Quantum Lower Bounds for Tripartite Versions of the Hidden Shift and the Set Equality Problems

Aleksandrs Belovs[1]
Faculty of Computing, University of Latvia, Raina 19, Riga, Latvia

Ansis Rosmanis[2]
Centre for Quantum Technologies, National University of Singapore,
Block S15, 3 Science Drive 2, Singapore

Abstract

In this paper, we study quantum query complexity of the following rather natural tripartite generalisations (in the spirit of the 3-sum problem) of the hidden shift and the set equality problems, which we call the 3-shift-sum and the 3-matching-sum problems.

The 3-shift-sum problem is as follows: given a table of $3\times n$ elements, is it possible to circularly shift its rows so that the sum of the elements in each column becomes zero? It is promised that, if this is not the case, then no 3 elements in the table sum up to zero. The 3-matching-sum problem is defined similarly, but it is allowed to arbitrarily permute elements within each row. For these problems, we prove lower bounds of $\Omega(n^{1/3})$ and $\Omega(\sqrt{n})$, respectively. The second lower bound is tight.

The lower bounds are proven by a novel application of the dual learning graph framework and by using representation-theoretic tools from [7].

2012 ACM Subject Classification Theory of computation → Quantum query complexity

Keywords and phrases Adversary Bound, Dual Learning Graphs, Quantum Query Complexity, Representation Theory

Digital Object Identifier 10.4230/LIPIcs.TQC.2018.3

Related Version `https://arxiv.org/abs/1712.10194`

Acknowledgements Part of this research was done while A.B. was visiting the Centre for Quantum Technologies at the National University of Singapore. A.B. would like to thank Miklos Santha for hospitality.

1 Introduction

One of the starting points of this paper was the following problem, posed by Aaronson and Ambainis [1]: construct a partial Boolean function with polylogarithmic quantum query complexity but whose randomised query complexity is $\omega(\sqrt{n})$, where n is the number of input variables. There are relatively many functions known with the required quantum query complexity and randomised query complexity $\Theta(\sqrt{n})$. For instance, one can take the forrelation problem of [1] with quantum query complexity 1, or the better-known hidden subgroup problem [13]. However, no function with polylogarithmic quantum query complexity

[1] This research is partially supported by the ERDF project number 1.1.1.2/I/16/113. Part of this work was done while supported by the ERC Advanced Grant MQC.
[2] This research is partially funded by the Singapore Ministry of Education and the National Research Foundation under grant R-710-000-012-135.

13th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2018).
Editor: Stacey Jeffery; Article No. 3; pp. 3:1–3:15
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

and randomised query complexity $\omega(\sqrt{n})$ is known. As shown in [10, 2], such a function would also yield a larger than 5/2-power separation between quantum and randomised query complexities for *total* Boolean functions.

Aaronson and Ambainis proposed a candidate function, which they call the k-fold forrelation. It has a very simple quantum $O(1)$-query algorithm, but it seems hard to lower bound its randomised query complexity. However, it is also possible to go in the opposite direction: find a function whose randomised query complexity is $\omega(\sqrt{n})$, and construct an efficient quantum algorithm computing this function. A potential candidate might be a modification of a function already known to be easy quantumly, preserving the hope the modification is still easy.

One particularly neat starting function, in our opinion, is the following *hidden shift problem*. Given two strings $x, y \in [q]^n$, the task is to distinguish two cases: in the positive case, x is a circular shift of y; in the negative case, all the input variables in x and y are distinct. This problem is equivalent to the hidden subgroup problem in the dihedral group [17], and its quantum query complexity is logarithmic.[3] It is also easy to see that its randomised query complexity is $\Theta(\sqrt{n})$.

In this paper we consider the following modification, which we call the *3-shift-sum* problem. We are given an input string $x \in [q]^{3n}$, which we treat as a $3 \times n$ table. In the positive case, it is possible to circularly shifts the rows of the table so that the sum of the elements in each column becomes divisible by q. In the negative case, no matter how we shift the rows, there is no column with the sum of its elements divisible by q. (In other words, there is no three elements, one from each row, whose sum is divisible by q.) This is a natural amalgamation of the hidden shift and the 3-sum problem, both studied quantumly.

It is easy to see that the randomised query complexity of this problem is $\Theta(n^{2/3})$. This raises the question of what its quantum query complexity is. Our first result is a simple proof that, unlike the hidden shift problem, the quantum query complexity of the 3-shift-sum problem is polynomial: $\Omega(n^{1/3})$. Thus, the 3-shift-sum problem fails to provide the desired separation.

Similarly as the 3-shift-sum problem is a tripartite version of the hidden shift problem, the *3-matching-sum* problem is a tripartite version of the *set equality* problem. In the set equality problem, the negative inputs are as in the hidden shift problem, but in a positive input, y is an arbitrary permutation of x, not necessary a circular shift. Unlike the hidden shift problem, the set equality problem has polynomial quantum query complexity: $\Theta(n^{1/3})$ [23, 25, 7]. In our tripartite version of it, the negative inputs are the same as in the 3-shift-sum problem, but for a positive input, there exists an arbitrary permutation of the elements within each row such that the sum of each column becomes divisible by q. Our second result is a complete characterisation of the quantum query complexity of this problem: it is $\Theta(\sqrt{n})$.

1.1 Techniques

Our main tool is the framework of dual learning graphs, which is "compiled" to the adversary lower bound.

---

[3] The canonical version of the Dihedral HSP also assumes that all the symbols in x are pairwise distinct and all the symbols in y are pairwise distinct. However, this condition is not relevant for the query complexity being logarithmic as easily follows from [14, Theorem 2]. Since it is not immediately obvious how to generalise this condition for the tripartite version, we omitted it from our definition of the hidden shift problem. Also note that this is the decision version of the problem, where it is not required to find the shift. The latter may be difficult if x contains repeated symbols.

The first version of the adversary method was developed by Ambainis [3]. This version, later known as the positive-weighted adversary, is easy to use, and it has found many applications, but it is also subject to some limitations: the certificate complexity [24, 26] and the property testing [15] barriers. The property testing barrier, which is relevant to our problems, states that, if any positive input differs from any negative input on at least ε fraction of the input variables, the positive-weighted adversary fails to prove a lower bound better than $\Omega(1/\varepsilon)$. In most cases $\varepsilon = \Omega(1)$, thus this only gives a trivial lower bound.

The next version of the bound, the negative-weighted adversary [15], is known to be tight [20], but it is also harder to apply. An application of the bound to the k-sum problem was obtained in [9]. This result was later stated in the framework of dual learning graphs [6], which we are about to describe.

Learning graphs is a model of computation introduced in [4, 5]. They are most naturally stated in terms of certificate structures, which describe where 1-certificates can be located in a positive input. Learning graphs capture quantum query complexity of certificate structures in the following sense. Let L be the learning graph complexity of a certificate structure $\mathcal{C}$. First, for *any* function with certificate structure $\mathcal{C}$, there exists a quantum algorithm solving it in $O(L)$ queries. Second, there exists *some* function with certificate structure $\mathcal{C}$ and quantum query complexity $\Omega(L)$. In general these functions are rather contrived, yet one example of them being natural are the following sum problems. A sum problem is a *total* function parametrised by a family $\mathcal{S}$ of $O(1)$-sized subsets of $[n]$. The task, given an input string $x \in [q]^n$, is to detect whether there exists $S \in \mathcal{S}$ such that $\sum_{i\in S} x_i$ is divisible by q. Note that our problems do not fall into this category, because every positive input is promised to have many such subsets.

While dual learning graphs give tight lower bounds for all of the above sum problems, in general, of course, they do not give lower bounds for all problems with a given certificate structure. For example, the learning graph complexity of the certificate structure corresponding to the hidden shift problem is $\Theta(n^{1/3})$, whereas its quantum query complexity is logarithmic. What about the 3-shift-sum problem? It turns out that dual learning graphs are still of help here, but in a slightly different way. The learning graph complexity of the corresponding certificate structure is $\Theta(\sqrt{n})$, yet we do not know whether it can be converted into a quantum query lower bound. However, a dual learning graph for a *different* certificate structure can be converted into, albeit not tight, but still a polynomial lower bound. This shows that dual learning graphs are more versatile than we thought.

