Adversary Lower Bound for the k-sum Problem

Aleksandrs Belovs* stiboh@gmail.com
Robert Špalek† spalek@google.com

Abstract

We prove a tight quantum query lower bound $\Omega(n^{k/(k+1)})$ for the problem of deciding whether there exist k numbers among n that sum up to a prescribed number, provided that the alphabet size is sufficiently large.

The technical version of the paper can be found at [9].

1 Introduction

Two main techniques for proving lower bounds on quantum query complexity are the polynomial method [6] developed by Beals *et al.* in 1998, and the adversary method [2] developed by Ambainis in 2000. Both techniques are incomparable. There are functions with adversary bound strictly larger than polynomial degree [3], as well as functions with the reverse relation.

One of the examples of the reverse relation is exhibited by the element distinctness function. The input to the function is a string of length n of symbols in an alphabet of size q, i.e., $x = (x_i) \in [q]^n$. We use notation $[q]$ to denote the set $\{1, \ldots, q\}$. The element distinctness function evaluates to 0 if all symbols in the input string are pairwise distinct, and to 1 otherwise.

The quantum query complexity of element distinctness is $\Theta(n^{2/3})$ with the algorithm given by Ambainis [5]. The tight lower bounds were given by Aaronson and Shi [1], Kutin [14] and Ambainis [4] using the polynomial method.

The adversary bound, however, fails for this function. The reason is that the function has 1-certificate complexity 2, and the so-called certificate complexity barrier [17, 18] states that the adversary method fails to achieve anything better than $\Omega(\sqrt{n})$ for any function with 1-certificate complexity bounded by a constant.

In 2006, a stronger, negative-weight version of the adversary bound was developed by Høyer *et al.* [12]. Later, the negative-weight adversary bound was proven to be optimal by Reichardt *et al.* [16, 15]. Despite this fact, it has almost never been used to prove lower bounds for explicit functions. Vast majority of lower bounds by the adversary method uses the old positive-weight version of this method. But since the only competing polynomial method is known to be non-tight, a better understanding of the negative-weight adversary method would be very beneficial. In the sequel, we consider the negative-weight adversary bound only, and we will omit the adjective "negative-weight".

In this paper we use the adversary method to prove a lower bound for the following variant of the knapsack packing problem. Let $\mathbb{G}$ be a finite Abelian group, and $t \in \mathbb{G}$ be its arbitrary element. For a positive integer k, the *k-sum problem* consists in deciding whether the input string $x_1, \ldots, x_n \in \mathbb{G}$ contains a subset of k elements that sums up to t. We assume that k is an arbitrary but fixed constant. The main result of the paper is the following

*Faculty of Computing, University of Latvia
†Google, Inc.

Theorem 1. *For a fixed k, the quantum query complexity of the k-sum problem is $\Omega(n^{k/(k+1)})$ provided that $|\mathbb{G}| \geq n^k$.*

Clearly, the 1-certificate complexity of the k-sum problem is k, hence, it is also subject to the certificate complexity barrier.

The result of Theorem 1 is tight thanks to the quantum algorithm based on quantum walks on the Johnson graph [5]. This algorithm was first designed to solve the k-distinctness problem. This problem asks for detecting whether the input string $x \in [q]^n$ contains k elements that are all equal. Soon enough it was realized that the same algorithm works for any function with 1-certificate complexity k [11], in particular, for the k-sum problem. The question of the tightness of this algorithm remained open for a long time. It has been known to be tight for $k = 2$ due to the lower bound for the element distinctness problem. Now we know that it is not optimal for the k-distinctness problem if $k > 2$ [8]. However, Theorem 1 shows that, for every k, quantum walk on the Johnson graph is optimal for some functions with 1-certificate complexity k. Finally, we note that the k-sum problem is also interesting because of its applications in quantum Merkle puzzles [10, 13].

In fact, we get Theorem 1 as a special case of a more general result we are about to describe. The following is a special case of a well-studied combinatorial object:

Definition 2. Assume T is a subset of $[q]^k$ of size q^{k-1}. We say that T is an *orthogonal array of length k* iff, for every index $i \in [k]$ and for every vector $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_k \in [q]$, there exists exactly one $x_i \in [q]$ such that $(x_1, \ldots, x_k) \in T$.

For $x = (x_i) \in [q]^n$ and $S \subseteq [n]$ let x_S denote the projection of x on S, i.e., the vector $(x_{s_1}, \ldots, x_{s_\ell})$ where $s_1, \ldots, s_\ell$ are the elements of S in the increasing order.

Assume each k-subset S of $[n]$ is equipped with an orthogonal array T_S. The *k-orthogonal array* problem consists in finding an element of any of the orthogonal arrays in the input string. More precisely, the input $x \in [q]^n$ evaluates to 1 iff there exists $S \subseteq [n]$ of size k such that $x_S \in T_S$. Consider the following two examples:

Example 3. Let $\mathbb{G}$ be a commutative group with q elements and $t \in \mathbb{G}$. $T = \{x \in \mathbb{G}^k : \sum_{i=1}^k x_i = t\}$ is an orthogonal array of length k. This choice corresponds to the k-sum problem of Theorem 1.

Example 4. $T = \{x \in [q]^2 : x_1 = x_2\}$ is an orthogonal array of length 2. This corresponds to the element distinctness problem from [7].

