THEORY OF COMPUTING, Volume 1 (2005), pp. 29–36
http://theoryofcomputing.org

Quantum Lower Bound for the Collision Problem with Small Range

Samuel Kutin

Received: December 8, 2004; published: April 14, 2005.

Abstract: We extend Aaronson and Shi's quantum lower bound for the r-to-one collision problem. An r-to-one function is one where every element of the image has exactly r preimages. The r-to-one collision problem is to distinguish between one-to-one functions and r-to-one functions over an n-element domain.

Recently, Aaronson and Shi proved a lower bound of $\Omega((n/r)^{1/3})$ quantum queries for the r-to-one collision problem. Their bound is tight, but their proof applies only when the range has size at least $3n/2$. We give a modified version of their argument that removes this restriction.

ACM Classification: F.1.2

AMS Classification: 81P68, 68Q17

Key words and phrases: quantum, query, oracle, collision, polynomial method, lower bound, small range

1 Introduction

How many quantum queries does it take to find a collision? A *collision* in a function is a pair of inputs that map to the same value. We consider the problem of finding a collision in an r-to-one function; i.e., a function where every element of the image has exactly r preimages. (We require that r be a divisor of n, the size of the input space.) The difficulty of this problem for a quantum computer has attracted much interest [1, 2, 4, 3, 6, 10].

In some cases, explicit information about a function may make it easier to find collisions. For example, if we know a function is periodic, we can find a collision using Shor's algorithm [11]. Rather than use such explicit information, we focus on a *black-box* model: our only access to the function is as a quantum oracle. Brassard, Høyer, and Tapp [6] use Grover's search [7] to find a collision in an r-to-one function in $O((n/r)^{1/3})$ quantum queries, an improvement over the $\Theta((n/r)^{1/2})$ classical queries needed. In this note, we are concerned with the matching lower bound.

For a lower bound, it is easier to consider a decision problem: the input function is guaranteed to be either one-to-one or r-to-one, and our task is to distinguish between these two cases. Aaronson [1] proved the first significant lower bound: $\Omega((n/r)^{1/5})$ queries.

More recently, Shi [10] proved a lower bound of $\Omega((n/r)^{1/3})$, given the additional condition that the size of the range of the function is at least $3n/2$. (In the case where the range is only n, Shi provides a lower bound of $\Omega((n/r)^{1/4})$). The proof is a novel application of the methods of Nisan and Szegedy [8] and Paturi [9] to the case where one cannot fully symmetrize the multivariate polynomials.

Our main result is a new version of this theorem, but without the additional constraint on the size of the range:

Theorem 1.1. *Let $n > 0$ and $r \geq 2$ be integers with $r \mid n$, and let a function from $[n]$ to $[n]$ be given as an oracle with the promise that it is either one-to-one or r-to-one. Then any quantum algorithm for distinguishing these two cases must evaluate the function $\Omega\left((n/r)^{1/3}\right)$ times.*

The argument is very similar to that of Aaronson and Shi. (See [2] for a combined version of [1] and [10].) As stated above, we remove the requirement that the range be at least $3n/2$. Our proof is conceptually simpler for other reasons:

1. The natural automorphism group on the set of functions from $[n]$ to $[N]$ is $S_n \times S_N$. Our argument symmetrizes with respect to the entire group.

2. For technical reasons, Shi introduces an additional decision problem called Half-r-to-one, where one must distinguish between r-to-one functions and functions that are r-to-one on half the domain and one-to-one on the other half. We avoid using this Half-r-to-one problem.

An independent approach

Independent of this work, Ambainis [4] gave an alternate proof of Theorem 1.1. His approach is more general: he shows that, given any lower bound for a symmetric function property with a restriction on the size of the range, we can remove that restriction.

Ambainis's work, together with Shi's paper, implies Theorem 1.1. It is worth noting another consequence of those two papers: Aaronson and Shi prove that, given a black-box function f on n inputs whose range has size $\Omega(n^2)$, it takes $\Omega(n^{2/3})$ queries to determine if f is one-to-one. Theorem 1.1 implies a similar result; the constant hidden in the $\Omega(n^2)$ term improves, but the dependence on n does not. Neither Aaronson and Shi [2] nor this paper gives a lower bound for element distinctness with small range.