Another interesting feature of our result is that it might be the simplest constructed example of the adversary bound surpassing the property testing barrier. Examples of the negative-weighted adversary breaking the certificate complexity and the property testing barriers were already obtained in [15]. But [15] did not cover the most interesting regime $\varepsilon = \Omega(1)$ of the property testing barrier. The sum problems of [6] are relatively simple examples of overcoming the certificate complexity barrier. An example for the $\varepsilon = \Omega(1)$ regime of the property testing barrier was constructed in [7], but the construction is quite technical. Our result gives a similar example by much simple means, comparable to that of [6].

Concerning the 3-matching-sum problem, our lower bound is an application of the technique developed for the set equality problem [7]. It is based on the representation theory of the symmetric group. Surprisingly, the technique can be used for the 3-matching-sum problem with essentially no modifications: our proof uses representation theory to a minimal extent, and mostly follows from combinatorial estimates involving the dual learning graph. This indicates that our technique has a potential to be used in proving lower bounds for other symmetric problems.

1.2 Results in property testing

In the property testing model, one is given some property (a set of positive inputs), and the task is to distinguish whether the input possesses the property, or is ε-far, in the relative Hamming distance, from any input that has the property.

Overcoming the property testing barrier automatically gives a lower bound for a property testing problem—that of testing whether the input is positive. But it is not always the most natural way to state the problem. We give an example of a lower bound for a problem that is most naturally stated in the setting of property testing.

The 3-shift-sum problem, as formulated above, must have relatively large q for the problem to be interesting. For instance, it is easy to see that for $q = 2$ there are almost no negative inputs. In our lower bound, we require that $q = \Omega(n^3)$. But it is possible to formulate a version of the problem that is interesting even when the input alphabet is Boolean. Define the set of positive inputs as before, and define the set of negative inputs as being at relative Hamming distance at least, say, $1/7$ to it. We prove a lower bound of $\Omega(n^{1/3})$ also for this version of the problem.

Although there is quite a number of quantum algorithms for property testing problems, there are not so many quantum lower bounds known. (An interested reader might consult a recent survey [19] for more information on the topic.) One of the main reasons, of course, is the property testing barrier for the positive-weighted adversary. Up to our knowledge, our result is the first property testing lower bound proven using the adversary method, which answers the problem mentioned in [19]. This shows yet another area of applications of dual learning graphs.

2 Preliminaries

For positive integers m and $\ell \geq m$, let $[m]$ denote the set $\{1, 2, ..., m\}$ and $[m..\ell]$ denote the set $\{m, m+1, \dots, \ell\}$. For P a predicate, we use 1_P to denote the variable that is 1 if P is true, and 0 otherwise.

For an $\mathcal{I} \times \mathcal{J}$-matrix A, $i \in \mathcal{I}$, and $j \in \mathcal{J}$, we denote by $A[\![i,j]\!]$ its (i,j)-th entry. For $\mathcal{I}' \subseteq \mathcal{I}$ and $\mathcal{J}' \subseteq \mathcal{J}$, $A[\![\mathcal{I}', \mathcal{J}']\!]$ denotes the corresponding submatrix. We use similar notation also for vectors. Next, $\|\cdot\|$ denotes the spectral norm (the largest singular value), and $\circ$ denotes the Hadamard (i.e., entry-wise) product of matrices. We often identify projectors with the spaces they project onto.

2.1 Adversary Bound

For background on quantum query complexity the reader may refer to [11]. In the paper, we only require the knowledge of the (negative-weighted) adversary bound for decision problems, which we are about to define.

Let $f\colon \mathcal{D} \to \{0,1\}$ with $\mathcal{D} \subseteq [q]^n$. An *adversary matrix* for f is a non-zero $f^{-1}(1) \times f^{-1}(0)$-matrix Γ. For any $j \in [n]$, the $f^{-1}(1) \times f^{-1}(0)$-matrix Δ_j is defined by

$$\Delta_j[\![x,y]\!] = \begin{cases} 0, & \text{if } x_j = y_j; \\ 1, & \text{if } x_j \neq y_j. \end{cases} \qquad (1)$$

▶ **Theorem 1** (Adversary bound, [15, 18, 9]). *In the above notation, the quantum query complexity of the function f is $\Theta\big(\mathrm{ADV}^{\pm}(f)\big)$, where $\mathrm{ADV}^{\pm}(f)$ is the optimal value of the semi-definite program*

$$\text{maximise} \quad \|\Gamma\| \tag{2a}$$
$$\text{subject to} \quad \|\Delta_j \circ \Gamma\| \le 1 \qquad \textit{for all } j \in [m], \tag{2b}$$

where the maximisation is over all adversary matrices Γ for f.

We can choose any adversary matrix Γ and scale it so that the condition $\|\Delta_j \circ \Gamma\| \le 1$ holds. Thus, we often use the condition $\|\Delta_j \circ \Gamma\| = \mathrm{O}(1)$ instead of $\|\Delta_j \circ \Gamma\| \le 1$.

Working with the matrix $\Delta_j \circ \Gamma$ might be cumbersome, so we do the following transformation instead. We write $\Gamma \overset{\Delta_j}{\longmapsto} \Gamma'$ if $\Gamma \circ \Delta_j = \Gamma' \circ \Delta_j$. In other words, we modify the entries of Γ with $x_j = y_j$. Now, from the fact [18] that $\gamma_2(\Delta_j) = \max_B \big\{\|\Delta_j \circ B\| : \|B\| \le 1\big\} \le 2$, we conclude that $\|\Delta_j \circ \Gamma\| \le 2\|\Gamma_j\|$, hence we can replace $\Delta_j \circ \Gamma$ with Γ' in (2b).

It is sometimes convenient [9] to allow several rows or columns corresponding to the same input x. We add labels to distinguish different rows corresponding to the same input.

2.2 Certificate Structures and Dual Learning Graphs

Let $f\colon \mathcal{D} \to \{0,1\}$ be a function with domain $\mathcal{D} \subseteq [q]^n$. For $x \in f^{-1}(1)$, a *certificate* for x is a subset $S \subseteq [n]$ such that $f(z) = 1$ for all $z \in \mathcal{D}$ satisfying $x_i = y_i$ for all $i \in S$. A *certificate structure* $\mathcal{C}$ is a collection of non-empty subsets of $2^{[n]}$. We say that f *has* certificate structure $\mathcal{C}$ if, for every $x \in f^{-1}(1)$, there exists $\mathcal{M} \in \mathcal{C}$ such that every $S \in \mathcal{M}$ is a certificate for x. It is natural to assume that all $\mathcal{M} \in \mathcal{C}$ are upward closed.

There are two formulations of the learning graph complexity: primal and dual. For the purposes of this paper, it is enough to state the dual one. A *dual learning graph* for a certificate structure $\mathcal{C}$ is a feasible solution to the following optimisation problem:

$$\text{maximise} \quad \sqrt{\sum\nolimits_{\mathcal{M}\in\mathcal{C}} \alpha(\mathcal{M},\emptyset)^2} \tag{3a}$$
$$\text{subject to} \quad \sum_{\mathcal{M}\in\mathcal{C}} \big(\alpha(\mathcal{M},S) - \alpha(\mathcal{M},S\cup\{j\})\big)^2 \le 1 \qquad \forall S \subseteq [n],\ \forall j \in [n]\setminus S; \tag{3b}$$
$$\alpha(\mathcal{M},S) = 0 \qquad \forall S \in \mathcal{M}; \tag{3c}$$
$$\alpha(\mathcal{M},S) \in \mathbb{R} \qquad \forall \mathcal{M} \in \mathcal{C},\ \forall S \subseteq [n]. \tag{3d}$$

The optimal value of this optimisation problem is called the learning graph complexity of $\mathcal{C}$.

