THEORY OF COMPUTING, Volume 1 (2005), pp. 37–46
http://theoryofcomputing.org

Polynomial Degree and Lower Bounds in Quantum Complexity: Collision and Element Distinctness with Small Range

Andris Ambainis

Received: November 23, 2004; published: May 13, 2005.

Abstract: We give a general method for proving quantum lower bounds for problems with small range. Namely, we show that, for any symmetric problem defined on functions $f : \{1,\ldots,N\} \to \{1,\ldots,M\}$, its polynomial degree is the same for all $M \geq N$. Therefore, if we have a quantum query lower bound for some (possibly quite large) range M which is shown using the polynomials method, we immediately get the same lower bound for all ranges $M \geq N$. In particular, we get $\Omega(N^{1/3})$ and $\Omega(N^{2/3})$ quantum lower bounds for collision and element distinctness with small range, respectively. As a corollary, we obtain a better lower bound on the polynomial degree of the two-level AND–OR tree.

ACM Classification: F.1.2

AMS Classification: 81P68, 68Q17

Key words and phrases: quantum computation, quantum query algorithms, quantum lower bounds, polynomials method, complexity of Boolean functions, element distinctness

1 Introduction

Quantum computing provides speedups for many search problems. The most famous example is Grover's algorithm [14], which computes OR of N variables with $O(\sqrt{N})$ queries. Other examples include counting [8], estimating mean and median [15, 19], finding collisions [7], determining element distinctness [11, 5], finding triangles in a graph [18] and verifying matrix products [12]. For many of these problems, we can also prove that known quantum algorithms are optimal or nearly optimal.

DOI: 10.4086/toc.2005.v001a003

In at least two cases, the lower bounds match the best known algorithm only with an additional "large range" assumption. For example, consider the collision problem [7, 2] which models collision-free hash functions. We have to distinguish if a function $f : \{1,\ldots,N\} \to \{1,\ldots,M\}$ is one-to-one or two-to-one. A quantum algorithm can solve the problem with $O(N^{1/3})$ queries (evaluations of f) [7], which is better than the $\Theta(N^{1/2})$ queries required classically. A lower bound by Aaronson and Shi [2] says that $\Omega(N^{1/3})$ quantum queries are required if $M \geq 3N/2$. If $M = N$, the lower bound becomes $\Omega(N^{1/4})$.

A similar problem exists for element distinctness. (Again, we are given $f : \{1,\ldots,N\} \to \{1,\ldots,M\}$ but f can be arbitrary and we have to determine if there are i, j, $i \neq j$, $f(i) = f(j)$.) If $M = \Omega(N^2)$, the lower bound is $\Omega(N^{2/3})$ [2], which matches the best algorithm [5]. But, if $M = N$, the lower bound is only $\Omega(\sqrt{N})$ or $\Omega(\sqrt{N \log N})$, depending on the model [11, 16].

Thus, it might be possible that a quantum algorithm could use the small M to decrease the number of queries. While unlikely, this cannot be ruled out. Remember that classically, sorting requires $\Omega(N \log_2 N)$ steps in the general case but only $O(N)$ steps if the items to be sorted are all from the set $\{1,\ldots,N\}$ (Bucket Sort, [13]).

In this paper, we show that the collision and element distinctness problems require $\Omega(N^{2/3})$ and $\Omega(N^{1/3})$ queries even if the range M is equal to N. Our result follows from a general result on the polynomial degree of Boolean functions.

We show that, for any symmetric property ϕ of functions $f : \{1,2,\ldots,N\} \to \{1,2,\ldots,M\}$, its polynomial degree is the same for all $M \geq N$. The polynomial degree of ϕ provides a lower bound for both classical and quantum query complexity. (This was first shown by Nisan and Szegedy [20] in the classical case and then extended to the quantum case by Beals et al. [6] for $M = 2$ and Aaronson [1, 2] for $M > 2$.) Thus, one can prove lower bounds on quantum query complexity of a function ϕ by lower-bounding the polynomial degree of ϕ. This is known as the *polynomials method* for proving quantum lower bounds [6, 10, 2].

