Adversary lower bounds for the collision and the set equality problems

Belovs, Aleksandrs; Rosmanis, Ansis

2017

Belovs, A., & Rosmanis, A. (2017). Adversary lower bounds for the collision and the set equality problems. *Quantum Information and Computation*, 18(3-4).

<https://hdl.handle.net/10356/137663>

© 2017 Rinton Press. All rights reserved. This paper was published in *Quantum Information and Computation* and is made available with permission of Rinton Press.

Downloaded on 26 Apr 2023 22:01:46 SGT

Adversary Lower Bounds for the Collision and the Set Equality Problems

Aleksandrs Belovs*

Ansis Rosmanis†

Abstract

We prove tight $\Omega(n^{1/3})$ lower bounds on the quantum query complexity of the COLLISION and the SET EQUALITY problems, provided that the size of the alphabet is large enough. We do this using the negative-weight adversary method. Thus, we reprove the result by Aaronson and Shi, as well as a more recent development by Zhandry.

1 Introduction

In theory, quantum query complexity is surprisingly well-understood. For each function, the (negative-weight) adversary method [14] gives a tight characterisation of its quantum query complexity [20]. The adversary bound is a semi-definite optimisation problem (SDP). It performs nicely under composition [20], and it was used to prove a strong direct product theorem [18]. There is no similar theory known for randomised query complexity.

In practice, however, it is quite hard to construct a feasible solution to the adversary bound SDP that would be close to the optimal. In many cases, the positive-weight version of the bound is used. It is the original version of the bound developed by Ambainis [3]. This version uses combinatorial reasoning, and it is easy to apply. Unfortunately, the positive-weight adversary is subject to some severe constraints like the certificate complexity barrier [26, 29] and the property testing barrier [14]. The latter one, for instance, states that, if each positive and each negative input differ in at least ε fraction of the input variables, no positive-weight adversary can prove a lower bound better than $\Omega(1/\varepsilon)$.

The proper negative-weight adversary method usually requires spectral analysis of the adversary matrix. It does not have many applications so far. It was used to prove tight lower bounds for formulae evaluation [21]. Also, a recent line of development [10, 9] resulted in a relatively general construction of adversary bounds superseding the certificate complexity barrier: For any family $\mathcal{S}$ of $O(1)$ -sized subsets of $[n]$, a relatively simple optimisation problem gives a tight lower bound on the problem of detecting whether the input string $x \in [q]^n$ contains a subset of elements that belongs to $\mathcal{S}$ and whose sum is divisible by q . The upper bound emerges from the model of non-adaptive learning graphs [8]. And the dual to this model gives a lower bound. In particular, this gives nearly tight lower bounds for the k -SUM and the TRIANGLE-SUM problems.

For some other functions, lower bounds on quantum query complexity were obtained using the polynomial method [5]. This method is known to be suboptimal [4], but it is not subject to the

*Faculty of Computing, University of Latvia.

†SPMS, Nanyang Technological University and Centre for Quantum Technologies.

limitations of the positive-weight adversary. This method works well for functions with a lot of symmetries. A notable example of its application is for the COLLISION and the SET EQUALITY problems.

In the COLLISION problem, one has to decide whether the input function $x: [2n] \rightarrow [q]$ is one-to-one or two-to-one, provided that one of the two cases holds. The SET EQUALITY problem is a special case of COLLISION with an extra promise that the function x is one-to-one on both subdomains $[1..n]$ and $[n+1..2n]$. The COLLISION problem was defined in [11], and a quantum $O(n^{1/3})$ -query algorithm for the problem was given (which also works for SET EQUALITY). The SET EQUALITY problem was defined by Shi [24] and conjectured to be as hard as COLLISION. By the property testing barrier, no positive-weight adversary can give more than a trivial lower bound for any of these problems. Consequently, all known lower bounds for these problems were obtained using the polynomial method.

First, Aaronson and Shi [2] proved an $\Omega(n^{1/3})$ lower bound for the COLLISION problem. The proof was later simplified by Kutin [16]. Next, Midrijānis showed an $\Omega((n/\log n)^{1/5})$ lower bound for SET EQUALITY using a combination of the positive-weight adversary and the polynomial methods [19]. The lower bound for the SET EQUALITY problem was used by Aaronson and Ambainis in their proof of the polynomial relation between the randomised and the quantum query complexity of partial permutation-invariant functions [1]. Recently, Zhandry [28] proved a tight $\Omega(n^{1/3})$ lower bound for the SET EQUALITY problem using rather complicated machinery from [27] based on the polynomial method. This automatically strengthened the results in [1].

In this paper, we extend the range of applications of the adversary bound, and use it to prove tight lower bounds for the COLLISION and the SET EQUALITY problems given that the size of the input alphabet, q , is at least $\Omega(n^2)$. Thus, we resolve Shi's conjecture affirmatively, which was done independently from Zhandry's work.

There are several reasons why better understanding of the adversary method can be beneficial. First, this method is always tight. Second, known constructions of the adversary lower bounds follow through duality to the upper bounds. This simplifies the construction of the lower bound. The construction in this paper is also based on duality, as explained in the beginning of Section 5. Interestingly, the proofs for the COLLISION and the SET EQUALITY problems in this paper are almost identical, showing that the adversary method can be easier adapted for a specific function, as soon as a lower bound for a similar function is obtained. This is in contrast to the polynomial method, as more than ten years separated Shi's and Zhandry's results. Also, we note that it is the first, to our knowledge, application of the adversary method that supersedes the property testing barrier. We hope that some of our ideas will be useful in proving lower bounds via the adversary method for further problems.

The paper is organised as follows. In Section 2, we define the adversary method and basics of representation theory. In Section 3, we give a general treatment of the COLLISION and the SET EQUALITY problems, without actually defining the lower bound. In Section 4, we describe the dual learning graph for these problems, which inspires our adversary construction. In Section 5, we construct the lower bound. Proofs of two technical results from this section are given in Section 6.

As a final note about organisation of our paper, let us address a possible question of why we prove lower bounds for both the COLLISION and the SET EQUALITY problems. Indeed, since the latter is a special case of the former, a lower bound for the former is superfluous. Our justification is as follows. First, COLLISION is a better-known problem than SET EQUALITY. And then, the proof for SET EQUALITY is slightly more involved and contains all the necessary ingredients for the

COLLISION lower bound. Thus, a reader mainly interested in the COLLISION problem can attain the proof by merely ignoring all parts on SET EQUALITY.

2 Preliminaries

For integers m and ℓ , let $[m..\ell]$ denote the set $\{m, m+1, \dots, \ell\}$, and we often use shorthand $[m]$ for $[1..m]$. For sets $\mathcal{I}$ and $\mathcal{J}$, let $\mathcal{I}^{\mathcal{J}}$ denote the set of functions from $\mathcal{J}$ to $\mathcal{I}$. We use shorthand $\mathcal{I}^m$ for $\mathcal{I}^{[m]}$ and $\mathcal{I}^{m..\ell}$ for $\mathcal{I}^{[m..\ell]}$. We may also think of $\mathcal{I}^m$ as strings of length m .

For finite sets $\mathcal{I}$ and $\mathcal{J}$, an $\mathcal{I} \times \mathcal{J}$ matrix is a matrix whose rows are labelled by the elements of $\mathcal{I}$ and whose columns are labelled by the elements of $\mathcal{J}$. For such a matrix A , and $i \in \mathcal{I}$ and $j \in \mathcal{J}$, we denote by $A[i, j]$ its (i, j) -th entry. We use similar notation for vectors. We use A^* for the conjugate operator. Let $\circ$ denote the Hadamard (entry-wise) product of matrices. We use $\mathbb{I}$ and $\mathbb{J}$ to denote the identity and the all-ones matrices, respectively, to which, if necessary, we attach a subscript to indicate the space on which the corresponding operator acts. If $\{b_1, \dots, b_m\}$ is a basis of $\mathbb{C}^m$, we say that a vector $v = \alpha_1 b_1 + \dots + \alpha_m b_m$ uses a basis vector b_i if the corresponding coefficient α_i is non-zero.

For a linear operator $A: V \rightarrow U$, its *singular-value decomposition* is $A = \sum_i \sigma_i u_i v_i^*$, where σ_i are positive real numbers, $\{v_i\}$ form an orthonormal system in V and $\{u_i\}$ in U . The names for σ_i , u_i , and v_i are *singular value*, *left-singular vector*, and *right-singular vector*, respectively. The span of the left-singular vectors is called the *image* and the span of the right-singular vectors the *coimage* of A . The *spectral norm*, $\|A\|$, is the largest singular value. Singular vectors corresponding to the largest singular value are called *principal*. Finally, by principal singular vectors of A on a subspace $\mathcal{H} \subseteq U$ we understand principal singular vectors of the operator $A\Pi$, where Π is the orthogonal projector on $\mathcal{H}$.

2.1 Adversary Method

We identify decision problems with Boolean-valued functions $f: \mathcal{D} \rightarrow \{0, 1\}$ with domain $\mathcal{D} \subseteq [q]^m$. We call the inputs in $f^{-1}(1)$ and $f^{-1}(0)$ *positive* and *negative*, respectively. We are interested in quantum query complexity of decision problems. For the definitions and basic properties refer to [12]. In the paper, we only require the knowledge of the (negative-weight) adversary bound that we are about to define.

Definition 1. An *adversary matrix*¹ for a decision problem f is a non-zero $f^{-1}(1) \times f^{-1}(0)$ matrix Γ . For any $i \in [m]$, the $f^{-1}(1) \times f^{-1}(0)$ matrix Δ_i is defined by

$$\Delta_i[x, y] = \begin{cases} 0, & x_i = y_i; \\ 1, & x_i \neq y_i. \end{cases}$$

Theorem 2 (Adversary bound [14, 17, 10]). *In notation of Definition 1, the quantum query complexity of the decision problem f is $\Theta(\text{ADV}^\pm(f))$, where $\text{ADV}^\pm(f)$ is the optimal value of the semi-definite program*

$$\text{maximise} \quad \|\Gamma\| \tag{1a}$$

$$\text{subject to} \quad \|\Delta_i \circ \Gamma\| \leq 1 \quad \text{for all } i \in [m], \tag{1b}$$

¹Compared to the general definition of an adversary matrix [14], we consider only a quarter of the matrix, as this quarter completely specifies the whole matrix (see [10] for details).

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

Note that we can choose any adversary matrix Γ and scale it so that the condition $\max_i \|\Delta_i \circ \Gamma\| \leq 1$ holds. Consequently, we often use a relaxed condition $\|\Delta_i \circ \Gamma\| = O(1)$ instead of $\|\Delta_i \circ \Gamma\| \leq 1$.

Precise evaluation of $\|\Delta_i \circ \Gamma\|$ may be hard, but we can upper bound $\|\Delta_i \circ \Gamma\|$ using the following trick first introduced in [6] and later used in [10, 9, 25]. If A is a matrix of the same dimensions as Δ_i , we use the notation $\Delta_i \diamond A$ for a matrix B satisfying $\Delta_i \circ B = \Delta_i \circ A$. Or, we write $A \xrightarrow{\Delta_i} B$. Note that B is not uniquely defined, and our task will be to choose one that fits our needs. Now, from the fact that

$$\gamma_2(\Delta_i) = \max_B \{ \|\Delta_i \circ B\| : \|B\| \leq 1 \} \leq 2$$

(see [17] for this and other facts about the γ_2 norm), it follows that

$$\|\Delta_i \circ A\| = \|\Delta_i \circ (\Delta_i \diamond A)\| \leq 2\|\Delta_i \diamond A\|.$$

Note that we can always choose $\Delta_i \diamond A = A$ and

$$\Delta_i \diamond (\alpha' A' + \alpha'' A'') = \alpha'(\Delta_i \diamond A') + \alpha''(\Delta_i \diamond A'').$$

In order to show that $\|\Delta_i \circ \Gamma\| = O(1)$, it suffices to show that $\|\Delta_i \diamond \Gamma\| = O(1)$ for some $\Delta_i \diamond \Gamma$. That is, it suffices to show that there is a way to modify the entries $\Gamma[x, y]$ of Γ with $x_i = y_i$ so that the spectral norm of the resulting matrix is bounded by a constant.

