

Lower Bounds on Quantum Query and Learning Graph Complexities

by

Ansis Rosmanis

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Computer Science

Waterloo, Ontario, Canada, 2014

© Ansis Rosmanis 2014

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

In this thesis we study the power of quantum query algorithms and learning graphs; the latter essentially being very specialized quantum query algorithms themselves. We almost exclusively focus on proving lower bounds for these computational models.

First, we study lower bounds on learning graph complexity. We consider two types of learning graphs: adaptive and, more restricted, non-adaptive learning graphs. We express both adaptive and non-adaptive learning graph complexities of Boolean-valued functions (i.e., decision problems) as semidefinite minimization problems, and derive their dual problems. For various functions, we construct feasible solutions to these dual problems, thereby obtaining lower bounds on the learning graph complexity of the functions. Most notably, we prove an almost optimal $\Omega(n^{9/7}/\sqrt{\log n})$ lower bound on the non-adaptive learning graph complexity of the TRIANGLE problem. We also prove an $\Omega(n^{1-2^{k-2}/(2^k-1)})$ lower bound on the adaptive learning graph complexity of the k -DISTINCTNESS problem, which matches the complexity of the best known quantum query algorithm for this problem.

Second, we construct optimal adversary lower bounds for various decision problems. Our main procedure for constructing them is to embed the adversary matrix into a larger matrix whose properties are easier to analyze. This embedding procedure imposes certain requirements on the size of the input alphabet. We prove optimal $\Omega(n^{1/3})$ adversary lower bounds for the COLLISION and SET EQUALITY problems, provided that the alphabet size is at least $\Omega(n^2)$. An optimal lower bound for COLLISION was previously proven using the polynomial method, while our lower bound for SET EQUALITY is new. (An optimal lower bound for SET EQUALITY was also independently and at about the same time proven by Zhandry using the polynomial method [arXiv, 2013].)

We compare the power of non-adaptive learning graphs and quantum query algorithms that only utilize the knowledge on the possible positions of certificates in the input string. To do that, we introduce a notion of a certificate structure of a decision problem. Using the adversary method and the dual formulation of the learning graph complexity, we show that, for every certificate structure, there exists a decision problem possessing this certificate structure such that its non-adaptive learning graph and quantum query complexities differ by at most a constant multiplicative factor. For a special case of certificate structures, we construct a relatively general class of problems having this property. This construction generalizes the adversary lower bound for the k -SUM problem derived recently by Belovs and Špalek [ACM ITCS, 2013].

We also construct an optimal $\Omega(n^{2/3})$ adversary lower bound for the ELEMENT DISTINCTNESS problem with minimal non-trivial alphabet size, which equals the length of the input. Due to the strict requirement on the alphabet size, here we cannot use the embedding procedure, and the construction of the adversary matrix heavily relies on the representation theory of the symmetric group. While an optimal lower bound for ELEMENT DISTINCTNESS using the polynomial method

had been proven for any input alphabet, an optimal adversary construction was previously only known for alphabets of size at least $\Omega(n^2)$.

Finally, we introduce the ENHANCED FIND-TWO problem and we study its query complexity. The ENHANCED FIND-TWO problem is, given n elements such that exactly k of them are marked, find two distinct marked elements using the following resources: (1) one initial copy of the uniform superposition over all marked elements, (2) an oracle that reflects across this superposition, and (3) an oracle that tests if an element is marked. This relational problem arises in the study of quantum proofs of knowledge. We prove that its query complexity is $\Theta(\min\{\sqrt{n/k}, \sqrt{k}\})$.

Acknowledgements

I would like to thank my supervisor John Watrous for his invaluable support, advice, and guidance throughout my graduate studies. I also thank him for trusting me to work on research problems that interested me.

I would like to thank Andris Ambainis, Aleksandrs Belovs, Troy Lee, Miklos Santha, Aarthi Sundaram, and Dominique Unruh for the collaboration on the research results presented in this thesis. I would also like to thank Chris Godsil, Robin Kothari, Hari Krovi, Rajat Mittal, Abel Molina, and Robert Špalek for fruitful discussions and useful comments regarding my work. I thank Andrew Childs, Richard Cleve, Frédéric Magniez, Ashwin Nayak, and John Watrous for being on my PhD thesis committee, reading this thesis, and providing me with useful comments that helped me to improve it.

I would like to thank the faculty and students at the Institute for Quantum Computing for invaluable group and personal meetings, which broadened my knowledge on the field and gave me many ideas for my research. I have spent a portion of my studies as a visitor to Laboratoire d'Informatique Algorithmique: Fondements et Applications in Paris, the University of Latvia in Riga, and the Centre for Quantum Technologies in Singapore, and I would like to thank these institutions for their hospitality.

I also thank Canada for its hospitality. Over the years of my graduate studies, I have met many people that have made my days here very colorful. You are too many to name, and I thank you all. I would also like to thank my family, my dad and my two sisters, for their support.