We call a solution to the dual learning graph for $\mathcal{C}$ any mapping $\alpha(\mathcal{M},S)$ satisfying (3d), where we implicitly assume (3c). A solution is feasible if it satisfies (3b). It is easy to see that any optimal solution $\alpha(\mathcal{M},S)$ to (3) is entry-wise non-negative and non-increasing in S. We will implicitly assume that any feasible solution satisfies these requirements.

Inspired by this optimisation problem, we define the norm of a solution α as $\|\alpha\| = \max_{S\subseteq[n]} \sqrt{\sum_{\mathcal{M}\in\mathcal{C}} \alpha(\mathcal{M},S)^2}$. It satisfies the usual axioms of a norm, although we will not use this fact. For $j \in [n]$, we define an operation ∂_j given by

$$\partial_j\alpha(\mathcal{M},S) = \begin{cases} \alpha(\mathcal{M},S) - \alpha(\mathcal{M},S\cup\{j\}), & \text{if } j \notin S; \\ 0, & \text{if } j \in S. \end{cases}$$

If α is a solution to the dual learning graph, so is $\partial_j\alpha$. Condition (3b) can be restated as $\|\partial_j\alpha\| \le 1$ for all $j \in [n]$. If $\alpha(\mathcal{M},S)$ is non-increasing in S, the objective value (3a) is given by $\|\alpha\|$.

Dual learning graphs have close connection to adversary matrices, which we discuss in Section 4.

2.3 Representation Theory

In this section, we introduce basic notions from the representation theory of finite groups with special emphasis on the symmetric group. For more background, the reader may refer to [12, 22] for general theory, and to [16, 21] for the special case of the symmetric group.

Assume G is a finite group. The *group algebra* $\mathbb{C}G$ is the complex vector space with the elements of G forming an orthonormal basis, where the multiplication law of G is extended to $\mathbb{C}G$ by linearity. A (left) *G-module*, also called a *representation* of G, is a complex vector space V with the left multiplication operation by the elements of $\mathbb{C}G$ satisfying the usual associativity and distributivity conditions. We can treat elements of $\mathbb{C}G$ as linear operators acting on V.

A *G-morphism* (or just morphism, if G is clear from the context) between two G-modules V and W is a linear operator $\theta\colon V \to W$ that commutes with all $\alpha \in \mathbb{C}G$: $\theta\alpha = \alpha\theta$, where the first α acts on V and the second one on W.

A G-module is called *irreducible* (or just irrep for irreducible representation) if it does not contain a non-trivial G-submodule. For any G-module V, one can define its *canonical* decomposition into the direct sum of *isotypic* subspaces, each spanned by all copies of a fixed irrep in V. Different isotypic subspaces in this decomposition are orthogonal. If an isotypic subspace contains at least one copy of the irrep, we say that V *uses* this irrep.

If G and H are finite groups, then the irreducible $G \times H$-modules are of the form $V \otimes W$ where V is an irreducible G-module and W is an irreducible H-module. And the corresponding group action is given by $(g,h)(v \otimes w) = gv \otimes hw$, with $g \in G$, $h \in H$, $v \in V$, and $w \in W$, which is extended by linearity.

We use the following results:

▶ **Lemma 2** (Schur's Lemma). *Assume $\theta\colon V \to W$ is a morphism between two irreducible G-modules V and W. Then, $\theta = 0$ if V and W are non-isomorphic; otherwise, θ is uniquely determined up to a scalar multiplier.*

▶ **Lemma 3** ([7]). *Let $\theta\colon V \to W$ be a morphism between two G-modules V and W. Then, there exists an irrep in V all consisting of right-singular vectors of θ of singular value $\|\theta\|$. (We call such right-singular vectors principal.)*

Let $\mathfrak{S}_L$ denote the *symmetric group* on a finite set L, that is, the group with the permutations of L as elements, and composition as the group operation. If m is a positive integer, $\mathfrak{S}_m$ denotes the isomorphism class of the symmetric groups $\mathfrak{S}_L$ with $|L| = m$. Representation theory of $\mathfrak{S}_m$ is closely related to *Young diagrams*, defined as follows.

A *partition* λ of an integer m is a non-increasing sequence $(\lambda_1, \dots, \lambda_t)$ of positive integers satisfying $\lambda_1 + \cdots + \lambda_t = m$. A partition $\lambda = (\lambda_1, \dots, \lambda_t)$ is often represented in the form of a Young diagram that consists, from top to bottom, of rows of $\lambda_1, \lambda_2, \dots, \lambda_t$ boxes aligned by the left side. We say that a partition *has k boxes below the first row* if $\lambda_1 = m - k$. For each partition λ of m, there exists an irreducible $\mathfrak{S}_m$-module S^λ, called the *Specht module*. All these modules are pairwise non-isomorphic, and give a complete list of all the irreps of $\mathfrak{S}_m$.

3 Formulation of the Problems and Easy Observations

In this section, we formulate the 3-shift-sum problem and define the closely related 3-matching-sum problem. We also sketch proofs of few simple observations about these problems.

Both the 3-shift-sum and the 3-matching-sum are partial Boolean functions defined on $[q]^{3n}$, with q and n positive integers. The $3n$ input variables are divided into three groups $A = [1..n]$, $B = [n+1..2n]$, and $C = [2n+1..3n]$. A *3-dimensional matching* is a partition μ of the set $[3n]$ into n triples, $\mu = \{T_1, \dots, T_n\}$, such that $|T_i \cap A| = |T_i \cap B| = |T_i \cap C| = 1$ for all i. This is a natural generalisation of the usual (2-dimensional) matching between sets A and B. We denote the set of 3-dimensional matchings by $M_{\rm M}$ (we omit n, assuming its value is clear from the context). We consider a special type of 3-dimensional matchings, we call *3-shifts*. A 3-shift is a matching $\mu = \{T_1, \dots, T_n\}$ such that there exist two numbers $b, c \in [n]$ such that $T_i = \bigl\{i,\ n+1+(i+b \bmod n),\ 2n+1+(i+c \bmod n)\bigr\}$ for all i. We denote the set of 3-shifts by $M_{\rm S}$.

We define the *3-shift-sum* and the *3-matching-sum* problems as follows. Let $M_{\rm Q}$ stand for $M_{\rm S}$ in 3-shift-sum and for $M_{\rm M}$ in 3-matching-sum. In a positive input $x \in [q]^{3n}$, there exists $\mu \in M_{\rm Q}$ such that $x_a + x_b + x_c$ is divisible by q for every triple $\{a,b,c\} \in \mu$. We say that x is *of the form* μ in this case. In a negative input $y \in [q]^{3n}$, we have $y_a + y_b + y_c \not\equiv 0 \pmod q$ for any choice of $a \in A$, $b \in B$, and $c \in C$. The task is to determine whether the input is positive or negative, provided that one of the two options holds. Since 3-shift-sum is a special case of 3-matching-sum, the latter is a harder problem.

3.1 Randomised and Quantum Complexity

Let us describe what we can immediately say about quantum and randomised query complexities of these problems. Neither result will be relevant later in the paper.

▶ **Proposition 4.** *The quantum query complexity of the 3-shift-sum and the 3-matching-sum problems is* $O(\sqrt{n})$.

Proof sketch. Consider a positive input x, and let $\mu \in M_{\rm Q}$ be its form. Take random subsets $A' \subseteq A$ and $B' \subseteq B$ of size approximately $\sqrt{n}$, and query all the variables in $A' \cup B'$. With high probability, there exists $T \in \mu$ that intersects both A' and B'. Now use Grover's search to find an element $c \in C$ satisfying $x_a + x_b + x_c \equiv 0 \pmod q$ for some $a \in A'$ and $b \in B'$. ◀

▶ **Proposition 5.** *The randomised query complexity of the 3-shift-sum and the 3-matching-sum problems is* $\Theta(n^{2/3})$.

The proof is totally standard, and it can be found in the full version of the paper [8].