2.2 Representation Theory

In this section, we introduce basic notions from representation theory of finite groups with special emphasis on the symmetric group. For more background, the reader may refer to [13, 23] for general theory, and to [15, 22] 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. The module V is equipped with an inner product invariant under the multiplication by the elements of G . A G -*morphism* (or just morphism, if G is clear from the context) between two G -modules V and W is a linear operator $\theta: V \rightarrow W$ satisfying $\theta(uv) = u\theta(v)$ for all $u \in \mathbb{C}G$ and $v \in V$.

A G -module is called *irreducible* (or just irrep for irreducible representation) if it does not contain a non-trivial G -submodule. An essential basic result in representation theory is the following

Lemma 3 (Schur's Lemma). *Assume $\theta: V \rightarrow W$ 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.*

Copies of non-isomorphic irreps in a fixed G -module V are orthogonal. 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 . 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.

Symmetric group. Let $\mathbb{S}_L$ 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, $\mathbb{S}_m$ denotes the isomorphism class of the symmetric groups $\mathbb{S}_L$ with $|L| = m$. Representation theory of $\mathbb{S}_m$ is closely related to *Young diagrams*, defined as follows.

A *partition* λ of an integer m is a non-increasing sequence $(\lambda_1, \dots, \lambda_k)$ of positive integers satisfying $\lambda_1 + \dots + \lambda_k = m$. We denote this by $\lambda \vdash m$, or write $m = |\lambda|$. A partition $\lambda = (\lambda_1, \dots, \lambda_k)$ is often represented in the form of a Young diagram that consists, from top to bottom, of rows of $\lambda_1, \lambda_2, \dots, \lambda_k$ boxes aligned by the left side. For a partition $\lambda = (\lambda_1, \dots, \lambda_k)$ of m and an integer $\ell \geq \lambda_1$, by (ℓ, λ) we denote the partition $(\ell, \lambda_1, \dots, \lambda_k)$ of $m + \ell$.

For each partition $\lambda \vdash m$, we assign an irreducible $\mathbb{S}_m$ -module S^λ , called the *Specht module*. All these modules are pairwise non-isomorphic, and give a complete list of all the irreps of $\mathbb{S}_m$. We describe these modules following the classical approach by Young (see [15, Chapter 3] or [13, §28] for more detail).

From now on, we assume that aforementioned $L = [m]$ for concreteness. Assume $\lambda \vdash m$. A *Young tableau* of shape λ is a Young diagram of λ with each box containing an integer from $[m]$, each integer used exactly once. We use $\text{sh}(t)$ for the shape of t . As an example, the following is a Young tableau of shape $(3, 2)$:

4	5	1
2	3	

For any Young tableau t , we define the group of its *row permutations* R_t and the group of its *column permutations* C_t to consist of the permutations in $\mathbb{S}_m$ that permute the elements within each row or column of t , respectively. In our example above, $R_t = \mathbb{S}_{\{1,4,5\}} \times \mathbb{S}_{\{2,3\}}$ and $C_t = \mathbb{S}_{\{2,4\}} \times \mathbb{S}_{\{3,5\}}$. In the representation theory of the symmetric group, the elements

$$R_t^+ = \sum_{\pi \in R_t} \pi, \quad C_t^- = \sum_{\rho \in C_t} (\text{sgn } \rho) \rho, \quad \text{and} \quad E_t = C_t^- R_t^+, \quad (2)$$

of the group algebra $\mathbb{CS}_m$ are widely used. (Here $\text{sgn } \rho$ stands for the sign of the permutation ρ .) Suitably scaled multiples $\tilde{R}_t^+$, $\tilde{C}_t^-$, and $\tilde{E}_t$ of these operators are non-zero idempotents: $(\tilde{R}_t^+)^2 = \tilde{R}_t^+$, $(\tilde{C}_t^-)^2 = \tilde{C}_t^-$, and $\tilde{E}_t^2 = \tilde{E}_t$. Note that, while $\tilde{R}_t^+$ and $\tilde{C}_t^-$ are orthogonal projectors, $\tilde{E}_t$, in general, is *not* an orthogonal projector.

For any Young tableau t , the left ideal $\mathbb{CS}_m E_t$ in the group algebra $\mathbb{CS}_m$ is an irrep of $\mathbb{S}_m$. The isomorphism class of $\mathbb{CS}_m E_t$ only depends on the shape of t , and we denote by S^λ an arbitrary representative from this class, where $\lambda = \text{sh}(t)$. There exists a non-zero vector $v \in S^\lambda$ such that $\tilde{E}_t v = v$ (in $\mathbb{CS}_m E_t$ one can take $v = E_t$). On the other hand, if $\mu \neq \lambda$ is another partition of m , then $\tilde{E}_t$ annihilates S^μ : $\tilde{E}_t v = 0$ for all $v \in S^\mu$.

It is not hard to see that $\pi R_t^+ = R_{\pi t}^+ \pi$, where πt is the Young tableau t with entry $i \in [m]$ replaced by $\pi(i)$. As a consequence, the module S^λ is spanned by the images of R_t^+ as t runs through all the Young tableau of shape λ : $S^\lambda = \text{span}\{R_t^+(S^\lambda) \mid \text{sh}(t) = \lambda\}$.

3 Analysis of the Problem

Consider the following three types of input strings $x = (x_1, \dots, x_{2n})$ in $[q]^{2n}$:

- (a) For each $a \in [q]$, there are either exactly two or none $i \in [2n]$ such that $x_i = a$.

- (b) For each $a \in [q]$, either there is no x_i equal to a , or there is unique $i \in [1..n]$ and unique $j \in [n + 1..2n]$ satisfying $x_i = x_j = a$.
- (c) For each $a \in [q]$, there is at most one $i \in [2n]$ such that $x_i = a$.

In the COLLISION problem, given $x \in [q]^{2n}$ satisfying (a) or (c), the task is to distinguish these two cases. We say that (a) forms the set of positive inputs, and (c) is the set of negative inputs. In the SET EQUALITY problem, the task is to distinguish (b) and (c) in a similar manner. Since any string in (b) also satisfies the requirements in (a), SET EQUALITY is a simpler problem than COLLISION. In the following, we will use subscripts CP and SE to denote relation to COLLISION and SET EQUALITY, respectively. To avoid unnecessary repetitions, we use notation Q that may refer to both CP and SE.

Positive inputs in (a) and (b) naturally give rise to the corresponding *matchings*. A matching μ on $[2n]$ is a decomposition

$$[2n] = \{\mu_{1,1}, \mu_{1,2}\} \cup \{\mu_{2,1}, \mu_{2,2}\} \cup \cdots \cup \{\mu_n, \mu_{n,2}\}$$

of the set $[2n]$ into n mutually disjoint pairs of elements. For concreteness, we will assume that $\mu_{i,j}$ s are sorted: $\mu_{i,1} < \mu_{i,2}$ for all $i \in [n]$ and $\mu_{1,1} < \mu_{2,1} < \cdots < \mu_n, \mu_{n,1}$. In particular, $\mu_{1,1} = 1$. Clearly, this assumption is without loss of generality. Let M_{CP} denote the set of all matchings on $[2n]$, and let M_{SE} denote the set of matchings μ on $[2n]$ such that $1 \leq \mu_{i,1} \leq n$ and $n + 1 \leq \mu_{i,2} \leq 2n$ for all i .

Embedding. Our aim is to construct adversary matrices Γ_{CP} and Γ_{SE} for the COLLISION and the SET EQUALITY problems, respectively. As described in Definition 1, the rows of the adversary matrices are labelled by the positive inputs and the columns by the negative inputs. We will use the trick initially used in [6], and embed the adversary matrix Γ_Q into a larger $|M_Q|q^n \times q^{2n}$ matrix $\tilde{\Gamma}_Q$. The columns of $\tilde{\Gamma}_Q$ are labelled by all possible inputs in $[q]^{2n}$. The rows of $\tilde{\Gamma}_Q$ are split into blocks corresponding to the matchings in M_Q . Inside the block corresponding to a matching $\mu \in M_Q$, there are all possible row labels $x \in [q]^{2n}$ such that $x[\mu_{i,1}] = x[\mu_{i,2}]$ for all i . We will label the rows by specifying both the label and the block, i.e., as (x, μ) .

A *legal* label of a row or a column of $\tilde{\Gamma}_Q$ is one that also features in Γ_Q . Besides legal labels, $\tilde{\Gamma}_Q$ also contains *illegal* labels. A column is illegal if its label contains two equal elements. A row corresponding to a matching μ is illegal if $x[\mu_{i,1}] = x[\mu_{j,1}]$ for some $i \neq j$. We obtain Γ_Q from $\tilde{\Gamma}_Q$ by removing all the illegal rows and columns.

This embedding is used because it is easier to work with $\tilde{\Gamma}_Q$ than with Γ_Q . We clearly have

$$\|\Delta_i \circ \Gamma_Q\| \leq \|\Delta_i \circ \tilde{\Gamma}_Q\| \tag{3}$$

because $\Delta_i \circ \Gamma_Q$ is a submatrix of $\Delta_i \circ \tilde{\Gamma}_Q$. If we could show that $\|\Gamma_Q\|$ is not much smaller than $\|\tilde{\Gamma}_Q\|$, that would allow us to use $\tilde{\Gamma}_Q$ instead of Γ_Q in Theorem 2. This is not true for every choice of $\tilde{\Gamma}_Q$, however. For instance, if $\tilde{\Gamma}_Q$ contains non-zero entries only in illegal rows or columns, then $\Gamma_Q = 0$. But for our specific choice of $\tilde{\Gamma}_Q$ this will be the case, as shown in Section 5.3. The condition $q = \Omega(n^2)$ is essential for the proof.

Space $\mathcal{H}$ and its subdivision. Let $\mathcal{H} = \mathbb{C}^q$, and let $\mathcal{M}_Q = \mathbb{C}^{M_Q}$ be the complex vector space with the elements of M_Q forming an orthonormal basis. Then, $\tilde{\Gamma}_Q$ can be considered as an operator from $\mathcal{H}^{\otimes 2n}$ to $\mathcal{M}_Q \otimes \mathcal{H}^{\otimes n}$ if we identify a basis element $(\mu, z) \in M_Q \times [q]^n$ with the row label x in the μ -block of $\tilde{\Gamma}_Q$ defined by $x[\mu_{i,a}] = z_i$.

Now we are going to define two bases of $\mathcal{H}$. The *standard* basis of $\mathcal{H}$ is the one used in Definition 1, in which the basis vectors correspond to the symbols of the input alphabet. The *e-basis* is an orthonormal basis $e_0, e_1, \dots, e_{q-1}$ satisfying $e_0[j] = 1/\sqrt{q}$ for all $j \in [q]$. The precise choice of the remaining basis vectors is irrelevant. Further in the text, we almost exclusively work in the *e-basis*. Let

$$\mathcal{H}_0 = \text{span}\{e_0\} \quad \text{and} \quad \mathcal{H}_1 = e_0^\perp = \text{span}\{e_1, \dots, e_{q-1}\}.$$

Let us agree on a notational convention that Π with arbitrary sub- and superscripts denotes the orthogonal projector onto the space denoted by $\mathcal{H}$ with the same sub- and superscripts. Thus, for instance, $\Pi_0 = e_0 e_0^* = \mathbb{J}_{\mathcal{H}}/q$ and $\Pi_1 = \mathbb{I}_{\mathcal{H}} - \mathbb{J}_{\mathcal{H}}/q$ are the projectors onto e_0 and its orthogonal complement, respectively. An important relation is

$$\Pi_1 \xrightarrow{\Delta} -\Pi_0, \tag{4}$$

where Δ acts on the only variable.

Similarly, for the space $\mathcal{H}^{\otimes m}$, the *e-basis* consists of all possible tensor products of the vectors in $\{e_i\}$ of length m . Vector e_0 in the tensor product is called the *zero component*. The *weight* of the basis vector is the number of non-zero components in the product. The space $\mathcal{H}^{\otimes m}$ can be decomposed as $\mathcal{H}^{\otimes m} = \bigoplus_{k=0}^m \mathcal{H}_k^{(m)}$, where

$$\mathcal{H}_k^{(m)} = \bigoplus_{c \in \{0,1\}^m, |c|=k} \mathcal{H}_{c_1} \otimes \dots \otimes \mathcal{H}_{c_m}$$

is the space spanned by all the basis elements of weight k .