For the financial support, I thank Mike and Ophelia Lazaridis Fellowship, David R. Cheriton Graduate Scholarship, and the US ARO.

Table of Contents

List of Tables	xi
List of Figures	xii
Introduction	1
I Preliminaries	8
1 Mathematical preliminaries	9
1.1 Linear algebra	9
1.2 Semidefinite programming	11
1.3 Linear group representations	12
1.3.1 Basics of group theory and the symmetric group	12
1.3.2 Basic definitions	13
1.3.3 Subrepresentations	14
1.3.4 Transporter bases of isotypical subspaces	15
1.3.5 Character theory	17
1.3.6 Group action and regular representations	18
1.3.7 Restricted and induced representations	19
1.3.8 Composition of representations	20
1.4 Representation theory of the symmetric group	20

1.4.1	Young diagrams and Young tableaux	21
1.4.2	Specht modules	23
1.4.3	Induction and restriction of representations	25
1.4.4	The orthogonal form of \mathcal{S}^λ	26
1.4.5	Decomposition of inner tensor products	29
1.4.6	Representation theory of the unitary group	30
1.5	Association schemes	31
1.5.1	Hamming scheme	31
1.5.2	Johnson scheme	32
2	Quantum query complexity	34
2.1	Computational problems	34
2.1.1	Common computational problems	35
2.1.2	Certificates for decision problems	37
2.2	Quantum query algorithm	40
2.2.1	Registers and states of the computation	40
2.2.2	Automorphisms of problems and symmetrization	43
2.2.3	Algorithms with an input register	46
2.2.4	Symmetries of the input register	47
2.3	Adversary bound	49
2.3.1	Intuition behind the bound	51
2.3.2	Simplification tools	52
2.3.3	Structure of adversary constructions	56
2.3.4	Limitations of positive-weights adversary bound	59
2.4	Span programs and learning graphs	60

II	Results	65
3	Lower bounds on learning graph complexity	66
3.1	Learning graph complexity as a semidefinite program	66
3.1.1	SDPs for adaptive learning graph complexity	66
3.1.2	SDPs for learning graph complexity of certificate structures	70
3.2	Learning graph complexity of certificate structures	72
3.2.1	Lower bounds for the k -subset and hidden shift certificate structures	73
3.2.2	Lower bound for the triangle certificate structure	74
3.3	Lower bounds on adaptive learning graph complexity	78
3.3.1	Adaptive learning graph complexity of the And function	78
3.3.2	Adaptive learning graph complexity of k -Distinctness	78
4	Adversary bounds using matrix embedding	83
4.1	Adversary bound for Element Distinctness	84
4.1.1	Construction of the adversary matrix	85
4.1.2	Bounding $\ \Delta_1 \circ \tilde{\Gamma}'\ $	88
4.1.3	Bounding $\ \Delta_n \circ \tilde{\Gamma}''\ $	88
4.1.4	Removal of illegal columns	89
4.2	Adversary lower bounds for the Collision and Set Equality problems	90
4.2.1	Preliminaries	90
4.2.2	Simple yet unsuccessful construction	92
4.2.3	Successful construction	93
4.2.4	Removal of illegal rows and columns	96
4.3	Adversary bounds for certificate structures	97
4.3.1	Outline of the lower bound	98
4.3.2	Common parts of the proofs	99
4.3.3	Comparison to adversary constructions of Sections 4.1 and 4.2	101
4.3.4	Boundedly generated certificate structures	102
4.3.5	General certificate structures	104

5	Adversary bound for Element Distinctness with small range	109
5.1	Preliminaries	109
5.2	Building blocks of Γ	111
5.2.1	Decomposition of U and V into irreps	111
5.2.2	Γ as a linear combination of transporters	113
5.3	Specification of Γ via $\Gamma_{1,2}$	114
5.3.1	Necessary and sufficient symmetries of $\Gamma_{1,2}$	114
5.3.2	Labeling of projectors and transporters	116
5.3.3	Decomposition of $\Gamma_{1,2}$ into projectors and transporters	116
5.4	Tools for estimating $\ \Delta_1 \circ \Gamma\ $	117
5.4.1	Division of $\Delta_1 \circ \Gamma$ into two parts	117
5.4.2	Commutativity with the action of Δ_i	118
5.4.3	Relations among irreps of $\mathbb{S}_{[3..n]} \times \mathbb{S}_\Sigma$ within an isotypical subspace	118
5.4.4	Relations among irreps of $\mathbb{S}_{[4..n]} \times \mathbb{S}_\Sigma$ within an isotypical subspace	119
5.4.5	Summing the permutations of $(\Delta_1 \circ \Gamma_{1,2})^*(\Delta_1 \circ \Gamma_{1,2})$	120
5.5	Construction of the optimal adversary matrix	121
5.5.1	Approximate action of Δ_i	122
5.5.2	Bounding $\ \Delta_1 \circ \Gamma'\ $	123
5.5.3	Bounding $\ \Delta_1 \circ \Gamma''\ $	125
6	Lower bound for the Enhanced Find-Two problem	127
6.1	Framework of the lower bound	129
6.1.1	Proof of Lemma 6.3	131
6.1.2	Proof of Lemma 6.4	131
6.2	Proof of Lemma 6.5	134
6.2.1	Decomposition of \mathcal{X}_Q into irreps	136
6.2.2	Necessary and sufficient conditions for the irrep $\mathcal{S}^{(n-1,1)}$	139
6.2.3	Solution for the irreps $\mathcal{S}^{(n-2,2)}$ and $\mathcal{S}^{(n-2,1,1)}$	141
6.2.4	Solution for the irrep $\mathcal{S}^{(n-1,1)}$	142