However, Ambainis's work gives a lower bound of $\Omega(n^{2/3})$ without any range restriction. Ambainis has also given a matching upper bound [3].

2 Preliminaries

2.1 Functions as quantum oracles

Let $n, N > 0$ be integers. Let $\mathcal{F}(n,N)$ be the set of functions from $[n]$ to $[N]$.

A function is given to us as a quantum oracle. We can perform a transformation O_f, which applies f to the contents of some of the quantum state:

$$O_f |i,j,z\rangle = |i, f(i)+j \pmod N, z\rangle \quad .$$

Here z is a placeholder for the unaffected portion of the quantum state.

The query complexity of a quantum algorithm is the number of times it calls O_f. We think of our algorithm as alternating between $T+1$ unitary operators and T applications of O_f.

Let $\delta_{i,j}(f)$ be 1 when $f(i) = j$ and 0 otherwise. Then, after T queries, the amplitude of each quantum base state is a degree-T polynomial in these $\delta_{i,j}(f)$. Hence, the acceptance probability $P(f)$ is a polynomial over $\delta_{i,j}$ of degree at most $2T$. The connection between quantum complexity and polynomial degree is due to Beals, et al. [5]; the application to functions using variables $\delta_{i,j}$ is due to Aaronson [1].

Note that this polynomial $P(f)$ is constrained to be in the interval $[0,1]$ whenever the $\delta_{i,j}$ correspond to a valid input; i.e.,

$$\begin{aligned} \forall i, j, \qquad & \delta_{i,j} \in \{0,1\} \quad , \\ \forall i, \qquad & \sum_j \delta_{i,j} = 1 \quad . \end{aligned} \tag{2.1}$$

The connection between polynomial degree and query complexity was first made by Nisan and Szegedy [8]. In their applications, they symmetrize over all permutations of the variables, reducing the multivariate polynomial to a univariate polynomial. They then apply results from approximation theory to prove a lower bound on the degree of the polynomial. Beals, et al. [5] follow the same approach.

A nice, general version of the approximation theory results was shown by Paturi [9]. Following Shi [10], we use a slight modification of Paturi's theorem:

Theorem 2.1 (Paturi). *Let $q(\alpha) \in \mathbb{R}[\alpha]$ be a polynomial of degree d. Let a and b be integers, $a < b$, and let $\xi \in [a,b]$ be a real number. If*

1. *$|q(i)| \le c_1$ for all integers $i \in [a,b]$, and*
2. *$|q(\lceil \xi \rceil) - q(\xi)| \ge c_2$ for some constant $c_2 > 0$,*

then

$$d = \Omega(\sqrt{(\xi - a + 1)(b - \xi + 1)}) \quad ,$$

where the hidden constant depends on c_1 and c_2.

Note that, if the conditions of the theorem are met for any ξ, we have $d = \Omega(\sqrt{b-a})$. If they are met for some $\xi \approx (a+b)/2$, then $d = \Omega(b-a)$.

In our setting, the automorphism group for the variables $\delta_{i,j}$ is $S_n \times S_N$. If we symmetrize with respect to this group, we do not immediately obtain a univariate polynomial. Hence, we will have to work harder to apply Theorem 2.1.

For $\sigma \in S_n$, $\tau \in S_N$, we define $\Gamma^{\sigma}_{\tau} : \mathcal{F}(n,N) \to \mathcal{F}(n,N)$ by

$$\Gamma^{\sigma}_{\tau}(f) = \tau \circ f \circ \sigma \quad .$$

Let $P(f)$ be an acceptance polynomial as above. We can write P as a sum $\sum_S C_S I_S(f)$, where S ranges over subsets of $[n] \times [N]$ and

$$I_S = \prod_{(i,j)\in S} \delta_{i,j} \quad .$$

By (2.1), we may assume that each pair $(i,j) \in S$ has a distinct value of i; we thus write

$$I_S = \prod_{k=1}^{t} \prod_{i \in S_k} \delta_{i,j_k} \quad , \tag{2.2}$$

where the sets S_k are disjoint. The degree of the monomial is $\sum_k |S_k|$.

