorem for the nonnegative adversary method, Ambainis obtains a separation between quantum query complexity and polynomial degree of $Q_\epsilon(f^d) = \Omega(\text{deg}(f^d)^{1.321})$. We have verified that this function indeed gives the largest separation between adversary bounds and polynomial degree over all functions on 4 or fewer variables.

In the next theorem, we construct an adversary matrix with negative weights which shows that $\text{ADV}^\pm(f) \geq 2.5135$. Using the composition theorem [Theorem 13](#) we obtain $\text{ADV}^\pm(f) \geq \text{ADV}(f)^{1.005}$ and improve the separation between quantum query complexity and polynomial degree to $Q_\epsilon(f^d) = \Omega(\text{deg}(f^d)^{1.3296})$.

Theorem 20 *Let $f : \{0, 1\}^4 \rightarrow \{0, 1\}$ be Ambainis’ function. Then $\text{ADV}^\pm(f) \geq 2.5135$.*

Proof. Following our own advice, to design a good adversary matrix the first thing we do is look at the automorphisms of the function. Notice that the element $g = (1432) \times ((01), \text{id}, \text{id}, \text{id}) \in S_4 \times S_{\{0,1\}^4}$ preserves the property of the bits being ordered, and thus also the function value. We are using cycle notation here, so for example, $(1432) \times ((01), \text{id}, \text{id}, \text{id}) \cdot 0000 = 0001$. Let G be the group generated by g . As g has order eight, this group is isomorphic to \mathbb{Z}_8 . The group G is both f -transitive and index transitive, thus we know that the uniform vector will be a principal eigenvector of our eventual adversary matrix, and that all $\Gamma \circ D_i$ will be unitarily equivalent.

We now construct our adversary matrix. If $f(x) = f(y)$ we set $\Gamma[x, y] = 0$. So now we just consider the “interesting” part of the adversary matrix with rows labeled by inputs which map to

¹The function was first described in this way by Laplante, Lee, and Szegedy [LLS06]. The function defined by Ambainis [Amb06] can be obtained from this function by exchanging the first and third input bits and negating the output.

one, and columns labeled by inputs which map to 0. To highlight the group structure in the matrix, we let the i^{th} row be $g^i(0000)$ and similarly let the i^{th} column be $g^i(0010)$.

It turns out there are four types of pairs (x, y) which are not equivalent under automorphism. We let $\Gamma[x, y] = a$ if (x, y) have Hamming distance one. If (x, y) have Hamming distance 3, they are also related by automorphism, and we set $\Gamma[x, y] = d$ in this case. There are two automorphism types for (x, y) pairs which have Hamming distance 2. If they differ on bits which are both sensitive or both not sensitive, we set $\Gamma[x, y] = b$; otherwise, we set $\Gamma[x, y] = c$. The adversary matrix then looks as follows:

	0010	0101	1011	0110	1101	1010	0100	1001
0000	a	c	d	b	d	c	a	b
0001	b	a	c	d	b	d	c	a
0011	a	b	a	c	d	b	d	c
0111	c	a	b	a	c	d	b	d
1111	d	c	a	b	a	c	d	b
1110	b	d	c	a	b	a	c	d
1100	d	b	d	c	a	b	a	c
1000	c	d	b	d	c	a	b	a

As we have remarked, all $\Gamma \circ D_i$ matrices are equivalent up to permutation, and it can be shown that they consist of four 4-by-4 disjoint blocks, each of these blocks being some permutation of rows and columns of the following matrix B :

$$B = \begin{pmatrix} c & b & d & d \\ b & c & d & a \\ d & d & c & b \\ d & a & b & c \end{pmatrix}. \quad (8)$$

The particular block B above is one of the four blocks of $\Gamma \circ D_1$ with columns indexed by zero-inputs 0010, 0100, 0101, 0110, and rows indexed by one-inputs 1000, 1110, 1111, 1100. A principal eigenvector of Γ is given by the uniform vector, which has eigenvalue $2(a + b + c + d)$. Thus our optimization problem is to maximize $2(a + b + c + d)$ while keeping $\|B\| \leq 1$. The optimal setting of the four variables can be found numerically by semidefinite programming and is the following:

	ADV	ADV $^\pm$
a	3/4	0.5788
b	1/2	0.7065
c	0	0.1834
d	0	-0.2120
λ	5/2	2.5135

(9)

The eigenvalues of $\Gamma \circ D_i$ are $\{1, 1, \frac{1}{4}, \frac{1}{4}\}$, and the eigenvalues of $\Gamma^\pm \circ D_i$ are $\{1, 1, -1, -0.2664\}$. Both spectral bounds are tight due to the existence of matching dual solutions; we, however, omit them here.

□

7.2 Breaking the certificate complexity barrier

We now consider a function on six bits. We will consider this function in two guises. We first define a partial function f to show that ADV^\pm can break the property testing barrier. We then extend this partial function to a total monotone function g which gives a larger separation between the ADV and ADV^\pm bounds, and also shows that ADV^\pm can break the certificate complexity barrier.

We define the partial function f on six bits as follows:

- The zero inputs of f are: 111000, 011100, 001110, 100110, 110010, 101001, 010101, 001011, 100101, 010011.
- The one inputs of f are: 000111, 100011, 110001, 011001, 001101, 010110, 101010, 110100, 011010, 101100.

Notice that f is defined on all inputs with Hamming weight three, and only on these inputs. This function is inspired by a function defined by Kushilevitz which appears in [NW95] and is also discussed by Ambainis [Amb06]. Kushilevitz's function has the same behavior as the above on inputs of Hamming weight three; it is additionally defined to be 0 on inputs with Hamming weight 0, 4, or 5, and to be 1 on inputs with Hamming weight 1, 2, or 6.

All zero inputs of f have Hamming distance at least 2 from any one input, thus the relative Hamming distance between any zero and one input is $\epsilon = 1/3$. In [Theorem 21](#) we show that $\text{ADV}^\pm(f) \geq 2 + 3\sqrt{5}/5 \approx 3.341$. This implies $\text{ADV}^\pm(f) \geq (1/\epsilon(f))^{1.098}$, and as both bounds compose we obtain $\text{ADV}^\pm(f^d) \geq (1/\epsilon(f^d))^{1.098}$. This shows that the property testing barrier does not apply to ADV^\pm as it does to ADV . The relative Hamming distance $\epsilon(f^d)$, however, goes to zero when d increases. We don't know of an asymptotic separation for constant ϵ .

We also consider a monotone extension of f to a total function, denoted g . It is additionally defined to be 0 on inputs with Hamming weight 0, 1, or 2, and to be 1 on inputs with Hamming weight 4, 5, or 6. Recall that the maxterms of a monotone Boolean function are the maximal, under subset ordering, inputs x which evaluate to 0, and similarly the minterms are the minimal inputs which evaluate to 1. The zero inputs of f become maxterms of g and the one inputs become minterms. Since f is defined on all inputs with Hamming weight three, g is a total function. The extended function g is at least as hard as its sub-function f , hence $\text{ADV}^\pm(g) \geq \text{ADV}^\pm(f)$. The 0-certificates of g are given by the location of 0's in the maxterms and the 1-certificates are given by the location of 1's in the minterms, thus $C_0(g) = C_1(g) = 3$. Both bounds compose thus $C_0(g^d) = C_1(g^d) = 3^d$.

Applying the composition theorem [Theorem 13](#) we obtain $\text{ADV}^\pm(g^d) \geq (C_0(g)C_1(g))^{0.549}$. As $\text{ADV}(h) \leq \sqrt{C_0(h)C_1(h)}$ for a total function h , we also conclude $\text{ADV}^\pm(g^d) \geq \text{ADV}(g^d)^{1.098}$.

Theorem 21 $\text{ADV}^\pm(f) \geq 2 + 3\sqrt{5}/5$.

