

Adversary Lower Bound for the k -sum Problem

Aleksandrs Belovs
University of Latvia
19 Raina Boulevard
Riga, Latvia
stiboh@gmail.com

Robert Špalek
Google Inc.
1600 Amphitheatre Parkway
Mountain View, CA, USA
rspalek@gmail.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.

Categories and Subject Descriptors

F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*Computations on discrete structures*; H.2.4 [Database Management]: Systems—*Query processing*

General Terms

Algorithms

Keywords

Quantum query complexity, Knapsack packing problem, Orthogonal arrays

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 $\{1, \dots, q\}$. 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 $\Omega(n^{2/3})$ 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] implies that for any function with 1-certificate complexity bounded by a constant, the adversary method fails to achieve anything better than $\Omega(\sqrt{n})$.

In 2006 a stronger version of the adversary bound was developed by Høyer *et al.* [12]. This is the negative-weight adversary lower bound defined in Section 2. Later it was proved to be optimal by Reichardt *et al.* [16, 15]. Although the negative-weight adversary lower bound is known to be tight, it has almost never been used to prove lower bounds for explicit functions. The vast majority of lower bounds by the adversary method used the old nonnegative-weight version of this method. But since the 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 $x_1, \dots, x_n \in \mathbb{G}$ 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

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 that, for a variable k , states the nonnegative-weight adversary cannot prove a better lower bound than $\Omega(\sqrt{kn})$.

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 was 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

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ITCS'13, January 9–12, 2012, Berkeley, California, USA.

Copyright 2013 ACM 978-1-4503-1859-4/13/01 ...\$15.00.

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 [9, 13].

Actually, 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 1. 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 $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k \in [q]$, there exists exactly one $x_i \in [q]$ such that $(x_1, \dots, x_k) \in T$.

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 $(x_{s_1}, \dots, x_{s_\ell})$ where s_1, \dots, s_ℓ are the elements of S in the increasing order.

Assume each subset S of $[n]$ of size k 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 1. 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 2. $T = \{x \in [q]^2 : x_1 = x_2\}$ is an orthogonal array of length 2. This corresponds to the element distinctness problem from [7].

THEOREM 2. *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 has queried no elements out of it. 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 2 is a formalization of this intuition.

Let us elaborate on the requirement on the size of the alphabet. It is easy to see that some requirement is necessary. Indeed, the k -sum problem can be solved in $O(\sqrt{n})$ queries if the size of \mathbb{G} is $O(1)$, using the Grover search to find up to k copies of every element of \mathbb{G} in the input string, and trying to construct t out of what is found. In some cases, e.g., when t is the identity element and k equals the order of the group, the problem becomes trivial if n is large enough.

The requirement on the size of the alphabet for the element distinctness problem is a subtle issue. The lower bounds by Aaronson and Shi [1] and Kutin [14] require the size of the alphabet to be at least $\Omega(n^2)$ that is the same that gives Theorem 2. However, later Ambainis [4] showed that the lower bound remains the same even if one allows the alphabet of size n . Reducing the alphabet size in Theorem 2 is one of our open problems.

2. ADVERSARY LOWER BOUND

In the paper we are interested in the quantum query complexity of solving the orthogonal array problem. For the definitions and main properties of quantum query complexity refer to, e.g., Ref. [10]. For the purposes of our paper, it suffices to work with the definition of the adversary bound we give in this section.

Compared to the original formulation of the negative-weighted adversary bound [12], our formulation has two differences. Firstly, in order to simplify notations we call an adversary matrix a matrix with rows labelled by positive inputs, and the columns—by the negative ones. It is a quarter of the original adversary matrix that completely specifies the latter. Secondly, due to technical reasons, we allow several rows to be labelled by the same positive input. All this is captured by the following definition and theorem.