Symmetry. Let $\mathbb{S}_{\text{CP}} = \mathbb{S}_{[2n]}$ be the symmetric group on $2n$ elements, and let $\mathbb{S}_{\text{SE}}$ be its subgroup $\mathbb{S}_{[1..n]} \times \mathbb{S}_{[n+1..2n]}$. The COLLISION and the SET EQUALITY problems are invariant under the permutations of the input variables in $\mathbb{S}_{\text{CP}}$ and $\mathbb{S}_{\text{SE}}$, respectively. Hence, we may assume that Γ_Q is also invariant under these permutations [14]. We extend this symmetry to $\tilde{\Gamma}_Q$ by requiring that, for each $\pi \in \mathbb{S}_Q$, and labels (x, μ) and y , we have

$$\tilde{\Gamma}_Q[(x, \mu), y] = \tilde{\Gamma}_Q[(\pi x, \pi \mu), \pi y], \tag{5}$$

where $(\pi x)_i = x_{\pi^{-1}(i)}$ and $\pi \mu = \{\{\pi(\mu_{1,1}), \pi(\mu_{1,2})\}, \dots, \{\pi(\mu_{n,1}), \pi(\mu_{n,2})\}\}$. Note that $\mathcal{H}_k^{(m)}$ is invariant under all the permutations in $\mathbb{S}_{[m]}$. Because of this symmetry, we may use the representation theory in the construction of $\tilde{\Gamma}_Q$.

4 Dual Learning Graph Perspective

Our lower bounds for the COLLISION and the SET EQUALITY problems are intrinsically based on the dual learning graph for these problems developed in [9]. In this section, we explain the dual learning graph and how it relates to the adversary bound.

The concept of learning graphs is based on *certificate structures*. Informally, they describe all possible dispositions of certificates in positive inputs. For COLLISION and SET EQUALITY, the corresponding certificate structure can be identified with the set of matchings M_Q . For each positive input x , there exists a matching $\mu \in M_Q$ that pairs up equal elements. That is, $x_a = x_b$ for all

$\{a, b\} \in \mu$. A subset $S \subseteq [2n]$ is a 1-certificate for x if and only if it contains a pair from μ as a subset (that is, $2^S \cap \mu \neq \emptyset$). Let us denote the latter relationship by $S \sim \mu$.

The (dual) learning graph complexity of the certificate structure is the optimal value of the following optimisation problem (the formulation is tailored to the case of COLLISION and SET EQUALITY):

$$\text{maximise} \quad \sqrt{\sum_{\mu \in M_Q} \alpha_\mu(\emptyset)^2} \quad (6a)$$

$$\text{subject to} \quad \sum_{\mu \in M_Q} (\alpha_\mu(S) - \alpha_\mu(S \cup \{j\}))^2 \leq 1 \quad \text{for all } S \subseteq [2n] \text{ and } j \in [2n] \setminus S; \quad (6b)$$

$$\alpha_\mu(S) = 0 \quad \text{whenever } S \sim \mu; \quad (6c)$$

$$\alpha_\mu(S) \in \mathbb{R} \quad \text{for all } S \subseteq [2n] \text{ and } \mu \in M_Q. \quad (6d)$$

A *dual learning graph* is any feasible solution to this optimisation problem. Let us note that one can also define the primal learning graph complexity, which is equal to the dual one, but we do not use it in this paper.

In some sense, learning graphs precisely capture quantum query complexity of certificate structures. First, for any decision problem with a given certificate structure, there exists a quantum algorithm for this problem with query complexity equal (up to a constant factor) to the learning graph complexity of the certificate structure. Second, for any certificate structure, there exists a decision problem possessing this certificate structure and whose quantum query complexity is equal (up to a constant factor) to the learning graph complexity of the certificate structure.

For our problems, we have the following

Proposition 4. *The dual learning graph complexity of the COLLISION and the SET EQUALITY problems is $\Omega(n^{1/3})$.*

Proof. Define the following potential solution to (6):

$$\alpha_\mu(S) = \begin{cases} \frac{1}{\sqrt{|M_Q|}} \max\{n^{1/3} - |S|, 0\}, & \text{if } S \not\sim \mu; \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

It is easy to see that the objective value is $n^{1/3}$. Let us prove its feasibility (up to a constant factor). Fix $S \subseteq [2n]$ and $j \notin S$. If $|S| \geq n^{1/3}$, the left-hand side of (6b) is 0, so let us further assume $|S| < n^{1/3}$. For each μ , the difference $\alpha_\mu(S) - \alpha_\mu(S \cup \{j\})$ can take the following values:

- If $S \sim \mu$, then $\alpha_\mu(S) = \alpha_\mu(S \cup \{j\}) = 0$, and the difference is 0.
- If $S \cup \{j\} \not\sim \mu$, then the difference is $1/\sqrt{|M_Q|}$.
- Finally, if $S \not\sim \mu$ and $S \cup \{j\} \sim \mu$, then the difference can be as large as $n^{1/3}/\sqrt{|M_Q|}$. However, this only happens if j is matched by μ to an element already in S , and at most $|S|/n \leq n^{-2/3}$ fraction of all matchings μ satisfy this condition.

Thus, the value of the left-hand side of (6b) is at most

$$|M_Q| \cdot \frac{1}{|M_Q|} + n^{-2/3} \cdot |M_Q| \cdot \frac{n^{2/3}}{|M_Q|} = O(1).$$

By scaling down the solution by a constant factor, we obtain a feasible solution with objective value $\Omega(n^{1/3})$. $\square$

Now we would like to convert this dual learning graph into an adversary lower bound. In order to get some intuition, let us first consider the following $(M_Q \times [q]^{2n}) \times [q]^{2n}$ matrices. Their rows are labelled by pairs (μ, x) , where $\mu \in M_Q$ and x is an *arbitrary* string in $[q]^{2n}$. And their columns are labelled by strings in $[q]^{2n}$. We treat such matrices as a $M_Q \times 1$ block matrices, where each block is an $[q]^{2n} \times [q]^{2n}$ matrix.

Let $\hat{\Gamma}$ be one such matrix whose μ -th block is given by $\sum_S \alpha_\mu(S) \Pi_S$. Here Π_S is defined as $\bigotimes_{j \in [2n]} \Pi_{s_j}$, where $s_j = 1$ if $j \in S$ and $s_j = 0$ otherwise, and Π_0 and Π_1 are defined right before relation (4). Using (4), we can define

$$\Delta_j \diamond \Pi_S = \begin{cases} -\Pi_{S \setminus \{j\}}, & \text{if } j \in S; \\ \Pi_S, & \text{otherwise.} \end{cases}$$

Since Π_S are pairwise orthogonal projectors, we have

$$\|\hat{\Gamma}\|^2 = \max_S \sum_{\mu \in M_Q} \alpha_\mu(S)^2 \quad \text{and} \quad \|\Delta_j \diamond \hat{\Gamma}\|^2 = \max_{S \not\ni j} \sum_{\mu \in M_Q} (\alpha_\mu(S) - \alpha_\mu(S \cup \{j\}))^2.$$

This is the intuition behind the expressions in (6a) and (6b).

Now we would like to switch from $\hat{\Gamma}$ to $\tilde{\Gamma}$. The simplest solution, adapted in [9], is to restrict $\hat{\Gamma}$ to the rows used in $\tilde{\Gamma}$ and scale each block up, so that the norm of the thus transformed matrix $\Pi_\emptyset$ is still 1 in each block. However, the norm of $\hat{\Gamma}$ (and consequently, $\Delta_j \diamond \hat{\Gamma}$) may grow after this transformation. For example, consider a simple case of $n = 1$ and two matrices

$$\Pi_\emptyset = \Pi_0 \otimes \Pi_0 \quad \text{and} \quad \Pi_{\{1\}} + \Pi_{\{2\}} = \Pi_1 \otimes \Pi_0 + \Pi_0 \otimes \Pi_1.$$

The norms of both matrices are 1. After the transformation, these matrices become

$$\Psi_0 = \Pi_0 \otimes e_0^* = e_0^* \otimes \Pi_0 \quad \text{and} \quad \Psi_1 = \Pi_1 \otimes e_0^* + e_0^* \otimes \Pi_1. \quad (8)$$

The norm of the first matrix is still 1, whereas the norm of the second is $\sqrt{2}$. This growth in norm occurs essentially because two initially orthogonal left-singular vectors $e_i \otimes e_0$ and $e_0 \otimes e_i$ collapse into one vector e_i , for any e_i orthogonal to e_0 .

Ideally, we would like the norm of $\Delta_j \diamond \hat{\Gamma}$ not to grow more than by a constant factor. The condition (6c) is essential for this. In the lower bound constructions of [9], the decision problems possessing a given certificate structure were chosen in such a way that the norm did not grow much for any choice of $\alpha_\mu(S)$ satisfying (6c). However, here we do not have a luxury of choosing the problems: we have to consider COLLISION and SET EQUALITY. We will see in Section 5.1 that the same simple restriction does not work for these problems. In order to control the growth of the norm, we will have to switch to different operators, which we describe in Section 5.2.

5 Construction of the Adversary Matrix

The matrix $\tilde{\Gamma}_Q$ is constructed as a linear combination

$$\tilde{\Gamma}_Q = \sum_k \alpha_k \bar{W}_{Q,k}, \quad (9)$$

where, for each k , $\bar{W}_{Q,k}$ is an operator from $\mathcal{H}_k^{(2n)}$ to $\mathcal{M}_Q \otimes \mathcal{H}_k^{(n)}$. The coefficients α_k are given by $\alpha_k = \max\{0, n^{1/3} - k\}$. We again assume that $\bar{W}_{Q,k}$ are invariant under the action of S_Q (in the sense of (5)).

In terms of Section 4, one should think of $\bar{W}_{Q,k}$ as corresponding to a combination of matrices Π_S with $|S| = k$. As mentioned in the previous section, the first intention is to use the techniques of [9] to construct the constituent matrices in (9). Unfortunately, this does not work, as we show in Section 5.1. Luckily, it is possible to modify the construction of the matrices so that (9) gives an optimal adversary matrix. We describe this in Section 5.2. Finally, in Section 5.3, we show how to transform $\tilde{\Gamma}_Q$ into a valid adversary matrix Γ_Q .

5.1 First Attempt

In this section, we define matrices $W_{Q,k}$ that may seem as the most natural choice for the decomposition (9). Unfortunately, they do not work well enough, so we will have to modify the construction in Section 5.2.

As described in Section 3, the matrix $\tilde{\Gamma}_Q$ can be decomposed into blocks corresponding to different matchings $\mu \in M_Q$. We first define one block of the matrix. Recall the operators $\Psi_0, \Psi_1: \mathcal{H}^{\otimes 2} \rightarrow \mathcal{H}$ defined in (8). For every $k \in [0..n]$ and every $\mu \in M_Q$, define the operator $W_k^\mu: \mathcal{H}^{\otimes 2n} \rightarrow \mathcal{H}^{\otimes n}$ by

$$W_k^\mu = \sum_{c \in \{0,1\}^n, |c|=k} \Psi_{c_1} \otimes \dots \otimes \Psi_{c_n}, \quad (10)$$

where, for $i \in [n]$, Ψ_{c_i} maps the $\mu_{i,1}$ -th and the $\mu_{i,2}$ -th multiplier in $\mathcal{H}^{\otimes 2n}$ to the i -th multiplier in $\mathcal{H}^{\otimes n}$. Note that the image of W_k^μ is contained in $\mathcal{H}_k^{(n)}$ and its coimage in $\mathcal{H}_k^{(2n)}$, that is, $W_k^\mu = \Pi_k^{(n)} W_k^\mu \Pi_k^{(2n)}$. The block of the matrix $W_{Q,k}$ corresponding to $\mu \in M_Q$ is defined by $\frac{1}{\sqrt{|M_Q|}} W_k^\mu$.

We have $W_{Q,k} = W_{Q,k} \Pi_k^{(2n)}$.