3.2 Certificate Structures

It is easy to describe the certificate structures $\mathcal{C}_{\rm S}$ and $\mathcal{C}_{\rm M}$ of the 3-shift-sum and the 3-matching-sum problems. For each $\mu \in M_{\rm Q}$, there is a corresponding $\mathcal{M}_\mu \in \mathcal{C}_{\rm Q}$ obtained as follows: a subset $S \subseteq [n]$ is in $\mathcal{M}_\mu$ if and only if there exists a triple $T \in \mu$ satisfying $T \subseteq S$.

The lower bound from the following proposition will be our main source of inspiration when constructing adversary bounds later in the paper.

▶ **Proposition 6.** *The learning graph complexity of the certificate structures* $\mathcal{C}_{\rm S}$ *and* $\mathcal{C}_{\rm M}$ *is* $\Theta(\sqrt{n})$.

Proof. The upper bound is similar to Proposition 4, and we omit the proof. The upper bound is stated here for completeness, and we do not use it further in the paper.

Let us prove the lower bound. For that we have to construct a feasible solution to the dual learning graph. For $\mathcal{M} \in \mathcal{C}_{\rm Q}$, define

$$\alpha(\mathcal{M}, S) = \frac{1}{\sqrt{|M_{\rm Q}|}} \max\bigl\{\sqrt{n} - |S|, 0\bigr\} \qquad \text{if } S \notin \mathcal{M}, \tag{4}$$

and as 0 otherwise. It is easy to see that the objective value (3a) is $\sqrt{n}$, and that (3c) holds.

It remains to check (3b). Fix S and j. If $|S| \geq \sqrt{n}$, then the left-hand side of (3b) is zero, so assume $|S| \leq \sqrt{n}$. We have the following contributions to the left-hand side of (3b):

- If $S \cup \{j\} \notin \mathcal{M}$, then the value of $\alpha(\mathcal{M}, S)$ changes by $\frac{1}{\sqrt{|M_Q|}}$ as $|S|$ increases by 1.
- If $\mu \in M_Q$ is taken uniformly at random, the probability is $O\big((|S|/n)^2\big) = O(1/n)$ that $S \notin \mathcal{M}_\mu$ but $S \cup \{j\} \in \mathcal{M}_\mu$. In this case, $\alpha(\mathcal{M}_\mu, S)$ changes by at most $\sqrt{\frac{n}{|M_Q|}}$.

Altogether we have:

$$\sum_{\mathcal{M} \in \mathcal{C}} \big(\alpha(\mathcal{M}, S) - \alpha(\mathcal{M}, S \cup \{j\})\big)^2 \leq |M_Q| \cdot \frac{1}{|M_Q|} + O\left(\frac{|M_Q|}{n}\right) \cdot \frac{n}{|M_Q|} = O(1).$$

Scaling down α by a constant factor, we get a feasible solution with objective value $\Omega(\sqrt{n})$. ◀

4 Basic Definitions

In this section we introduce our basic notation, and describe a procedure of converting a solution to the dual learning graph into an adversary matrix. This is a general procedure from [6] tailored for the special case of the 3-shift-sum and the 3-matching-sum problems. This procedure does not immediately result in good adversary matrices for these problems, but we are able modify it in Sections 5 and 6 so that it works. Let again M_Q stand for either M_S or M_M.

4.1 Fourier Basis

Let $\mathcal{H} = \mathbb{C}^{\mathbb{Z}_q}$ and $e_0, \ldots, e_{q-1}$ be the Fourier basis of $\mathcal{H}$. Recall that it is an orthonormal basis given by $e_i[\![j]\!] = \frac{1}{\sqrt{q}}\omega^{ij}$, where $\omega = \mathrm{e}^{2\pi\mathrm{i}/q}$. For m a positive integer, the Fourier basis of $\mathcal{H}^{\otimes m}$ is given by tensor products $e_{a_1} \otimes \cdots \otimes e_{a_m}$. A component e_{a_i} in this tensor product is called non-zero if $a_i \neq 0$. The weight of the Fourier basis element is the number of non-zero components.

We define two projectors in $\mathcal{H}$: $\Pi_0 = e_0 e_0^*$ and $\Pi_1 = I - \Pi_0 = \sum_{i=1}^{q-1} e_i e_i^*$. All the entries of Π_0 are equal to $1/q$. Important relations are $\Pi_0 \overset{\Delta}{\longmapsto} \Pi_0$ and $\Pi_1 \overset{\Delta}{\longmapsto} -\Pi_0$, where Δ is as in (1) and acts on the sole variable. For two sets $R \subseteq T$, we define a projector Π_R^T in the space $\mathcal{H}^T$ by $\Pi_R^T = \bigotimes_{j \in T} \Pi_{1_{j \in R}}$. As R ranges over all subsets of T, this gives an orthogonal decomposition of $\mathcal{H}^T$. By the above relations:

$$\Pi_R^T \overset{\Delta_j}{\longmapsto} \Pi_R^T \quad \text{if} \quad j \notin R \qquad \text{and} \qquad \Pi_R^T \overset{\Delta_j}{\longmapsto} -\Pi_{R \setminus \{j\}}^T \quad \text{if} \quad j \in R. \tag{5}$$

If $\mathcal{A}$ is a collection of subsets of T, we can define projector $\Pi_{\mathcal{A}}^T = \sum_{R \in \mathcal{A}} \Pi_R^T$. We clearly have $\Pi_{\mathcal{A}} \Pi_{\mathcal{B}} = \Pi_{\mathcal{A} \cap \mathcal{B}}$. We will use this construction only for some special cases, in particular, for a positive integer k, we define $\Pi_k^T = \sum_{R \subseteq T, |R| = k} \Pi_R^T$.

4.2 Basic Operators

Let $\mu = \{T_1, \ldots, T_n\}$ be a 3-dimensional matching. Let $\mathcal{P}^\mu$ denote the set of positive inputs of form μ. We use $\mathcal{P}$ for the set of pairs (μ, x) with $\mu \in M_Q$ and $x \in \mathcal{P}^\mu$. Think of $\mathcal{P}$ as the set of positive inputs with additional labels so that some inputs x can appear multiple times. We use $\mathcal{N}$ for the set of negative inputs, and $\mathcal{U} = [q]^{3n}$ for the set of all strings. Similarly to the proof of Proposition 5, $\mathcal{U}$ will be close to $\mathcal{N}$, and we use the former as a proxy for the latter.

Now assume T is a triple of elements. Think of it as an element of a 3-dimensional matching μ. Denote

$$P^T = \{(a,b,c) \in [q]^T \mid a+b+c \equiv 0 \pmod q\}. \tag{6}$$

Thus, $\mathcal{P}^\mu$ is the Cartesian product $\prod_{T\in\mu} P^T$. For $R \subseteq T$, define $\Psi^T_R = \sqrt{q}\,\Pi^T_R [\![P^T, [q]^T]\!]$, where the factor $\sqrt{q}$ is introduced to account for the reduced number of rows. For $S \subseteq [3n]$, let

$$\Psi^\mu_S = \bigotimes\nolimits_{T\in\mu} \Psi^T_{S\cap T} = q^{n/2}\,\Pi^{[3n]}_S [\![\mathcal{P}^\mu, \mathcal{U}]\!].$$

As for $\Pi^T_{\mathcal{A}}$, we will use $\Psi^\mu_{\mathcal{A}} = \sum_{S\in\mathcal{A}} \Psi^\mu_S$ for a family $\mathcal{A}$ of subsets of $[3n]$. Again, $\Psi^\mu_{\mathcal{A}}\Pi^{[3n]}_{\mathcal{B}} = \Psi^\mu_{\mathcal{A}\cap\mathcal{B}}$. Using (5), we have

$$\Psi^\mu_S \overset{\Delta_j}{\longmapsto} \Psi^\mu_S \quad \text{if} \quad j \notin S \qquad \text{and} \qquad \Psi^\mu_S \overset{\Delta_j}{\longmapsto} -\Psi^\mu_{S\setminus\{j\}} \quad \text{if} \quad j \in S. \tag{7}$$

4.3 From Dual Learning Graphs to Adversary Matrices

Now we explain how to convert a solution α to the dual learning graph (3) into a $\mathcal{P}\times\mathcal{U}$-matrix $G(\alpha)$. In [6], the adversary matrix Γ was obtained by restricting $\Gamma = G(\alpha)[\![\mathcal{P}, \mathcal{N}]\!]$. It is convenient to allow all the columns corresponding to $\mathcal{U}$, and restrict them to $\mathcal{N}$ only at the very end.