Our result means that, if we have a quantum lower bound for a symmetric property ϕ shown by the polynomials method for some range size M, we also have the same quantum lower bound for all $M \geq N$. As particular cases, we get lower bounds on the collision and element distinctness problems with small range. Since many quantum lower bounds are shown using the polynomials method, our result may have other applications.

A corollary of our lower bound on element distinctness with small range is that the polynomial degree of the two-level AND–OR tree on N^2 variables is $\Omega(N^{2/3})$. This improves over the previously known lower bound of $\Omega(\sqrt{N \log N})$ by Shi [21].

Related work. The $\Omega(N^{1/3})$ lower bound for the collision problem with small range was independently discovered by the author of this paper and Kutin [17], at about the same time, with completely different proofs. Kutin [17] takes the proof of the $\Omega(N^{1/3})$ lower bound for the collision problem with a large range [2] and changes it so that it works for all $M \geq N$. Our result is more general because it applies to any symmetric property and any lower bound shown by the polynomials method. On the other hand, Kutin's proof has the advantage that it also simplifies the lower bound for the collision problem with large range by Aaronson and Shi [2].

2 Preliminaries

2.1 Quantum query model

Let $[k]$ denote the set $\{1,\ldots,k\}$. Let $\mathcal{F}(N,M)$ be the set of all $f:[N]\to[M]$. We are given a function $f\in\mathcal{F}(N,M)$ by an oracle that answers queries. In one query, we can give i to the oracle and it returns $f(i)$ to us.

We would like to know whether f has a certain property (for example, whether f is one-to-one). More formally, we would like to compute a partial function $\phi:\mathcal{F}'\to\{0,1\}$, where $\mathcal{F}'\subseteq\mathcal{F}(N,M)$. In particular, we are interested in the following two properties:

Problem 1: Collision. $\phi(f)=1$ if the input function f is one-to-one. $\phi(f)=0$ if f is two-to-one (i.e., if, for every $k\in[M]$, there are either zero or two $x\in[N]$ satisfying $f(x)=k$). $\phi(f)$ is undefined for all other f.

Problem 2: Element distinctness. $\phi(f)=1$ if the input function f is one-to-one. $\phi(f)=0$ if there exist i,j, $i\neq j$, $f(i)=f(j)$.

A quantum algorithm with T queries is a sequence of unitary transformations

$$U_0\to O_f\to U_1\to O_f\to\cdots\to U_{T-1}\to O_f\to U_T.$$

The U_j's can be arbitrary unitary transformations that do not depend on $f(1),\ldots,f(N)$. O_f is a query (oracle) transformation. To define O_f, we represent basis states as $|i,b,z\rangle$ where i consists of $\lceil\log N\rceil$ bits, b is $\lceil\log M\rceil$ bits and z consists of all other bits. Then, O_f maps $|i,b,z\rangle$ to $|i,(b+f(i)) \bmod M,z\rangle$.

The computation starts with a state $|0\rangle$. Then, we apply U_0, O_f, ..., O_f, U_T and measure the final state. The result of the computation is the rightmost bit of the state obtained by the measurement.

The quantum algorithm computes ϕ with error ε if, for every $f\in\mathcal{F}(N,M)$ such that $\phi(f)$ is defined, the probability that the rightmost bit of $U_TO_fU_{T-1}\cdots O_fU_0|0\rangle$ equals $\phi(f)$ is at least $1-\varepsilon$. (Throughout this paper, ε is an arbitrary but fixed value, with $0<\varepsilon<1/2$.)

2.2 Polynomial lower bound

We can describe a function $f:[N]\to[M]$ by $N\times M$ Boolean variables y_{ij} which are 1 if $f(i)=j$ and 0 otherwise. Let $y=(y_{11},\ldots,y_{NM})$.

Definition 2.1. We say that a polynomial P ε-approximates the property ϕ if

1. $\phi(f)=1$ implies $1-\varepsilon\le P(y)\le 1$ for $y=(y_{11},\ldots,y_{NM})$ corresponding to f;
2. $\phi(f)=0$ implies $0\le P(y)\le\varepsilon$ for $y=(y_{11},\ldots,y_{NM})$ corresponding to f;
3. If $\phi(f)$ is undefined, then $0\le P(y)\le 1$ for the corresponding y.