Now, if we define $\tilde{\Gamma}_Q$ as in (9) with $\bar{W}_{Q,k} = W_{Q,k}$, we obtain the same matrix we would have obtained using the construction in Section 4 with $\alpha_\mu(S)$ given by (7). One can see that $W_{Q,k}$ thus constructed satisfy the symmetry (5). Because of this, $\|\Delta_i \circ \tilde{\Gamma}_Q\|$ is the same for all $i \in [2n]$. Therefore, it suffices to estimate $\|\Delta_1 \circ \tilde{\Gamma}_Q\|$. For that, we use the following simple decomposition

$$\Pi_k^{(2n)} = \Pi_0 \otimes \Pi_k^{(2n-1)} + \Pi_1 \otimes \Pi_{k-1}^{(2n-1)}. \quad (11)$$

It is not hard to check that the μ -th block of $W_{Q,k}(\Pi_0 \otimes \Pi_k^{(2n-1)})$ is given by $\frac{1}{\sqrt{|M_Q|}}(X_k^\mu + Y_k^\mu)$ and the μ -th block of $W_{Q,k}(\Pi_1 \otimes \Pi_{k-1}^{(2n-1)})$ is given by $\frac{1}{\sqrt{|M_Q|}}Z_k^\mu$, where (with the same order of multipliers as in (10)):

$$\begin{aligned} X_k^\mu &= \Psi_0 \otimes \sum_{c \in \{0,1\}^{2..n}, |c|=k} \Psi_{c_2} \otimes \dots \otimes \Psi_{c_n}, \\ Y_k^\mu &= (e_0^* \otimes \Pi_1) \otimes \sum_{c \in \{0,1\}^{2..n}, |c|=k-1} \Psi_{c_2} \otimes \dots \otimes \Psi_{c_n}, \\ Z_k^\mu &= (\Pi_1 \otimes e_0^*) \otimes \sum_{c \in \{0,1\}^{2..n}, |c|=k-1} \Psi_{c_2} \otimes \dots \otimes \Psi_{c_n}. \end{aligned} \quad (12)$$

Here we used that $\mu(1,1) = 1$, and also (8). Thus, if we define $X_{Q,k}$, $Y_{Q,k}$, and $Z_{Q,k}$ similarly to $W_{Q,k}$, we get the following decomposition:

$$W_{Q,k} = X_{Q,k} + Y_{Q,k} + Z_{Q,k}. \quad (13)$$

Again, one can see that $X_{Q,k}$, $Y_{Q,k}$, and $Z_{Q,k}$ are symmetric under the action of $\mathbb{S}'_Q$, where $\mathbb{S}'_{\text{CP}} = \mathbb{S}_{[2..2n]}$ and $\mathbb{S}'_{\text{SE}} = \mathbb{S}_{[2..n]} \times \mathbb{S}_{[n+1..2n]}$. Using (4), it is reasonable to define $\Delta_1 \diamond X_{Q,k} = X_{Q,k}$, $\Delta_1 \diamond Y_{Q,k} = Y_{Q,k}$, and $\Delta_1 \diamond Z_{Q,k} = -X_{Q,k-1}$, so that

$$\tilde{\Gamma}_Q \xrightarrow{\Delta_1} \sum_k (\alpha_{k-1} - \alpha_k) X_{Q,k-1} + \sum_k \alpha_k Y_{Q,k}. \quad (14)$$

This construction restates that of Section 4, where the first and the second terms of this relation correspond to the second and the third bullets in the proof of Proposition 4, respectively. In order to maintain a meaningful lower bound, we would need the norms of $W_{Q,k}$ and $X_{Q,k}$ to be $O(1)$, and the norm of $Y_{Q,k}$ to be $O(\sqrt{k/n})$. However, in reality, it is not hard to show that

$$\|W_{Q,k}\| = \Theta(2^{k/2}), \quad \|X_{Q,k}\| = \Theta(2^{k/2}), \quad \text{and} \quad \|Y_{Q,k}\| = \Theta(2^{k/2} \sqrt{k/n}). \quad (15)$$

This growth by the factor of $2^{k/2}$ can be interpreted as the $\sqrt{2}$ growth in (8) taken to the k -th power. And this construction fails to give anything better than the trivial lower bound.

There is an explanation for this failure. The construction above only used that the learning graph complexity of the COLLISION problem is $\Omega(n^{1/3})$. On the other hand, the learning graph complexity of the hidden shift problem is also $\Omega(n^{1/3})$ [9, Proposition 12]. Thus, if the current construction had worked, we would also have proven an $\Omega(n^{1/3})$ lower bound for the hidden shift problem contradicting the fact that the query complexity of this problem is logarithmic. Thus, in order to obtain an optimal solution, we have to use again the structure of the problem.

5.2 Successful Construction

Our aim is to get rid of the $2^{k/2}$ factor in (15) while preserving an analogue of (14). Recall that $W_k^\mu = W_k^\mu \Pi_k^{(2n)}$. As we will show in Corollary 9, for m a positive integer, $\mathcal{H}_k^{(m)}$, as an $\mathbb{S}_m$ -module, uses irreps whose Young diagrams have at most k boxes below the first row. Let us define $\bar{\mathcal{H}}_k^{(m)}$ as the subspace of $\mathcal{H}_k^{(m)}$ spanned by the irreps of $\mathbb{S}_m$ having *exactly* k boxes below the first row, i.e., of the form $(m-k, \lambda)$ with $\lambda \vdash k$. We restrict each W_k^μ to this subspace, or, more precisely, we define

$$\bar{W}_{Q,k} = W_{Q,k} \bar{\Pi}_{Q,k}, \quad (16)$$

where $\bar{\Pi}_{Q,k}$ is the orthogonal projector on one of the following subspaces:

$$\bar{\mathcal{H}}_{\text{CP},k} = \bar{\mathcal{H}}_k^{(2n)} \quad \text{or} \quad \bar{\mathcal{H}}_{\text{SE},k} = \sum_{\ell=0}^k \bar{\mathcal{H}}_{k-\ell}^{(n)} \otimes \bar{\mathcal{H}}_\ell^{(n)}. \quad (17)$$

Here, for $\bar{\mathcal{H}}_{\text{SE},k}$, the first and the second multiplier reside in the first n and the second n copies of $\mathcal{H}$ in $\mathcal{H}^{\otimes 2n}$, respectively.

While applying Δ_1 in Section 5.1, we used (11), which effected to

$$\Pi_k^{(m)} = \Pi_0 \otimes \Pi_k^{(m-1)} + \Pi_1 \otimes \Pi_{k-1}^{(m-1)}.$$

Now we would like to have a similar decomposition for $\bar{\Pi}_k^{(m)}$. Unfortunately, this time there is a non-zero error term

$$\Phi_k^{(m)} = \bar{\Pi}_k^{(m)} - \Pi_0 \otimes \bar{\Pi}_k^{(m-1)} - \Pi_1 \otimes \bar{\Pi}_{k-1}^{(m-1)}. \quad (18)$$

Note that $\Phi_k^{(m)}$ is a normal operator with image (and coimage) contained in $\mathcal{H}_k^{(m)}$ and it is symmetric with respect to $\mathbb{S}'_{m-1} = \mathbb{S}_{[2..m]}$. Luckily, in Section 6.1, we will be able to bound it as follows:

Lemma 5. *If $k < m/3$, then $\|\Phi_k^{(m)}\| = O(1/\sqrt{m})$. Moreover, the image of $\Phi_k^{(m)}$ only uses irreps of $\mathbb{S}'_{m-1}$ with exactly $k-1$ boxes below the first row.*

With $\Phi_k^{(m)}$ as in (18), let us define

$$\begin{aligned}\bar{\Pi}'_{\text{CP},k} &= \bar{\Pi}_k^{(2n-1)}, & \Phi_{\text{CP},k} &= \Phi_k^{(2n)}, \\ \bar{\Pi}'_{\text{SE},k} &= \sum_{\ell=0}^k (\bar{\Pi}_{k-\ell}^{(n-1)} \otimes \bar{\Pi}_\ell^{(n)}), & \Phi_{\text{SE},k} &= \sum_{\ell=0}^{k-1} (\Phi_{k-\ell}^{(n)} \otimes \bar{\Pi}_\ell^{(n)}).\end{aligned}\tag{19}$$

Note that $\Phi_{Q,k}$ acts on $\mathcal{H}^{\otimes 2n}$ while $\bar{\Pi}'_{Q,k}$ acts on $\mathcal{H}^{\otimes (2n-1)}$. From (18), we have

$$\bar{\Pi}_{Q,k} = \Pi_0 \otimes \bar{\Pi}'_{Q,k} + \Pi_1 \otimes \bar{\Pi}'_{Q,k-1} + \Phi_{Q,k}.\tag{20}$$

With $X_{Q,k}$, $Y_{Q,k}$ and $Z_{Q,k}$ as in Section 5.1, let

$$\bar{X}_{Q,k} = X_{Q,k}(\Pi_0 \otimes \bar{\Pi}'_{Q,k}), \quad \bar{Y}_{Q,k} = Y_{Q,k}(\Pi_0 \otimes \bar{\Pi}'_{Q,k}), \quad \text{and} \quad \bar{Z}_{Q,k} = Z_{Q,k}(\Pi_1 \otimes \bar{\Pi}'_{Q,k-1}),$$

so that from (13) and (20) we get

$$\bar{W}_{Q,k} = W_{Q,k} \bar{\Pi}_{Q,k} = \bar{X}_{Q,k} + \bar{Y}_{Q,k} + \bar{Z}_{Q,k} + W_{Q,k} \Phi_{Q,k}.$$

We define the action of Δ_1 on these operators by

$$\bar{X}_{Q,k} \xrightarrow{\Delta_1} \bar{X}_{Q,k}, \quad \bar{Y}_{Q,k} \xrightarrow{\Delta_1} \bar{Y}_{Q,k}, \quad W_{Q,k} \Phi_{Q,k} \xrightarrow{\Delta_1} W_{Q,k} \Phi_{Q,k}, \quad \text{and} \quad \bar{Z}_{Q,k} \xrightarrow{\Delta_1} -\bar{X}_{Q,k-1}.$$

It is not hard to check that this definition satisfies the requirements of Section 2.1. Thus, for $\tilde{\Gamma}_Q$ as defined in (9), we have

$$\tilde{\Gamma}_Q \xrightarrow{\Delta_1} \sum_k (\alpha_{k-1} - \alpha_k) \bar{X}_{Q,k-1} + \sum_k \alpha_k \bar{Y}_{Q,k} + \sum_k \alpha_k W_{Q,k} \Phi_{Q,k}.$$

So far we have merely constructed an analogue of (14). The main difference between this construction and the one in Section 5.1 is given by the following result, which we prove in Section 6.2 (note the difference with (15)):

Lemma 6. *In the above notations, we have:*

$$(a) \quad \|\bar{X}_{Q,k}\| \leq 1, \quad (b) \quad \|\bar{Y}_{Q,k}\| = O(\sqrt{k/n}), \quad (c) \quad \|W_{Q,k} \Phi_{Q,k}\| = O(1/\sqrt{n}).$$

With this result, it is not hard to show that $\alpha_k = \max\{n^{1/3} - k, 0\}$ is a good choice for the values of α_k in the decomposition (9). Indeed, for different k , all the operators $\bar{X}_{Q,k}$ are mutually orthogonal, and the same is true for $\bar{Y}_{Q,k}$ and $W_{Q,k} \Phi_{Q,k}$. Hence, the following conditions ensure that $\|\Delta_1 \diamond \tilde{\Gamma}_Q\| = O(1)$:

$$|\alpha_{k-1} - \alpha_k| \leq 1, \quad |\alpha_k| \leq \sqrt{n/k}, \quad \text{and} \quad |\alpha_k| \leq \sqrt{n}$$

for all k . Our choice $\alpha_k = \max\{n^{1/3} - k, 0\}$ satisfies these conditions, giving us

$$\|\tilde{\Gamma}_Q\| \geq \|\alpha_0 \bar{W}_{Q,0}\| = \alpha_0 = n^{1/3}.$$

5.3 Removal of Illegal Rows and Columns

So far we have only constructed the matrix $\tilde{\Gamma}_Q$ in which the actual adversary matrix Γ_Q is embedded. We obtain Γ_Q by deleting all the illegal rows and columns of $\tilde{\Gamma}_Q$. By (3), we already have that $\|\Gamma_Q \circ \Delta_i\| = O(1)$ for all i . It remains to show that $\|\Gamma_Q\|$ is not much smaller than $\|\tilde{\Gamma}_Q\|$, that is, not much smaller than α_0 . Let us assume that $q \in \Omega(n^2)$, so that a constant ratio of rows and columns of $\tilde{\Gamma}_Q$ remain in Γ_Q (that is, they are legal).