2.2 Some special functions

We now define a collection of functions which are a-to-one on part of the domain, and b-to-one on the rest of the domain. (These will enable us to interpolate between one-to-one and r-to-one functions.)

Fix $N \geq n > 0$. We say that a triple (m,a,b) of integers is *valid* if $0 \leq m \leq n$, $a \mid m$, and $b \mid (n-m)$. For any such valid triple, we have a function $f_{m,a,b} \in \mathcal{F}(n,N)$, given by

$$f_{m,a,b} = \begin{cases} \lceil i/a \rceil & 1 \leq i \leq m \ , \\ N - \lfloor (n-i)/b \rfloor & m < i \leq n \ . \end{cases}$$

So $f_{m,a,b}$ is a-to-one on m points, and b-to-one on the remaining $n-m$ points. (Since $N \geq n$, the two parts of the range do not overlap.)

Note that our $f_{m,a,b}$ plays the same role as Aaronson and Shi's $f_{m,g}$, with $a = g$ and $b = 2$.

We now examine the behavior of $f_{m,a,b}$ after we symmetrize by all of $S_n \times S_N$.

Lemma 2.2. *Let $P(f)$ be a degree-d polynomial in $\delta_{i,j}$. For a valid triple (m,a,b), define $Q(m,a,b)$ by*

$$Q(m,a,b) = \mathbf{E}_{\sigma,\tau}\left[P\left(\Gamma^{\sigma}_{\tau}(f_{m,a,b})\right)\right] \quad .$$

Then Q is a degree-d polynomial in m,a,b.

The key new step in this paper lies in the proof of Lemma 2.2. To show that the expected value $Q(m,a,b)$ is a polynomial, we break down S_N into a union of disjoint events A_U. We then write $Q(m,a,b)$ as a sum over all U, and we show that each term in the sum is a polynomial in m, a, and b.

Definition 2.3. For integers k,ℓ, let $\ell^{\underline{k}}$ denote the falling power $\ell(\ell-1)\cdots(\ell-k+1)$.

Proof of Lemma 2.2. It suffices to prove the lemma in the case where P is a monomial I_S. We write I_S in the form (2.2); then $d = |S|$. We write $s_k = |S_k|$.

For each subset $U \subseteq [t]$, let A_U be the following event: for each $k \in U$, $\tau^{-1}(j_k) \leq m/a$; for each $k \notin U$, $\tau^{-1}(j_k) \geq N-(n-m)/b+1$.

Clearly the events A_U are disjoint. If $I_S(\Gamma_\tau^\sigma(f_{m,a,b}))$ is nonzero, then every $\tau^{-1}(j_k)$ must lie in the range of $f_{m,a,b}$, so some event A_U must occur. Hence, we write

$$Q(m,a,b) = \sum_{U \subseteq [t]} \Pr(A_U) Q_U(m,a,b) \ ,$$

where

$$Q_U(m,a,b) = \mathbf{E}_{\sigma,\tau}\left[I_S(\Gamma_\tau^\sigma(f_{m,a,b})) \mid A_U\right] \ .$$

Choose some U, and let $u = |U|$. Then $\Pr(A_U)$ is given by

$$\Pr(A_U) = \frac{\left(\frac{m}{a}\right)^{\underline{u}} \left(\frac{n-m}{b}\right)^{\underline{t-u}}}{N^{\underline{t}}} \ ,$$

which is a rational function in m,a,b. The numerator has degree t, and the denominator is $a^u b^{t-u}$.

Also,

$$Q_U(m,a,b) = \frac{1}{n^{\underline{d}}} \prod_{k \in U} a^{\underline{s_k}} \prod_{k \notin U} b^{\underline{s_k}} \ .$$

This is a polynomial in a,b of degree d; furthermore Q_U is divisible by $a^u b^{t-u}$.

Hence, for each U, $\Pr(A_U)Q_U$ is a degree-d polynomial in m,a,b. Therefore $Q(m,a,b)$ is itself a degree-d polynomial. This concludes the lemma. □

3 Main Proof

We are now ready to prove Theorem 1.1.