Conclusion	144
References	146
APPENDICES	152
A Proofs of lemmas in Section 4.2	153
A.1 Proof of Lemma 4.4	153
A.2 Proof of Lemma 4.3	156
A.3 Proof of Lemma A.3	160
B Necessary conditions on the adversary matrix for Element Distinctness with small range	163
B.1 Action of Δ_i on Π_λ^λ and transporters	163
B.2 Necessary conditions for $\ \Delta_1 \circ \Gamma\ = O(1)$	164
B.3 Proof of Claim 5.4	167

List of Tables

1.1	The number of standard tableaux of shape λ for the last seven λ in the lexicographical order.	23
5.1	Available operators for the construction of Γ . We distinguish three cases: both λ and ν are the same below the first row (label “ \surd_0 ”), λ has one box more below the first row than ν (label “ \surd_1 ”), λ has two boxes more below the first row than ν (labels “ \surd_2 ” and “ \surd_2 ”).	113

List of Figures

1.1	The Young diagram corresponding to the partition $(5, 3, 3, 2)$	21
1.2	Hook lengths of the boxes of the Young diagrams $(5, 3, 3, 2)$ (left) and $(m - 3, 3)$ (right).	21
1.3	A standard (left) and a non-standard (right) Young tableau of shape $(5, 3, 3, 2)$	22
1.4	A tabloid of shape $(2, 2, 1)$	23
1.5	The last letter order of the standard tableaux of shape $(3, 2)$	27
1.6	The generating matrices of the orthogonal form $\omega^{(3,2)}$. (Here we have omitted entries 0.)	27
2.1	The circuit diagram of the standard quantum oracle. For the sake of conciseness, in circuit diagrams we write \mathcal{O}_x instead of $\mathcal{O}(x)$	41
2.2	The circuit diagram of a generic quantum query algorithm.	42
2.3	Unitary transformations of the symmetrized algorithm $\bar{\mathcal{A}}$: (top) the transformation between the queries t and $t + 1$, if $t \in [T - 1]$, or the initial transformation, if $t = 0$; (bottom) the final transformation (i.e., $t = T$). The controlled unitary transformations are controlled by $\gamma \in G$	46
2.4	The oracle controlled by the input.	47
2.5	The circuit diagram of a generic quantum query algorithm with the input and symmetrization registers.	48
2.6	Embedding the adversary matrix Γ into a larger matrix $\tilde{\Gamma}$	55
2.7	Decomposition of W_k as the sum of three matrices. Wiggly arrows connect matrices that cancel out after the application of Δ_i	58
2.8	The matrix W_k corresponds to Young diagrams λ with k boxes below the first row. Wiggly arrows connect Young diagrams λ whose corresponding matrices $\Delta_i \circ W_\lambda$ have a potential for cancellation.	59

4.1	The decomposition of W_k into blocks. Here, by abuse of notation, we write $W_{\widehat{\mu},k}$ instead of $W_{\mu,k}$.	87
4.2	The decomposition of $\Delta_1 \circ \tilde{\Gamma}$ for the sake of estimating its norm. The top part of the matrix on the right consists of $n - 1$ blocks, the bottom part to $\binom{n-1}{2}$ blocks.	87
5.1	Symmetries of $\Gamma_{1,2}$ for $n = 5$ and $\Sigma = \{a, b, c, d, e\}$. With respect to the bijection f , the order of rows and columns matches. The solid arrows show that U^τ and V^τ act symmetrically on $\Gamma_{1,2}$ (here we use $\tau = (aeb)(cd) \in \mathbb{S}_\Sigma$), and so do U_π and V_π for $\pi \in \mathbb{S}_{[3..n]}$ (here we use $\pi = (354)$). However, as shown by the dash-dotted arrows, $U_{(12)}$ acts as the identity on the rows, while $V_{(12)}$ transposes the columns.	115

Introduction

Quantum query complexity and learning graphs

In quantum computation, one of the main questions that we are interested in is: What is the quantum circuit complexity of a given computational problem? This question is hard to answer, and so we consider an alternative question: What is the quantum query complexity of the problem? For many problems, it is seemingly easier to upper and lower bound the number of times an algorithm requires to access the input rather than to bound the number of elementary quantum operations required by the algorithm. Nonetheless, the study of the quantum query complexity can give us great insights for the quantum circuit complexity. For example, a query-efficient algorithm for SIMON'S PROBLEM [Sim97] helped Shor to develop a time-efficient algorithm for factoring [Sho97]. On the other hand, $\Omega((n/\log n)^{1/5})$ and $\Omega(n^{1/2})$ lower bounds on the (bounded-error) quantum query complexity of the SET EQUALITY [Mid04] and the INDEX ERASURE [AMRR11] problems, respectively, ruled out certain approaches for constructing time-efficient quantum algorithms for the GRAPH ISOMORPHISM problem.