A polynomial P approximates f if it ε-approximates f for some fixed $\varepsilon<1/2$.

The polynomial P is allowed to take any value if y does not correspond to any f. (This happens if for some $i\in[N]$ there is no or there is more than one $j\in[M]$ with $y_{ij}=1$.)

Lemma 2.2 ([1, 2]). *If a quantum algorithm computes ϕ with error ε using T queries then there is a polynomial $P(y_{11},\ldots,y_{NM})$ of degree at most $2T$ that ε-approximates ϕ.*

A lower bound on the number of queries can be then shown by proving that such a polynomial P does not exist. For the collision and element distinctness problems, we have

Theorem 2.3 ([22, 2]).

1. *If a polynomial P approximates the collision property for $M \geq \frac{3N}{2}$, the degree of P is $\Omega(N^{1/3})$;*
2. *If a polynomial P approximates the element distinctness property for $M = \Omega(N^2)$, the degree of P is $\Omega(N^{2/3})$;*

Note. More precisely, Shi [22, 2] proved that any polynomial approximating another problem, the *half two-to-one* problem, has degree $\Omega(N^{1/3})$. He then used that to deduce that $\Omega(N^{1/3})$ and $\Omega(N^{2/3})$ quantum queries are needed for the collision problem (when $M \geq \frac{3N}{2}$) and the element distinctness problem (when $M = \Omega(N^2)$). His proof can be easily modified to show a lower bound on the degree of polynomials approximating the collision and element distinctness problems.

By Theorem 2.3, $\Omega(N^{1/3})$ and $\Omega(N^{2/3})$ queries are required to solve the collision problem and element distinctness problem if the range M is sufficiently large. Previously, only weaker lower bounds of $\Omega(N^{1/4})$ [2] and $\Omega(\sqrt{N\log N})$ [16] were known if $M = N$.

3 Results

We call a property ϕ *symmetric* if, for any $\pi \in S_N$ and $\sigma \in S_M$,

$$\phi(f) = \phi(\sigma f \pi).$$

That is, $\phi(f)$ should remain the same if we permute the input set $\{1,\ldots,N\}$ before applying f or permute the output set $\{1,\ldots,M\}$ after applying f. The collision and element distinctness properties are both symmetric.

Our main result is

Theorem 3.1. *Let $\phi : \mathcal{F}' \to \{0,1\}$ be a symmetric property defined on a set of functions $\mathcal{F}' \subseteq \mathcal{F}(N,M)$. Let ϕ' be the restriction of ϕ to $f : [N] \to [N]$. Then, the minimum degree of a polynomial (in y_{ij}, $i \in [N]$, $j \in [M]$) approximating ϕ is equal to the minimum degree of a polynomial (in y_{ij}, $i \in [N]$, $j \in [N]$) approximating ϕ'.*

Theorems 2.3 and 3.1 imply that $\Omega(N^{1/3})$ and $\Omega(N^{2/3})$ queries are needed to solve the collision and element distinctness problems, even if $M = N$. (For $M < N$, these problems do not make sense because they both involve f being one-to-one as one of the cases.)

The proof of Theorem 3.1 is in two steps.

1. We describe a different way to describe an input function f by variables z_1, ..., z_M instead of $y_{11},\ldots,y_{NM}$. We prove that a polynomial of degree k in z_1, ..., z_M exists if and only if a polynomial of degree k in y_{11}, ..., y_{NM} exists.

2. We show that a polynomial $Q(z_1,\ldots,z_M)$ for $M > N$ exists if and only $Q(z_1,\ldots,z_N)$ exists.

The first step can be useful on its own. The representation of f by y_{11}, ..., y_{NM} gave the lower bounds of [2]. The new representation by z_1, ..., z_N might yield new lower bounds that are easier to prove using this approach.

3.1 New polynomial representation

We introduce variables z_1, ..., z_M, with $z_j = |f^{-1}(j)|$ (equivalently, $z_j = |\{i \mid y_{ij} = 1\}|$). We say that a polynomial Q in $z_1,\ldots,z_M$ approximates ϕ if it satisfies requirements similar to Definition 2.1. ($Q \in [1-\varepsilon,1]$ if $\phi(f) = 1$, $Q \in [0,\varepsilon]$ if $\phi(f) = 0$, and $Q \in [0,1]$ if $z_1,\ldots,z_M$ correspond to $f \in \mathcal{F}(N,M)$ for which $\phi(f)$ is not defined.)