Definition 2. Let f be a function $f: \mathcal{D} \rightarrow \{0, 1\}$ with domain $\mathcal{D} \subseteq [q]^n$. Let $\tilde{\mathcal{D}}$ be a set of pairs (x, a) with the property that the first element of each pair belongs to \mathcal{D} , and $\tilde{\mathcal{D}}_i = \{(x, a) \in \tilde{\mathcal{D}} : f(x) = i\}$ for $i \in \{0, 1\}$. An *adversary matrix* for the function f is a non-zero real $\tilde{\mathcal{D}}_1 \times \tilde{\mathcal{D}}_0$ matrix Γ . And, for $i \in [n]$, let Δ_i denote the $\tilde{\mathcal{D}}_1 \times \tilde{\mathcal{D}}_0$ matrix defined by

$$\Delta_i[(x, a), (y, b)] = \begin{cases} 0, & x_i = y_i; \\ 1, & \text{otherwise.} \end{cases}$$

THEOREM 3 (ADVERSARY BOUND). *In the notations of Definition 2, $Q_2(f) = \Omega(\text{Adv}^\pm(f))$, where*

$$\text{Adv}^\pm(f) = \sup_{\Gamma} \frac{\|\Gamma\|}{\max_{i \in n} \|\Gamma \circ \Delta_i\|} \quad (1)$$

with the maximization over all adversary matrices for f , $\|\cdot\|$ is the spectral norm, and $Q_2(f)$ is the quantum query complexity of f .

PROOF. In the original negative-weight adversary bound paper [12], Eq. (1) is proven when Γ is a real symmetric $\mathcal{D} \times \mathcal{D}$ matrix with the property $\Gamma[x, y] = 0$ if $f(x) \neq f(y)$, and Δ_i are modified accordingly. We describe a reduction from the adversary matrix in our definition, Γ , to the adversary matrix Γ' in the definition of [12]. Also, let Δ'_i be the $\mathcal{D} \times \mathcal{D}$ matrix with $\Delta'_i[x, y] = 1$ if $x_i \neq y_i$, and 0, otherwise.

At first, define $\bar{\Gamma}$ as

$$\bar{\Gamma} = \begin{pmatrix} 0 & \Gamma^* \\ \Gamma & 0 \end{pmatrix}.$$

Note that $\|\bar{\Gamma}\| = \|\Gamma\|$, and the spectrum of $\bar{\Gamma}$ is symmetric. Also, for all i , $\|\bar{\Gamma} \circ \bar{\Delta}_i\| = \|\Gamma \circ \Delta_i\|$, where $\bar{\Delta}_i$ is defined similarly to $\bar{\Gamma}$.

Let $\delta = (\delta_{x,a})$ be the normalized eigenvalue $\|\Gamma\|$ eigenvector of $\bar{\Gamma}$. For all $x, y \in \mathcal{D}$, let:

$$\delta'_x = \sqrt{\sum_{a:(x,a) \in \tilde{\mathcal{D}}} \delta_{x,a}^2}$$

and

$$\Gamma'[(x, a), (y, b)] = \frac{1}{\delta'_x \delta'_y} \sum_{\substack{a:(x,a) \in \tilde{\mathcal{D}} \\ b:(y,b) \in \tilde{\mathcal{D}}}} \delta_{x,a} \delta_{y,b} \bar{\Gamma}[(x, a), (y, b)].$$

Then it is easy to see that $\delta' = (\delta'_x)$ satisfies $\|\delta'\| = 1$ and $(\delta')^* \Gamma' \delta' = \delta^* \bar{\Gamma} \delta = \|\Gamma\|$, hence, $\|\Gamma'\| \geq \|\Gamma\|$.

And vice versa, if ε' is such that $\|\varepsilon'\| = 1$ and $(\varepsilon')^*(\Gamma' \circ \Delta'_i)\varepsilon' = \|\Gamma' \circ \Delta'_i\|$, let $\varepsilon_{x,a} = \delta_{x,a}\varepsilon'_x/\delta'_x$. Again, $\|\varepsilon\| = 1$ and $\varepsilon^*(\Gamma \circ \Delta_i)\varepsilon = (\varepsilon')^*(\Gamma' \circ \Delta'_i)\varepsilon'$, hence, $\|\Gamma' \circ \Delta'_i\| \leq \|\Gamma \circ \Delta_i\| = \|\Gamma \circ \Delta_i\|$.