Theorem 5. *For a fixed k and any choice of the orthogonal arrays T_S, the quantum query complexity of the k-orthogonal array problem is $\Omega(n^{k/(k+1)})$ provided that $q \geq n^k$. The constant behind big-Omega depends on k, but not on n, q, or the choice of T_S.*

The orthogonal array condition specifies that even if an algorithm has queried $k-1$ elements out of any k-tuple, it has the same information whether this k-tuple is a 1-certificate as if it had queried no elements at all. Because of this, the search for a k-tuple as a whole entity is the best the quantum algorithm can do. Our proof of Theorem 5 is a formalization of this intuition.

2 Techniques used

If the size of the alphabet is big enough then almost all inputs are negative, i.e., they do not contain elements of T_S. We use this observation to approximate the set of negative inputs by the set of all possible input strings $[q]^n$. This may seem faulty, because in this case any positive

input figures as a negative one, hence, it is easy to come up with an adversary matrix having arbitrary large objective value. We do not fall victim to this, because the building blocks we construct our adversary matrix from are robust to this approximation.

We split the set of the positive inputs into $\binom{n}{k}$ subsets labelled by k-subsets of $[n]$. The inputs in the block labelled by S consists of all inputs x satisfying $x_S \in T_S$. Some inputs appear more than once, but it is fine.

By extending the set of negative inputs, we bring independence to their values. We can describe the matrices as the elements of the Hamming association scheme, i.e, as various tensor products featuring two $q \times q$ matrices E_0 and E_1 defined by

$$E_0[\![x,y]\!] = \frac{1}{q}, \qquad E_1[\![x,y]\!] = \begin{cases} 1-1/q, & x=y; \\ -1/q, & x \neq y. \end{cases}$$

These are orthogonal projectors, and Δ_j from the adversary bound transforms E_1 in the jth position into $-E_0$. By using these properties, we were able to construct the adversary matrix.

3 Summary

The polynomial-based proof of the lower bound for the element distinctness problem is rather circuitous: The element distinctness problem is reduced to the collision problem, and the lower bound is proven for the latter. We give a direct proof for a more general problem using the adversary method. To our knowledge, this is the first application of the negative-weighted adversary that does not rely on the composition results. We hope these techniques will be useful in proving lower bounds for other functions, such as set equality, k-distinctness, and others.

Acknowledgements

A.B. would like to thank Andris Ambainis, Troy Lee and Ansis Rosmanis for valuable discussions. We are grateful to Kassem Kalach for informing about the applications of the k-sum problem in Merkle puzzles, and for reporting on some minor errors in the early version of the paper.

A.B. is supported by the European Social Fund within the project "Support for Doctoral Studies at University of Latvia" and by FET-Open project QCS.

References

[1] S. Aaronson and Y. Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM*, 51(4):595–605, 2004.

[2] A. Ambainis. Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64(4):750–767, 2002, `arXiv:quant-ph/0002066`.

[3] A. Ambainis. Polynomial degree vs. quantum query complexity. In *Proc. of 44th IEEE FOCS*, pages 230–239, 2003, `arXiv:quant-ph/0305028`.

[4] A. Ambainis. Polynomial degree and lower bounds in quantum complexity: Collision and element distinctness with small range. *Theory of Computing*, 1:37–46, 2005, `arXiv:quant-ph/0305179`.

[5] A. Ambainis. Quantum walk algorithm for element distinctness. *SIAM Journal on Computing*, 37:210–239, 2007, arXiv:quant-ph/0311001.

[6] 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, arXiv:quant-ph/9802049.

[7] A. Belovs. Adversary lower bound for element distinctness. 2012, arXiv:1204.5074.

[8] A. Belovs. Learning-graph-based quantum algorithm for k-distinctness. 2012, arXiv:1205.1534.

[9] A. Belovs and R. Špalek. Adversary lower bound for the k-sum problem. 2012, arXiv:1206.6528.

[10] G. Brassard, P. Høyer, K. Kalach, M. Kaplan, S. Laplante, and L. Salvail. Merkle puzzles in a quantum world. In *CRYPTO 2011*, pages 391–410. Springer, 2011, arXiv:1108.2316.

[11] A. Childs and J. Eisenberg. Quantum algorithms for subset finding. *Quantum Information & Computation*, 5(7):593–604, 2005, arXiv:quant-ph/0311038.

[12] P. Høyer, T. Lee, and R. Špalek. Negative weights make adversaries stronger. In *Proc. of 39th ACM STOC*, pages 526–535, 2007, arXiv:quant-ph/0611054.

[13] K. Kalach. Personal communication, 2012.

[14] S. Kutin. Quantum lower bound for the collision problem with small range. *Theory of Computing*, 1(1):29–36, 2005.

[15] T. Lee, R. Mittal, B. Reichardt, R. Špalek, and M. Szegedy. Quantum query complexity of the state conversion problem. In *Proc. of 52nd IEEE FOCS*, pages 344–353, 2011, arXiv:1011.3020.

[16] B. Reichardt. Reflections for quantum query algorithms. In *Proc. of 22nd ACM-SIAM SODA*, pages 560–569, 2011, arXiv:1005.1601.

[17] R. Špalek and M. Szegedy. All quantum adversary methods are equivalent. *Theory of Computing*, 2:1–18, 2006, arXiv:quant-ph/0409116.

[18] S. Zhang. On the power of ambainis lower bounds. *Theoretical Computer Science*, 339(2):241–256, 2005, arXiv:quant-ph/0311060.