We have $\Gamma_Q = \sum_k \alpha_k \check{W}_{Q,k}$, where $\check{W}_{Q,k}$ is obtained from $\bar{W}_{Q,k}$ by deleting all the illegal rows and columns. In particular, since $\bar{W}_{Q,0} = W_{Q,0}$ is the matrix of all entries equal and $\|\bar{W}_{Q,0}\| = 1$, we have $\|\check{W}_{Q,0}\| = \Omega(1)$ and its principal right-singular vector is the all-ones vector $\mathbf{1}$ of length $q!/(q-2n)!$. All that is left to show is that $\check{W}_{Q,k}\mathbf{1} = 0$ whenever $k \neq 0$.

The coimage of $\check{W}_{Q,k}$ is contained in $\check{\mathcal{H}}_{Q,k}$. Let $\mathcal{L}$ be the domain of $\check{W}_{Q,k}$, which is spanned by the standard basis vectors corresponding to the legal negative inputs. Note that $\mathcal{L}$ is a submodule of $\mathcal{H}^{\otimes 2n}$, which is the domain of $\bar{W}_{Q,k}$. Hence, the coimage of $\check{W}_{Q,k}$ is contained in the span of the irreps of $\mathbb{S}_Q$ used in $\bar{\mathcal{H}}_{Q,k}$. The vector $\mathbf{1}$ is contained in the trivial irrep of $\mathbb{S}_Q$ ($S^{(2n)}$ for COLLISION and $S^{(n)} \otimes S^{(n)}$ for SET EQUALITY), and hence it is orthogonal to the coimage of $\check{W}_{Q,k}$. Thus, $\check{W}_{Q,k}\mathbf{1} = 0$ for $k > 0$. This gives us the main theorem of the paper.

Theorem 7. *For both $Q \in \{\text{CP}, \text{SE}\}$, let*

$$\tilde{\Gamma}_Q = \sum_{k=0}^{n^{1/3}} (n^{1/3} - k) \bar{W}_{Q,k},$$

where $\bar{W}_{Q,k}$ is defined in (16), and let Γ_Q be obtained from $\tilde{\Gamma}_Q$ by removing all the illegal rows and columns. Given that $q \in \Omega(n^2)$, Γ_{CP} and Γ_{SE} are adversary matrices for COLLISION and SET EQUALITY, respectively, giving an $\Omega(n^{1/3})$ lower bound on the quantum query complexity of both problems.

6 Proofs

In this section, we present the proofs of Lemmata 5 and 6. Throughout the whole section, we will work in the e -basis. For the group algebra $\mathbb{CS}_m$, we will use the standard basis consisting of the permutations $\mathbb{S}_m$.

Recall the operator $E_t = C_t^- R_t^+$ from Section 2.2 with the property that the submodule $\mathbb{CS}_m E_t$ of the module $\mathbb{CS}_m$ is isomorphic to S^λ , where λ is the shape of t . Its scaled version $\tilde{E}_t$ is an idempotent. Moreover, there exists a non-zero $v \in S^\lambda$ with $\tilde{E}_t v = v$, whereas $\tilde{E}_t v = 0$ for all $v \in S^\mu$ with $\mu \neq \lambda$. We start with constructing a similar operator for the whole space $\tilde{\mathcal{H}}_k^{(m)}$. Let $F \in \mathbb{CS}_m$ be defined by

$$F = \frac{1}{2^k} (\varepsilon - (a_1, b_1)) (\varepsilon - (a_2, b_2)) \cdots (\varepsilon - (a_k, b_k)). \quad (21)$$

Here, $a_1, b_1, \dots, a_k, b_k$ are some distinct fixed elements of $[m]$, ε is the identity element of $\mathbb{S}_m$, and (a_i, b_i) denotes the transposition of a_i and b_i . Note that F is an orthogonal projector.

Lemma 8. *Let $\lambda \vdash m$. If $\lambda_1 = m - k$, then there exists a non-zero vector $v \in S^\lambda$ such that $Fv = v$. If $\lambda_1 > m - k$, then $Fv = 0$ for all $v \in S^\lambda$.*

Proof. We start with the first statement. Let $\ell = \lambda_2$, and let t be a Young tableau of shape λ with $a_1, \dots, a_k$, in this order, being the first k entries in the first row, and $b_1, \dots, b_k$ being the entries in

the remaining rows, so that $b_1, \dots, b_\ell$ form the second row. Since S^λ is isomorphic to $\mathbb{C}\mathbb{S}_m E_t$, we will work with the latter from now on.

Let us define

$$v = (\varepsilon - (a_{\ell+1}, b_{\ell+1})) \cdots (\varepsilon - (a_k, b_k)) E_t \in \mathbb{C}\mathbb{S}_m E_t. \quad (22)$$

Then, clearly, $\frac{1}{2}(\varepsilon - (a_i, b_i))v = v$ for $i > \ell$. On the other hand, $\frac{1}{2}(\varepsilon - (a_i, b_i))v = v$ for all $i \leq \ell$ as well since for them $(a_i, b_i) \in C_t$. Hence, we can conclude that $Fv = v$.

It remains to show that $v \neq 0$. Open the brackets in (22). The result is $\sum_{\pi} (\text{sgn } \pi) \pi E_t$, where π runs through all possible products of the transpositions $(a_{\ell+1}, b_{\ell+1}), \dots, (a_k, b_k)$. Let $\rho = (a_{\ell+1}, b_{\ell+1}) \cdots (a_k, b_k)$. Since the coefficient of the basis vector ε in E_t is 1, the vector ρE_t uses the basis vector ρ . Note that ρ maps all $b_{\ell+1}, \dots, b_k$ to the elements of $[m]$ lying outside of the first ℓ columns of t . On the other hand, any permutation σ used in E_t maps $b_{\ell+1}, \dots, b_k$ to the elements in the first ℓ columns of t . Hence, $\pi\sigma$ maps some of $b_{\ell+1}, \dots, b_k$ to an element in the first ℓ columns of t , unless π is a product of all the $k - \ell$ transpositions, i.e., unless $\pi = \rho$. Thus, πE_t does not use ρ unless $\pi = \rho$. Hence, the coefficient of ρ is non-zero in v , and, in particular, $v \neq 0$.

For the second statement, recall that any S^λ is spanned by the images of R_t^+ as t runs through the Young tableau of shape λ . Fix t of shape λ . Since $\lambda_1 > m - k$, at least one pair $\{a_1, b_1\}, \dots, \{a_k, b_k\}$ is contained in the first row of t , and hence F is zero on the image of R_t^+ . Thus, F is zero on the whole S^λ . $\square$

Corollary 9. *The $\mathbb{S}_m$ -module $\mathcal{H}_k^{(m)}$ only uses irreps S^λ with $\lambda_1 \geq m - k$.*

Proof. Assume towards contradiction that $\mathcal{H}_k^{(m)}$ contains a copy of irrep S^λ with $\lambda_1 < m - k$. Then, by Lemma 8, there exist distinct elements $a_1, b_1, \dots, a_{k+1}, b_{k+1} \in [m]$ and a non-zero vector $v \in \mathcal{H}_k^{(m)}$ such that $F''v = v$, where $F'' = \frac{1}{2^{k+1}}(\varepsilon - (a_1, b_1))(\varepsilon - (a_2, b_2)) \cdots (\varepsilon - (a_{k+1}, b_{k+1}))$. On the other hand, for every e -basis vector $e_{c_1} \otimes \cdots \otimes e_{c_m}$ of weight k , there exists $i \in [k+1]$ such that $c_{a_i} = c_{b_i} = 0$. Hence, F'' nullifies this basis vector, and, consequently, is zero on the whole $\mathcal{H}_k^{(m)}$. This contradiction finishes the proof. $\square$

Corollary 10. *If $v \in \mathcal{H}_k^{(m)}$ is such that $Fv = v$, then $v \in \bar{\mathcal{H}}_k^{(m)}$.*

Proof. By Lemma 8, v is orthogonal to all copies of S^λ in $\mathcal{H}_k^{(m)}$ with $\lambda_1 > m - k$. And Corollary 9 then only leaves the possibility that v is contained in $\bar{\mathcal{H}}_k^{(m)}$. $\square$

Before we proceed with the proofs, let us introduce the following piece of notation. For $a \in [m]$ and $z \in [0..q-1]$, $\overset{a}{(z)}$ denotes the component e_z residing in the a -th copy of $\mathcal{H}$ in the tensor product $\mathcal{H}^{\otimes m}$. We might omit the upper index a when $z = 0$ and when it is clear in which copies of $\mathcal{H}$ zero components e_0 must reside. A juxtaposition of such components denotes the tensor product. For instance, given $m = 6$,

$$\overset{4}{(3)}\overset{2}{(0)}\overset{5}{(3)}\overset{1}{(7)}(0)(0) = e_7 \otimes e_0 \otimes e_0 \otimes e_3 \otimes e_3 \otimes e_0.$$

6.1 Proof of Lemma 5

Let us for brevity denote $\Phi = \Phi_k^{(m)}$, and recall that

$$\Phi = \bar{\Pi}_k^{(m)} - \Pi_0 \otimes \bar{\Pi}_k^{(m-1)} - \Pi_1 \otimes \bar{\Pi}_{k-1}^{(m-1)}.$$

Recall that Φ is symmetric with respect to $\mathbb{S}'_{m-1} = \mathbb{S}_{[2..m]}$, and it is zero on the orthogonal complement of $\mathcal{H}_k^{(m)}$. By the branching rule [22, Section 2.8] applied to the term $\bar{\Pi}_k^{(m)}$ in the equation above, the image of Φ only uses irreps S^λ of $\mathbb{S}'_{m-1}$ corresponding to partitions $\lambda \vdash m-1$ with $\lambda_1 = m-k$ and $\lambda_1 = m-k-1$.

Let us start with the latter case: $\lambda_1 = m-k-1$. Fix $\lambda \vdash m-1$ with $\lambda_1 = m-k-1$. By Schur's lemma, there exists a copy of S^λ in $\mathcal{H}_k^{(m)}$ consisting of principal right-singular vectors of Φ on the isotypic subspace corresponding to λ . Fix k disjoint pairs $\{a_1, b_1\}, \dots, \{a_k, b_k\} \subset [2..m]$, and let them specify F as in (21). By Lemma 8, there is a vector v in this irrep that satisfies $Fv = v$.

Consider the vector v in the e -basis. Since F is an orthogonal projector, v only uses those basis vectors that are not nullified by F . The operator F nullifies any basis vector unless it has a non-zero (i.e., different from e_0) component in positions specified by each pair $\{a_i, b_i\}$. Since each basis vector in $\mathcal{H}_k^{(m)}$ has exactly k non-zero components, the first component of every basis vector used by v must be e_0 . Hence, $v \in \mathcal{H}_0 \otimes \bar{\mathcal{H}}_k^{(m-1)}$. By Corollary 10, we also have $v \in \bar{\mathcal{H}}_k^{(m)}$. Thus

$$\Phi v = \bar{\Pi}_k^{(m)} v - \Pi_0 \otimes \bar{\Pi}_k^{(m-1)} v - \Pi_1 \otimes \bar{\Pi}_{k-1}^{(m-1)} v = v - v - 0 = 0,$$

and Φ is zero on the isotypic subspace corresponding to λ . This proves the second statement of Lemma 5.

Now let us consider the remaining case: $\lambda_1 = m-k$. We use $\bar{\mathcal{H}}_k^{(m)}$ to denote the subspace of $\mathcal{H}_k^{(m)}$ spanned by the irreps having exactly $k-1$ boxes below the first row. In this notation, the subspace of $\mathcal{H}_k^{(m)}$ spanned by the irreps S^λ of $\mathbb{S}'_{m-1}$ with $\lambda_1 = m-k$ is given by $\mathcal{H}_0 \otimes \bar{\mathcal{H}}_k^{(m-1)} \oplus \mathcal{H}_1 \otimes \bar{\mathcal{H}}_{k-1}^{(m-1)}$. We treat these two cases separately in the two claims below. If Φ has norm $O(1/\sqrt{m})$ on both of them, then so it does on their direct sum, thus proving Lemma 5.

Claim 11. Φ has norm $O(1/\sqrt{m})$ on $\mathcal{H}_1 \otimes \bar{\mathcal{H}}_{k-1}^{(m-1)}$.