This means that Γ' provides at least as good adversary lower bound as Γ . \square

3. PROOF

In this section we prove Theorem 2 using the adversary lower bound, Theorem 3. The idea of our construction is to embed the adversary matrix Γ into a slightly larger matrix $\tilde{\Gamma}$ with additional columns. Then $\Gamma \circ \Delta_i$ is a submatrix of $\tilde{\Gamma} \circ \Delta_i$, hence, $\|\Gamma \circ \Delta_i\| \leq \|\tilde{\Gamma} \circ \Delta_i\|$. (In this section we use Δ_i to denote all matrices defined like in Definition 2, with the size and the labels of the rows and columns clear from the context.) It remains to prove that $\|\tilde{\Gamma}\|$ is large, and that $\|\Gamma\|$ is not much smaller than $\|\tilde{\Gamma}\|$.

The proof is organized as follows. In Section 3.1 we define $\tilde{\Gamma}$ in dependence on parameters α_m , in Section 3.2 we analyze its norm, in Sections 3.3 and 3.4 calculate $\|\tilde{\Gamma} \circ \Delta_i\|$, in Section 3.5 optimize α_m s, and, finally, in Section 3.6 prove that the norm of the true adversary matrix Γ is not much smaller than the norm of $\tilde{\Gamma}$.

3.1 Adversary matrix

Matrix $\tilde{\Gamma}$ consists of $\binom{n}{k}$ matrices $\tilde{G}_{s_1, \dots, s_k}$ stacked one on another for all possible choices of $S = \{s_1, \dots, s_k\} \subset [n]$:

$$\tilde{\Gamma} = \begin{pmatrix} \tilde{G}_{1,2,\dots,k} \\ \tilde{G}_{1,2,\dots,k-1,k+1} \\ \dots \\ \tilde{G}_{n-k+1,n-k+2,\dots,n} \end{pmatrix}. \quad (2)$$

Each \tilde{G}_S is a $q^{n-1} \times q^n$ matrix with rows indexed by inputs $(x_1, \dots, x_n) \in [q]^n$ such that $x_S \in T_S$, and columns indexed by all possible inputs $(y_1, \dots, y_n) \in [q]^n$.

We say a column with index y is *illegal* if $y_S \in T_S$ for some $S \subseteq [n]$. After removing all illegal columns, \tilde{G}_S will represent the part of $\tilde{\Gamma}$ with the rows indexed by the inputs having an element of the orthogonal array on S . Note that some positive inputs appear more than once in $\tilde{\Gamma}$. More specifically, an input x appears as many times as many elements of the orthogonal arrays it contains.

This construction may seem faulty, because there are elements of $[q]^n$ that are used as labels of both rows and columns in $\tilde{\Gamma}$, and hence, it is trivial to construct a matrix $\tilde{\Gamma}$ such that the value in (1) is arbitrarily large. But we design $\tilde{\Gamma}$ in a specifically restrictive way so that it still is a good adversary matrix after the illegal columns are removed.

Let J_q be the $q \times q$ all-ones matrix. Assume e_0, \dots, e_{q-1} is an orthonormal eigenbasis of J_q with $e_0 = 1/\sqrt{q}(1, \dots, 1)$ being the eigenvalue q eigenvector. Consider the vectors of the following form:

$$v = e_{v_1} \otimes e_{v_2} \otimes \dots \otimes e_{v_n}, \quad (3)$$

where $v_i \in \{0, \dots, q-1\}$. These are eigenvectors of the Hamming Association Scheme on $[q]^n$. For a vector v from (3), the *weight* $|v|$ is defined as the number of non-zero entries in (v_1, \dots, v_n) . Let $E_k^{(n)}$, for $k = 0, \dots, n$, be the orthogonal projector on the space spanned by the vectors from (3) having weight k . These are the projectors on the eigenspaces of the association scheme. Let us denote $E_i = E_i^{(1)}$ for $i = 0, 1$.

