Lower bounds for randomized and quantum query complexity using Kolmogorov arguments

Sophie Laplante[*]
LRI, Université Paris-Sud
laplante@lri.fr

Frédéric Magniez[*]
LRI, CNRS
magniez@lri.fr

Abstract

We prove a very general lower bound technique for quantum and randomized query complexity, that is easy to prove as well as to apply. To achieve this, we introduce the use of Kolmogorov complexity to query complexity. Our technique generalizes the weighted, unweighted methods of Ambainis, and the spectral method of Barnum, Saks and Szegedy. As an immediate consequence of our main theorem, it can be shown that adversary methods can only prove lower bounds for boolean functions f in $O(\min(\sqrt{nC_0(f)}, \sqrt{nC_1(f)}))$, where C_0, C_1 is the certificate complexity, and n is the size of the input. We also derive a general form of the ad hoc weighted method used by Høyer, Neerbek and Shi to give a quantum lower bound on ordered search and sorting.

1. Introduction

1.1. Overview

In this paper, we study lower bounds for randomized and quantum query complexity. In the query model, the input is accessed using oracle queries, and the query complexity of an algorithm is the number of calls to the oracle. Since it is difficult to obtain lower bounds on time directly, the query model is often used to prove concrete lower bounds, in classical as well as quantum computation.

The two main tools for proving lower bounds of randomized query complexity, the polynomial method [7] and the adversary method [2], were successfully extended to quantum computation. In the randomized setting, the adversary method is most often applied using Yao's minimax principle [18]. Using a different approach, which introduces the notion of quantum adversaries, Ambainis developed a general scheme in which it suffices to analyze the combinatorial properties of the function in order to obtain a quantum lower bound. Recently, Aaronson [1] brought these combinatorial properties back to randomized computation, using Yao's minimax principle.

The most general method for proving lower bounds in quantum query complexity is the semidefinite programming method of Barnum, Saks and Szegedy [6]. This method is in fact an exact characterization of the query complexity. However, the method is so general as to be very difficult to apply to obtain concrete lower bounds. Barnum, Saks and Szegedy gave a weaker method derived from the semidefinite programming approach, using weight matrices and their largest eigenvalue. This spectral method can be thought of as a generalization of Ambainis' unweighted method. Other generalizations of Ambainis' unweighted method have been previously introduced [5, 3]. All of them use a weight function on the instances. The difficulty in applying these methods is finding a good weight function on the instances. Høyer, Neerbek and Shi [14] were the first to use such weight assignments to prove lower bounds for searching in ordered lists and sorting. Their *ad hoc* method, though similar in many respects, does not fall into setting of the weighted method of Ambainis [3].

This paper presents a new, very general adversary technique (**Theorem 1**) to prove lower bounds in quantum and randomized query complexity. We believed that this technique is simpler to prove and to apply. It is based on the framework of Kolmogorov complexity. This framework has proven to be very useful for proving negative results in other models of computation, for example for number of rounds and length of advice in random-self-reductions in [12, 4]. The techniques we use here are an adaptation of those techniques to the framework of query complexity. We expect that this framework will not only prove to be useful for negative results in other quantum models of computation, for instance, communication complexity, but also for finer analysis of query complexity in terms of the number of rounds of queries.

[*] LRI, UMR 8623 CNRS and Université Paris–Sud, 91405 Orsay, France. Partially supported by the EU 5th framework programs RESQ IST-2001-37559 and RAND-APX IST-1999-14036, and by ACI Cryptologie CR/02 02 0040 and ACI Sécurité Informatique 03 511 grants of the French Research Ministry.

The proof of Theorem 1 is in two parts. The first part (**Divergence Lemma**) shows how fast the computations can diverge when they start on different inputs. This part depends on the model of computation (randomized or quantum). The quantum case of this lemma was first proven by Ambainis [2]. The second part (**Query Information Lemma**) does not depend on the model of computation. It establishes the relationship between the Kolmogorov complexity of individual positions of the input, and the probability that a given algorithm makes a query to this position. Whereas Aaronson [1] used a different approach to prove a version of Ambainis' method for randomized algorithms, here we use the same framework to establish lower bounds for both quantum and randomized query complexities (QQC and RQC).

We show that our method encompasses all previous adversary methods, including the quantum and randomized weighted methods [3, 1] (**Theorem 2**) and the spectral method [6] (**Theorem 3**). As an immediate consequence of our main theorem (observed by Troy Lee), our method can only prove lower bounds for boolean functions in $O(\min(\sqrt{nC_0(f)}, \sqrt{nC_1(f)}))$, where C_0 and C_1 is the certificate complexity of negative and positive instances, respectively, of f, and n is the size of the input (**Theorem 4**). Prior to our work, it was known [3] that the unweighted Ambainis method [2, Theorem 5.1] could not prove bounds better than $\Omega(\sqrt{C_0(f)C_1(f)})$; Szegedy [17] also proved independently that the semidefinite programming method could not prove lower bounds better than $O(\min(\sqrt{nC_0(f)}, \sqrt{nC_1(f)}))$, and Zhang [19] proved the same thing for Ambainis' weighted method.

We also give a generalization (**Theorem 5**) of the *ad hoc* proofs of Høyer, Neerbek and Shi [14] as a corollary of our method. For this we introduce a new distance scheme. This new scheme separates the quantum part from the combinatorial part of these *ad hoc* proofs. Using it, we prove the lower bound of [14] using only combinatorial arguments. We end the paper by giving some applications of our method to prove lower bounds for some graph properties: bipartiteness (**Theorem 7**) and connectivity (**Theorem 6**). The lower bound on connectivity was proven in [11], and the one on bipartiteness by Dürr (personal communication) and independently in [19]. We reprove it here to illustrate the simplicity of our method.

1.2. Main result

The conditional Kolmogorov complexity $\mathsf{K}(a|b)$ (defined formally in Section 2.1) is the length of the shortest program which prints a given b as input. Our main result is stated in terms of $\mathsf{K}(i|x, A)$ and $\mathsf{K}(i|y, A)$, where x, y are inputs for which $f(x) \neq f(y)$, i is an index into the inputs where $x_i \neq y_i$, and A is an algorithm for f.

Theorem 1. *There exists a constant $C > 0$ such that the following holds. Let Σ be a finite set, let $n \geq 1$ be an integer, and let $S \subseteq \Sigma^n$ and S' be sets. Let $f : S \to S'$. Let A be an algorithm that for all $x \in S$ computes f, with bounded error ε and at most T queries to the input. Then for every $x, y \in S$ with $f(x) \neq f(y)$:*

1. *If A is a quantum algorithm then*
$$T \geq C \times \frac{1-2\sqrt{\varepsilon(1-\varepsilon)}}{\sum_{i:x_i\neq y_i} \sqrt{2^{-\mathsf{K}(i|x,A)-\mathsf{K}(i|y,A)}}};$$
2. *If A is a randomized algorithm then*
$$T \geq C \times \frac{1-2\varepsilon}{\sum_{i:x_i\neq y_i} \min\left(2^{-\mathsf{K}(i|x,A)}, 2^{-\mathsf{K}(i|y,A)}\right)}.$$

We briefly describe the intuition behind the proof of Theorem 1. Consider an algorithm that purports to compute f, presented with two inputs x, y that lead to different outputs. The algorithm must query those positions where x and y differ with average probability of the order of $\frac{1}{T}$, or it will not successfully compute the function. On the other hand, the queries that are made with high average probability can be described succinctly given the input and the algorithm, using the Shannon-Fano code. If we exhibit a pair of strings x, y for which there is no succinct description of any of the positions where x and y differ, then the number of queries must be large.

The same reasoning can be applied to classical and to quantum computing; the only difference is how fast two different input states cause the outputs to diverge to different outcomes.

To conclude the introduction we give a very simple application, for Grover search.