Proof. Fix $k-1$ disjoint pairs $\{a_1, b_1\}, \dots, \{a_{k-1}, b_{k-1}\} \subset [2..m]$, and let $a_k = 1$ for notational convenience. Define an orthogonal projector

$$F' = \frac{1}{2^{k-1}} (\varepsilon - (a_1, b_1)) (\varepsilon - (a_2, b_2)) \cdots (\varepsilon - (a_{k-1}, b_{k-1})). \quad (23)$$

By Schur's Lemma and Lemma 8, there exists a principal right-singular vector $v \in \mathcal{H}_1 \otimes \bar{\mathcal{H}}_{k-1}^{(m-1)}$ of $\Phi(\Pi_1 \otimes \bar{\Pi}_{k-1}^{(m-1)})$ such that $F'v = v$.

As $z = (z_1, \dots, z_k)$ runs through $[1..q-1]^k$, the following vectors

$$w_z = \left[\frac{a_1 \ b_1}{(z_1)(0) - (0)(z_1)} \otimes \frac{a_2 \ b_2}{(z_2)(0) - (0)(z_2)} \otimes \cdots \otimes \frac{a_{k-1} \ b_{k-1}}{(z_{k-1})(0) - (0)(z_{k-1})} \right] \otimes \frac{a_k}{(z_k)(0)(0) \cdots (0)}$$

form an orthonormal basis of the image of F' in $\mathcal{H}_1 \otimes \bar{\mathcal{H}}_{k-1}^{(m-1)}$. Hence, we can write $v = \sum_z \beta_z w_z$.

For brevity, let $u_z \in \mathcal{H}^{\otimes(2k-2)}$ denote the unit vector in the square brackets in the above definition of w_z . Then we can decompose $w_z = w'_z + w''_z$, with

$$\begin{aligned} w'_z &= \frac{1}{m-2k+2} u_z \otimes \sum_{b_k \in [m] \setminus \{a_1, b_1, \dots, a_{k-1}, b_{k-1}, a_k\}} \left[\frac{a_k \ b_k}{(z_k)(0) - (0)(z_k)} \right] \otimes (0) \cdots (0), \\ w''_z &= \frac{1}{m-2k+2} u_z \otimes \sum_{c \in [m] \setminus \{a_1, b_1, \dots, a_{k-1}, b_{k-1}\}} \frac{c}{(z_k)(0)(0) \cdots (0)}. \end{aligned}$$

Note that $w'_z \in \bar{\mathcal{H}}_k^{(m)}$ by Lemma 8. In addition, v is orthogonal to $\mathcal{H}_0 \otimes \bar{\mathcal{H}}_k^{(m-1)}$, therefore

$$\Phi v = \bar{\Pi}_k^{(m)} v - v = \sum_z \beta_z \left(w'_z + \bar{\Pi}_k^{(m)} w''_z \right) - \sum_z \beta_z \left(w'_z + w''_z \right) = - \left(\mathbb{I}_{\mathcal{H}}^{\otimes m} - \bar{\Pi}_k^{(m)} \right) \sum_z \beta_z w''_z.$$

We have $\|w''_z\| = 1/\sqrt{m-2k+2}$ and the vectors w''_z are mutually orthogonal. Thus

$$\|\Phi v\|^2 \leq \left\| \sum_z \beta_z w''_z \right\|^2 = \sum_z \beta_z^2 \|w''_z\|^2 = \frac{1}{m-2k+2} \sum_z \beta_z^2 = \frac{1}{m-2k+2} \|v\|^2,$$

and the norm of Φ on $\mathcal{H}_1 \otimes \bar{\mathcal{H}}_{k-1}^{(m-1)}$ is at most $1/\sqrt{m-2k+2} = O(1/\sqrt{m})$. $\square$

Claim 12. Φ has norm $O(1/\sqrt{m})$ on $\mathcal{H}_0 \otimes \bar{\mathcal{H}}_k^{(m-1)}$.

Proof. By Schur's Lemma there exist $\lambda \vdash m-1$ with $\lambda_1 = m-k$ and a copy of the irrep S^λ in $\mathcal{H}_0 \otimes \bar{\mathcal{H}}_k^{(m-1)}$ consisting of principal right-singular vectors of $\Phi(\Pi_0 \otimes \bar{\Pi}_k^{(m-1)})$. Let t be a Young tableau of shape λ with the entries in $[2..m]$. Recall the operator $E_t = C_t^- R_t^+$ from (2). There exists a non-zero vector v in this irrep that is also in the image of E_t . Let $d_1, \dots, d_\ell$ be the entries in the columns of t of height 1. We have $\ell \geq m-2k+1$. Also, for convenience, let $d_{\ell+1} = 1$.

Since C_t does not affect $d_1, \dots, d_\ell$, and R_t uses all the symmetric group on these elements, the vector v is symmetric with respect to permuting them. Also, v is in the image of C_t^- , which means it only uses basis vectors such that, for entries of each column of t , at most one corresponding component is e_0 . Consequently, these basis vectors have at most one non-zero component in positions $d_1, \dots, d_\ell$, and the vector v is of the form

$$v = u_0 \otimes \binom{d_1}{0} \binom{d_2}{0} \cdots \binom{d_\ell}{0} \otimes \binom{d_{\ell+1}}{0} + \sum_{z \in [1..q-1]} u_z \otimes \left(\sum_{j=1}^{\ell} \binom{d_j}{z} \binom{d_{\ell+1}}{0} \right) \otimes \binom{d_{\ell+1}}{0}$$

for some vectors $u_0, u_1, \dots, u_{q-1} \in \mathcal{H}^{\otimes(m-\ell-1)}$. The vector v can be decomposed as $v' + v''$ with

$$\begin{aligned} v' &= u_0 \otimes \binom{d_1}{0} \binom{d_2}{0} \cdots \binom{d_{\ell+1}}{0} + \sum_{z \in [1..q-1]} u_z \otimes \left(\sum_{j=1}^{\ell+1} \binom{d_j}{z} \binom{d_{\ell+1}}{0} \right), \\ v'' &= - \sum_{z \in [1..q-1]} u_z \otimes \binom{d_{\ell+1}}{z} \binom{d_{\ell+1}}{0} \cdots \binom{d_{\ell+1}}{0}. \end{aligned}$$

Note the intentional resemblance between v and v' , and note that

$$\|v\|^2 \geq \sum_{z \in [1..q-1]} \ell \|u_z\|^2 = \ell \|v''\|^2.$$

We claim that v' is in the image of $E_{t'}$, where t' is obtained from t by appending the element 1 to the first row. This means that v' is contained in $S^{\text{sh}(t')}$. In particular, v' is orthogonal to $\bar{\mathcal{H}}_k^{(m)}$. Because v is orthogonal to both $\mathcal{H}_0 \otimes \bar{\mathcal{H}}_k^{(m-1)}$ and $\mathcal{H}_1 \otimes \bar{\mathcal{H}}_{k-1}^{(m-1)}$, we have

$$\|\Phi v\|^2 = \|\bar{\Pi}_k^{(m)} v\|^2 = \|\bar{\Pi}_k^{(m)} v''\|^2 \leq \|v''\|^2 \leq \|v\|^2 / \ell.$$

Hence, the norm of Φ on $\mathcal{H}_0 \otimes \bar{\mathcal{H}}_k^{(m-1)}$ is at most $1/\sqrt{m-2k+1} = O(1/\sqrt{m})$.

Now let us prove our claim. By linearity, it suffices to prove it when $v = E_t e$ for some basis vector e with e_0 as the first component. Let R_1 be the set of entries in the first row of t , and R' be the group of entry permutations within the rows of t except the first one. Denote by $A \subset R_1$ the set of entries in R_1 that correspond to non-zero components in e . Then, $E_t e$ is proportional to

$$\sum_{\substack{S \subseteq R_1: |S|=|A| \\ |S \cap \{d_1, \dots, d_\ell\}| \leq 1}} \sum_{\substack{\pi: A \rightarrow S \\ \pi \text{ is a bijection}}} \sum_{\pi' \in R'} \sum_{\rho \in C_t} (\text{sgn } \rho) \rho \pi' \pi e,$$

where πe is defined in the obvious way, and $E_{t'} e$ is proportional to

$$\sum_{\substack{S \subseteq R_1 \cup \{d_{\ell+1}\}: |S|=|A| \\ |S \cap \{d_1, \dots, d_{\ell+1}\}| \leq 1}} \left[\sum_{\substack{\pi: A \rightarrow S \\ \pi \text{ is a bijection}}} \sum_{\pi' \in R'} \sum_{\rho \in C_t} (\text{sgn } \rho) \rho \pi' \pi e \right].$$

The only difference between the two expressions above is that there appeared new subsets S in the outermost sum satisfying $S \cap \{d_1, \dots, d_{\ell+1}\} = \{d_{\ell+1}\}$. Since neither R' nor C_t affect $\{d_1, \dots, d_{\ell+1}\}$, the expression in the brackets with $S \cap \{d_1, \dots, d_{\ell+1}\} = \{d_{\ell+1}\}$ can be obtained from the same expression for $(S \setminus \{d_{\ell+1}\}) \cup \{d_\ell\}$ by exchanging the d_ℓ -th and the $d_{\ell+1}$ -th copies of the space $\mathcal{H}$ in the tensor product $\mathcal{H}^{\otimes m}$. This proves that v' , defined via the same vectors $u_0, u_1, \dots, u_{q-1}$ as v , is proportional to $E_{t'} e$. $\square$

6.2 Proof of Lemma 6

Recall that the matrix $W_{Q,k}$ is a block column matrix consisting of $|M_Q|$ blocks $\frac{1}{\sqrt{|M_Q|}} W_k^\mu$, one for every matching $\mu \in M_Q$. So, let us focus on W_k^μ given by (10). Although W_k^μ is defined on the whole $\mathcal{H}^{\otimes 2n}$, in reality it maps $\mathcal{H}_k^{(2n)}$ to $\mathcal{H}_k^{(n)}$ and is zero on the orthogonal complement.

A convenient way to illustrate the action of W_k^μ is as follows. Consider the graph G_μ on $2n$ vertices with n edges given by the pairs in μ . The basis vectors of $\mathcal{H}_k^{(2n)}$ correspond to the labellings of the vertices with components in $\{e_0, e_1, \dots, e_{q-1}\}$ that have exactly k non-zero components. Fix such a labelling and let e be the corresponding basis vector. If there is an edge connecting two non-zero components, then $W_k^\mu e = 0$. Otherwise, $W_k^\mu e$ is a basis vector of $\mathcal{H}_k^{(n)}$, which corresponds to the following labelling of the edges: an edge connecting two components e_0 is labeled by e_0 and an edge connecting e_0 and a non-zero component e_i is labeled by e_i .

Lemma 13. *Recall the operator F from (21). For any vector $v \in \mathcal{H}_k^{(2n)}$ satisfying $Fv = v$ and any matching μ , we have $\|W_k^\mu v\| \leq \|v\|$, $\|X_k^\mu v\| \leq \|v\|$, and $\|Y_k^\mu v\| \leq \|v\|$, where W_k^μ , X_k^μ , and Y_k^μ are defined in (10) and (12).*

Proof. We prove the result for W_k^μ , the proofs for X_k^μ and Y_k^μ being similar. Let, for brevity, $W = W_k^\mu$, and let $A_i = \{a_i, b_i\}$ be the pairs from the definition of F . Similarly as in the proof of Claim 11, the following vectors give an orthonormal basis of the image of F as z runs through $[1..q-1]^k$:

$$w_z = \frac{\begin{pmatrix} a_1 & b_1 \\ z_1 & 0 \end{pmatrix} - \begin{pmatrix} a_1 & b_1 \\ 0 & z_1 \end{pmatrix}}{\sqrt{2}} \otimes \frac{\begin{pmatrix} a_2 & b_2 \\ z_2 & 0 \end{pmatrix} - \begin{pmatrix} a_2 & b_2 \\ 0 & z_2 \end{pmatrix}}{\sqrt{2}} \otimes \dots \otimes \frac{\begin{pmatrix} a_k & b_k \\ z_k & 0 \end{pmatrix} - \begin{pmatrix} a_k & b_k \\ 0 & z_k \end{pmatrix}}{\sqrt{2}} \otimes (0)(0) \dots (0).$$

However, contrary to Claim 11, the images of these vectors under W are not orthogonal, thus we have to adopt a different proof strategy.