These are $q \times q$ matrices. All entries of E_0 are equal to $1/q$, and the entries of E_1 are given by

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

Elements of S in \tilde{G}_S should be treated differently from the rest of the elements. For them, we define a $q^{k-1} \times q^k$ matrix F_S . It has rows labelled by the elements of T_S and columns by the elements of $[q]^k$, and is defined as follows.

Definition 3. Let

$$E_{<k}^{(k)} = I - E_k^{(k)} = \sum_{i=0}^{k-1} E_i^{(k)} = \sum_{\substack{u=e_{u_1} \otimes \dots \otimes e_{u_k} \\ |u|<k}} uu^*$$

be the projector onto the subspace spanned by the vectors of less than maximal weight. Let F_S be \sqrt{q} times the submatrix of $E_{<k}^{(k)}$ consisting of only the rows from T_S .

Finally, we define $\tilde{\Gamma}$ as in (2) with \tilde{G}_S defined by

$$\tilde{G}_S = \sum_{m=0}^{n-k} \alpha_m F_S \otimes E_m^{(n-k)}, \quad (4)$$

where F_S acts on the elements in S and E_m acts on the remaining $n - k$ elements. Coefficients α_m will be specified later.

3.2 Norm of $\tilde{\Gamma}$

LEMMA 1. Let $\tilde{\Gamma}$ be like in (2) with \tilde{G}_S defined as in (4). Then

- (a) $\|\tilde{\Gamma}\| = \Omega(\alpha_0 n^{k/2})$,
- (b) $\|\tilde{\Gamma}\| = O(\max_m \alpha_m n^{k/2})$.

PROOF. Fix a subset S and denote $T = T_S$ and $F = F_S$. Recall that $E_{<k}^{(k)}$ is the sum of uu^* over all $u = e_{u_1} \otimes \dots \otimes e_{u_k}$ with at least one u_j equal to 0, and F is the restriction of $E_{<k}^{(k)}$ to the rows in T .

For $u = e_{u_1} \otimes \dots \otimes e_{u_k}$ and ℓ such that $u_\ell = 0$, let $u^{(\ell)}$ denote the \sqrt{q} multiple of u restricted to the elements in T . The reason for the superscript is that we consider the following process of obtaining $u^{(\ell)}$: we treat T as $[q]^{k-1}$ by erasing the ℓ -th element in any string of T , then $u^{(\ell)}$ coincides on this set with u with the ℓ -th term removed.

In this notation, the contribution from uu^* to F equals $u^{(\ell_u)}u^*$, where ℓ_u is any position in u containing e_0 . In general, we do not know how $u^{(\ell)}$ s relate for different ℓ . However, we know that, for a fixed ℓ , they are all orthogonal; and for any ℓ , $(e_0^{\otimes k})^{(\ell)}$ is the vector $1/\sqrt{q^{k-1}}(1, \dots, 1)$.

Let us start with proving (a). We estimate $\|\tilde{\Gamma}\|$ from below as $w^* \tilde{\Gamma} w'$, where w and w' are unit vectors with all elements equal. In other words, $\|\tilde{\Gamma}\|$ is at least the sum of all its entries divided by $\sqrt{\binom{n}{k} q^{2n-1}}$. In order to estimate the sum of the entries of $\tilde{\Gamma}$, we rewrite (4) as

$$\tilde{G}_S = \alpha_0 e_0^{\otimes(n-1)} (e_0^{\otimes n})^* + \sum_{u,v} \alpha_{|v|} (u^{(\ell_u)} \otimes v)(u \otimes v)^*, \quad (5)$$

where the summation is over all u and v such that at least one of them contains an element different from e_0 . The sum

of all entries in the first term of (5) is $\alpha_0 q^{n-1/2}$. The sum of each column in each of $(u^{(\ell_u)} \otimes v)(u \otimes v)^*$ is zero because at least one of $u^{(\ell_u)}$ or v sums up to zero. By summing over all $\binom{n}{k}$ choices of S , we get that $\|\tilde{\Gamma}\| \geq \alpha_0 \sqrt{\binom{n}{k}} = \Omega(\alpha_0 n^{k/2})$.