Example 1. Fix n and a quantum algorithm A for Grover search for instances of length n. Let z be a binary string of length $\log n$, with $\mathsf{K}(z|A) \geq \log n$. Let j be the integer between 0 and $n-1$ whose binary expansion is z. Consider x, the all 0's string, and let y be everywhere 0 except at position $i = j+1$, where it is 1. Then $\mathsf{K}(i|x, A) \geq \log n - O(1)$ and $\mathsf{K}(i|y, A) \geq 0$, therefore, $\mathsf{QQC}(\textsc{Search}) = \Omega(\sqrt{n})$.

2. Preliminaries

2.1. Kolmogorov complexity

We use a few standard results in Kolmogorov complexity and information theory in this paper. We briefly review these here. The reader is invited to consult standard textbooks such as [15] for more background on Kolmogorov complexity, and [8] for more on information theory. We denote the length of a finite string x by $|x|$. We assume that the Turing machine's alphabet is the same finite alphabet as the alphabet used to encode instances of the function under consideration. Letters x, y typically represent instances; i is an index into the representation of the instance; and p, q are probability distributions. Programs are denoted P, and the output of a Turing machine M on input x is written $M(x)$. When there are multiple inputs, we assume that a standard encoding of tuples is used.

Definition 1. *Let M be a Turing machine. Let x and y be finite strings.*

1. *The* Kolmogorov complexity of x given y with respect to M *is denoted $\mathsf{C}_M(x|y)$, and defined as follows:*

$$\mathsf{C}_M(x|y) = \min(|P| \textit{ such that } M(P, y) = x).$$

2. *A set of strings is* prefix-free *if no string is a prefix of another in the set.*
3. *The* prefix-free Kolmogorov complexity of x given y with respect to M *is denoted $\mathsf{K}_M(x|y)$, and defined as follows:*

$$\mathsf{K}_M(x|y) = \min(|P| \textit{ such that } M(P, y) = x),$$

where P is taken in some fixed prefix-free set.

In the rest of the paper M is some fixed universal Turing machine, and we will write C and K instead of C_M and K_M. When y is the empty string, we write $\mathsf{K}(x)$ instead of $\mathsf{K}(x|y)$.

Proposition 1. *There exists a constant $c \geq 0$ such that for every finite string σ,*

$$\mathsf{K}(x|\sigma) \leq \mathsf{K}(x) + c, \textit{ and}$$

$$\mathsf{K}(x) \leq \mathsf{K}(\sigma) + K(x|\sigma) + c.$$

Proposition 2 (Kraft's inequality). *Let S be any prefix-free set of finite strings. Then $\sum_{x \in S} 2^{-|x|} \leq 1$.*

Proposition 3 (Shannon's coding theorem). *Consider a source $\mathcal{S}$ of finite strings where x occurs with probability $p(x)$. Then for any code for $\mathcal{S}$, the average code length is bounded below by the entropy of the source, that is, if x is encoded by the code word $c(x)$ of length $|c(x)|$, $H(\mathcal{S}) = \sum_{x:p(x)\neq 0} p(x)\log(\frac{1}{p(x)}) \leq \sum_{x:p(x)\neq 0} p(x)|c(x)|$.*

Lemma 1. *Let $\mathcal{S}$ be a source as above. Then for any fixed finite string σ, there exists a string x such that $p(x) \neq 0$ and $\mathsf{K}(x|\sigma) \geq \log(\frac{1}{p(x)})$.*

Proof. By Shannon's coding theorem,

$$H(\mathcal{S}) = \sum_{x:p(x)\neq 0} p(x)\log(\tfrac{1}{p(x)}) \leq \sum_{x:p(x)\neq 0} p(x)\mathsf{K}(x|\sigma),$$

because $\mathsf{K}(x|\sigma)$ is the length of an encoding of x. Therefore there exists x such that $p(x) \neq 0$ and $\mathsf{K}(x) \geq \log(\frac{1}{p(x)})$. □

The Shannon-Fano code is a prefix-free code that encodes each word x with $p(x) \neq 0$, using $\lceil \log(\frac{1}{p(x)}) \rceil$ bits. We will write $\log(\frac{1}{p(x)})$ to simplify notation. The code can easily be computed given a description of the probability distribution. This allows us to write the following proposition, where $\mathsf{K}(x|\mathcal{S})$ means the prefix-free Kolmogorov complexity of x given a finite description of $\mathcal{S}$.

Proposition 4 (Shannon-Fano code). *There exists a constant $c \geq 0$, such that for every source $\mathcal{S}$ as above, for all x such that $p(x) \neq 0$, $\mathsf{K}(x|\mathcal{S}) \leq \log(\frac{1}{p(x)}) + c$.*

We shall also use the following bound on conditional Kolmogorov complexity.

Proposition 5. *There is a constant $c \geq 0$ such that for any three strings x, y, z,*

$$\begin{aligned} \mathsf{K}(z|x) \quad \geq \quad & \mathsf{K}(x,y) - \mathsf{K}(x) - \mathsf{K}(y|z,x) + \\ & \mathsf{K}(z|x,y,\mathsf{K}(x,y)) - c. \end{aligned}$$

Proof. Using [15, Theorem 3.9.1, page 232], there is a constant $c_1 \geq 0$ such that

$$|\mathsf{K}(a,b) - \mathsf{K}(a) - \mathsf{K}(b|a,\mathsf{K}(a))| \leq c_1.$$

Substituting x, y for a and z for b:

$$\begin{aligned} \mathsf{K}(x,y) + \mathsf{K}(z|x,y,\mathsf{K}(x,y)) - c_1 &\leq \mathsf{K}(x,y,z) \\ &\leq \mathsf{K}(x) + \mathsf{K}(z|x) + \mathsf{K}(y|z,x) + c_2, \end{aligned}$$

which gives the result. □

2.2. Query models

The quantum query model was implicitly introduced by Deutsch, Jozsa, Simon and Grover [9, 10, 16, 13], and explicitly by Beals, Buhrman, Cleve, Mosca and de Wolf [7]. In this model, as in its classical counterpart, we pay for accessing the oracle, but unlike the classical case, the machine can use the power of quantum parallelism to make queries in superposition. Access to the input $x \in \Sigma^n$, where Σ is a finite set, is achieved by way of a query operator O_x. The *query complexity* of an algorithm is the number of calls to O_x.

The state of a computation is represented by a register R composed of three subregisters: the *query register* $i \in \{0, \dots, n\}$, the *answer register* $z \in \Sigma$ and the *work register* w. We denote a register using the ket notation $|R\rangle = |i\rangle|z\rangle|w\rangle$, or simply $|i, z, w\rangle$. In the quantum (resp., randomized) setting, the state of the computation is a complex (resp., non-negative real) combination of all possible values of the registers. Let $\mathcal{H}$ denote the corresponding finite-dimensional vector space. We denote the state of the computation by a vector $|\psi\rangle \in \mathcal{H}$ over the basis $(|i, z, w\rangle)_{i,z,w}$. Furthermore, the state vectors are unit length for the ℓ_2 norm in the quantum setting, and for the ℓ_1 norm in the randomized setting.

A *T-query algorithm* A is specified by a $(T{+}1)$-uple $(U_0, U_1, \ldots, U_T)$ of matrices. When A is quantum (resp., randomized), the matrices U_i are unitary (resp., stochastic). The computation takes place as follows. The *query operator* is the unitary (resp., stochastic) matrix O_x that satisfies $O_x|i,z,w\rangle = |i, z \oplus x_i, w\rangle$, for every i, z, w, where by convention $x_0 = 0$. Initially the state is set to some fixed value $|0,0,0\rangle$. Then the sequence of transformations $U_0, O_x, U_1, O_x, \ldots, U_{T-1}, O_x, U_T$ is applied.