If $\mu \in M_{\mathrm{Q}}$, let us for brevity write $\alpha(\mu, S)$ for $\alpha(\mathcal{M}_\mu, S)$. The matrix $G(\alpha)$ is defined block-wise by $G(\alpha)[\![\mathcal{P}^\mu, \mathcal{U}]\!] = G^\mu(\alpha) = \sum_{S\subseteq[n]} \alpha(\mu,S)\Psi^\mu_S$. Eq. (7) gives the following important relation:

$$G(\alpha) \overset{\Delta_j}{\longmapsto} G(\partial_j\alpha). \tag{8}$$

4.4 Extended Matrices

Eq. (8) gives one connection between $G(\alpha)$ and the optimisation problem in (3). Here we give another one. For that, we define an extended version $\widetilde{G}(\alpha)$ of $G(\alpha)$.

Let $\widetilde{\mathcal{U}} = M_{\mathrm{Q}} \times \mathcal{U}$. We use $\widetilde{\mathcal{U}}^\mu$ to denote $\{\mu\} \otimes \mathcal{U}$. The $\widetilde{\mathcal{U}} \times \mathcal{U}$-matrix $\widetilde{G}(\alpha)$ is defined block-wise: $\widetilde{G}(\alpha)[\![\mathcal{U}^\mu, \mathcal{U}]\!] = \widetilde{G}^\mu(\alpha) = \sum_{S\subseteq[n]} \alpha(\mu,S)\Pi^{[3n]}_S$. Clearly, $G(\alpha) = q^{n/2}\widetilde{G}(\alpha)[\![\mathcal{P}, \mathcal{U}]\!]$.

Using that $\{\Pi^{[3n]}_S\}$ is a decomposition of $\mathcal{H}^{3n}$ into orthogonal subspaces, we get

$$\widetilde{G}(\alpha)^*\widetilde{G}(\alpha) = \sum_{\mu\in M_{\mathrm{Q}}} \bigl(\widetilde{G}^\mu(\alpha)\bigr)^*\widetilde{G}^\mu(\alpha) = \sum_{S\subseteq[3n]} \left[\sum_{\mu\in M_{\mathrm{Q}}} \alpha(\mu,S)^2\right]\Pi^{[3n]}_S.$$

As $\|A\| = \sqrt{\|A^*A\|}$ for any matrix A, we obtain another important relation:

$$\|\widetilde{G}(\alpha)\| = \|\alpha\|. \tag{9}$$

Of course it also holds for $\partial_j\alpha$. If $\|\widetilde{G}(\alpha)\|$ and $\|G(\alpha)\|$ were close, then any feasible solution α would give an adversary matrix $\Gamma = G(\alpha)[\![\mathcal{P}, \mathcal{N}]\!]$ with value $\|\alpha\|$. It is easy to lower bound $\|\Gamma\|$ in terms of $\|\alpha\|$, see Lemma 8 below, but, in general, $\|G(\partial_j\alpha)\|$ will be much larger than $\|\widetilde{G}(\partial_j\alpha)\|$. In particular, this is the case when α is the solution from Proposition 6. Our main challenge in the coming sections will be to find ways to reduce $\|G(\partial_j\alpha)\|$.

4.5 Reducing Extended Matrices

Here we will give a finer relation between $G(\alpha)$ and $\widetilde{G}(\alpha)$ than the trivial relation $G(\alpha) = q^{n/2}\widetilde{G}(\alpha)[\![\mathcal{P},\mathcal{U}]\!]$. For $\mathcal{M}_\mu \in \mathcal{C}_{\mathrm{Q}}$, we define

$$\Pi^\mu = \sum_{S\subseteq[3n], S\notin\mathcal{M}_\mu} \Pi_S^{[3n]} \qquad \text{and} \qquad \Psi^\mu = \sum_{S\subseteq[3n], S\notin\mathcal{M}_\mu} \Psi_S^\mu. \tag{10}$$

By condition (3c), we have $\widetilde{G}^\mu(\alpha) = \Pi^\mu \widetilde{G}^\mu(\alpha)$, and, thus, $G^\mu(\alpha) = \Psi^\mu \widetilde{G}^\mu(\alpha)$. If we define a linear operator $\Psi_{\mathrm{Q}} \colon \mathcal{H}^{\widetilde{\mathcal{U}}} \to \mathcal{H}^{\mathcal{P}}$ by $\Psi_{\mathrm{Q}} = \bigoplus_{\mu\in M_{\mathrm{Q}}} \Psi^\mu$, we get

$$G(\alpha) = \Psi_{\mathrm{Q}} \widetilde{G}(\alpha). \tag{11}$$

In the light of discussion after (9), it would help if we could upper bound the norm of Ψ_{Q}. Unfortunately, its norm is exponential. Indeed, we can write $\Psi^\mu = \bigotimes_{T\in\mu} \Psi^T_{\le 2}$,where $\Psi^T_{\le 2} = \sum_{R\subset T, R\ne T} \Psi^T_R$. We prove its basic properties in the next claim, where we also study the operator $\Psi^T_{\le 1} = \sum_{R\subset T, |R|\le 1} \Psi^T_R$.

▶ **Claim 7.** *We have the following estimates*

(a) $\left\|\Psi^T_{\le 2}\right\| = \sqrt{3}$,

(b) $\left\|\Psi^T_{\le 2}(\Pi_0\otimes I_{\mathcal{H}}\otimes I_{\mathcal{H}})\right\| = \left\|\Psi^T_{\le 2}(I_{\mathcal{H}}\otimes \Pi_0\otimes I_{\mathcal{H}})\right\| = \left\|\Psi^T_{\le 2}(I_{\mathcal{H}}\otimes I_{\mathcal{H}}\otimes \Pi_0)\right\| = 1$,

(c) $\left\|\Psi^T_{\le 1}\right\| = 1$,

(d) $(\Psi^T_\emptyset)^*\Psi^T_{\le 2} = \Pi^T_\emptyset$, *and* $\|\Psi^T_\emptyset\| = 1$.

The proof of the claim can be found in the full version of the paper. In the next sections, we will use points (b) and (c) of this claim to upper bound the norm of $G(\partial_j\alpha)$ using (11).

4.6 Restricting from $\mathcal{U}$ to $\mathcal{N}$

Finally, we give a general way of bounding the norm of $\Gamma = G(\alpha)[\![\mathcal{P},\mathcal{N}]\!]$ in terms of α. For our upcoming application in Section 6, we prove a slightly more general result. Note that the bound is related to the objective value (3a) of α.

▶ **Lemma 8.** *Let α be a solution to the dual learning graph of $\mathcal{C}_{\mathrm{Q}}$, and V is an arbitrary linear operator in $\mathbb{C}^{\mathcal{U}}$ satisfying $\Pi^{[3n]}_\emptyset V = \Pi^{[3n]}_\emptyset$. Then, $\left\|(G(\alpha)V)[\![\mathcal{P},\mathcal{N}]\!]\right\| \ge \sqrt{\frac{|\mathcal{N}|}{|\mathcal{U}|}\sum_{\mu\in M_{\mathrm{Q}}}\alpha(\mu,\emptyset)^2}$.*

An easy proof of this lemma can be found in the full version of the paper.

5 Lower Bound for the 3-Shift-Sum Problem

The goal is to prove a quantum query lower bound for the 3-shift-sum problem.