In order to prove (b), express F_S as $\sum_{\ell=1}^k F_S^{(\ell)}$ with $F_S^{(\ell)} = \sum_{u \in U_\ell} u^{(\ell)} u^*$. Here $\{U_\ell\}$ is an arbitrary decomposition of all u such that U_ℓ contains only u with e_0 in the ℓ -th position. Define $\tilde{G}_S^{(\ell)}$ as in (4) with F_S replaced by $F_S^{(\ell)}$, and $\tilde{\Gamma}^{(\ell)}$ as in (2) with \tilde{G}_S replaced by $\tilde{G}_S^{(\ell)}$.

Since all $u^{(\ell)}$ s are orthogonal for a fixed ℓ , we get that

$$(\tilde{G}^{(\ell)})^* \tilde{G}^{(\ell)} = \sum_{u \in U_\ell, v} \alpha_{|v|}^2 (u \otimes v)(u \otimes v)^*,$$

thus $\|(\tilde{G}^{(\ell)})^* \tilde{G}^{(\ell)}\| = \max_m \alpha_m^2$. By the triangle inequality,

$$\|\tilde{\Gamma}^{(\ell)}\|^2 = \left\| \left(\tilde{\Gamma}^{(\ell)} \right)^* \tilde{\Gamma}^{(\ell)} \right\| = \left\| \sum_S (\tilde{G}_S^{(\ell)})^* \tilde{G}_S^{(\ell)} \right\| \leq \binom{n}{k} \max_m \alpha_m^2.$$

Since $\tilde{\Gamma} = \sum_{\ell=1}^k \tilde{\Gamma}^{(\ell)}$, another application of the triangle inequality finishes the proof of (b). \square

3.3 Action of Δ_1

The adversary matrix is symmetric in all input variables and hence it suffices to only consider the entry-wise multiplication by Δ_1 . Precise calculation of $\|\tilde{\Gamma} \circ \Delta_1\|$ is very tedious, but one can get an asymptotically tight bound using the following trick. Instead of computing $\tilde{\Gamma} \circ \Delta_1$ directly, we arbitrarily map $\tilde{\Gamma} \xrightarrow{\Delta_1} \tilde{\Gamma}_1$ such that $\tilde{\Gamma}_1 \circ \Delta_1 = \tilde{\Gamma} \circ \Delta_1$, and use the inequality $\|\tilde{\Gamma}_1 \circ \Delta_1\| \leq 2\|\tilde{\Gamma}_1\|$ that holds thanks to $\gamma_2(\Delta_1) \leq 2$ [15]. In other words, we change arbitrarily the entries with $x_1 = y_1$. We use the mapping

$$E_0 \xrightarrow{\Delta_1} E_0, \quad E_1 \xrightarrow{\Delta_1} -E_0, \quad I \xrightarrow{\Delta_1} 0. \quad (6)$$

The projector $E_{<k}^{(k)}$ is mapped by Δ_1 as

$$E_{<k}^{(k)} = I - E_k^{(k)} \xrightarrow{\Delta_1} E_0 \otimes E_1^{\otimes(k-1)}.$$

It follows that

$$F \xrightarrow{\Delta_1} e_0^* \otimes E_1^{\otimes(k-1)} = \sum_{\substack{u=e_{u_1} \otimes \dots \otimes e_{u_k} \\ u_0=0, |u|=k-1}} u^{(1)} u^*, \quad (7)$$

where $u^{(1)}$ is defined like in the proof of Lemma 1.