We say that the algorithm A *ε-computes* a function $f : S \to S'$, for some sets $S \subseteq \Sigma^n$ and S', if the observation of the last bits of the work register equals $f(x)$ with probability at least $1 - \varepsilon$, for every $x \in S$. Then $\mathsf{QQC}(f)$ (resp., $\mathsf{RQC}(f)$) is the minimum query complexity of quantum (resp., randomized) query algorithms that ε_0-compute f, where ε_0 is a fixed positive constant no greater than $\frac{1}{3}$.

3. Proof of the main theorem

This section is devoted to the proof of the main theorem. We prove Theorem 1 in two main steps. Lemma 2 shows how fast the computations diverge when they start on different individual inputs, in terms of the query probabilities. This lemma depends on the model of computation. Lemma 3 establishes the relationship between the Kolmogorov complexity of individual positions of the input, and the probability that a given algorithm makes a query to this position. This lemma is independent of the model of computation. Theorem 1 follows immediately by combining these two lemmas.

In the following two lemmas, let A be an ε-bounded error algorithm for f that makes at most T queries to the input. Let $p_t^x(i)$ be the probability that A queries x_i at query t on input x, and let $\overline{p}^x(i) = \frac{1}{T}\sum_{t=1}^{T} p_t^x(i)$ be the average query probability over all the time steps up to time T. We assume henceforth without loss of generality that $\overline{p}^x(i) > 0$. (For example, we start by uniformly querying all positions and reverse the process.)

Lemma 2 (Divergence Lemma). *For every input $x, y \in S$ such that $f(x) \neq f(y)$ the following holds.*

1. *For quantum algorithms:*

$$2T \sum_{i : x_i \neq y_i} \sqrt{\overline{p}^x(i)\overline{p}^y(i)} \geq 1 - 2\sqrt{\varepsilon(1-\varepsilon)}.$$

2. *For randomized algorithms:*

$$2T \sum_{i : x_i \neq y_i} \min\left(\overline{p}^x(i), \overline{p}^y(i)\right) \geq 1 - 2\varepsilon.$$

We defer the proof of Lemma 2 to the end of this section.

The next lemma relates the query probabilities to the Kolmogorov complexity of the strings. In this lemma and the results that follow, we assume that a finite description of the algorithm is given. Using the knowledge of A, we may assume without loss of generality that the function f that it computes is also given, as is the length n of the inputs. With additional care, the additive constants in all of the proofs can be made very small by adding to the auxiliary information made available to the description algorithms, those constant-size programs that are described within the proofs.

Lemma 3 (Query Information Lemma). *There exists an absolute constant $c \geq 0$ such that for every input $x \in S$ and position $i \in \{1, \ldots n\}$,*

$$\mathsf{K}(i|x, A) \leq \log\left(\tfrac{1}{\overline{p}^x(i)}\right) + c.$$

Proof. We describe the program that prints i given x and A. Given x, use A and x to compute the probabilities $\overline{p}^x(i)$. This can be done in a finite number of steps because the number of queries is bounded by T. The program includes a hard coded copy of the encoding of i under the Shannon-Fano code for this probability distribution. Decode this and print i. □

From these two lemmas we derive the main theorem.

Proof of Theorem 1. By Lemma 3, there is a constant $c \geq 0$ such that for any algorithm that makes at most T queries, and any x, y, i,

$$\overline{p}^x(i) \leq 2^{-\mathsf{K}(i|x,A)+c} \quad \text{and} \quad \overline{p}^y(i) \leq 2^{-\mathsf{K}(i|y,A)+c}.$$

This is true in particular for all those i where $x_i \neq y_i$. Combining this with Lemma 2 concludes the proof of the main theorem with $C = 2^{-c-1}$. □

We now give the proof of Lemma 2. The proof of the quantum case is very similar to the proofs found in many papers which give quantum lower bounds on query complexity. To our knowledge, the randomized case is new despite the simplicity of its proof. Whereas Aaronson [1] used a different approach to prove a version of Ambainis' method for randomized algorithms, our lemma allows us to use the same framework to establish lower bounds for both quantum and randomized query complexities.

Proof of Lemma 2. Let $|\psi_t^x\rangle$ be the state of the ε-bounded error algorithm A just before the tth oracle query, on input x. By convention, $|\psi_{T+1}^x\rangle$ is the final state. When A is a quantum algorithm $|\psi_t^x\rangle$ is a unit vector for the ℓ_2 norm; otherwise it is a probabilistic distribution, that is, a nonnegative and unit vector for the ℓ_1 norm. Observe that the ℓ_1 distance is the total variation distance.

First we prove the quantum case. Initially, the starting state of A does not depend on the input, thus before the first question we have $|\psi_1^x\rangle = |\psi_1^y\rangle$, so $\langle\psi_1^x|\psi_1^y\rangle = 1$. At the end of the computation, if the algorithm is correct with probability ϵ, then $|\langle\psi^x_{T+1}|\psi^y_{T+1}\rangle| \leq 2\sqrt{\epsilon(1-\epsilon)}$. At each time step, we consider how much the two states can diverge.

Claim 1.

$$|\langle\psi^x_t|\psi^y_t\rangle - \langle\psi^x_{t+1}|\psi^y_{t+1}\rangle| \leq 2\sum_{i:x_i\neq y_i}\sqrt{p^x_t(i)p^y_t(i)}.$$

The proof of Claim 1 can be found in Appendix A.

Over T time steps, the two states diverge as follows. The proof uses only Claim 1 and the Cauchy-Schwartz inequality.

$$\begin{aligned}
1-2\sqrt{\varepsilon(1-\varepsilon)} &\leq |\langle\psi^x_1|\psi^y_1\rangle - \langle\psi^x_{T+1}|\psi^y_{T+1}\rangle| \\
&\leq \sum_{t=1}^{T}|\langle\psi^x_t|\psi^y_t\rangle - \langle\psi^x_{t+1}|\psi^y_{t+1}\rangle| \\
&\leq \sum_{t=1}^{T}2\sum_{i:x_i\neq y_i}\sqrt{p^x_t(i)p^y_t(i)} \\
&\leq 2\sum_{i:x_i\neq y_i}\sqrt{\sum_{t=0}^{T-1}p^x_t(i)\sum_{t=0}^{T-1}p^y_t(i)} \\
&= 2T\sum_{i:x_i\neq y_i}\sqrt{\overline{p}^x(i)\overline{p}^y(i)}.
\end{aligned}$$

Now we prove the randomized case. We use the *ket* notation for real-valued normalized vectors, for consistency in notation. Again, initially $|\psi^x_1\rangle = |\psi^y_1\rangle$. At the end of the computation, if the algorithm is correct with probability ϵ, then $\parallel |\psi^x_{T+1}\rangle - |\psi^y_{T+1}\rangle \parallel_1 \geq 1-2\epsilon$. At each time step, the distribution states now diverge according the following claim.

Claim 2.

$$\begin{aligned}
&\parallel |\psi^x_{t+1}\rangle - |\psi^y_{t+1}\rangle \parallel_1 \\
&\quad\leq \parallel |\psi^x_t\rangle - |\psi^y_t\rangle \parallel_1 + 2\sum_{i:x_i\neq y_i}\min\left(p^x_t(i), p^y_t(i)\right).
\end{aligned}$$

The proof of Claim 2 can be found in Appendix A. We now conclude the proof.

$$\begin{aligned}
1-2\varepsilon &\leq \sum_{t=1}^{T}\parallel |\psi^x_{t+1}\rangle - |\psi^y_{t+1}\rangle \parallel_1 - \parallel |\psi^x_t\rangle - |\psi^y_t\rangle \parallel_1 \\
&\leq \sum_{t=1}^{T}2\sum_{i:x_i\neq y_i}\min\left(p^x_t(i), p^y_t(i)\right) \\
&\leq 2T\sum_{i:x_i\neq y_i}\min\left(\overline{p}^x(i), \overline{p}^y(i)\right).
\end{aligned}$$

□

4. Comparison with previous adversary methods

In this section, we reprove, as a corollary of Theorem 1, the previously known adversary lower bounds. Our framework also allows us to obtain somewhat stronger statements for free.