Currently, two main techniques for proving lower bounds on quantum query complexity are the *polynomial method* developed by Beals, Buhrman, Cleve, Mosca, and de Wolf [BBC⁺01], and the *adversary method* originally developed by Ambainis [Amb02] in what later became known as the *positive-weights adversary method*. The adversary method was later strengthened by Høyer, Lee, and Špalek [HLŠ07] by allowing negative weights in the adversary matrix. In recent results [Rei11, LMR⁺11], Lee, Mittal, Reichardt, Špalek, and Szegedy showed that, unlike the polynomial method [Amb03], the general (i.e., strengthened) adversary method can give optimal lower bounds for all function-evaluation problems.

The optimality of the adversary method was proven by, first, expressing the adversary bound as a semidefinite program and, then, showing that each feasible solution of its dual program yields a quantum query algorithm whose query complexity equals the objective value of the program. Soon afterwards, Belovs introduced the computational model of *learning graph* [Bel12d], which can be translated into such a feasible solution and, thus, in turn, into a quantum query algorithm. In a series of works that followed [BL11, Zhu12, LMS12, Bel12c, LMS13], learning graphs and

their generalizations improved upon previously best-known quantum query algorithms for various query problems. The learning graph complexity of a problem is the minimum among complexities of all learning graphs for the problem.

Motivation, results, and relevance

Collision and Set Equality. Once it was proven that the adversary method can always give optimal bounds, a natural question arose: How to use it effectively? A good starting point is to consider problems for which we do not know yet how to construct adversary bounds that would match lower bounds obtained by other methods. Because, if one knows what bound is attainable, one is more likely to succeed at attaining it. For about a decade, ELEMENT DISTINCTNESS and COLLISION were prime examples of such problems. Given an input string $z \in \Sigma^n$, the ELEMENT DISTINCTNESS problem is to decide whether each character of z is unique, and the COLLISION problem is its special case given a promise that each character of z is either unique or appears in z exactly twice.

The quantum query complexity of these two problems is known. Brassard, Høyer, and Tapp first gave an $O(n^{1/3})$ quantum query algorithm for COLLISION [BHT98]. Aaronson and Shi then gave a matching $\Omega(n^{1/3})$ lower bound for COLLISION via the polynomial method, requiring that $|\Sigma| \geq 3n/2$ [AS04] (Aaronson [Aar02] gave the first non-trivial lower bound, $\Omega(n^{1/5})$, which was then improved by Shi [Shi02]). Due to a particular reduction from COLLISION to ELEMENT DISTINCTNESS, their lower bound also implied an $\Omega(n^{2/3})$ lower bound for ELEMENT DISTINCTNESS, requiring that $|\Sigma| = \Omega(n^2)$. Subsequently, Kutin (for COLLISION) and Ambainis (for both) removed these requirements on the alphabet size [Kut05, Amb05]. Finally, Ambainis gave an $O(n^{2/3})$ quantum query algorithm for ELEMENT DISTINCTNESS based on a quantum walk [Amb07], thus improving the best previously known $O(n^{3/4})$ upper bound [BDH⁺05].

The first of these problems “to fall” was ELEMENT DISTINCTNESS: Belovs gave an $\Omega(n^{2/3})$ adversary bound for the problem when $|\Sigma| = \Omega(n^2)$ [Bel12b]. (Due to the *certificate complexity barrier* [Zha05, ŠS06], the positive-weights adversary method fails to give a better lower bound than $\Omega(n^{1/2})$.) As hoped, this new insight on the usage of the adversary method turned out to be very useful, and only a few months later Belovs and Špalek gave a tight $\Omega(n^{k/(k+1)})$ adversary bound for the k -SUM problem [BŠ13], improving over the best previously known lower bound. The k -SUM problem is, given a constant k and assuming that Σ is an additive group, to decide whether there exist k numbers among n that sum up to 0. Similarly to ELEMENT DISTINCTNESS, their lower bound also required that the alphabet size is sufficiently large, in particular, $|\Sigma| = \Omega(n^k)$.