3.4 Norm of $\tilde{\Gamma}_1$

LEMMA 2. Let $\tilde{\Gamma}$ be like in (2) with \tilde{G}_T defined as in (4), and map $\tilde{\Gamma} \xrightarrow{\Delta_1} \tilde{\Gamma}_1$ and $\tilde{G}_T \xrightarrow{\Delta_1} (\tilde{G}_T)_1$ using (6) and (7). Then

$$\|\tilde{\Gamma}_1\| = O(\max_m (\max(\alpha_m m^{(k-1)/2}, (\alpha_m - \alpha_{m+1}) n^{k/2}))).$$

PROOF. We have $\|\tilde{\Gamma}_1\|^2 = \|\tilde{\Gamma}_1^* \tilde{\Gamma}_1\| = \|\sum_S (\tilde{G}_S)_1^* (\tilde{G}_S)_1\|$. Decompose the set of all possible k -tuples of indices into $\mathcal{S}_1 \cup \mathcal{S}_2$, where \mathcal{S}_1 are k -tuples containing 1 and \mathcal{S}_2 are k -tuples that don't contain 1. We upper-bound the contribution of \mathcal{S}_1 to $\|\tilde{\Gamma}_1\|^2$ by $\max_m \alpha_m^2 \binom{m+k-1}{k-1}$ and the contribution of \mathcal{S}_2 by $\max_m (\alpha_m - \alpha_{m+1})^2 k \binom{n-1}{k}$, and apply the triangle inequality.

Let $v = e_{v_1} \otimes \dots \otimes e_{v_n}$ with $|v| = m + k - 1$, and let $S \in \mathcal{S}_1$. Then, by (7),

$$(\tilde{G}_S)_1 v = \begin{cases} \alpha_m v^{(1)}, & v_1 = 0 \text{ and } |v_S| = k - 1, \\ 0, & \text{otherwise.} \end{cases}$$

Here $v^{(1)} = e_{v_2} \otimes \dots \otimes e_{v_n}$ is v with the first term removed and $v_S = \bigotimes_{s \in S} e_{v_s}$.

For different v , these are orthogonal vectors, and hence v is an eigenvector of $(\tilde{G}_S)_1^* (\tilde{G}_S)_1$ of eigenvalue α_m^2 if $v_1 = 0$ and $|v_S| = k - 1$, and of eigenvalue 0 otherwise. For every v with $v_1 = 0$ and $|v| = m + k - 1$, there are $\binom{m+k-1}{k-1}$ sets $S \in \mathcal{S}_1$ such that $(\tilde{G}_S)_1 v \neq 0$. Thus, the contribution of \mathcal{S}_1 is as claimed.

Now consider an $S \in \mathcal{S}_2$, that means $1 \notin S$.

$$\begin{aligned} \tilde{G}_S &= \sum_{m=0}^{n-k} \alpha_m F_S \otimes E_m^{(n-k)} \\ &= \sum_{m=0}^{n-k} \alpha_m F_S \\ &\quad \otimes (E_0 \otimes E_m^{(n-k-1)} + E_1 \otimes E_{m-1}^{(n-k-1)}) \\ &\xrightarrow{\Delta_1} \sum_{m=0}^{n-k} \alpha_m F_S \otimes E_0 \otimes (E_m^{(n-k-1)} - E_{m-1}^{(n-k-1)}) \\ (\tilde{G}_S)_1 &= \sum_{m=0}^{n-k} (\alpha_m - \alpha_{m+1}) F_S \otimes E_0 \otimes E_m^{(n-k-1)}. \end{aligned}$$

Therefore $(\tilde{G}_S)_1$ is of the same form as \tilde{G}_S , but with coefficients $(\alpha_m - \alpha_{m+1})$ instead of α_m and on one dimension less. We get the required estimate from Lemma 1(b).

Since $k = O(1)$, we get the claimed bound. \square

3.5 Optimization of α_m

To maximize the adversary bound, we maximize $\|\tilde{\Gamma}\|$ while keeping $\|\tilde{\Gamma}_1\| = O(1)$. That means, we choose the coefficients $\{\alpha_m\}$ to maximize $\alpha_0 n^{k/2}$ (Lemma 1) so that, for every m , $\alpha_m \leq m^{(1-k)/2}$ and $\alpha_m \leq \alpha_{m+1} + n^{-k/2}$ (Lemma 2).