To obtain the previously known adversary methods as a corollary of Theorem 1, we must give a lower bound on terms $\mathsf{K}(i|x,A)$ and $\mathsf{K}(i|y,A)$. To this end, we apply Proposition 5, and give a lower bound on $\mathsf{K}(x,y)$, and upper bounds on $\mathsf{K}(x|i,y)$ and $\mathsf{K}(y|i,x)$. The lower bound on $\mathsf{K}(x,y)$ is obtained by applying Lemma 1, a consequence of Shannon's coding theorem, for an appropriate distribution. The upper bounds on $\mathsf{K}(x|i,y)$ and $\mathsf{K}(y|i,x)$ are obtained using the Shannon-Fano code, for appropriate distributions.

The following lemma is the general formulation of the sketch above.

Lemma 4. *There exists a constant $C > 0$ such that the following holds. Let Σ be a finite set, let $n \geq 1$ be an integer, and let $S \subseteq \Sigma^n$. Let q be a probability distribution on S^2, let p be a probability distribution on S and let $\{p'_{x,i} : x \in S, 1 \leq i \leq n\}$ be a family of probability distributions on S. Assume that whenever $q(x,y) \neq 0$ then $p(x)$, $p(y)$, $p'_{y,i}(x)$ and $p'_{x,i}(y)$ are non-zero, for every i such that $x_i \neq y_i$. Then for every finite string σ, there exist $x, y \in S$ with $q(x,y) \neq 0$, such that*

$$\begin{aligned}
&\frac{1}{\sum_{i:x_i\neq y_i}\sqrt{2^{-\mathsf{K}(i|x,\sigma)-\mathsf{K}(i|y,\sigma)}}} \\
&\quad\geq C \times \min_{i:x_i\neq y_i}\left(\frac{\sqrt{p(x)p'_{x,i}(y)\ p(y)p'_{y,i}(x)}}{q(x,y)}\right),
\end{aligned}$$

and

$$\begin{aligned}
&\frac{1}{\sum_{i:x_i\neq y_i}\min\left(2^{-\mathsf{K}(i|x,\sigma)}, 2^{-\mathsf{K}(i|y,\sigma)}\right)} \\
&\quad\geq C \times \min_{i:x_i\neq y_i}\left(\max\left(\frac{p(x)p'_{x,i}(y)}{q(x,y)}, \frac{p(y)p'_{y,i}(x)}{q(x,y)}\right)\right).
\end{aligned}$$

Proof. In this proof, $c_1, \ldots, c_5$ are some appropriate nonnegative constants. By Lemma 1, there exists a pair (x,y) such that $q(x,y) \neq 0$ and

$$\mathsf{K}(x,y|\sigma,p,p') \geq \log(\tfrac{1}{q(x,y)}),$$

where p' stands for a complete description of all the $p'_{x,i}$.

Fix x and y so that this holds. By using the Shannon-Fano code (Proposition 3),

$$\mathsf{K}(x|p) \leq \log(\tfrac{1}{p(x)}) + c_1$$

and

$$\mathsf{K}(y|x,i,p'_{x,i}) \leq \log(\tfrac{1}{p'_{i,x}(y)}) + c_1,$$

for any i such that $x_i \neq y_i$. By Proposition 5,

$$\begin{array}{rl}\mathsf{K}(i|x,\sigma) & \\ \geq & \mathsf{K}(i|x,\sigma,p,p') - c_3 \\ \geq & \mathsf{K}(x,y|\sigma,p,p') - \mathsf{K}(x|p) - \mathsf{K}(y|i,x,p'_{x,i}) + \\ & \mathsf{K}(i|x,y,\mathsf{K}(x,y),\sigma,p,p') - c_4 \\ \geq & \log(\frac{1}{q(x,y)}) - \log(\frac{1}{p(x)}) - \log(\frac{1}{p'_{x,i}(y)}) + \\ & \mathsf{K}(i|x,y,\mathsf{K}(x,y),\sigma,p,p') - c_5 \\ = & \log(\frac{p(x)p'_{x,i}(y)}{q(x,y)}) + \mathsf{K}(i|x,y,\mathsf{K}(x,y),\sigma,p,p') - c_5.\end{array}$$

Similarly,

$$\begin{array}{rcl}\mathsf{K}(i|y,\sigma) & \geq & \log(\frac{p(y)p'_{y,i}(x)}{q(x,y)}) + \\ & & \mathsf{K}(i|x,y,\mathsf{K}(x,y),\sigma,p,p') - c_5\end{array}$$

This concludes the proof of the lemma using Kraft's inequality (Proposition 2) and letting $C = 2^{-c_5}$. □

4.1. Ambainis' weighted scheme

Theorem 2 (Ambainis' weighted method). *Let Σ be a finite set, let $n \geq 1$ be an integer, and let $S \subseteq \Sigma^n$ and S' be sets. Let $f : S \to S'$. Consider a weight scheme as follows:*

- *Every pair $(x,y) \in S^2$ is assigned a non-negative weight $w(x,y)$ such that $w(x,y) = 0$ whenever $f(x) = f(y)$.*
- *Every triple (x,y,i) is assigned a non-negative weight $w'(x,y,i)$ such that $w'(x,y,i) = 0$ whenever $x_i = y_i$ or $f(x) = f(y)$.*

For all x, i, let

$$\begin{array}{rcl} wt(x) & = & \sum_y w(x,y) \textit{and} \\ v(x,i) & = & \sum_y w(x,y,i). \end{array}$$

If $w'(x,y,i)w'(y,x,i) \geq w^2(x,y)$ for all x,y,i such that $x_i \neq y_i$, then

$$\mathsf{QQC}(f) = \Omega\left(\min_{\substack{x,y,i\\ w(x,y)\neq 0, x_i \neq y_i}} \left(\sqrt{\frac{wt(x)wt(y)}{v(x,i)v(y,i)}}\right)\right).$$

Furthermore, if $w'(x,y,i), w'(y,x,i) \geq w(x,y)$ for all x,y,i such that $x_i \neq y_i$, then

$$\mathsf{RQC}(f) = \Omega\left(\min_{\substack{x,y,i\\ w(x,y)\neq 0, x_i \neq y_i}} \left(\max\left(\frac{wt(x)}{v(x,i)}, \frac{wt(y)}{v(y,i)}\right)\right)\right).$$

The relation in Ambainis' original statement is implicit in this formulation, since it corresponds to the non-zero-weight pairs. A weaker version of the randomized case was proven independently by Aaronson [1] using a completely different approach. We show that Theorem 2 follows from Theorem 1.

Proof. We derive probability distributions q, p, p' from the weight schemes as follows. Let $W = \sum_{x,y} w(x,y)$. Define

$$\begin{array}{rcl} q(x,y) & = & \frac{w(x,y)}{W}, \\ p(x) & = & \frac{wt(x)}{W}, \\ p'_{x,i}(y) & = & \frac{w'(y,x,i)}{v(x,i)}, \quad \text{for any } x,y,i. \end{array}$$

It is easy to check that by construction and hypothesis, these distributions satisfy the conditions of Lemma 4. Rearranging and simplifying the terms allows us to conclude. □

We conclude this section by sketching the proof of the unweighted version of Ambainis' adversary method, as it affords a simpler combinatorial proof, that does not require Lemma 4. To simplify notation we omit additive constants and the usual auxiliary strings including A.

Let $R \subseteq S \times S$, be a relation on pairs of instances, where $(x,y) \in R \Longrightarrow f(x) \neq f(y)$, and let R_i be the restriction of R to pairs x, y for which $x_i \neq y_i$. Viewing the relation R as a bipartite graph, let l, l', m, m' be as follows.