Example 3.2. A polynomial $Q(z_1,\ldots,z_M)$ approximates the collision property if:

1. $Q(z_1,\ldots,z_M) \in [1-\varepsilon,1]$ if N of the variables z_1, ..., z_M are 1 and the remaining $M-N$ variables are 0;
2. $Q(z_1,\ldots,z_M) \in [0,\varepsilon]$ if $\frac{N}{2}$ of the variables z_1, ..., z_M are 2 and the remaining $M-\frac{N}{2}$ variables are 0;
3. $Q(z_1,\ldots,z_M) \in [0,1]$ if $z_1,\ldots,z_M$ are non-negative integers and $z_1+\cdots+z_M = N$.

Example 3.3. A polynomial $Q(z_1,\ldots,z_M)$ approximates element distinctness if:

1. $Q(z_1,\ldots,z_M) \in [1-\varepsilon,1]$ if N of the variables z_1, ..., z_M are 1 and the remaining $M-N$ variables are 0;
2. $Q(z_1,\ldots,z_M) \in [0,\varepsilon]$ if $z_1,\ldots,z_M$ are non-negative integers, $z_1+\cdots+z_M = N$, and $z_i > 1$ for some i.

In both cases, there is no restriction on $Q(z_1,\ldots,z_M)$ when $z_1+\cdots+z_M \neq N$ because such z_1, ..., z_M do not correspond to any $f : [N] \to [M]$.

Lemma 3.4. *Let $\phi : \mathcal{F}' \to \{0,1\}$, $\mathcal{F}' \subseteq \mathcal{F}(N,M)$ be symmetric. Then, the following two statements are equivalent:*

(1) There exists a polynomial Q of degree at most k in $z_1,\ldots,z_M$ approximating ϕ;

(2) There exists a polynomial P of degree at most k in y_{11}, ..., y_{NM} approximating ϕ.

Proof. To see that (1) implies (2), we substitute $z_j = y_{1j}+y_{2j}+\ldots+y_{Nj}$ into Q and obtain a polynomial in y_{ij} with the same approximation properties. Next, we show that (2) implies (1).

Let $P(y_{11},\ldots,y_{NM})$ be a polynomial approximating ϕ. We define $Q(z_1,\ldots,z_M)$ as follows. Let S be the set of all $y = (y_{11},\ldots,y_{NM})$ corresponding to functions $f : [N] \to [M]$ with the property that, for every $i \in [M]$ the number of j with $f(j) = i$ is exactly z_i. We define $Q(z_1,\ldots,z_M)$ as the expectation of $P(y_{11},\ldots,y_{NM})$ when $y = (y_{11},\ldots,y_{NM})$ is picked uniformly at random from S. (An equivalent way to define Q is to fix one function f with this property and to define Q as the expectation of of

$P(y_{11},\ldots,y_{NM})$, for $y=(y_{11},\ldots,y_{NM})$ corresponding to the function $f\pi$, with π being a random element of S_N.)

Since ϕ is symmetric, we have $\phi(f)=\phi(f\pi)$. Therefore, if $P(y_{11},\ldots,y_{NM})$ approximates ϕ, then $Q(z_1,\ldots,z_M)$ also approximates ϕ.

It remains to prove that Q is a polynomial of degree at most k in $z_1,\ldots,z_M$. Let

$$I=y_{i_1j_1}y_{i_2j_2}\cdots y_{i_kj_k}$$

be a monomial of P. It suffices to prove that each $E[I]$ is a polynomial of degree at most k because $E[P]$ is the sum of $E[I]$ over all I.