Regarding the COLLISION problem, the hope was that a tight adversary bound for it would help to prove the same lower bound for the closely related SET EQUALITY problem, which is a special case of COLLISION given an extra promise that each character of the first half (and,

thus, the second half) of the input string is unique. Shi conjectured SET EQUALITY to be as hard as COLLISION [Shi02]. Unfortunately, the $\Omega(n^{1/3})$ lower bound for COLLISION obtained via the polynomial method did not generalize to SET EQUALITY. The best known lower bound for SET EQUALITY was given by Midrijānis, who showed an $\Omega((n/\log n)^{1/5})$ lower bound using a combination of the positive-weights adversary and the polynomial methods [Mid04]. Due to the *property testing barrier* [HLŠ07], on its own, the positive-weights adversary fails to give a better lower bound for COLLISION and SET EQUALITY than the trivial $\Omega(1)$.

In this thesis, we construct tight $\Omega(n^{1/3})$ adversary bounds for both COLLISION and SET EQUALITY, assuming that $|\Sigma| = \Omega(n^2)$. This work was done in collaboration with Aleksandrs Belovs, and it appears in Ref. [BR14]. Independently and at about the same time, the $\Omega(n^{1/3})$ lower bound for SET EQUALITY was also proven by Zhandry [Zha13] using machinery from Ref. [Zha12] based on the polynomial method. (Zhandry’s lower bound does not require any assumptions on the alphabet size.) Thus, Shi’s conjecture is resolved affirmatively. The lower bound for the SET EQUALITY problem was used by Aaronson and Ambainis in their proof of the polynomial relation between the randomized and quantum query complexities of partial, permutation-invariant functions [AA11]. By improving the lower bound, we automatically improve the exponent in their result, as explained in their paper.

Interestingly, our adversary constructions for the COLLISION and SET EQUALITY problems are almost identical, suggesting that the adversary method can be easier adopted 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, to the best of our knowledge, our application of the adversary method is the first that supersedes the property testing barrier.

Element Distinctness with small alphabet. When the adversary bound was strengthened by allowing negative weights in the adversary matrix, it was not immediately clear how to take advantage of these negative weights. Unlike with positive weights only, when the weight corresponding to a pair of inputs generally indicated how hard it is to distinguish between the two inputs, it was not clear how to interpret the sign of the weight. One of the first applications of the adversary method that truly exploited negative weights was for the INDEX ERASURE problem [AMRR11]. In this purely quantum problem, given an injective function, one is asked to generate a uniform superposition over its image.

In Ref. [HLŠ07], Høyer, Lee, and Špalek also showed that, without loss of generality, one can assume that the adversary matrix respects symmetries of the problem—this is known as the *automorphism principle*. Ambainis, Magnin, Rötteler, and Roland built the adversary matrix for INDEX ERASURE using the symmetries given by the automorphism principle, in particular, they expressed it as a linear combination of projectors on certain irreducible representations of the symmetric group. From this viewpoint, the difference between the general and the positive-

weighted adversary methods is that, in the former, one can construct the adversary matrix as any real linear combination of these projectors (in the latter, only certain linear combinations are permitted). And one does not even attempt to calculate what are the entries of the adversary matrix, and which of them are positive and which negative. Seeing the adversary matrix as a linear combination of such projectors also highly simplified the evaluation of the spectral norms of the adversary matrix and its entry-wise matrix products with the difference matrices Δ_i , the norms that are essential to the adversary bound.

In this thesis, we develop a similar, representation-theory-inspired construction for the ELEMENT DISTINCTNESS problem with minimal non-trivial alphabet size, that is, $|\Sigma| = n$. This work appears in Ref. [Ros14]. Even though a tight adversary bound for ELEMENT DISTINCTNESS was already given by Belovs, the present result has a potential importance. For example, for lower bounding the quantum query complexity of the k -DISTINCTNESS problem, which, given a constant k , asks to decide whether the input string contains some character at least k times. Belovs' adversary bound for ELEMENT DISTINCTNESS (as well as the adversary bounds for k -SUM, COLLISION, and SET EQUALITY mentioned above) uses the technique of embedding the adversary matrix in a larger matrix, and this technique has certain limitations:

- It requires that a random string in Σ^n is a negative input of the problem with a high probability. This requirement, in turn, imposes restrictions on the size of the alphabet: for example, $|\Sigma| = \Omega(n^2)$ for ELEMENT DISTINCTNESS, COLLISION, and SET EQUALITY, and $|\Sigma| = \Omega(n^k)$ for k -SUM.
- It seems to require that, with a high probability, a random negative input is “hard”. However, the hardest negative inputs for k -DISTINCTNESS, for example, seem to be the ones in which each character appears $k - 1$ times, and a randomly chosen negative input is such only with a minuscule probability. This might be a reason why an $\Omega(n^{2/3})$ adversary bound for k -DISTINCTNESS [Špa13] based on the technique of the embedding does not narrow the gap to the best known upper bound, $O(n^{1-2^{k-2}/(2^k-1)})$ [Bel12c]. (The $\Omega(n^{2/3})$ lower bound was already known previously via the reduction from ELEMENT DISTINCTNESS attributed to Aaronson in Ref. [Amb07].)