- m is a lower bound on the degree of all $x \in X$,
- m' is a lower bound on the degree of all $y \in Y$,
- for any fixed x and $i, 1 \leq i \leq n$, the number of y adjacent to x for which $x_i \neq y_i$ is at most l,
- for any fixed y and $i, 1 \leq j \leq n$, the number of x adjacent to y for which $x_i \neq y_i$ is at most l'.

We make the following observations.

1. $|R| \geq \max\{m|X|, m'|Y|\}$, so $\exists x,y\ \mathsf{K}(x,y) \geq \max\left(\log(m|X|), \log(m'|Y|)\right)$.
2. $\forall x \in X, \mathsf{K}(x) \leq \log(|X|)$ and $\mathsf{K}(y) \leq \log(|Y|)$, for all $y \in Y$.
3. $\forall x,y,i$ with $(x,y) \in R_i, \mathsf{K}(y|i,x) \leq \log(l)$ and similarly, $\mathsf{K}(x|i,y) \leq \log(l')$.

For any i with $x_i \neq y_i$, by Proposition 5,

$$\begin{array}{rcl}\mathsf{K}(i|x) & \geq & \mathsf{K}(x,y) - \mathsf{K}(x) - \mathsf{K}(y|i,x) + \\ & & \mathsf{K}(i|x,y,\mathsf{K}(x,y)) \\ & \geq & \log(m|X|) - \log(|X|) - \log(l) + \\ & & \mathsf{K}(i|x,y,\mathsf{K}(x,y)) \\ & = & \log(\frac{m}{l}) + \mathsf{K}(i|x,y,\mathsf{K}(x,y))\end{array}$$

The same proof works to show that $\mathsf{K}(i|y) \geq \log(\frac{m'}{l'}) + \mathsf{K}(i|x,y,\mathsf{K}(x,y))$. By Theorem 1 and Kraft's inequality,

$$\mathsf{QQC}(f) = \Omega\left(\sqrt{\frac{mm'}{ll'}}\right).$$

4.2. Spectral lower bound

We now show how to prove the spectral lower bound of Barnum, Saks ans Szegedy [6] as a corollary of Theorem 1. Recall that for any matrix Γ, $\lambda(\Gamma)$ is the largest eigenvalue of Γ.

Theorem 3 (Barnum-Saks-Szegedy spectral method). *Let Σ be a finite set, let $n \geq 1$ be an integer, and let $S \subseteq \Sigma^n$ and S' be sets. Let $f : S \to S'$. Let Γ be an arbitrary $S \times S$ nonnegative real symmetric matrix that satisfies $\Gamma(x, y) = 0$ whenever $f(x) = f(y)$. For $i = 1, \dots, n$ let Γ_i be the matrix:*

$$\Gamma_i(x,y) = \begin{cases} 0, & \textit{if } x_i = y_i; \\ \Gamma(x,y), & \textit{otherwise.} \end{cases}$$

Then

$$\mathsf{QQC}(f) = \Omega\left(\frac{\lambda(\Gamma)}{\max_i \lambda(\Gamma_i)}\right).$$

Proof. Let $|\alpha\rangle$ (resp., $|\alpha_i\rangle$) be the unit eigenvector of Γ (resp., Γ_i) with nonnegative entries and whose eigenvalue is $\lambda(\Gamma)$ (resp., $\lambda(\Gamma_i)$). We define the probability distributions q, p, p' as follows. Let $W = \sum_{x,y} w(x, y)$, and define

$$\begin{aligned} q(x,y) &= \frac{\Gamma(x,y)\langle x|\alpha\rangle\langle y|\alpha\rangle}{\langle\alpha|\Gamma|\alpha\rangle}, \\ p(x) &= \langle x|\alpha\rangle^2, \\ p'_{i,x}(y) &= \frac{\Gamma_i(x,y)\langle y|\alpha_i\rangle}{\langle x|\Gamma_i|\alpha_i\rangle}, \qquad \text{for any } x, y, i. \end{aligned}$$

By construction these distributions satisfy the conditions of Lemma 4, which suffices to conclude. □

5. Certificate complexity and adversary techniques

Let f be a boolean function. For any positive instance $x \in \Sigma^n$ of f ($f(x)$=1), a *positive certificate* for $f(x)$ is the smallest subset of indices $I \subseteq [n]$ of x, such that for any y with $x_i = y_i$ for all $i \in I$, $f(y)$=1.

The 1-*certificate complexity* of f, denoted $C_1(f)$, is the size of the largest positive certificate for $f(x)$, over all positive instances x. The 0-*certificate complexity* is defined similarly for negative instances x of f ($f(x) = 0$).

Prior to our work, it was known that the best possible bound that could be proven using the unweighted adversary technique [2, Theorem 5.1] is $O(\sqrt{C_0(f)C_1(f)})$. Independently, Szegedy [17] showed that the best possible lower bound using the spectral method is $O(\min(\sqrt{nC_0(f)}, \sqrt{nC_1(f)}))$, and Zhang [19] proved the same for Ambainis' weighted method.

The following lemma, due to Troy Lee, results in a very simple proof of the fact that our method, and hence, all the known variants of the adversary method, cannot prove lower bounds larger than $\min(\sqrt{nC_0(f)}, \sqrt{nC_1(f)})$.

Lemma 5. *There exists a constant $c \geq 0$ such that the following holds. Let Σ be a finite set, let $n \geq 1$ be an integer, and let $S \subseteq \Sigma^n$ be a set. Let $f : S \to \{0, 1\}$. For every $x, y \in S$ with $f(x) = 0$ and $f(y) = 1$, there is an i with $x_i \neq y_i$ for which $\mathsf{K}(i|x, f) \leq \log(C_0(f)) + c$, and similarly, there is a j with $x_j \neq y_j$ such that $\mathsf{K}(j|y, f) \leq \log(C_1(f)) + c$.*

Proof. Let I be the lexicographically smallest certificate for $f(x)$. Since $f(x) \neq f(y)$, x and y must differ on some $i \in I$. To describe i given x, it suffices to give an index into I, which requires $\log(C_0(f)) + c$ bits. The same can also be done with x and y reversed. □

Theorem 4. *Let Σ be a finite set, let $n \geq 1$ be an integer, and let $S \subseteq \Sigma^n$ be a set. Let $f : S \to \{0, 1\}$. Then any quantum query lower bound for f given by Theorem 1 is in $O(\min(\sqrt{nC_0(f)}, \sqrt{nC_1(f)}))$.*

Proof. Let A be a quantum algorithm that computes f with bounded error by making at most T queries to the input. Since a description of f can be obtained from a description of A, $\mathsf{K}(i|x, A) \leq \mathsf{K}(i|x, f) + O(1)$. Therefore, the lower bound given by Theorem 1 is $O\left(\frac{1}{\sum_{i:x_i \neq y_i} \sqrt{2^{-\mathsf{K}(i|x,f)-\mathsf{K}(i|y,f)}}}\right)$, where $f(x) \neq f(y)$. This is $O(\min(\sqrt{nC_0(f)}, \sqrt{nC_1(f)}))$ by Lemma 5. □

6. Applications

6.1. A general method for distance schemes

We generalize the technique of Høyer, Neerbek and Shi [14], which they used to prove lower bounds on ordered search and sorting. Though their technique is similar, it does not appear to be a special case of the weighted adversary method.

Here, we restrict ourselves to those weight functions that take values of the form $\frac{1}{d}$, for integer values d. Therefore, instead of a weight function, we consider an integer function D, which may be thought of as a distance function on pairs of instances (even though it is not strictly speaking a distance measure in general). We will define the *load* of an instance x, to be the maximum number of instances y at any given distance d from x. This will allow us to bound the complexity of printing y, given x and d. (In the case of ordered search, the load will be 1 for all instances.)