We can assume that i_ℓ for $\ell\in\{1,\ldots,k\}$ are all distinct. (If the monomial I contains two variables y_{ij} with the same i, j, one of them is redundant because $y_{ij}^2=y_{ij}$. If I contains y_{ij}, $y_{ij'}$, $j\neq j'$, then $y_{ij}y_{ij'}=0$ because $f(i)$ cannot be equal j and j' at the same time. Then, $I=0$.) We have

$$E[I]=\Pr[y_{i_1j_1}=1]\prod_{\ell=2}^{k}\Pr\left[y_{i_\ell j_\ell}=1\mid y_{i_1j_1}\cdots y_{i_{\ell-1}j_{\ell-1}}=1\right] .$$

There are N variables y_{ij_1}. Out of them, z_{j_1} variables are equal to 1 and each y_{ij_1} is equally likely to be 1. Therefore,

$$\Pr[y_{i_1j_1}=1]=\frac{z_{j_1}}{N} .$$

Furthermore, let s_ℓ be the number of $\ell'<\ell$ such that $j_\ell=j_{\ell'}$. Then,

$$\Pr\left[y_{i_\ell j_\ell}=1\mid y_{i_1j_1}\cdots y_{i_{\ell-1}j_{\ell-1}}=1\right]=\frac{z_{j_\ell}-s_\ell}{N-\ell-1}$$

because, once we have set $y_{i_1j_1}=1$, …, $y_{i_{\ell-1}j_{\ell-1}}=1$, we have also set all other y_{i_1j}, …, $y_{i_{\ell-1}j}$ to 0. Then, we have $N-\ell-1$ variables y_{ij_ℓ} which are not set yet and, out of them, $z_{j_\ell}-s_\ell$ must be 1.

Therefore, $E[I]$ is a product of k terms, each of which is a linear function of $z_1,\ldots,z_M$. This means that $E[I]$ is a polynomial in $z_1,\ldots,z_M$ of degree k. This completes the proof of the lemma. □

3.2 Lower bound for properties with small range

We now finish the proof of Theorem 3.1. Obviously, the minimum degree of a polynomial approximating ϕ' is at most the minimum degree of a polynomial approximating ϕ (because we can take a polynomial approximating ϕ and obtain a polynomial approximating ϕ' by restricting it to variables y_{ij}, $j\in[N]$). In the other direction, we can take a polynomial P' approximating ϕ' and obtain a polynomial Q' in $z_1,\ldots,z_N$ approximating ϕ' by Lemma 3.4. We then construct a polynomial Q in $z_1,\ldots,z_M$ of the same degree approximating ϕ. After that, using Lemma 3.4 in the other direction gives us a polynomial P in $y_{11},\ldots,y_{NM}$ approximating ϕ.

It remains to construct Q from Q'. For that, we can assume that Q' is symmetric w.r.t. permuting $z_1,\ldots,z_N$. (Otherwise, replace Q' by the expectation of $Q'(z_{\pi(1)},\ldots,z_{\pi(N)})$, where π is a uniformly random permutation of $\{1,2,\ldots,N\}$.) Since Q' is symmetric, it is a sum of elementary symmetric polynomials

$$Q'_{c_1,\ldots,c_l}=\sum_{i_1,\ldots,i_l\in[N]} z_{i_1}^{c_1}z_{i_2}^{c_2}\cdots z_{i_l}^{c_l} .$$

Let Q be the sum of elementary symmetric polynomials in $z_1,\ldots,z_M$ with the same coefficients.

We claim that Q approximates ϕ. To see this, consider an input function $f:[N]\to[M]$. There are at most N values $j\in\{1,\ldots,M\}$ such that there exists $i\in\{1,\ldots,N\}$ with $f(i)=j$. This means that, out of M variables $z_1,\ldots,z_M$ corresponding to f, at most N are nonzero.

Consider a permutation $\pi\in S_M$ that maps all $i\in[M]$ with $z_i\neq 0$ to $\{1,\ldots,N\}$. Let $f'=\pi f$. Since ϕ is symmetric, $\phi(f)=\phi(f')$. Since f' is a function from $[N]$ to $[N]$, Q' correctly approximates ϕ on f'. Since $Q(z_1,\ldots,z_N,0,\ldots,0)=Q'(z_1,\ldots,z_N)$, Q also correctly approximates ϕ on f'. Since Q is symmetric w.r.t. permutations of $z_1,\ldots,z_M$, Q approximates ϕ on the input function f as well. This completes the proof of Theorem 3.1.