If μ contains a pair A_i for some i , then $Wv = 0$, so we will further assume this is not the case. Construct the graph G' from the graph G_μ by adding k edges $a_1b_1, \dots, a_kb_k$. Unlike the edges of G_μ , these new edges will not be labeled, as they do not correspond to copies of $\mathcal{H}$ in $\mathcal{H}^{\otimes n}$. This is a simple graph of maximal degree 2, and it is a collection of even-length cycles and odd-length paths. The set of vertices of each cycle is a union of at least two pairs from $A_1, \dots, A_k$. A path starts and ends outside of $\bigcup_i A_i$, but the set of its internal vertices is a union (possibly empty) of some of the pairs from $A_1, \dots, A_k$. Let $\mathcal{C}$ denote the set of connected components of G' , which from now on, to avoid confusion, we will call *connected parts*. For each $C \in \mathcal{C}$, we identify C with the set of its vertices. We decompose the space $\mathcal{H}^{\otimes 2n}$ as $\mathcal{H}^{\otimes 2n} = \bigotimes_{C \in \mathcal{C}} \mathcal{H}^{\otimes C}$, where $\mathcal{H}^{\otimes C}$ corresponds to the vertices in C . The above discussion implies that each w_z can be decomposed as $w_z = \bigotimes_{C \in \mathcal{C}} w_{z,C}$ with $w_{z,C} \in \mathcal{H}^{\otimes C}$ of norm 1.

Our goal is to prove that $\|WF\| \leq 1$. Let $\Xi_F: \mathcal{H}^{\otimes 2n} \rightarrow \mathcal{H}^{\otimes 2n}$ be the orthogonal projector onto the span of the basis vectors used by some w_z . These are the basis vectors having exactly one non-zero component in (positions corresponding to) each A_i . Let $\Xi_W: \mathcal{H}^{\otimes 2n} \rightarrow \mathcal{H}^{\otimes 2n}$ denote the orthogonal projector onto the span of the basis vectors *not* nullified by W , i.e., having at most one non-zero component in each pair of μ . Since the image of F is contained in the coimage of Ξ_F , and the coimage of W is contained in the image of Ξ_W , we have $WF = W\Xi_W\Xi_F F$.

Note that both Ξ_W and Ξ_F are diagonal $(0,1)$ -matrices in the e -basis. Let $\Xi = \Xi_W\Xi_F$, which is the orthogonal projector onto the span of all the basis vectors that are used any w_z and not mapped to 0 by W . These basis vectors can be described in terms of G' . In each cycle, vertices labeled by zero and non-zero components alternate. In each path, each edge a_ib_i has exactly one non-zero component (as its endpoint), and each edge of G_μ has at most one non-zero component. (See Figure 1 for an example of such graph.) The conditions for distinct $C \in \mathcal{C}$ are independent, so Ξ can be decomposed as $\Xi = \bigotimes_{C \in \mathcal{C}} \Xi_C$ with orthogonal projectors $\Xi_C: \mathcal{H}^{\otimes C} \rightarrow \mathcal{H}^{\otimes C}$.

Figure 1: An example of the graph G' for $n = 12$ and $k = 8$, which has $K = 2$ cycles. The edges of μ are represented by straight lines and the edges a_ib_i by wobbly curves. For readability, labels e_z are shown as z instead, and the labels of edges are in italics.

We have $\|WF\| = \|W\Xi F\| \leq \|W\Xi\| \|\Xi F\|$. We bound both multipliers separately in the two claims below, from which the lemma follows. Let K denote the number of cycles in G' .

Claim 14. *We have $\|\Xi F\| \leq 1/\sqrt{2^K}$.*

Proof. First, w_z form an orthonormal basis of the image of the projector F . Next, distinct w_z use disjoint sets of basis vectors, hence Ξ maps them to mutually orthogonal vectors. Thus, it suffices to show that $\|\Xi w_z\| \leq 1/\sqrt{2^K}$ for each z . We have

$$\|\Xi w_z\| = \left\| \bigotimes_{C \in \mathcal{C}} \Xi_C w_{z,C} \right\| = \prod_{C \in \mathcal{C}} \|\Xi_C w_{z,C}\|. \quad (24)$$

If C is a path, we conclude $\|\Xi_C w_{z,C}\| \leq 1$ from Ξ_C being a projector. Now assume C is a cycle. As the length of C is at least 4, the basis vectors of $\mathcal{H}^{\otimes C}$ used by $w_{z,C}$ appear in $w_{z,C}$ with coefficients $1/\sqrt{2}^{|C|/2} \leq 1/2$ by the absolute value. But, since zero and non-zero components in a cycle must alternate, at most 2 of these basis vectors are in the image of Ξ_C . Hence, $\|\Xi_C w_{z,C}\| \leq 1/\sqrt{2}$. Multiplying over all $C \in \mathcal{C}$, we get the claim. $\square$

Claim 15. *We have $\|W\Xi\| \leq \sqrt{2^K}$.*

Proof. Consider the matrix M corresponding to $W\Xi$ written in the e -basis. Its rows are labelled by the basis vectors of $\mathcal{H}_k^{(n)}$, and its columns by the basis vectors used by Ξ . Because each basis vector used by Ξ is mapped to a basis vector of $\mathcal{H}_k^{(n)}$, each column of M contains exactly one 1, and the remaining entries are all 0. It suffices to prove that each row of M contains at most 2^K non-zero entries.

Fix a labelling of the edges of G_μ , and we can assume the corresponding row is non-zero. Consider a connected part $C \in \mathcal{C}$. If C is a path, then there is a unique way to extend the labelling of the edges of G_μ to the vertices of C . Indeed, one of the edges of G_μ in the path must have label e_0 . Both its endpoints are then e_0 , and the whole labelling of the vertices of C is uniquely determined. If C is a cycle, then there are exactly two possibilities to extend the labelling of the edges of G_μ to the vertices in C , because zero and non-zero components must alternate. This proves that there are exactly 2^K entries equal to 1 in this row, and the remaining entries are 0. $\square$

This ends the proof of Lemma 13 for $W = W_k^\mu$. The operators X_k^μ and Y_k^μ are similar to W_k^μ with the difference that they map more basis vectors to 0, and the bounds above hold for them as well. $\square$

Lemma 16. *Recall the operator F' from (23). For any vector $v \in \mathcal{H}_k^{(2n)}$ satisfying $F'v = v$ and any matching μ , we have $\|W_k^\mu v\| \leq \sqrt{3}\|v\|$, where W_k^μ is defined in (10).*

Proof. This is a modification of the proof of Lemma 13, and we adopt the notation from that proof. We define $A_1, \dots, A_{k-1}$ as before. The orthonormal basis of the image of F' consists of the vectors of the form

$$w_{a_i,z} = \frac{\begin{smallmatrix} a_1 & b_1 \\ (z_1) & (0) \end{smallmatrix} - \begin{smallmatrix} a_1 & b_1 \\ (0) & (z_1) \end{smallmatrix}}{\sqrt{2}} \otimes \dots \otimes \frac{\begin{smallmatrix} a_i & b_i \\ (z_i) & (z_k) \end{smallmatrix} - \begin{smallmatrix} a_i & b_i \\ (z_k) & (z_i) \end{smallmatrix}}{\sqrt{2}} \otimes \dots \otimes \frac{\begin{smallmatrix} a_{k-1} & b_{k-1} \\ (z_{k-1}) & (0) \end{smallmatrix} - \begin{smallmatrix} a_{k-1} & b_{k-1} \\ (0) & (z_{k-1}) \end{smallmatrix}}{\sqrt{2}} \otimes (0)(0) \dots (0),$$

where $z_i < z_k$, and

$$w_{c,z} = \frac{\begin{smallmatrix} a_1 & b_1 \\ (z_1) & (0) \end{smallmatrix} - \begin{smallmatrix} a_1 & b_1 \\ (0) & (z_1) \end{smallmatrix}}{\sqrt{2}} \otimes \dots \otimes \frac{\begin{smallmatrix} a_{k-1} & b_{k-1} \\ (z_{k-1}) & (0) \end{smallmatrix} - \begin{smallmatrix} a_{k-1} & b_{k-1} \\ (0) & (z_{k-1}) \end{smallmatrix}}{\sqrt{2}} \otimes \begin{smallmatrix} c \\ (z_k) \end{smallmatrix} (0) \dots (0).$$

As before, if μ contains a pair A_i for some i , then $W_k^\mu v = 0$, so we assume this is not the case. The graph G' , the set of connected parts $\mathcal{C}$, and the projector Ξ are all also defined as before. There is still a decomposition $w = \bigotimes_{C \in \mathcal{C}} w_C$ for each $w = w_{i,z}$. And the projector Ξ still maps different $w_{i,z}$ s to orthogonal vectors.

The main difference is in the description of the basis of the image of Ξ in terms of G' . In each cycle, zero and non-zero components still alternate. However, there is now one *special path*: the one containing z_k . This path has a *special place*: either a non-zero component outside of $\bigcup_i A_i$,

Figure 2: Two types of special paths. Special places are denoted by one (left) or two (right) hollow circles.

or a pair A_i with two non-zero components. Starting from the special place, zero and non-zero components alternate. (See Figure 2 for an example.) All the other paths are as before.

Another difference is that the conditions on the basis vectors in Ξ are no longer independent for different $C \in \mathcal{C}$ (one path being special prevent other paths from being special). Hence, no longer Ξ decomposes into a tensor product. However, there is a simple remedy. Let $P \in \mathcal{C}$ be a path, and Ξ^P be the projector onto the span of the basis vectors of the image of Ξ that have P as their special path. Now the conditions become independent, and we regain the decomposition $\Xi^P = \bigotimes_{C \in \mathcal{C}} \Xi_C^P$. Also note that all the basis vectors used by given $w_{i,z}$ share the same special path. So, for a path $P \in \mathcal{C}$, let F^P denote the projector onto the span of $w_{i,z}$ with special path P .

Let again K be the number of cycles in G' . We have the following analogues of Claims 14 and 15, respectively. Here, $|P|$ denotes the number of vertices in path P , that is, its length plus 1.

Claim 17. *For a path $P \in \mathcal{C}$, we have $\|\Xi^P F^P\| \leq 1/\sqrt{2^{K+\max\{0, |P|-4\}}/2}$.*

Proof. The projector Ξ^P maps distinct $w_{i,z}$ to orthogonal vectors, hence it suffices to prove that $\|\Xi^P w\| \leq 1/\sqrt{2^{K+\max\{0, |P|-4\}}/2}$ for any $w = w_{i,z}$ with special path P .

As in (24), we have $\|\Xi^P w\| = \prod_{C \in \mathcal{C}} \|\Xi_C^P w_C\|$. And we still use the bounds $\|\Xi_C^P w_C\| \leq 1$ and $\|\Xi_C^P w_C\| \leq 1/\sqrt{2}$ when C is a non-special path or a cycle, respectively.

Now consider $\Xi_P^P w_P$. If $|P| = 2$, then the bound $\|\Xi_P^P w_P\| \leq \|w_P\|$ suffices. So let us assume $|P| \geq 4$. Then, w_P uses $2^{(|P|-2)/2}$ basis vectors, and it is easy to check that at most 2 of them are used in Ξ^P . Indeed, the special position is uniquely determined, and, starting from it, zero and non-zero components must alternate. The factor 2 comes from the two possible arrangements of non-zero components in the special pair A_i (if there is one). Thus, $\|\Xi_P^P w_P\| \leq \sqrt{2^{(|P|-4)/2}} \|w_P\|$. Multiplying over all $C \in \mathcal{C}$, we get the required result. $\square$

Claim 18. *For a path $P \in \mathcal{C}$, we have $\|W \Xi^P\| \leq \sqrt{2^{K-1}(|P|+2)}$.*

Proof. Define the matrix M for $W \Xi^P$ as in the proof of Claim 15. Still, each column of W contains exactly one 1, the remaining entries all being 0. Fix a labelling of the edges of G_μ such that the corresponding row of M is non-zero.