For every $r \in [n]$, $\alpha_0 \leq \alpha_r + r n^{-k/2} \leq r^{(1-k)/2} + r n^{-k/2}$. The expression on the right-hand side achieves its minimum, up to a constant, $\alpha_0 = 2 n^{k(1-k)/(2(k+1))}$ for $r = n^{k/(k+1)}$. This corresponds to the following solution:

$$\alpha_m = \max \left\{ 2 - \frac{m}{n^{k/(k+1)}}, 0 \right\} n^{\frac{k(1-k)}{2(k+1)}} \quad (8)$$

With this choice of α_m , $\|\tilde{\Gamma}\| = \Omega(\alpha_0 n^{k/2}) = \Omega(n^{k/(k+1)})$.

3.6 Constructing Γ from $\tilde{\Gamma}$

The matrix $\tilde{\Gamma}$ gives us the desired ratio of norms of $\tilde{\Gamma}$ and $\tilde{\Gamma} \circ \Delta_i$. Unfortunately, $\tilde{\Gamma}$ cannot directly be used as an adversary matrix, because it contains illegal columns y with $f(y) = 1$, that is, y that contain an element of the orthogonal array on $S \subset [n] : |S| = k$, i.e., $y_S \in T_S$. We show that after removing the illegal columns it is still good enough.

LEMMA 3. Let Γ be the sub-matrix of $\tilde{\Gamma}$ with the illegal columns removed. Then $\|\Gamma \circ \Delta_1\| \leq \|\tilde{\Gamma} \circ \Delta_1\|$, and $\|\Gamma\|$ is still $\Omega(\alpha_0 n^{k/2})$ when $q \geq n^k$.

PROOF. We estimate $\|\Gamma\|$ from below by $w^* \Gamma w'$ using unit vectors w, w' with all elements equal. Recall Equation (5):

$$\tilde{G}_T = \alpha_0 e_0^{\otimes(n-1)} (e_0^{\otimes n})^* + \sum_{u,v} \alpha_{|v|} (u^{(\ell_u)} \otimes v) (u \otimes v)^*,$$

where the summation is over all u and v such that at least one of them contains an element different from e_0 . The sum of each column in each of $(u^{(\ell_u)} \otimes v) (u \otimes v)^*$ still is zero because at least one of $u^{(\ell_u)}$ or v sums up to zero. Therefore the contribution of the sum is zero regardless of which columns have been removed.

By summing over all $\binom{n}{k}$ choices of S , we get

$$\|\Gamma\| \geq w^* \Gamma w' = \sqrt{\binom{n}{k}} \alpha_0 (e_0^{\otimes n})^*_{L'} w',$$

where e_L denotes the sub-vector of e restricted to L , and L is the set of legal columns. Since both e_0 and w' are unit vectors with all elements equal, and w' is supported on L , $(e_0^{\otimes n})^*_{L'} w' = \sqrt{|L|/q^n}$.

Let us estimate the fraction of legal columns. The probability that a uniformly random input $y \in [q]^n$ contains an orthogonal array at any given k -tuple S is $\frac{1}{q}$. By the union bound, the probability that there exists such S is at most $\binom{n}{k} \frac{1}{q}$. Therefore the probability that a random column is legal is $\frac{|L|}{q^n} \geq 1 - \binom{n}{k} \frac{1}{q}$, which is $\Omega(1)$ when $q \geq n^k$. \square

Thus, with the choice of α_m from (8), we have $\text{Adv}^\pm(f) = \Omega(\alpha_0 n^{k/2}) = \Omega(n^{k/(k+1)})$. This finishes the proof of Theorem 2.

4. OPEN PROBLEMS

Our technique relies crucially on the n^k lower bound on the alphabet size. Can one relax this bound in some special cases? For example, element distinctness is nontrivial when $q \geq n$, but our lower bound only holds for $q \geq n^2$.