▶ **Theorem 9.** *Assume $q \ge 2n^3$. Then the quantum query complexity of the 3-shift-sum problem is $\Omega(n^{1/3})$.*

The main idea behind the lower bound is to use Claim 7(c). In order to do that, we perform a transition to a different certificate structure $\mathcal{C}'_{\mathrm{S}}$. For each $\mu \in M_{\mathrm{S}}$, there is a corresponding $\mathcal{M}'_\mu \in \mathcal{C}'_{\mathrm{S}}$ obtained as follows: a subset $S \subseteq [3n]$ is in $\mathcal{M}'_\mu$ if and only if there exists a triple $T \in \mu$ satisfying $|T \cap S| \ge 2$. Note that this is *not* the certificate structure for the 3-shift-sum problem. Rather it is the certificate structure of a problem one might call the 3-shift-equal problem. The input is a $3\times n$-matrix. In the positive case, there exist circular shifts of rows such that the elements in each column become equal. In the negative case, any two elements from two different rows are different.

▶ **Proposition 10.** *The learning graph complexity of the certificate structure* $\mathcal{C}'_{\mathrm{s}}$ *is* $\Omega(n^{1/3})$.

Proof. The proof is similar to that of Proposition 12 from [6] for the hidden shift problem. We have $|\mathcal{C}'_{\mathrm{s}}| = n^2$. Define

$$\alpha(\mathcal{M}, S) = \frac{1}{n}\max\Bigl\{n^{1/3} - |S|, 0\Bigr\} \qquad \text{if } S \notin \mathcal{M}, \tag{12}$$

and as 0 otherwise. It is easy to see that the objective value (3a) is $n^{1/3}$, and that (3c) holds.

Fix S and j, and let us check (3b). If $|S| \geq n^{1/3}$, then the left-hand side of (3b) is zero, so assume $|S| \leq n^{1/3}$. There are n^2 choices of $\mathcal{M} \in \mathcal{C}'_{\mathrm{s}}$. If $S \cup \{j\} \notin \mathcal{M}$, then the value of $\alpha(\mathcal{M}, S)$ changes by $1/n$ as the size of S increases by 1. Also, there are at most $|S|n \leq n^{4/3}$ choices of $\mathcal{M}$ such that $S \notin \mathcal{M}$ but $S \cup \{j\} \in \mathcal{M}$. For each of them, the value of $\alpha(\mathcal{M}, S)$ changes by at most $n^{-2/3}$. Thus,

$$\sum_{\mathcal{M}\in\mathcal{C}} \bigl(\alpha(\mathcal{M}, S) - \alpha(\mathcal{M}, S \cup \{j\})\bigr)^2 \leq n^2 \cdot \frac{1}{n^2} + n^{4/3} \cdot n^{-4/3} = O(1). \qquad ◀$$

5.1 Regular Version

In this section we prove Theorem 9. Let α'_{s} be the feasible solution (12) for the $\mathcal{C}'_{\mathrm{s}}$ certificate structure. It is also a feasible solution for the $\mathcal{C}_{\mathrm{s}}$ certificate structure. As in Section 4, we define the adversary matrix by $\Gamma = G(\alpha'_{\mathrm{s}})[\![\mathcal{P}, \mathcal{N}]\!]$. By Lemma 8, we get $\|\Gamma\| = \Omega(n^{1/3})$ if we prove that $|\mathcal{N}| = \Omega(|\mathcal{U}|)$. But that is easy: for a uniformly random triple $(a, b, c) \in [q]^3$, the probability that $a + b + c$ is divisible by q is $1/q$. There are n^3 possible triples having one element in each of A, B, and C. Hence, by the union bound, a uniformly random input in $[q]^{3n}$ is negative with probability at least $1 - n^3/q \geq 1/2$. That is, $|\mathcal{N}| \geq q^{3n}/2$.

Now let us prove that $\|\Gamma \circ \Delta_j\| = O(1)$. By (8) and using that Γ is a submatrix of $G(\alpha'_{\mathrm{s}})$, it suffices to prove that $\|G(\partial_j \alpha'_{\mathrm{s}})\| = O(1)$.

Following (10), let us define an analogue of Ψ_{s} for our new certificate structure $\mathcal{C}'_{\mathrm{s}}$ by $\Psi'^{\mu} = \sum_{S\subseteq[3n], S\notin\mathcal{M}'_{\mu}} \Psi^{\mu}_S$ and $\Psi'_{\mathrm{s}} = \bigoplus_{\mu\in M_{\mathrm{s}}} \Psi'^{\mu}$. Similarly to (11), we get $G(\partial_j \alpha'_{\mathrm{s}}) = \Psi'_{\mathrm{s}} \widetilde{G}(\partial_j \alpha'_{\mathrm{s}})$. We have $\|\widetilde{G}(\partial_j \alpha'_{\mathrm{s}})\| = O(1)$ by (9) and Proposition 10. It suffices to prove that $\|\Psi'_{\mathrm{s}}\| = O(1)$. But it is easy to see that $\Psi'^{\mu} = \bigotimes_{T\in\mu} \Psi^T_{\leq 1}$, and, by Claim 7, $\|\Psi'^{\mu}\| = 1$, hence, $\|\Psi'_{\mathrm{s}}\| = 1$.

5.2 Property Testing Version

In this section, we prove a quantum lower bound for the property testing version of the 3-shift-sum problem. Unlike the original version of the 3-shift-sum problem, this problem makes sense even for $q = 2$, so, for concreteness, we will define it for Boolean alphabet, however, similar results also hold for larger alphabet sizes.

An input is a string in $\{0, 1\}^{3n}$. For a positive input x, there exists $\mu \in M_{\mathrm{s}}$ such that $x_a \oplus x_b \oplus x_c = 0$ for every triple $\{a, b, c\} \in \mu$. Here $\oplus$ stands for xor. The negative inputs are defined as being at relative Hamming distance at least ε to the set of positive inputs.

▶ **Theorem 11.** *For* $\varepsilon \leq \frac{1}{7}$, *the property testing version of the 3-shift-sum problem requires* $\Omega(n^{1/3})$ *quantum queries to solve.*

The construction is identical to that in Section 5.1. The proof of $\|\Gamma \circ \Delta_j\| = O(1)$ is identical. In this part of the proof only $\mathcal{P}$ and $\mathcal{U}$ are used, which are the same in the regular and the property testing versions of the problem, and the size of the alphabet is never used.

The only place where the size of the alphabet is used is in lower bounding $\|\Gamma\|$, where it is proven that $|\mathcal{N}| = \Omega(|\mathcal{U}|)$. If we prove this for this version of the problem, we will be done.

Recall that we treat x as an $3\times n$-matrix. Fix the last two rows. The input x is negative if its first row is at relative Hamming distance at least $\frac{3}{7}$ from the xor of any of n^2 circular shifts of the last two rows. A simple application of the Chernoff and the union bounds shows that this is the case with probability $1-o(1)$.

6 Lower Bound for the 3-Matching-Sum Problem

The goal of this section is to prove the following theorem:

▶ **Theorem 12.** *Assume $q \geq 2n^3$. Then the quantum query complexity of the 3-matching-sum problem is $\Omega(\sqrt{n})$.*

Let α_{M} be the feasible solution (4) to the dual learning graph of $\mathcal{C}_{\mathrm{M}}$ from Proposition 6. We will obtain an adversary matrix to the 3-matching-sum problem multiplying $G(\alpha_{\mathrm{M}})$ by a suitably chosen projector V. We define it using symmetries of the problem.

The group $\mathfrak{S}_n$ acts on the set $[q]^n$ in the natural way: $\pi \in \mathfrak{S}_n$ maps $x = (x_1,\dots,x_n)$ to $\pi x = (x_{\pi^{-1}(1)},\dots,x_{\pi^{-1}(n)})$, and by linearity we extend this action to $\mathcal{H}^n$, the latter thus becoming an $\mathfrak{S}_n$-module.