Here we construct an adversary bound for ELEMENT DISTINCTNESS in the most general setting, only assuming that it satisfies the symmetries given by the automorphism principle, which is without loss of generality. Due to the optimality of the general adversary method, we know that one can construct a tight adversary bound for k -DISTINCTNESS that satisfies these symmetries, and the hope is that our construction for ELEMENT DISTINCTNESS might give insights in how to do that. We also hope that, due to similarities between ELEMENT DISTINCTNESS and k -SUM, this construction might help to reduce the required alphabet size in the $\Omega(n^{k/(k+1)})$ lower bound for k -SUM.

Other techniques utilizing representation theory. As evidenced by our adversary bounds for COLLISION, SET EQUALITY, and ELEMENT DISTINCTNESS, and the adversary bound for INDEX ERASURE in Ref. [AMRR11], the representation theory of the symmetric group is a powerful toolkit for handling symmetries of a problem when studying its quantum query complexity. Recently, Belovs also used the representation theory of the symmetric group when studying the junta learning problem [Bel14]. In this thesis, we present yet another of its applications when we study the query complexity of the ENHANCED FIND-TWO problem. This work was done in collaboration with Andris Ambainis and Dominique Unruh, and it appears in Ref. [ARU14], where it is explained how ENHANCED FIND-TWO arises in the study of quantum proofs of knowledge.

The ENHANCED FIND-TWO problem is defined as follows. Given n elements such that exactly k of them are marked, the problem is to find two distinct marked elements using the following resources: (1) one initial copy of the uniform superposition over all marked elements, (2) an oracle that reflects across this superposition, and (3) an oracle that tests if an element is marked.

There are two reasons why the adversary method cannot address ENHANCED FIND-TWO. First of all, this is a relational problem, not a function evaluation; namely, every valid input has multiple correct solutions. And, second, the proof of the adversary bound does not address non-standard oracles (in this case, the reflection oracle (2)). Nonetheless, by borrowing some ideas from the proof of the adversary bound and using the representation theory of the symmetric group, we prove that the query complexity of ENHANCED FIND-TWO is $\Theta(\min\{\sqrt{n/k}, \sqrt{k}\})$, the upper bound coming, essentially, from the Grover’s search [Gro96, BBHT98].

Learning graph complexity. There are two general types of learning graphs, non-adaptive and adaptive, the latter being more powerful, yet more complex. Further generalizations of learning graphs have been considered—most notably, in the best known quantum query algorithm for the k -DISTINCTNESS problem [Bel12c]—but they are problem-specific.

The original adversary bound for ELEMENT DISTINCTNESS was inspired by an $\Omega(n^{2/3})$ lower bound on its non-adaptive learning graph complexity [Bel12a]. In a collaboration with Aleksandrs Belovs, we introduce a general method for giving lower bounds on the non-adaptive learning graph complexity. Our work appears in Ref. [BR13a]. Our method is based on expressing the non-adaptive learning graph complexity as a semidefinite program and constructing *dual (non-adaptive) learning graphs*, that is, feasible solutions to the dual semidefinite program.

We use this method to show that the non-adaptive learning graph complexity of the TRIANGLE problem, which is to decide if an n -vertex graph contains a triangle, is at least $\Omega(n^{9/7}/\sqrt{\log n})$. This is almost optimal because, recently before this result, a non-adaptive learning graph for TRIANGLE of complexity $O(n^{9/7})$ was given by Lee, Magniez, and Santha [LMS13]. Thus we prove that, if one wants to improve upon their algorithm for TRIANGLE, one has to look beyond the model of non-adaptive learning graph. Very recently this has been done by Le Gall [LG14], who gave an $O(n^{5/4})$ quantum query algorithm for TRIANGLE based on quantum walks. (The

previous two improvements— $O(n^{35/27})$ [Bel12d] and $O(n^{9/7})$ [LMS13]—were obtained using non-adaptive learning graphs. An $O(n^{9/7})$ quantum query algorithm based on *nested quantum walks* was also later discovered by Jeffery, Kothari, and Magniez [JKM13].)

In our work, we also show that non-adaptive learning graphs are exactly as powerful as quantum query algorithms that only look for complete 1-certificates in the input, disregarding any additional structure of the problem. To formalize this statement, we introduce the *certificate structure* of a problem, which describes possible positions of 1-certificates in an input of the problem. The non-adaptive learning graph complexity of the problem depends only on its certificate structure. For every certificate structure, we construct a problem having this certificate structure whose non-adaptive learning graph and quantum query complexities are the same. (For lower bounding its quantum query complexity, we construct an optimal adversary matrix based on a dual learning graph.) For a special case of certificate structures generated by certificates of bounded size, one can choose this problem to be the corresponding CERTIFICATE-SUM problem. This generalizes the adversary bound for the k -SUM problem from Ref. [BS13].