Proof of Theorem 1.1. Let $\mathcal{A}$ be an algorithm which distinguishes one-to-one from r-to-one in T queries, and let $P(f)$ be the corresponding acceptance probability. $P(f)$ is a polynomial in $\delta_{i,j}$ of degree at most $2T$. Let $Q(m,a,b)$ be formed from P as in Lemma 2.2, and let $d = \deg Q$; we have $d \leq 2T$.

For any σ,τ, we know that $\Gamma_\tau^\sigma(f_{m,a,b})$ is a valid function. If $a = b$, this function is a-to-one. We conclude the following:

1. $0 \leq Q(m,a,b) \leq 1$ whenever (m,a,b) is a valid triple.
2. $0 \leq Q(m,1,1) \leq 1/3$ for any m.
3. $2/3 \leq Q(m,r,r) \leq 1$ for any m such that $r \mid m$.

The remainder of the proof consists of proving that $\deg Q = \Omega\left((n/r)^{1/3}\right)$. We will take $M \approx m/2$, and we will examine either the univariate polynomial $Q(M,1,rx)$ or $Q(M,rx,r)$ (depending on the value of $Q(M,1,r)$). If this polynomial remains bounded for large values of x, we can apply Theorem 2.1. Otherwise, we can use Theorem 2.1 on the first argument to Q. Either way, we get a lower bound on d.

For simplicity of exposition, we begin with the case $r = 2$. Let $M = 2\lfloor n/4 \rfloor$. We ask: is $Q(M,1,2) \geq 1/2$? In other words: does our algorithm accept (at least half the time) an input which is one-to-one on half the domain, and two-to-one on the other half?

Case I: $Q(M,1,2) \geq 1/2$. Let $g(x) = Q(M,1,2x)$, and let k be the least positive integer for which $|g(k)| \geq 2$. Then we have $g(x)$ between -2 and 2 for all positive integers $x < k$, and $g(1) - g(1/2) \geq 1/6$ by assumption. Let $c = 2k$. By Theorem 2.1, we have

$$d = \Omega(\sqrt{k}) = \Omega(\sqrt{c}) \quad . \tag{3.1}$$

Now, we consider the polynomial $h(i) = Q(n-ci,1,c)$. For any integer i in the range $0 \leq i \leq \lfloor n/c \rfloor$, the triple $(n-ci,1,c)$ is valid, so $0 \leq h(i) \leq 1$. But

$$\left| h\left(\frac{n-M}{c}\right) \right| = |Q(M,1,c)| = |g(k)| \geq 2 \quad .$$

We conclude, by Theorem 2.1, that

$$d = \Omega(n/c) \quad . \tag{3.2}$$

Case II: $Q(M,1,2) < 1/2$. Now, let $g(x) = Q(M,2x,2)$. Let k be the least positive integer for which $|g(k)| \geq 2$, and let $c = 2k$. We have $g(1) - g(1/2) \geq 1/6$; as in Case I, we obtain (3.1) using Theorem 2.1.

Next, we consider $h(i) = Q(ci,c,2)$. For any integer i in the range $0 \leq i \leq \lfloor n/c \rfloor$, the triple $(ci,c,2)$ is valid (both n and c are even), so $0 \leq h(i) \leq 1$. But $|h(M/c)| = |g(k)| \geq 2$. Again, as in Case I, we obtain (3.2) using Theorem 2.1.

In either case, we use (3.1) and (3.2) to obtain $d = \Omega(n^{1/3})$. We could divide into cases (depending on whether $c \geq n^{2/3}$), or we could simply square (3.1) and multiply by (3.2) to obtain $d^3 = \Omega(n)$.

For general r, the setup is almost identical: we let $M = r\left\lfloor \frac{n}{2r} \right\rfloor$ and split into cases based on whether $Q(M,1,r) \geq 1/2$.

Case I: $Q(M,1,r) \geq 1/2$. Let $g(x) = Q(M,1,rx)$, let k be the least positive integer for which $|g(k)| \geq 2$, and let $c = rk$. We have $g(1) - g(1/r) \geq 1/6$, so Theorem 2.1 yields

$$d = \Omega(\sqrt{k}) = \Omega(\sqrt{c/r}) \quad . \tag{3.3}$$

Next, we let $h(i) = Q(n-ci,1,c)$. As in the $r = 2$ analysis above, we conclude (3.2).