Similarly, the group $\mathfrak{G} = \mathfrak{S}_A\times\mathfrak{S}_B\times\mathfrak{S}_C$ acts on $\mathcal{U}$: a group element $(\pi_A,\pi_B,\pi_C)\in\mathfrak{G}$ acts on $x = (x_A,x_B,x_C)\in\mathcal{U}$ by mapping it to $(\pi_A x_A,\pi_B x_B,\pi_C x_C)$. This action renders $\mathbb{C}^{\mathcal{U}}$ a $\mathfrak{G}$-module. Let $\mathfrak{G}$ act on $\mu\in M_{\mathrm{M}}$ by mapping each triple $(a_1,a_2,a_3)\in\mu$ to $(\pi_A(a_1),\pi_B(a_2),\pi_C(a_3))$. Together with the action on inputs of length $3n$, this gives us an action of $\mathfrak{G}$ on $\mathcal{P}$, hence, $\mathbb{C}^{\mathcal{P}}$ is also a $\mathfrak{G}$-module.

The 3-matching-sum problem is invariant under this action of $\mathfrak{G}$: positive inputs are mapped to positive inputs, and negative inputs are mapped to negative inputs. This means that $\mathbb{C}^{\mathcal{N}}$ is a $\mathfrak{G}$-submodule of $\mathbb{C}^{\mathcal{U}}$. It is easy to see that α_{M} is symmetric with respect to $\mathfrak{G}$, hence, $G(\alpha_{\mathrm{M}})$ is symmetric with respect to $\mathfrak{G}$. In other words, it commutes with any element of $\mathfrak{G}$, or, different still, it is a $\mathfrak{G}$-morphism.

Let T be a finite set. It is easy to see that Π^T_k is a $\mathfrak{S}_T$-submodule of Π^T. From [7], the module Π^T_k only contains irreps with at most k boxes below the first row. Denote by $\bar\Pi^T_k$ the projector onto the span of all irreps with *exactly* k boxes below the first row. In particular, $\bar\Pi^T_0 = \Pi^T_0$.

In order to simplify statements of some results, in particular Lemma 15, let us assume there is a cutting point K such that

$$\alpha(\mu,S) = 0 \quad \text{whenever } |S| > K. \tag{13}$$

For α_{M}, we take $K = \lfloor\sqrt{n}\rfloor$. Define the projectors $V^T = \sum_{k=0}^K \bar\Pi^T_k$ and $V = V^A\otimes V^B\otimes V^C$. Note that $(\Pi_{k_A}\otimes\Pi_{k_B}\otimes\Pi_{k_C})V = (\bar\Pi_{k_A}\otimes\bar\Pi_{k_B}\otimes\bar\Pi_{k_C})$ for all k_A,k_B,k_C between 0 and K.

The adversary matrix Γ is obtained as $\Gamma = \big(G(\alpha_{\mathrm{M}})V\big)[\![\mathcal{P},\mathcal{N}]\!]$.

We know from Section 5.1 that $|\mathcal{N}| = \Omega(|\mathcal{U}|)$. Also $\Pi^{[3n]}_\emptyset V = \bar\Pi^{[3n]}_\emptyset = \Pi^{[3n]}_\emptyset$. Hence, by Lemma 8 and Proposition 6, we have $\|\Gamma\| = \Omega(\sqrt{n})$.

It remains to prove that $\|\Delta_j\circ\Gamma\| = O(1)$ for any j, which we do in the remaining part of this section. Due to symmetry, $\|\Delta_j\circ\Gamma\|$ is the same for all j, so it suffices to consider $j=1$. Note that $\Delta_1\circ\Gamma$ is a $\mathfrak{G}'$-morphism, where $\mathfrak{G}' = \mathfrak{S}_{[2..n]}\otimes\mathfrak{S}_B\otimes\mathfrak{S}_C$.

We have to understand how Δ_1 acts on V, or, in particular, how it acts on $\bar\Pi^{[n]}_k$. For the usual projector, $\Pi^{[n]}_k$, we have the identity $\Pi^{[n]}_k = \Pi_0\otimes\Pi^{[2..n]}_k + \Pi_1\otimes\Pi^{[2..n]}_{k-1}$. Ref. [7] gives the following analogue of this identity for $\bar\Pi^{[n]}_k$, where one should think of $\Phi^{[n]}_k$ as an error term.

▶ **Lemma 13.** *Let $\Phi^{[n]}_k = \bar\Pi^{[n]}_k - \Pi_0\otimes\bar\Pi^{[2..n]}_k - \Pi_1\otimes\bar\Pi^{[2..n]}_{k-1}$. If $k < n/3$, then $\|\Phi^{[n]}_k\| = O(1/\sqrt{n})$.*

Define $\Phi^A = \sum_{k=1}^K \Phi^A_k$. It is easy to see that $\Pi^A_k \Phi^A_k \Pi^A_k = \Phi^A_k$, hence, $\|\Phi^A_k\| = O(1/\sqrt{n})$. Let $\Phi = \Phi^A \otimes V^B \otimes V^C$ and $V' = \Pi_0 \otimes V^{[2..n]} \otimes V^B \otimes V^C$.

From Lemma 13, we get the following variant of relation (8), whose proof is relatively straightforward and can be found in the full version of the paper.

▶ **Lemma 14.** *Let α be a solution to $\mathcal{C}_{\mathrm{M}}$ satisfying (13). Then, $G(\alpha)V \overset{\Delta_1}{\longmapsto} G(\partial_1\alpha)V' + G(\alpha)\Phi$.*

Applying Lemma 14 to $G(\alpha_{\mathrm{M}})$, we obtain $G(\alpha_{\mathrm{M}})V \overset{\Delta_1}{\longmapsto} G(\partial_1\alpha_{\mathrm{M}})V' + G(\alpha_{\mathrm{M}})\Phi$. Denote $W = I^A \otimes V^B \otimes V^C$, where I^A is the identity operator on $\mathcal{H}^A$. Note that $V' = WV'$ and $\Phi = W\Phi$. Also, $\|V'\| = 1$ and $\|\Phi\| = O(1/\sqrt{n})$. Thus, since Γ is a submatrix of $G(\alpha_{\mathrm{M}})V$, it suffices to prove that

$$\big\|G(\partial_1\alpha_{\mathrm{M}})W\big\| = O(1) \qquad \text{and} \qquad \big\|G(\alpha_{\mathrm{M}})W\big\| = O(\sqrt{n}), \tag{14}$$

which is reasonable since $\|\partial_1\alpha_{\mathrm{M}}\| = O(1)$ and $\|\alpha_{\mathrm{M}}\| = O(\sqrt{n})$. We prove this using the following somewhat technical estimate on the norm of $G(\alpha)W$, whose proof can be found in the full version of the paper.

▶ **Lemma 15.** *Let α be a solution to the dual learning graph of $\mathcal{C}_{\mathrm{M}}$ satisfying (13) and symmetric with respect to $\mathfrak{S}_B \times \mathfrak{S}_C$. Then $\|G(\alpha)W\| \le \max_{k_B,k_C\in[0..K]} \Lambda_{k_B,k_C}(\alpha)$, where $\Lambda_{k_B,k_C}(\alpha)$ is defined in the following way. Let $R_B = [n+1..n+2k_B]$ and $R_C = [2n+1..2n+2k_C]$. Let $L(\mu,R_B,R_C)$ be the number of triples in the matching μ that intersect both R_B and R_C. Then*

$$\Lambda_{k_B,k_C}(\alpha) = \sqrt{\sum_{\mu\in M_{\mathrm{M}}} 3^{L(\mu,R_B,R_C)} \max_{S\subseteq A\cup R_B\cup R_C} \alpha(\mu,S)^2}. \tag{15}$$

Now we show how to use Lemma 15 to prove the estimates in (14). The exponential term in (15) might be somewhat of a concern, but we prove that the fraction of matchings with large $L(\mu,R_B,R_C)$ decreases even faster.