In this thesis we also present work on adaptive learning graph complexity. This work was done in collaboration with Troy Lee, Miklos Santha, and Aarthi Sundaram, and it is currently unpublished. As for non-adaptive learning graphs before, we express adaptive learning graph complexity as a semidefinite program, and consider its dual program. This way, we show that the adaptive learning graph complexity of the k -DISTINCTNESS problem is $\Omega(n^{1-2^{k-2}/(2^k-1)})$. Since there is a way how to construct adversary bounds from non-adaptive dual learning graphs, the hope is that something similar can be done in the adaptive case and this lower bound on the adaptive learning graph complexity of k -DISTINCTNESS might help to improve lower bounds on its quantum query complexity.

Organization of the thesis

This thesis is divided into two parts: Part I, in which we introduce the necessary preliminaries, and Part II, in which we describe the original results. Some of the more technical proofs are also left to appendices.

Part I – Preliminaries. In Chapter 1 we introduce basic mathematical concepts and notation that we use in this thesis. Here we also introduce basics of the representation theory of the symmetric group, and obtain some results necessary in the later chapters. In Chapter 2 we introduce quantum query algorithms and describe the process of their symmetrization, as well as the basic idea behind the adversary method. Here we also introduce learning graphs.

Part II – Results. In Chapter 3 we study the non-adaptive and adaptive learning graph complexities. We first express these complexities as semidefinite programs, then we obtain their dual

programs, and then we use the dual programs to obtain lower bounds on the learning graph complexity of various problems. In Chapter 4 we construct adversary bounds based on the technique of embedding adversary matrices into larger matrices. We first recall Belovs' construction for ELEMENT DISTINCTNESS, which pioneered this approach, in Section 4.1. Then, in Section 4.2, we construct adversary bounds for COLLISION and SET EQUALITY and, in Section 4.3, we construct adversary bounds for certificate structures based on dual non-adaptive learning graphs. In Chapter 5 we construct an adversary bound for ELEMENT DISTINCTNESS with minimal non-trivial alphabet size. Finally, in Chapter 6 we prove tight bounds for the ENHANCED FIND-TWO problem.

Part I

Preliminaries

Chapter 1

Mathematical preliminaries

Let \mathbb{N} , \mathbb{Z} , \mathbb{R} , and \mathbb{C} denote the set of positive integers, integers, real numbers, and complex numbers, respectively. For a complex number a , let $|a|$ and \bar{a} denote, respectively, its absolute value and its complex conjugate. For integers ℓ and $m \geq \ell$, let $[\ell..m]$ denote the set $\{\ell, \ell + 1, \dots, m\}$, and let $[m]$ be a shorthand for $[1..m]$ (sometimes we still use the latter notation). The *power set* $2^{[n]}$ is the set of all 2^n subsets of $[n]$.

Let \sqcup denote the disjoint union of sets, which we associate with the concept of a *decomposition* of a set. Throughout the thesis, we use “:=” instead of “=” in equations that define or assign a value to the term on the left hand side of “:=”. We use “ \cong ” to denote isomorphism.

We call every finite nonempty set denoted by Σ an *alphabet*. Given a string x over Σ , we use both x_i and $x[[i]]$ to denote its i -th entry. For a string $x \in \Sigma^n$ and $S \subseteq [n]$, let $x_S \in \Sigma^{|S|}$ denote the *projection* (or *restriction*) of x on S , i.e., the string $(x[[i_1]], \dots, x[[i_\ell]])$ indexed by the elements i_1, \dots, i_ℓ of S . Suppose Σ is equipped with a total ordering ‘ $<$ ’. For two strings x and y over Σ , we say that x *comes before* y in the *lexicographical order* and we write $x < y$ if, for i being the first position at which x and y differ, $x[[i]] < y[[i]]$.

1.1 Linear algebra

We assume that the reader is familiar with basic concepts of linear algebra like finite-dimensional Hilbert spaces, eigenvalues and eigenvectors, singular value decompositions, direct sums and tensor products, and normal, unitary, Hermitian, and positive semidefinite operators just to name a few. (For a reference, see, for example, [Bha97, Chapter I]). The main purpose of this section is to introduce the notation that we will be using throughout the thesis.

Every *Hilbert space* in this thesis is assumed to be finite-dimensional. Let \mathbb{R}^d and \mathbb{C}^d denote the d -dimensional real and complex Hilbert spaces, respectively. We think of elements $v \in \mathbb{C}^n$

as column vectors. Unless stated otherwise, the Hilbert spaces we consider are thought to be complex.

Suppose \mathcal{X} and \mathcal{Y} are Hilbert spaces. Let $L(\mathcal{X}, \mathcal{Y})$ be the space of *linear maps* from \mathcal{X} to \mathcal{Y} , and let $L(\mathcal{X}) := L(\mathcal{X}, \mathcal{X})$ be the space of *linear operators* on \mathcal{X} . Let $GL(\mathcal{X})$ and $U(\mathcal{X})$ be, respectively, the set of invertible and unitary linear operators on \mathcal{X} . $GL(\mathcal{X})$ is called the *general linear group* of \mathcal{X} , and its subgroup $U(\mathcal{X})$ is called the *unitary group* of \mathcal{X} . Let $\text{Tr}: L(\mathcal{X}) \rightarrow \mathbb{C}$ denote the trace and $\text{Tr}_{\mathcal{Y}}: L(\mathcal{X} \otimes \mathcal{Y}) \rightarrow L(\mathcal{X})$ denote the partial trace. The equality $\text{Tr}(AB) = \text{Tr}(BA)$ is called the *cyclic property* of the trace.