Case II: $Q(M,1,r) < 1/2$. Now, let $g(x) = Q(M,rx,r)$, let k be the least integer for which $|g(k)| \geq 2$, and let $c = rk$. We have $g(1) - g(1/r) \geq 1/6$; as in Case I, we obtain (3.3) using Theorem 2.1.

Next, we take $h(i) = Q(ci,c,r)$. For any integer i in the range $0 \leq i \leq \lfloor n/c \rfloor$, the triple (ci,c,r) is valid; note that $n - ci$ must be a multiple of r. But $|h(M/c)| = |g(k)| \geq 2$. So, as in the $r = 2$ analysis, we get (3.2).

In either case, we square (3.3) and multiply by (3.2) to obtain $d^3 = \Omega(n/r)$ as desired. □

Acknowledgments

The author thanks László Babai, Vincent Nesme, Natacha Portier, and the referees for helpful comments.

References

[1] * SCOTT AARONSON: Quantum lower bound for the collision problem. In *Proc. of the 34th ACM STOC*, pp. 635–642, 2002. [STOC:509907.509999, arXiv:quant-ph/0111102]. 1, 1, 2.1, 2

[2] * SCOTT AARONSON AND YAOYUN SHI: Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM*, 51(4):595–605, 2004. Based on [10] and [1]. [JACM:1008735]. 1, 1, 1

[3] * ANDRIS AMBAINIS: Quantum walk algorithm for element distinctness. In *Proc. of the 45th IEEE FOCS*, pp. 22–31, 2004. [FOCS:10.1109/FOCS.2004.54, arXiv:quant-ph/0311001]. 1, 1

[4] * ANDRIS AMBAINIS: Quantum lower bounds for collision and element distinctness with small range. *Theory of Computing*, 1(3), 2005. To appear. [ToC:v001/a003, arXiv:quant-ph/0305179]. 1, 1

[5] * BOB BEALS, HARRY BUHRMAN, RICHARD CLEVE, MICHELE MOSCA, AND RONALD DE WOLF: Quantum lower bounds by polynomials. In *Proc. of the 39th IEEE FOCS*, pp. 352–361, 1998. [FOCS:10.1109/SFCS.1998.743485, arXiv:quant-ph/9802049]. 2.1, 2.1

[6] * GILLES BRASSARD, PETER HØYER, AND ALAIN TAPP: *Quantum Cryptanalysis of Hash and Claw-Free Functions*, volume 1380 of *Lecture Notes in CS*, pp. 163–169. Springer-Verlag, 1998. [LATIN:11bhjthw46dxl2qa, arXiv:quant-ph/9805082]. 1

[7] * LOV K. GROVER: A fast quantum mechanical algorithm for database search. In *Proc. of the 28th ACM STOC*, pp. 212–219, 1996. [STOC:237814.237866, arXiv:quant-ph/9605043]. 1

[8] * NOAM NISAN AND MÁRIÓ SZEGEDY: On the degree of boolean functions as real polynomials. *Computational Complexity*, 4:301–313, 1994. [STOC:129757]. 1, 2.1

[9] * RAMAMOHAN PATURI: On the degree of polynomials that approximate symmetric boolean functions. In *Proc. of the 24th ACM STOC*, pp. 468–474, 1992. [STOC:129758]. 1, 2.1

[10] * YAOYUN SHI: Quantum lower bounds for the collision and the element distinctness problems. In *Proc. of the 43th IEEE FOCS*, pp. 513–519, 2002. [FOCS:10.1109/SFCS.2002.1181975, arXiv:quant-ph/0112086]. 1, 1, 2.1, 2

[11] * PETER W. SHOR: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997. [SICOMP:29317]. 1

AUTHOR

Samuel Kutin
Center for Communications Research
805 Bunn Drive
Princeton, NJ 08540
kutin@idaccr.org
http://www.kutin.com

ABOUT THE AUTHOR

Samuel (Sandy) Kutin received his Ph.D. from the University of Chicago in 2002 under Partha Niyogi and László Babai. His research interests include Boolean function complexity, computational learning theory, and quantum computing. Sandy and his family live in Princeton, where he spends his free time petting his cats, playing games, and participating in the National Puzzlers' League. He is proud to support Theory of Computing.