3.3 Lower bound on the polynomial degree of the AND–OR tree

As a by-product, our result provides a better lower bound on the polynomial degree of a well-studied Boolean function.

This Boolean function is the two-level AND–OR tree on N^2 variables. Let $x_1,\ldots,x_{N^2}\in\{0,1\}$ be the variables. We split them into N groups, with the i^{th} group consisting of $x_{(i-1)N+1}$, $x_{(i-1)N+2}$, …, x_{iN}. The AND–OR function $g(x_1,\ldots,x_{N^2})$ is defined as

$$g(x_1,\ldots,x_{N^2})=\bigwedge_{i=1}^{n}\bigvee_{j=(i-1)N+1}^{iN}x_j\ .$$

A polynomial $p(x_1,\ldots,x_{N^2})$ approximates g if $0\le p(x_1,\ldots,x_{N^2})\le\varepsilon$ whenever $g(x_1,\ldots,x_{N^2})=0$ and $1-\varepsilon\le p(x_1,\ldots,x_{N^2})\le 1$ whenever $g(x_1,\ldots,x_{N^2})=1$ (similarly to Definition 2.1).

It has been an open problem to determine the minimum degree of a polynomial approximating the two-level AND–OR tree. The best lower bound is $\Omega(\sqrt{N\log N})$ by Shi [21], while the best upper bound is $O(N)$. (Curiously, the quantum query complexity of this problem is known. It is $\Theta(N)$, as shown by [9, 3]. If the polynomial degree is $o(N)$, this would be the second example of a Boolean function with a gap between the polynomial degree and quantum query complexity, with the first example being the iterated functions in [4].) We show

Theorem 3.5. *Any polynomial approximating g has degree $\Omega(N^{2/3})$.*

Proof. Consider the element distinctness problem for $M=N$. An instance of this problem, $f\in\mathcal{F}(N,N)$ can be described by N^2 variables $y_{11},\ldots,y_{NN}$ (as shown in Section 2.2).

The values of the function, $f(1)$, $f(2)$, …, $f(N)$, are all distinct if and only if, for each $j\in[N]$, there exists $i\in[N]$ with $f(i)=j$. This, in turn, is equivalent to saying that, for each $i\in[N]$, one of the variables $y_{1i},y_{2i},\ldots,y_{Ni}$ is equal to 1.

Assume we have a polynomial $P(x_1,\ldots,x_{N^2})$ of degree d approximating the two-level AND–OR tree function g. Consider the polynomial $Q(y_{11},\ldots,y_{NN})$ obtained from P by replacing $x_{(i-1)N+j}$ with y_{ji}. If the N values $f(i)$ are all distinct, then, for each $j\in\{1,\ldots,N\}$, there exists i such that $f(i)=j$. Therefore, one of the variables $y_{1j},\ldots,y_{Nj}$ is 1 and the OR of those variables is also 1. This means that the AND–OR function $g(x_1,\ldots,x_{N^2})$ is equal to 1. If the values $f(i)$ are not all distinct, then there exists $j\in[N]$ such that there is no i with $f(i)=j$. Then, $y_{1i},y_{2i},\ldots,y_{Ni}$ are all 0, implying that $g(x_1,\ldots,x_{N^2})=0$ for the corresponding assignment x_1, …, x_{N^2}.

This means that Q approximates the element distinctness property, in the sense of Section 2.2. Since degree $\Omega(N^{2/3})$ is required to approximate element distinctness, $d = \Omega(N^{2/3})$. □

4 Conclusion

We have shown that, for any symmetric property of functions $f : [N] \to [M]$, its polynomial degree is the same for all $M \geq N$. Thus, if we prove a lower bound for the degree for some large M, this immediately implies the same bound for $M = N$. Since the polynomial degree is a lower bound for quantum query complexity, this can be used to show quantum lower bounds. As particular cases of our result, we get that the collision problem has degree $\Omega(N^{1/3})$ and that the element distinctness problem has degree $\Omega(N^{2/3})$, even if $M = N$. This implies $\Omega(N^{1/3})$ and $\Omega(N^{2/3})$ quantum lower bounds on these problems for $M = N$.