Proof. To design an adversary matrix for f , we again first consider its automorphisms. The automorphism group G of f contains a subgroup isomorphic to A_5 , the alternating group on five elements. As we have listed the zero and one instances of the function, one can easily see that the permutation (12345), in cycle notation, is an automorphism. This automorphism fixes the sixth

bit. It turns out that for every $1 \leq i \leq 6$ there is an automorphism of f of order 5 which fixes the i^{th} bit. Here are some examples: (25643), (15643), (12456), (16235), (14362), (12345). Taking the closure of these elements gives a group isomorphic to A_5 . This group is f -transitive and index transitive, thus we know that the uniform eigenvector will be a principal eigenvector of our eventual adversary matrix, and that all $\Gamma \circ D_i$ will have the same spectral norm.

Any two pairs (x, y) and (x', y') with the same Hamming distance are related by automorphism, thus the (x, y) entry of Γ will only depend on Hamming distance. As all valid inputs to f have Hamming weight three, the Hamming distance between x and y is even and is either two, four, or six. We label the matrix entries a, b, c respectively for Hamming distances two, four, six.

	000	100	110	011	001	010	101	110	011	101
	111	011	001	001	101	110	010	100	010	100
111000	c	b	a	a	b	b	a	a	a	a
011100	b	c	b	a	a	a	b	a	a	a
001110	a	b	c	b	a	a	a	b	a	a
100110	a	a	b	c	b	a	a	a	b	a
110010	b	a	a	b	c	a	a	a	a	b
101001	b	a	a	a	a	c	a	b	b	a
010101	a	b	a	a	a	a	c	a	b	b
001011	a	a	b	a	a	b	a	c	a	b
100101	a	a	a	b	a	b	b	a	c	a
010011	a	a	a	a	b	a	b	b	a	c

Notice that all rows have the same sum, thus the uniform vector is an eigenvector with eigenvalue $6a + 3b + c$.

From this ordering of rows and columns, one can easily read off the matrix $\Gamma \circ D_6$. This is a block diagonal matrix with each block equal, up to permutation, to a matrix B :

$$B = \begin{pmatrix} c & b & a & a & b \\ b & c & b & a & a \\ a & b & c & b & a \\ a & a & b & c & b \\ b & a & a & b & c \end{pmatrix}.$$

Our optimization problem then becomes: maximize $6a + 3b + c$ while keeping the spectral norm of B at most one. As B is a sum of circulants, its eigenvalues will be $c + b\omega^k + a\omega^{2k} + a\omega^{3k} + b\omega^{4k}$, for $0 \leq k \leq 4$, where ω is a 5th root of unity. An optimal solution is obtained by setting $a = (1 + \sqrt{5})/5$, $b = (1 - \sqrt{5})/5$, $c = 1/5$. This makes the eigenvalues of B equal to $\{1, 1, 1, -1, -1\}$, while $6a + 3b + c = 2 + 3\sqrt{5}/5$. \square

8 Conclusion

Breaking the certificate complexity and property testing barriers opens the possibility that ADV^\pm can prove better lower bounds where we know ADV cannot. Salient examples are element distinct-

ness, the collision problem, and triangle finding. For element distinctness, the best bound provable by the standard adversary method is $O(\sqrt{n})$ while the polynomial method is able to prove a tight lower bound of $\Omega(n^{2/3})$ [AS04]. For the collision problem, the adversary method is only able to prove a constant lower bound while the polynomial method again proves a tight lower bound of $\Omega(n^{1/3})$ [AS04]. Finally, for the problem of determining if a graph contains a triangle, the best bound provable by the adversary method is $O(n)$ and the best known algorithm is $O(n^{1.3})$ [MSS05].

It is also interesting to determine what limitations our new adversary method might face. The only limitation we are aware of is that the square of ADV^\pm is a lower bound on formula size. This is probably not a limitation, however, as [FGG07, CRŠZ07] have recently taken a major step towards proving the conjecture of Laplante, Lee, and Szegedy that the square of bounded-error quantum query complexity is in general a lower bound on formula size.