More formally, for any non-negative integer function D on pairs (x, y), we define the *right load* $l_{\mathbf{R}}(x, i)$ to be the maximum over all values d, of the number of y such that $D(x, y) = d$ and $x_i \neq y_i$. The *left load* $l_{\mathbf{L}}(y, i)$ is defined similarly, inverting x and y.

Theorem 5. *Let Σ be a finite set, let $n \geq 1$ be an integer, and let $S \subseteq \Sigma^n$ and S' be sets. Let $f : S \to S'$. Let D be a non-negative integer function on S^2 such that $D(x, y) = 0$ whenever $f(x) = f(y)$. Let $W = \sum_{x,y:D(x,y)\neq 0} \frac{1}{D(x,y)}$. Then*

$$\mathsf{QQC}(f) = \Omega\left(\frac{W}{|S|} \min_{\substack{x,y \\ D(x,y)\neq 0, \\ x_i = y_i}} \left(\frac{1}{\sqrt{l_{\mathbf{R}}(x,i) l_{\mathbf{L}}(y,i)}} \right) \right),$$

$$\mathsf{RQC}(f) = \Omega\left(\frac{W}{|S|} \min_{\substack{x,y \\ D(x,y)\neq 0, \\ x_i = y_i}} \left(\max\left(\frac{1}{l_{\mathbf{R}}(x,i)}, \frac{1}{l_{\mathbf{L}}(y,i)} \right) \right) \right).$$

Proof. We use a variation on Lemma 4. We define probability distributions $q(x, y) = \frac{1}{D(x,y)\times W}$ whenever $D(x, y) \neq 0$ and $q(x, y) = 0$ otherwise; $p(x) = \frac{1}{|S|}$. Fix σ to be the string containing a description of A and D, where D is a complete description of the distance function, and where we assume that A includes a description of f, hence $l_{\mathbf{R}}(x, i), l_{\mathbf{L}}(y, i)$ and are also given.

We give an upper bound on the terms $\mathsf{K}(y|x, i)$ and $\mathsf{K}(x|y, i)$ directly, using left and right loads. Given x, i and some integer $d > 0$, there are at most $l_{\mathbf{R}}(x, i)$ instances y such that $D(x, y) = d$ and $x_i \neq y_i$. Therefore

$$\mathsf{K}(y|x, i, \sigma) \leq \log(D(x, y)) + \log(l_{\mathbf{R}}(x, i)) + c,$$

where $c \geq 0$ is some constant, The same is true for $\mathsf{K}(x|y, i)$:

$$\mathsf{K}(y|x, i, \sigma) \leq \log(D(x, y)) + \log(l_{\mathbf{L}}(y, i)) + c.$$

Now, we conclude following the same sketch as the proof of Lemma 4. □

We reprove some of the lower bounds of Høyer, Neerbek and Shi. The distance schemes we use are exactly the ones of [14]. Whereas they did not separate the quantum part from the combinatorial part in their proofs, here we only need to evaluate the combinatorial objects $l_{\mathbf{R}}$ and $l_{\mathbf{L}}$ to get the results.

Corollary 1. $\mathsf{QQC}(\text{ORDERED SEARCH}) = \Omega(\log n)$ *and* $\mathsf{RQC}(\text{ORDERED SEARCH}) = \Omega(\log n)$.

Proof. Fix $\Sigma = \{0, 1\}$. We only consider the set of instances S of length n of the form $0^{a-1}1^{n-a}$. Note that $|S| = n$. Define distance for pairs $(x, y) \in S^2$ as $D(x, y) = b - a$, and $D(x, y) = 0$ for all other instances, where $x = 0^{a-1}1^{n-a}$ and $y = 0^{b-1}1^{n-b}$ with $1 \leq a < b \leq n$. The inverse distance has total weight $W = \Theta(n \log n)$. Furthermore, for every x, y, i such that $D(x, y) \neq 0$ and $x_i \neq y_i$, $l_{\mathbf{R}}(x, i) = l_{\mathbf{L}}(y, i) = 1$. The result follows by Theorem 5. □

A lower bound for sorting [14] in the comparison model can also be obtained by applying Theorem 5.

Corollary 2. $\mathsf{QQC}(\text{SORTING}) = \Omega(n \log n)$ *and* $\mathsf{RQC}(\text{SORTING}) = \Omega(n \log n)$.

Proof. Fix $\Sigma = \{0, 1\}$. An input is an $n \times n$ comparison matrix M_σ defined by $(M_\sigma)_{i,j} = 1$ if $\sigma(i) < \sigma(j)$, and $(M_\sigma)_{i,j} = 0$ otherwise, where σ is some permutation of $\{1, \ldots, n\}$. (In the usual array representation, the element of rank r in the array would be stored at position $\sigma^{-1}(r)$.) The set S of inputs is $\{M_\sigma : \sigma \in S_n\}$.

We consider pairs of instances $M_\sigma, M_{\sigma^{(k,d)}}$, where $\sigma^{(k,d)}$ is obtained from σ by changing the value of the element of rank $k + d$ to a value that immediately precedes the element of rank k in σ. This changes the rank of the d elements of intermediate rank, incrementing their rank by one.

More formally, define $\sigma^{(k,d)} = (k, k+1, \ldots, k+d) \circ \sigma$, for $d \neq 0$. For every pair of permutations σ, τ we let $D(M_\sigma, M_\tau) = d$ if there exists k, d such that $\tau = \sigma^{(k,d)}$, and $D(M_\sigma, M_\tau) = 0$ otherwise. Observe that whenever $\tau = \sigma^{(k,d)}$, the comparison matrices M_σ and M_τ differ only in entries $(\sigma^{-1}(k + d), \sigma^{-1}(i)) = (\tau^{-1}(k), \tau^{-1}(i + 1))$ and $(\sigma^{-1}(i), \sigma^{-1}(k + d)) = (\tau^{-1}(i + 1), \tau^{-1}(k))$, for $k \leq i \leq k + d - 1$.

Then for every $\sigma, \tau, (i, j)$ such that $D(M_\sigma, M_\tau) \neq 0$ and $(M_\sigma)_{i,j} \neq (M_\tau)_{i,j}$, $l_{\mathbf{R}}(\sigma, (i, j)) = l_{\mathbf{L}}(\sigma, (i, j)) = 2$. This is because given σ, i, j, d, either $i = \sigma^{-1}(k + d)$ or $j = \sigma^{-1}(k + d)$, so there are two possible values for (k, d). Similarly, $l_{\mathbf{R}}(\tau, (i, j)) = l_{\mathbf{L}}(\tau, (i, j)) = 2$. The inverse distance has total weight $W = \Theta((n!)n \log n)$ and the size of S is $|S| = (n!)$. Applying Theorem 5, we conclude the proof. □

6.2. Graph properties

Theorem 1 provides a simple and intuitive method to prove lower bounds for specific problems. We illustrate this by giving lower bounds for two graph properties: connectivity, and bipartiteness. These are direct applications of Theorem 1 in that we analyze directly the complexity $\mathsf{K}(i|x, A)$ without defining relations or weights or distributions: we only need to consider a "typical" hard pair of instances. In this section, we omit additive and multiplicative constants that result from using small, constant-size programs, as well as the constant length auxiliary string A to simplify the proofs.

6.2.1. Graph connectivity

Theorem 6 ([11]). *In the adjacency matrix model,*

$$\mathsf{QQC}(\text{GRAPHCONNECTIVITY}) = \Omega(n^{3/2}),$$

where n is the number of vertices in the graph.