A part of our result is a new representation for polynomials describing properties of functions $f : [N] \to [M]$. This new description might be useful for proving new quantum lower bounds. We conclude with two open problems.

Modified element distinctness problem. Say we are given $f : [N] \to [N]$ and we are promised that either f is one-to-one or there are i, j, k such that $f(i) = f(j) = f(k)$. We would like to know which of these two is the case. What is the quantum query complexity of this problem?

The problem is quite similar to element distinctness in which we have to distinguish one-to-one function from one having $f(i) = f(j)$ for some i, j with $i \neq j$. The known $O(N^{2/3})$ quantum algorithm still applies, but the $\Omega(N^{2/3})$ quantum lower bound of [2] (by a reduction from the collision problem) breaks down. The best lower bound that we can prove is $\Omega(N^{1/2})$ by a reduction from Grover's search. Improving this bound to $\Omega(N^{2/3})$ is an open problem.

This problem is also similar to element distinctness if we look at it in our new z_1, ..., z_M representation. For element distinctness, a polynomial Q must satisfy $Q(1,\ldots,1) \in [1-\varepsilon, 1]$ and $Q(z_1,\ldots,z_N) \in [0,\varepsilon]$ if $z_1 + \cdots + z_N = N$ and $z_i \geq 2$ for some i. For our new problem, we must have $Q(1,\ldots,1) \in [1-\varepsilon, 1]$ and $Q(z_1,\ldots,z_N) \in [0,\varepsilon]$ if $z_1 + \cdots + z_N = N$ and $z_i \geq 3$ for some i. In the first case, degree $\Omega(N^{2/3})$ is needed [2]. In the second case, no such lower bound is known.

Polynomial degree vs. quantum query complexity for symmetric properties. Let ϕ be a symmetric property of functions $f : [N] \to [M]$. Let $\deg(\phi)$ be the minimum degree of a polynomial that ε-approximates f and $Q_2(\phi)$ be the minimum number of queries in a quantum query algorithm computing ϕ with error at most ε. Is it true that these two quantities are polynomially related: $Q_2(\phi) = O(\deg^c(\phi))$ for some constant c?

This open problem was first proposed by Aaronson [1, 2], regarding properties which are only symmetric with respect to permuting inputs to f: $\phi(f) = \phi(f\pi)$ for any $\pi \in S_N$. It remains open both in this case and in the case of properties having the more general symmetry considered in this paper ($\phi(f) = \phi(\sigma f \pi)$, for all $\pi \in S_N$ and $\sigma \in S_M$). It is known that $Q_2(\phi) = O(\deg^2(\phi))$ if $M = 2$.

References

[1] * S. AARONSON: Quantum lower bound for the collision problem. In *Proceedings of STOC'02*, pp. 635–642, 2002. [STOC:509907.509999, arXiv:quant-ph/0111102]. 1, 2.2, 4, 2

[2] * S. AARONSON AND Y. SHI: Quantum lower bounds for the collision and the element distinctness problems. *Journal of ACM*, 51(4):595–605, 2004. Earlier versions in [1] and [22]. [JACM:1008731.1008735]. 1, 2.2, 2.3, 2.2, 3, 4

[3] * A. AMBAINIS: Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64:750–767, 2002. [JCSS:10.1006/jcss.2002.1826, arXiv:quant-ph/0002066]. 3.3

[4] * A. AMBAINIS: Polynomial degree vs. quantum query complexity. In *Proceedings of FOCS'03*, pp. 230–239, 2003. [FOCS:10.1109/SFCS.2003.1238197, arXiv:quant-ph/0305028]. 3.3

[5] * A. AMBAINIS: Quantum walk algorithm for element distinctness. In *Proceedings of FOCS'04*, pp. 22–31, 2004. [FOCS:10.1109/FOCS.2004.54, arXiv:quant-ph/0311001]. 1

[6] * R. BEALS, H. BUHRMAN, R. CLEVE, M. MOSCA, AND R. DE WOLF: Quantum lower bounds by polynomials. *Journal of ACM*, 48:778–797, 2001. Earlier version at FOCS'98. [JACM:502090.502097, arXiv:quant-ph/9802049]. 1