▶ **Lemma 16.** *Assume $|R_B|, |R_C| \le 2\sqrt{n}$. Then, $\Pr_\mu\big[L(\mu,R_B,R_C) = \ell\big] \le 8^\ell/\ell!$, where the probability is over uniformly random $\mu \in M_{\mathrm{M}}$.*

Proof. Fix ℓ elements in each R_B and R_C. The probability that these elements are mutually matched by a random μ is $\binom{n}{\ell}^{-1}$. Hence, by the union bound, the probability that for a randomly chosen μ there are ℓ (or more) elements in R_B matched to elements in R_C is at most $\binom{|R_B|}{\ell}\binom{|R_C|}{\ell}\Big/\binom{n}{\ell} \le \frac{(2\sqrt{n})^{2\ell}}{\ell!(n/2)^\ell} \le \frac{8^\ell}{\ell!}$, where we have assumed that n is large enough so that $\ell \le 2\sqrt{n} < n/2$. ◀

▶ **Claim 17.** *We have $\big\|G(\alpha_{\mathrm{M}})W\big\| = O(\sqrt{n})$.*

Proof. We apply Lemma 15. By (4), we have $\alpha_{\mathrm{M}}(\mu,S)^2 \le n/|M_{\mathrm{M}}|$ for all μ and S. Hence, $\Lambda_{k_B,k_C}(\alpha_{\mathrm{M}}) \le \sqrt{n}\sqrt{\mathbb{E}_{\mu\in M_{\mathrm{M}}}\big[3^{L(\mu,R_B,R_C)}\big]}$. And, using Lemma 16:

$$\mathbb{E}_{\mu\in M_{\mathrm{M}}}\big[3^{L(\mu,R_B,R_C)}\big] \le \sum_{\ell=0}^{\infty} 3^\ell \cdot \frac{8^\ell}{\ell!} = \mathrm{e}^{24} = O(1), \tag{16}$$

which gives the required bound. ◀

▶ **Claim 18.** *We have $\big\|G(\partial_1\alpha_{\mathrm{M}})W\big\| = O(1)$.*

Proof. We apply Lemma 15. By the analysis in Proposition 6, we see that

$$\max_{S\subseteq A\cup R_B\cup R_C} \partial_1\alpha_{\text{M}}(\mu,S)^2 \le \frac{1}{|M_{\text{M}}|}\begin{cases} n, \text{ if 1 is matched by } \mu \text{ to elements in both } R_B \text{ and } R_C;\\ 1, \text{ otherwise.}\end{cases}$$

Let us call the event in the first case above $Z(\mu)$. Then,

$$\Lambda_{k_A,k_B}(\partial_1\alpha_{\text{M}})^2 \le \mathbb{E}_{\mu\in M_{\text{M}}}\Big[3^{L(\mu,R_B,R_C)}\Big] + \Pr_\mu[Z(\mu)]\cdot 3n\,\mathbb{E}_{\mu\in M_{\text{M}}}\Big[3^{L(\mu,R_B,R_C)-1}\;\Big|\; Z(\mu)\Big].$$

The first term is $O(1)$ by (16). For the second term, it is easy to see that $\Pr_\mu[Z(\mu)] = |R_B||R_C|/n^2 = O(1/n)$, and the conditioned expectation the same as in (16), because we can remove the triple containing 1 from consideration thus reducing to the same problem with n, $|R_B|$, and $|R_C|$ smaller by 1. This gives the required bound of $O(1)$. ◀

7 Open Problems

The obvious open problem is to resolve the quantum query complexity of the 3-shift-sum problem. So far we only have an $\Omega(n^{1/3})$ lower bound and an $O(n^{1/2})$ upper bound, however, it is not clear how to improve on either of them.

For the 3-matching-sum problem, we have proven matching upper and lower bounds of $\Theta(\sqrt{n})$. An interesting problem is to generalise this to the k-matching-sum problem for arbitrary k. The main limitation seems to be the norm of the error term in Lemma 13.

Some other open problems can be formulated. What functions with randomised query complexity $\omega(\sqrt{n})$ could potentially have poly-logarithmic quantum query complexity? Or, can a relatively general result be proven that excludes some of such functions? For what other problems can the dual learning graph framework be useful?

References

1 Scott Aaronson and Andris Ambainis. Forrelation: A problem that optimally separates quantum from classical computing. In *Proc. of 47th ACM STOC*, pages 307–316, 2015.
2 Scott Aaronson, Shalev Ben-David, and Robin Kothari. Separations in query complexity using cheat sheets. In *Proc. of 48th ACM STOC*, pages 863–876, 2016.
3 Andris Ambainis. Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64(4):750–767, 2002.
4 Aleksandrs Belovs. Span programs for functions with constant-sized 1-certificates. In *Proc. of 44th ACM STOC*, pages 77–84, 2012.
5 Aleksandrs Belovs. *Applications of the Adversary Method in Quantum Query Algorithms*. PhD thesis, University of Latvia, 2014.
6 Aleksandrs Belovs and Ansis Rosmanis. On the power of non-adaptive learning graphs. *Computational Complexity*, 23(2):323–354, 2014.
7 Aleksandrs Belovs and Ansis Rosmanis. Adversary lower bounds for the collision and the set equality problems. *Quantum Information & Computation*, 18(3&4):198–222, 2018.
8 Aleksandrs Belovs and Ansis Rosmanis. Quantum lower bounds for tripartite versions of the hidden shift and the set equality problems. Full version. Avaiable at arXiv:1712.10194, 2018.
9 Aleksandrs Belovs and Robert Špalek. Adversary lower bound for the k-sum problem. In *Proc. of 4th ACM ITCS*, pages 323–328, 2013.
10 Shalev Ben-David. A super-Grover separation between randomized and quantum query complexities, 2015. `arXiv:1506.08106`.

11 Harry Buhrman and Ronald de Wolf. Complexity measures and decision tree complexity: a survey. *Theoretical Computer Science*, 288:21–43, 2002.
12 Charles W. Curtis and Irving Reiner. *Representation theory of finite groups and associative algebras.* AMS, 1962.
13 Mark Ettinger, Peter Høyer, and Emanuel Knill. The quantum query complexity of the hidden subgroup problem is polynomial. *Information Processing Letters*, 91(1):43–48, 2004.
14 Aram W. Harrow, Cedric Yen-Yu Lin, and Ashley Montanaro. Sequential measurements, disturbance and property testing. In *Proc. of 28th ACM-SIAM SODA*, pages 1598–1611, 2017.
15 Peter Høyer, Troy Lee, and Robert Špalek. Negative weights make adversaries stronger. In *Proc. of 39th ACM STOC*, pages 526–535, 2007.
16 Gordon James and Adalbert Kerber. *The Representation Theory of the Symmetric Group*, volume 16 of *Encyclopedia of Mathematics and its Applications.* Addison-Wesley, 1981.
17 Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM Journal on Computing*, 35:170–188, 2005.
18 Troy Lee, Rajat Mittal, Ben W. Reichardt, Robert Špalek, and Mario Szegedy. Quantum query complexity of state conversion. In *Proc. of 52nd IEEE FOCS*, pages 344–353, 2011.
19 Ashley Montanaro and Ronald de Wolf. A survey of quantum property testing. *Theory of Computing Graduate Surveys*, 7:1–81, 2016.
20 Ben W. Reichardt. Span programs and quantum query complexity: The general adversary bound is nearly tight for every Boolean function. In *Proc. of 50th IEEE FOCS*, pages 544–551, 2009.
21 Bruce E. Sagan. *The symmetric group: representations, combinatorial algorithms, and symmetric functions*, volume 203 of *Graduate Texts in Mathematics.* Springer, 2001.
22 Jean-Pierre Serre. *Linear Representations of Finite Groups*, volume 42 of *Graduate Texts in Mathematics.* Springer, 1977.
23 Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. In *Proc. of 43th IEEE FOCS*, pages 513–519, 2002.
24 Robert Špalek and Mario Szegedy. All quantum adversary methods are equivalent. *Theory of Computing*, 2:1–18, 2006.
25 Mark Zhandry. A note on the quantum collision and set equality problems. *Quantum Information & Computation*, 15(7&8):557–567, 2015.
26 Shengyu Zhang. On the power of Ambainis lower bounds. *Theoretical Computer Science*, 339(2):241–256, 2005.