Proof. We construct one negative and one positive instance of graph connectivity, using the incompressibility method, using the ideas of [11]. Let S be an incompressible string of length $\log(n-1)! + \log\binom{n}{2}$, chopped into two pieces S_1 and S_2 of length $\log(n-1)!$ and $\log\binom{n}{2}$, respectively. We think of S_1 as representing a hamilton cycle $C = (0, \pi(0)\cdots\pi(n-1), 0)$ through the n vertices, and S_2 as representing a pair of distinct vertices s, t. Let G contain the cycle C and let H be obtained from G by breaking the cycle into two cycles at s and t, that is, $H = G\setminus\{(\pi(s),\pi(s+1)),(\pi(t),\pi(t+1))\}\cup\{(\pi(s),\pi(t+1)),(\pi(s+1),\pi(t))\}$.

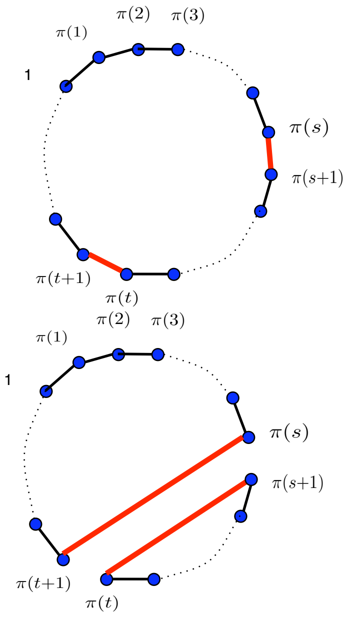

Figure 1. Graphs G, H for the graph lower bounds.

We show that for the four edges e where G and H differ, $\mathsf{K}(e|G)+\mathsf{K}(e|H) \geq 3\log n - 4$. Let e_-, e'_- be the edges removed from G, and e_+, e'_+ be the edges added to G. Observe that up to an additive constant, $\mathsf{K}(e_+|G) = \mathsf{K}(e'_+|G)$ and $\mathsf{K}(e_-|H) = \mathsf{K}(e'_-|H)$. Let e_- be one of the edges removed from G, w.l.o.g., $e_- = (\pi(s), \pi(s+1))$.

$$\begin{array}{rcl} \log(n-1)! + \log\binom{n}{2} & \leq & \mathsf{K}(S) \\ & \leq & \mathsf{K}(G) + \mathsf{K}(s|G) + \mathsf{K}(t|G) \\ & \leq & \mathsf{K}(G) + \mathsf{K}(e_-|G) + \log n \\ \mathsf{K}(e_-|G) & \geq & \log\binom{n}{2} - \log n = \log\frac{n-1}{2} \end{array}$$

Assume w.l.o.g. that the smallest cycle of H contains $\pi(s)$, and let l be its length.

$$\begin{array}{l} \log(n-1)! + \log\binom{n}{2} \\ \quad \leq \; \mathsf{K}(S) \\ \quad \leq \; \mathsf{K}(H) + \mathsf{K}(e_-|H) + \mathsf{K}(\pi(t),\pi(t+1)|H) \\ \quad \leq \; \log\frac{(n-1)!}{(n-l+1)!} + \log(n-l-1)! + \mathsf{K}(e_-|H) + \\ \qquad \log l + \log(n-l) \end{array}$$

$$\mathsf{K}(e_-|H) \quad \geq \quad 2\log n + \log(n-l) - \log(l) \geq 2\log n.$$

For the added edges, e_+, e'_+, consider w.l.o.g. $e_+ = (\pi(s), \pi(t+1))$. Since S is incompressible, $\mathsf{K}(e_+|G) \geq \mathsf{K}(s,t|G) \geq \log\binom{n}{2}$. Furthermore, $\mathsf{K}(S) \leq \mathsf{K}(H) + \mathsf{K}(e_+|H) + \mathsf{K}(e'_+|H)$, and $\mathsf{K}(e'_+|H) \leq \log n$, so $\mathsf{K}(e_+|H) \geq \log\binom{n}{2} - \log n = \log\frac{n-1}{2}$. The same proof shows that $\mathsf{K}(e'_+|H) \geq \log\frac{n-1}{2}$.

□

6.2.2. Bipartiteness The following lower bound was proven by Dürr (personal communication) and independently in [19].

Theorem 7. *In the adjacency matrix model,*

$$\mathsf{QQC}(\textsc{Bipartiteness}) = \Omega(n^{3/2}),$$

where n is the number of vertices in the graph.

Proof. The proof is similar to the one of Theorem 6 except that we construct G to be an even cycle on $n = 2m$ vertices, and H will be composed of two odd cycles on the same vertex set (see Figure 1).

Let S be an incompressible string of length $\log(n-1)! + \log(\binom{n}{2} - 1)$, chopped into two pieces S_1 and S_2 of length $\log(n-1)!$ and $\log(\binom{n}{2} - 1)$, respectively. We think of S_1 as representing a hamilton cycle $C = (0, \pi(0)\cdots\pi(n-1), 0)$ through the n vertices, and S_2 as representing a pair of distinct vertices s, t, with $s \not\equiv t \pmod 2$. Let G contain the cycle C and let H be obtained from G by breaking the cycle into two odd cycles at s and t, that is, $H = G\setminus\{(\pi(s),\pi(s+1)),(\pi(t),\pi(t+1))\}\cup\{(\pi(s),\pi(t+1)),(\pi(s+1),\pi(t))\}$.

The same analysis as Thaorem 6 yields the lower bound $\mathsf{QQC}(\textsc{Bipartiteness}) = \Omega(n^{3/2})$, as claimed. □

7. Acknowledgements

We wish to thank Troy Lee, Christoph Dürr for many useful discussions, and Andris Ambainis for his helpful answers to our questions.

References

[1] S. Aaronson. Lower bounds for local search by quantum arguments. In *Proceedings of 36th ACM Symposium on Theory of Computing*, 2004. To appear. Also in quant-ph/0307149.

[2] A. Ambainis. Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64:750–767, 2002.

[3] A. Ambainis. Polynomial degree vs. quantum query complexity. In *Proceedings of 44th IEEE Symposium on Foundations of Computer Science*, pages 230–239, 2003.

[4] L. Babai and S. Laplante. Stronger separations for random-self-reducibility, rounds, and advice. In *IEEE Conference on Computational Complexity*, pages 98–104, 1999.

[5] H. Barnum and M. Saks. A lower bound on the quantum query complexity of read-once functions. Technical Report quant-ph/0201007, arXiv, 2002.

[6] H. Barnum, M. Saks, and M. Szegedy. Quantum decision trees and semidefinite programming. In *Proceedings of the 18th IEEE Conference on Computational Complexity*, pages 179–193, 2003.

[7] 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.

[8] Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. Wiley-Interscience, 1991.

[9] D. Deutsch. Quantum theory, the Church-Turing principle and the universal quantum computer. In *Proceedings of the Royal Society of London A*, volume 400, pages 97–117, 1985.

[10] D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. In *Proceedings of the Royal Society A*, volume 439, 1992.

[11] C. Dürr, M. Heiligman, P. Høyer, and M. Mhalla. Quantum query complexity of some graph problems. In *Proceedings of 31st International Colloquium on Automata, Languages and Programming*, 2004. To appear. Also in quant-ph/0401091.

[12] J. Feigenbaum, L. Fortnow, S. Laplante, and A. V. Naik. On coherence, random-self-reducibility, and self-correction. *Computational Complexity*, 7(2):174–191, 1998.

[13] L. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of 28th ACM Symposium on Theory of Computing*, pages 212–219, 1996.

[14] P. Høyer, J. Neerbek, and Y. Shi. Quantum complexities of ordered searching, sorting, and element distinctness. *Algorithmica*, 34(4):429–448, 2002.

[15] M. Li and P. Vitányi. An introduction to Kolmogorov complexity and its applications. In *Graduate Texts in Computer Science*. Springer, 1997. Second edition.

[16] D. Simon. On the power of quantum computation. *SIAM Journal on Computing*, 26(5):1474–1483, 1997.