Acknowledgements

We would like to thank Aram Harrow and Umesh Vazirani for interesting discussions on the topics of this paper, and Ronald de Wolf for many valuable comments on an earlier draft.

References

- [Amb02] A. Ambainis. Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64(4):750–767, 2002. Earlier version in STOC’00.
- [Amb06] A. Ambainis. Polynomial degree vs. quantum query complexity. *Journal of Computer and System Sciences*, 72(2):220–238, 2006. Earlier version in FOCS’03.
- [AS04] 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.
- [BBC⁺01] 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.
- [Bha97] R. Bhatia. *Matrix Analysis*. Springer-Verlag, 1997.
- [BS04] 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.
- [BSS03] H. Barnum, M. Saks, and M. Szegedy. Quantum decision trees and semidefinite programming. In *Proc. of 18th IEEE Complexity*, pages 179–193, 2003.
- [BW02] H. Buhrman and R. de Wolf. Complexity measures and decision tree complexity: A survey. *Theoretical Computer Science*, 288(1):21–43, 2002.

- [CRŠZ07] A. M. Childs, B. W. Reichardt, R. Špalek, and S. Zhang. Every NAND formula of size N can be evaluated in time $O(N^{1/2+\epsilon})$ on a quantum computer. [quant-ph/0703015](https://arxiv.org/abs/quant-ph/0703015), 2007.
- [FGG07] E. Farhi, J. Goldstone, and S. Gutmann. A quantum algorithm for the Hamiltonian NAND tree. [quant-ph/0702144](https://arxiv.org/abs/quant-ph/0702144), 2007.
- [Gro96] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proc. of 28th ACM STOC*, pages 212–219, 1996.
- [HLŠ05] P. Høyer, T. Lee, and R. Špalek. Tight adversary bounds for composite functions. [quant-ph/0509067](https://arxiv.org/abs/quant-ph/0509067), 2005.
- [HLŠ06] P. Høyer, T. Lee, and R. Špalek. Source codes of semidefinite programs for ADV^\pm . <http://www.ucw.cz/~robert/papers/adv/>, 2006.
- [HŠ05] P. Høyer and R. Špalek. Lower bounds on quantum query complexity. *EATCS Bulletin*, 87:78–103, October, 2005.
- [KN97] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [KW88] M. Karchmer and A. Wigderson. Monotone connectivity circuits require super-logarithmic depth. In *Proc. of 20th ACM STOC*, pages 539–550, 1988.
- [LLS06] S. Laplante, T. Lee, and M. Szegedy. The quantum adversary method and classical formula size lower bounds. *Computational Complexity*, 15:163–196, 2006. Earlier version in Complexity’05.
- [LM04] S. Laplante and F. Magniez. Lower bounds for randomized and quantum query complexity using Kolmogorov arguments. In *Proc. of 19th IEEE Complexity*, pages 294–304, 2004.
- [MSS05] F. Magniez, M. Santha, and M. Szegedy. Quantum algorithms for the triangle problem. In *Proc. of 16th ACM-SIAM SODA*, pages 1109–1117, 2005.
- [NC00] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [NW95] N. Nisan and A. Wigderson. A note on rank vs. communication complexity. *Combinatorica*, 15(4):557–566, 1995.
- [sed06] SeDuMi version 1.1. available from <http://sedumi.mcmaster.ca/>, 2006.
- [Sho97] P. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997. Earlier version in FOCS’94.

- [Slo] N. Sloane. On-line encyclopedia of integer sequences. <http://www.research.att.com/~njas/sequences/>.
- [ŠS06] R. Špalek and M. Szegedy. All quantum adversary methods are equivalent. *Theory of Computing*, 2(1):1–18, 2006. Earlier version in ICALP’05.
- [Zha05] S. Zhang. On the power of Ambainis’s lower bounds. *Theoretical Computer Science*, 339(2–3):241–256, 2005. Earlier version in ICALP’04.