Consider a connected part $C \in \mathcal{C}$. Again, if C is a non-special path or a cycle, then there are exactly 1 or 2 ways, respectively, to extend the labelling of the edges of G_μ to the vertices in C . Now assume C is the special path P . All the edges of G_μ used by P must have non-zero labels. There are at most $(|P|+2)/2$ ways to extend this labelling to the labelling of the vertices in P : for each choice of the special position, there is at most one labelling. Hence, there are at most $2^{K-1}(|P|+2)$ entries equal to 1 in each row of M , and the norm of M is at most $\sqrt{2^{K-1}(|P|+2)}$. $\square$

Now we can finish the proof of Lemma 16. It suffices to prove that $\|WF'\| \leq \sqrt{3}$. We have $WF' = \sum_P WF^P = \sum_P W \Xi^P F^P$, where the sum is over all paths P in $\mathcal{C}$. Note that

W maps basis vectors with distinct special paths to orthogonal vectors. This means that $\|WF'\| = \max_P \|W\Xi^P F^P\|$. Given any path $P \in \mathcal{C}$, by the above two claims,

$$\|W\Xi^P F^P\| \leq \|W\Xi^P\| \|\Xi^P F^P\| \leq \sqrt{\frac{(|P|+2)/2}{2^{\max\{0, |P|-4\}/2}}} \leq \sqrt{3}. \quad \square$$

Now we are ready to prove Lemma 6. Let us start with point (a) stating that $\|\bar{X}_{Q,k}\| \leq 1$, where $\bar{X}_{Q,k} = X_{Q,k}(\Pi_0 \otimes \bar{\Pi}'_{Q,k})$. This matrix is symmetric with respect to $\mathbb{S}'_Q$. Hence, Schur's lemma implies that there exists an irreducible $\mathbb{S}'_Q$ -module all consisting of right-singular vectors of $\bar{X}_{Q,k}$ of singular value $\|\bar{X}_{Q,k}\|$.

By the definition of $\bar{\Pi}'_{Q,k}$, the module is isomorphic to either $S^{(2n-1-k, \lambda)}$ with $\lambda \vdash k$ for COLLISION, or $S^{(n-1-\ell, \lambda)} \otimes S^{(n-k+\ell, \lambda')}$ with $\lambda \vdash \ell$ and $\lambda' \vdash k - \ell$ for SET EQUALITY. By Lemma 8, in both cases, there exists a non-zero vector v in the module satisfying $Fv = v$ for some choice of $a_1, b_1, \dots, a_k, b_k \in [2..2n]$ (in the case of SET EQUALITY, one has to take the tensor product of two vectors obtained by two applications of Lemma 8). By Lemma 13, $\|X_k^\mu v\| \leq \|v\|$ for all μ . Hence $\|X_{Q,k}v\| \leq \|v\|$, and $\|\bar{X}_{Q,k}\| \leq 1$.

Consider point (b) now. Similarly as for (a), we get a right-singular vector v of singular value $\|\bar{Y}_{Q,k}\|$ such that $Fv = v$. Note that, if μ matches 1 with an element outside $\{a_1, b_1, \dots, a_k, b_k\}$, then, $Y_{Q,k}^\mu v = 0$, because both these components are e_0 for all basis vectors used in v . Otherwise, we still get $\|Y_{Q,k}^\mu v\| \leq 1$ by Lemma 13. The latter case only holds for an $O(k/n)$ fraction of all matchings, hence $\|\bar{Y}_{Q,k}\| = O(\sqrt{k/n})$ (cf. the third bullet in the proof of Proposition 4).

Now, let us prove (c). We start with $W_{CP,k}\Phi_{CP,k}$. Let Ξ denote the projector onto the subspace of $\mathcal{H}_k^{(2n)}$ spanned by the irreps of $\mathbb{S}'_{CP} = \mathbb{S}_{[2..2n]}$ with exactly $k-1$ boxes below the first row. From Lemma 5, we know that $\|\Phi_{CP,k}\| = O(1/\sqrt{n})$ and $\Phi_{CP,k} = \Xi\Phi_{CP,k}$. So, it suffices to prove that $\|W_{CP,k}\Xi\| = O(1)$. This matrix is symmetric with respect to $\mathbb{S}'_{CP}$; hence, it has an irrep isomorphic to S^λ for some $\lambda \vdash 2n-1$ consisting of principal right-singular vectors. Moreover, by the definition of Ξ , λ has exactly $k-1$ boxes below the first row. By Lemma 8, $W_{CP,k}\Xi$ has a principal right-singular vector that satisfies $F'v = v$ with F' as in (23). Lemma 16 then implies that $\|W_{CP,k}\Xi v\| = \|W_{CP,k}v\| \leq \sqrt{3}\|v\|$; hence, $\|W_{CP,k}\Phi_{CP,k}\| = O(1/\sqrt{n})$.

Let us now consider the case of $W_{SE,k}\Phi_{SE,k}$. Again, all these matrices are symmetric with respect to $\mathbb{S}'_{SE} = \mathbb{S}_{[2..n]} \times \mathbb{S}_{[n+1..2n]}$. Thus, there exist $\lambda \vdash n-1$, $\lambda' \vdash n$, and an irrep isomorphic to $S^\lambda \otimes S^{\lambda'}$ that consists solely of principal right-singular vectors of $W_{SE,k}\Phi_{SE,k}$. By the definition of $\Phi_{SE,k}$ from (19) and using Lemma 5, we have that λ has $k-\ell-1$ and λ' has ℓ boxes below the first row for some $\ell \in [0..k-1]$. By a double application of Lemma 8, there exists a principal right-singular vector v satisfying $F'v = v$. We have

$$W_{SE,k}\Phi_{SE,k}v = W_{SE,k} \left[\sum_{\ell=0}^{k-1} (\Phi_{k-\ell}^{(n)} \otimes \bar{\Pi}_\ell^{(n)}) \right] v = W_{SE,k} (\Phi_{k-\ell}^{(n)} \otimes \bar{\Pi}_\ell^{(n)}) v.$$

By Lemma 5, $\|(\Phi_{k-\ell}^{(n)} \otimes \bar{\Pi}_\ell^{(n)})v\| = O(1/\sqrt{n})\|v\|$. Moreover, since F' , as a group algebra element, commutes with $\Phi_{k-\ell}^{(n)} \otimes \bar{\Pi}_\ell^{(n)}$, the vector $(\Phi_{k-\ell}^{(n)} \otimes \bar{\Pi}_\ell^{(n)})v$ is also invariant under the action of F' . Applying Lemma 16, we get that $\|W_{SE,k}\Phi_{SE,k}v\| = O(1/\sqrt{n})\|v\|$, thus proving Lemma 6.

7 Discussion

In the paper, we prove tight lower bounds on the COLLISION and SET EQUALITY problems using the adversary method. This is done by restricting the construction from [9] on some specified irreducible representations of the symmetric group.

One oddity of this result is that it heavily uses the symmetry of the problem that is more typical for the polynomial method. It is an interesting question whether representation theory of the symmetric group can be avoided in the proof. In particular, consider the problem of distinguishing a 1-to-1 input from a 2-to-1 input whose matching, as defined in Section 3, belongs to some specified set of matchings M . Is it possible to characterise the quantum query complexity of this problem by a simple optimisation problem involving M ?

Another open problem is whether the results in this paper can be combined with the results in [9] in order to prove tight lower bounds for the k -DISTINCTNESS problem. In particular, is the algorithm in [7] optimal?

Acknowledgments

A.R. would like to thank Andris Ambainis, Chris Godsil, Robin Kothari, Laura Mančinska, and Robert Špalek for fruitful discussions. A.B. was supported by the ERC Advanced Grant MQC and FP7 FET Proactive project QALGO. A part of this research was performed while A.R. was at the University of Waterloo supported by Mike and Ophelia Lazaridis Fellowship, David R. Cheriton Graduate Scholarship, and the US ARO. A.R. also acknowledges the support of the Singapore National Research Foundation under NRF RF Award No. NRF-NRFF2013-13.

References

- [1] S. Aaronson and A. Ambainis. The need for structure in quantum speedups. In *Proc. of 2nd ACM ITCS*, pages 338–352, 2011. [arXiv:0911.0996](#).
- [2] 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.
- [3] A. Ambainis. Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64(4):750–767, 2002. Earlier: *STOC'00*, [arXiv:quant-ph/0002066](#).
- [4] A. Ambainis. Polynomial degree vs. quantum query complexity. In *Proc. of 44th IEEE FOCS*, pages 230–239, 2003. [arXiv:quant-ph/0305028](#).
- [5] 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. Earlier: *FOCS'98*, [arXiv:quant-ph/9802049](#).
- [6] A. Belovs. Adversary lower bound for element distinctness. [arXiv:1204.5074](#), 2012.
- [7] A. Belovs. Learning-graph-based quantum algorithm for k -distinctness. In *Proc. of 53rd IEEE FOCS*, pages 207–216, 2012. [arXiv:1205.1534](#).

- [8] A. Belovs. Span programs for functions with constant-sized 1-certificates. In *Proc. of 44th ACM STOC*, pages 77–84, 2012. [arXiv:1105.4024](#).
- [9] A. Belovs and A. Rosmanis. On the power of non-adaptive learning graphs. *Computational Complexity*, 23(2):323–354, 2014. Earlier: *CCC'13*, [arXiv:1210.3279](#).
- [10] A. Belovs and R. Špalek. Adversary lower bound for the k -sum problem. In *Proc. of 4th ACM ITCS*, pages 323–328, 2013. [arXiv:1206.6528](#).
- [11] G. Brassard, P. Høyer, and A. Tapp. Quantum cryptanalysis of hash and claw-free functions. In *Proc. of 3rd LATIN*, volume 1380 of *LNCS*, pages 163–169. Springer, 1998. [arXiv:quant-ph/9705002](#).
- [12] H. Buhrman and R. de Wolf. Complexity measures and decision tree complexity: a survey. *Theoretical Computer Science*, 288:21–43, 2002.
- [13] C. W. Curtis and I. Reiner. *Representation theory of finite groups and associative algebras*. AMS, 1962.
- [14] 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](#).
- [15] G. James and A. Kerber. *The Representation Theory of the Symmetric Group*, volume 16 of *Encyclopedia of Mathematics and its Applications*. Addison-Wesley, 1981.
- [16] S. Kutin. Quantum lower bound for the collision problem with small range. *Theory of Computing*, 1(1):29–36, 2005.
- [17] T. Lee, R. Mittal, B. W. Reichardt, R. Špalek, and M. Szegedy. Quantum query complexity of state conversion. In *Proc. of 52nd IEEE FOCS*, pages 344–353, 2011. [arXiv:1011.3020](#).
- [18] T. Lee and J. Roland. A strong direct product theorem for quantum query complexity. *Computational Complexity*, 22(2):429–462, 2013. Earlier: *CCC'12*, [arXiv:1104.4468](#).
- [19] G. Midrijānis. A polynomial quantum query lower bound for the set equality problem. In *Proc. of 31st ICALP*, volume 3142 of *LNCS*, pages 996–1005. Springer, 2004. [arXiv:quant-ph/0401073](#).
- [20] B. W. Reichardt. Reflections for quantum query algorithms. In *Proc. of 22nd ACM-SIAM SODA*, pages 560–569, 2011. [arXiv:1005.1601](#).
- [21] B. W. Reichardt and R. Špalek. Span-program-based quantum algorithm for evaluating formulas. *Theory of Computing*, 8:291–319, 2012. Earlier: *STOC'08*, [arXiv:0710.2630](#).
- [22] B. E. Sagan. *The symmetric group: representations, combinatorial algorithms, and symmetric functions*, volume 203 of *Graduate Texts in Mathematics*. Springer, 2001.
- [23] J.-P. Serre. *Linear Representations of Finite Groups*, volume 42 of *Graduate Texts in Mathematics*. Springer, 1977.
- [24] Y. Shi. Quantum lower bounds for the collision and the element distinctness problems. In *Proc. of 43th IEEE FOCS*, pages 513–519, 2002. [arXiv:quant-ph/0112086](#).

- [25] R. Špalek. Adversary lower bound for the orthogonal array problem. `arXiv:1304.0845`, 2013.
- [26] R. Špalek and M. Szegedy. All quantum adversary methods are equivalent. *Theory of Computing*, 2:1–18, 2006. Earlier: *ICALP'05*, `arXiv:quant-ph/0409116`.
- [27] M. Zhandry. How to construct quantum random functions. In *Proc. of 53rd IEEE FOCS*, pages 679–687, 2012. `ePrint:2012/182`.
- [28] M. Zhandry. A note on the quantum collision and set equality problems. *Quantum Information & Computation*, 15(7&8):557–567, 2015. `arXiv:1312.1027`.
- [29] S. Zhang. On the power of Ambainis lower bounds. *Theoretical Computer Science*, 339(2):241–256, 2005. `arXiv:quant-ph/0311060`.