[17] M. Szegedy. On the quantum query complexity of detecting triangles in graphs. Technical Report quant-ph/0310107, arXiv archive, 2003.

[18] A. Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of 18th IEEE Symposium on Foundations of Computer Science*, pages 222–227, 1977.

[19] S. Zhang. On the power of Ambainis's lower bounds. In *Proceedings of 31st International Colloquium on Automata, Languages and Programming*, 2004. To appear. Also in quant-ph/0311060.

A. Proofs of claims

Proof of Claim 1. Let

$$\begin{aligned} |\psi_t^x\rangle &= \sum_{i,z,w} \alpha_{i,z,w} |i,z,w\rangle, \text{ and} \\ |\psi_t^y\rangle &= \sum_{i,z,w} \beta_{i,z,w} |i,z,w\rangle. \end{aligned}$$

After the tth query is made, the states $|\psi_t'^x\rangle = O_x|\psi_t^x\rangle$ and $|\psi_t'^y\rangle = O_y|\psi_t^y\rangle$ are

$$\begin{aligned} |\psi_t'^x\rangle &= \sum_{i,z,w} \alpha_{i,z,w} |i,z \oplus x_i,w\rangle, \text{ and} \\ |\psi_t'^y\rangle &= \sum_{i,z,w} \beta_{i,z,w} |i,z \oplus y_i,w\rangle. \end{aligned}$$

Now, since the inner product is invariant under unitary transformations, we get

$$\langle \psi_{t+1}^x | \psi_{t+1}^y \rangle = \langle \psi_t'^x | \psi_t'^y \rangle,$$

and therefore,

$$\begin{aligned} &|\langle \psi_t^x | \psi_t^y \rangle - \langle \psi_{t+1}^x | \psi_{t+1}^y \rangle| \\ &\quad= |\sum_{i,z,w} \overline{\alpha_{i,z,w}} \beta_{i,z,w} - \sum_{i,z,w} \overline{\alpha_{i,z\oplus x_i,w}} \beta_{i,z\oplus y_i,w}| \\ &\quad= |\sum_{\substack{i,z,w \\ x_i \neq y_i}} \overline{\alpha_{i,z,w}} \beta_{i,z,w} - \overline{\alpha_{i,z\oplus x_i,w}} \beta_{i,z\oplus y_i,w}| \\ &\quad\leq \sum_{i:x_i\neq y_i} \left(|\sum_{z,w} \overline{\alpha_{i,z,w}} \beta_{i,z,w}| + |\sum_{z,w} \overline{\alpha_{i,z\oplus x_i,w}} \beta_{i,z\oplus y_i,w}| \right) \\ &\quad\leq 2 \sum_{i:x_i\neq y_i} \sqrt{\left(\sum_{z,w} |\alpha_{i,z,w}|^2 \right) \left(\sum_{z,w} |\beta_{i,z,w}|^2 \right)} \\ &\quad\leq 2 \sum_{i:x_i\neq y_i} \sqrt{p_t^x(i) p_t^y(i)} \end{aligned}$$

□

Proof of Claim 2. Let us write the distributions using the same formalism as above, that is,

$$\begin{aligned} |\psi_t^x\rangle &= \sum_{i,z,w} \alpha_{i,z,w} |i,z,w\rangle, \text{ and} \\ |\psi_t^y\rangle &= \sum_{i,z,w} \beta_{i,z,w} |i,z,w\rangle. \end{aligned}$$

Note that now, the vectors are unit for the ℓ_1 norm. After the tth query is made, the states $|\psi_t'^x\rangle = O_x|\psi_t^x\rangle$ and $|\psi_t'^y\rangle = O_y|\psi_t^y\rangle$ are

$$\begin{aligned} |\psi_t'^x\rangle &= \sum_{i,z,w} \alpha_{i,z,w} |i,z \oplus x_i,w\rangle, \text{ and} \\ |\psi_t'^y\rangle &= \sum_{i,z,w} \beta_{i,z,w} |i,z \oplus y_i,w\rangle. \end{aligned}$$

Now, since the ℓ_1 distance does not increase under stochastic matrices, we get

$$\| \, |\psi_{t+1}^x\rangle - |\psi_{t+1}^y\rangle \, \|_1 \leq \| \, |\psi_t'^x\rangle - |\psi_t'^y\rangle \, \|_1,$$

and therefore,

$$\begin{aligned} &\| \, |\psi_{t+1}^x\rangle - |\psi_{t+1}^y\rangle \, \|_1 \\ &\quad= \| \sum_{i,z,w} (\alpha_{i,z,w} |i,z\oplus x_i,w\rangle - \beta_{i,z,w} |i,z\oplus y_i,w\rangle) \, \|_1 \\ &\quad= \sum_i \| \sum_{z,w} (\alpha_{i,z,w} |i,z\oplus x_i,w\rangle - \beta_{i,z,w} |i,z\oplus y_i,w\rangle) \, \|_1 \; . \end{aligned}$$

We now bound each term of the last sum separately. Fix any i. If $x_i = y_i$ then

$$\begin{aligned} &\| \sum_{z,w} (\alpha_{i,z,w} |i,z\oplus x_i,w\rangle - \beta_{i,z,w} |i,z\oplus y_i,w\rangle) \, \|_1 \\ &\quad= \| \sum_{z,w} (\alpha_{i,z,w} |i,z,w\rangle - \beta_{i,z,w} |i,z,w\rangle) \, \|_1 \; . \end{aligned}$$

If $x_i \neq y_i$ then,

$$\begin{aligned} &\| \sum_{z,w} (\alpha_{i,z,w} |i,z\oplus x_i,w\rangle - \beta_{i,z,w} |i,z\oplus y_i,w\rangle) \, \|_1 \\ &\quad\leq \| \sum_{z,w} (\alpha_{i,z,w} |i,z\oplus y_i,w\rangle - \beta_{i,z,w} |i,z\oplus y_i,w\rangle) \, \|_1 + \\ &\qquad \| \sum_{z,w} (\alpha_{i,z,w} |i,z\oplus x_i,w\rangle - \alpha_{i,z,w} |i,z\oplus y_i,w\rangle) \, \|_1 \\ &\quad\leq \| \sum_{z,w} (\alpha_{i,z,w} |i,z,w\rangle - \beta_{i,z,w} |i,z,w\rangle) \, \|_1 + \\ &\qquad 2 \, \| \sum_{z,w} \alpha_{i,z,w} |i,z,w\rangle \, \|_1 \\ &\quad= \| \sum_{z,w} (\alpha_{i,z,w} |i,z,w\rangle - \beta_{i,z,w} |i,z,w\rangle) \, \|_1 + 2 p_t^x(i). \end{aligned}$$

In the same way we can prove that

$$\begin{aligned} &\| \sum_{z,w} (\alpha_{i,z,w} |i,z\oplus x_i,w\rangle - \beta_{i,z,w} |i,z\oplus y_i,w\rangle) \, \|_1 \\ &\quad\leq \| \sum_{z,w} (\alpha_{i,z,w} |i,z,w\rangle - \beta_{i,z,w} |i,z,w\rangle) \, \|_1 + 2 p_t^y(i). \end{aligned}$$

We group together these upper bounds and conclude

$$\begin{aligned} &\| \, |\psi_{t+1}^x\rangle - |\psi_{t+1}^y\rangle \, \|_1 \\ &\quad\leq \sum_i \| \sum_{z,w} (\alpha_{i,z,w} |i,z,w\rangle - \beta_{i,z,w} |i,z,w\rangle) \, \|_1 + \\ &\qquad 2 \sum_{i:x_i\neq y_i} \min\left(p_t^x(i), p_t^y(i)\right) \\ &\quad= \| \, |\psi_t^x\rangle - |\psi_t^y\rangle \, \|_1 + 2 \sum_{i:x_i\neq y_i} \min\left(p_t^x(i), p_t^y(i)\right). \end{aligned}$$

□