A tight $\Omega(n^{2/3})$ lower bound for element distinctness was originally proved by the polynomial method [1] by reduction via the collision problem. The k -collision problem is to decide whether a given function is $1 : 1$ or $k : 1$, provided that it is of one of the two types. One can use an algorithm for element distinctness to solve the 2-collision problem, and thus the tight $\Omega(n^{1/3})$ lower bound for collision implies in [1] implies a tight lower bound for element distinctness. Unfortunately, the reduction doesn't go in both directions and hence our result doesn't imply any nontrivial adversary bound for k -collision. The simpler nonnegative-weight adversary bound is limited to $O(1)$ due to the property testing barrier. Roughly speaking, if every 0-input differs from every 1-input in at least an ε -fraction of the input, the nonnegative-weight adversary bound is limited by $O(\frac{1}{\varepsilon})$. How does an explicit negative-weight adversary matrix for an $\omega(1)$ lower bound look like?

The recent learning graph-based algorithm for k -distinctness [8] solves the problem in $O(n^{1-2^{k-2}/(2^k-1)})$ quantum queries, which is less than $O(n^{k/(k+1)})$ but more than the $\Omega(n^{2/3})$ lower bound by reduction from 2-distinctness. k -distinctness is easier than the k -sum problem considered in our paper because one can obtain nontrivial information about the solution from partial solutions, i.e., from ℓ -tuples of equal numbers for $\ell < k$. Can one use our technique to prove an $\omega(n^{2/3})$ lower bound for k -distinctness?

The k -sum problem is very structured in the sense that all k -tuples of the input variables, and all possible values seen on a $(k-1)$ -tuple, are equal with respect to the function. The symmetry of this problem helped us to design a symmetric adversary matrix. The nonnegative-weight adversary bound gives nontrivial lower bounds for most problems, by simply putting most of the weight on hard-to-distinguish input pairs, regardless of whether the problem is symmetric or not. Can one use our technique to improve the best known lower bounds for some non-symmetric problems, for example, to prove an $\omega(\sqrt{n})$ lower bound for graph collision, $\omega(n)$ for triangle finding, or $\omega(n^{3/2})$ for verification of matrix products?

Acknowledgments

A.B. would like to thank Andris Ambainis, Troy Lee and Ansis Rosmanis for valuable discussions. We are grateful to Kassem Kalach for informing us 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, as well as to the anonymous reviewers for their suggestions on improving the presentation 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.

5. 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.
- [3] A. Ambainis. Polynomial degree vs. quantum query complexity. In *Proceedings of the 44th IEEE FOCS*, pages 230–239, 2003.
- [4] A. Ambainis. Quantum lower bounds for collision and element distinctness with small range. *Theory of Computing*, 1:37–46, 2005.
- [5] A. Ambainis. Quantum walk algorithm for element distinctness. *SIAM Journal on Computing*, 37:210–239, 2007.
- [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.
- [7] A. Belovs. Adversary lower bound for element distinctness. Technical Report 1204.5074, arXiv, 2012.
- [8] A. Belovs. Learning-graph-based quantum algorithm for k -distinctness. In *In Proc. of 53rd FOCS*, pages 207–216, 2012.
- [9] G. Brassard, P. Høyer, K. Kalach, M. Kaplan, S. Laplante, and L. Salvail. Merkle puzzles in a quantum world. In *CRYPTO 2011*, volume 6841, page 391. Springer, 2011.
- [10] H. Buhrman and R. de Wolf. Complexity measures and decision tree complexity: a survey. *Theoretical Computer Science*, 288:21–43, 2002.
- [11] A. Childs and J. Eisenberg. Quantum algorithms for subset finding. *Quantum Information & Computation*, 5(7):593–604, 2005.
- [12] P. Høyer, T. Lee, and R. Špalek. Negative weights

- make adversaries stronger. In *Proceedings of the 39th ACM STOC*, pages 526–535, 2007.
- [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 *Proceedings of 52nd IEEE FOCS*, pages 344–353, 2011.
- [16] B. Reichardt. Reflections for quantum query algorithms. In *Proceedings of 22nd ACM-SIAM SODA*, pages 560–569, 2011.
- [17] R. Špalek and M. Szegedy. All quantum adversary methods are equivalent. *Theory of Computing*, 2:1–18, 2006.
- [18] S. Zhang. On the power of Ambainis lower bounds. *Theoretical Computer Science*, 339(2):241–256, 2005.