[7] * G. BRASSARD, P. HØYER, AND A. TAPP: Quantum algorithm for the collision problem. *SIGACT News*, 28:14–19, 1997. [arXiv:quant-ph/9705002]. 1

[8] * G. BRASSARD, P. HØYER, AND A. TAPP: Quantum counting. In *Proceedings of ICALP'98*, pp. 820–831, 1998. [ICALP:ap2mrf08d8gcqppe, arXiv:quant-ph/9805082]. 1

[9] * H. BUHRMAN, R. CLEVE, AND A. WIGDERSON: Quantum vs. classical communication and computation. In *Proceedings of STOC'98*, pp. 63–68, 1998. [STOC:276698.276713]. 3.3

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

[11] * H. BUHRMAN, C. DURR, M. HEILIGMAN, P. HØYER, F. MAGNIEZ, M. SANTHA, AND R. DE WOLF: Quantum algorithms for element distinctness. In *16th IEEE Annual Conference on Computational Complexity (CCC'01)*, pp. 131–137, 2001. [CCC:10.1109/CCC.2001.933880, arXiv:quant-ph/0007016]. 1

[12] * H. BUHRMAN AND R. ŠPALEK: Quantum verification of matrix products. [arXiv:quant-ph/0409035]. 1

[13] * T. CORMEN, C. LEISERSON, R. RIVEST, AND C. STEIN: *Introduction to Algorithms, 2nd edition*. The MIT Press and McGraw-Hill Book Company, 2001. 1

[14] * L. GROVER: A fast quantum mechanical algorithm for database search. In *Proceedings of STOC'96*, pp. 212–219, 1996. [STOC:237814.237866, arXiv:quant-ph/9605043]. 1

[15] * L. GROVER: A framework for fast quantum mechanical algorithms. In *Proceedings of STOC'98*, pp. 53–62, 1998. [STOC:276698.276712, arXiv:quant-ph/9711043]. 1

[16] * P. HOYER, J. NEERBEK, AND Y. SHI: Quantum lower bounds of ordered searching, sorting and element distinctness. *Algorithmica*, 34:429–448, 2002. Earlier version at ICALP'01. [Algorithmica:25gl9elr5rxr3q6a, arXiv:quant-ph/0102078]. 1, 2.2

[17] * S. KUTIN: Quantum lower bound for the collision problem. *Theory of Computing*, 1(2):29–36, 2005. [ToC:v001/a002, arXiv:quant-ph/0304162]. 1

[18] * F. MAGNIEZ, M. SANTHA, AND M. SZEGEDY: Quantum algorithms for the triangle problem. In *Proceedings of SODA'05*, 2005. [arXiv:quant-ph/0310134]. 1

[19] * A. NAYAK AND F. WU: The quantum query complexity of approximating the median and related statistics. In *Proceedings of STOC'99*, pp. 384–393, 1999. [STOC:301250.301349, arXiv:quant-ph/9804066]. 1

[20] * N. NISAN AND M. SZEGEDY: On the degree of boolean functions as real polynomials. *Computational Complexity*, 4:301–313, 1994. 1

[21] * Y. SHI: Approximating linear restrictions of boolean functions. Manuscript. 1, 3.3

[22] * Y. SHI: Quantum lower bounds for the collision and the element distinctness problems. In *Proceedings of FOCS'02*, pp. 513–519, 2002. [FOCS:10.1109/SFCS.2002.1181975, arXiv:quant-ph/0112086]. 2.3, 2.2, 2

AUTHOR

Andris Ambainis
Department of Combinatorics and Optimization
Faculty of Mathematics
200 University Avenue West
Waterloo, ON N2L 3G1, Canada
ambainis@math.uwaterloo.ca
http://www.math.uwaterloo.ca/~ambainis

ABOUT THE AUTHOR

Andris Ambainis was born in Daugavpils, Latvia. After undergraduate studies at the University of Latvia, he received his Ph. D. from the University of California, Berkeley in 2001, supervised by Umesh Vazirani. Andris joined the University of Waterloo in August 2004. His research interests include quantum algorithms, quantum complexity theory, quantum cryptography as well as the classical theory of computation.