Given a vector $v \in \mathcal{X}$, let $\|v\|$ denote its *Euclidean norm*. And, given a linear map $A \in L(\mathcal{X}, \mathcal{Y})$, let $\|A\| := \max_{v \in \mathcal{X}} \|Av\|/\|v\|$ denote its *spectral norm*, which equals the largest singular value of A . The spectral norm is the only matrix norm considered in this thesis. Let v^* and A^* denote the conjugate transpose of v and A respectively. We think of v^* as a row vector. Let $\langle v, w \rangle := v^*w$ denote the inner product of two vectors u and v . The *Cauchy–Schwarz inequality* states that $|\langle v, w \rangle| \leq \|v\| \cdot \|w\|$. Two maps $A, B \in L(\mathcal{X}, \mathcal{Y})$ are *orthogonal* if both $A^*B = 0$ and $AB^* = 0$, and we write $A \perp B$ (note: our requirement for orthogonality is stronger than $\text{Tr}(A^*B) = 0$).

Let $A \succeq 0$ denote that an operator A is positive semidefinite. Let $A \succeq B$ stand for $A - B \succeq 0$; one calls \succeq the *semidefinite inequality*. For a normal operator A , the *support* of A is the space spanned by all the eigenvectors of A corresponding to non-zero eigenvalues.

Let $\dim \mathcal{X}$ denote the dimension of a Hilbert space \mathcal{X} , and suppose $d = \dim \mathcal{X}$. We use both $\mathbb{I}_{\mathcal{X}}$ and \mathbb{I}_d denote the identity operator on \mathcal{X} . Given a subspace $\mathcal{Y} \subseteq \mathcal{X}$, let $\Pi_{\mathcal{Y}} \in L(\mathcal{X})$ denote the projector on \mathcal{Y} . Note that $\Pi_{\mathcal{Y}}$ is essentially equal to $\mathbb{I}_{\mathcal{Y}}$, but we use the notation “ Π ” instead of “ \mathbb{I} ” when we want to stress that we are considering a subspace of some larger space.

Let $\mathcal{X} \oplus \mathcal{Y}$ and $\mathcal{X} \otimes \mathcal{Y}$ denote the direct sum and the tensor product of spaces \mathcal{X} and \mathcal{Y} , respectively. Suppose \mathcal{X} and \mathcal{Y} are such that $\mathcal{Y} \subseteq \mathcal{X}$. Let $\mathcal{X} \ominus \mathcal{Y}$ denote the *orthogonal complement* of \mathcal{Y} in \mathcal{X} , so that $\mathcal{Y} \oplus (\mathcal{X} \ominus \mathcal{Y}) = \mathcal{X}$. Given an operator $A \in L(\mathcal{X})$ such that $Av \in \mathcal{Y}$ for all $v \in \mathcal{Y}$, one says that \mathcal{Y} is *stable* under A . Given that \mathcal{Y} is stable under A , define $A|_{\mathcal{Y}} := A\Pi_{\mathcal{Y}} \in L(\mathcal{Y})$, which we call a *restriction* or a *reduction* of A to \mathcal{Y} .

With every finite set X , we associate an $|X|$ -dimensional Hilbert space denoted by \mathbb{C}^X , where, with every element $x \in X$, we associate a unit vector $\mathbf{x} \in \mathbb{C}^X$ so that $\{\mathbf{x}: x \in X\}$ is an orthonormal basis of \mathbb{C}^X , called the *standard basis* of \mathbb{C}^X . Note: we use the bold font for vectors of the standard basis. We call $\{\mathbf{xy}^*: (x, y) \in X \times Y\}$ the standard basis of $L(\mathbb{C}^Y, \mathbb{C}^X)$.

Given orthonormal bases of \mathcal{X} and \mathcal{Y} , we also think of a map $A \in L(\mathcal{Y}, \mathcal{X})$ as a $\dim \mathcal{X} \times \dim \mathcal{Y}$ matrix that one obtains by expressing A in these bases. If $\mathcal{X} := \mathbb{C}^X$ and $\mathcal{Y} := \mathbb{C}^Y$ for some finite sets X and Y and A is expressed in the standard basis, we may also say that A is an $X \times Y$ matrix, and the rows and columns of A are *labeled* by X and Y , respectively. For a matrix A , we denote its (x, y) -th entry by $A[x, y]$. Note that, if A is expressed in the standard basis, then

$A[x, y] = \mathbf{x}^* A \mathbf{y}$. Similarly, given a vector v expressed in some orthonormal basis, we denote its x -th entry by $v[x]$ or v_x (which equals $\mathbf{x}^* v$ in the standard basis).

Let \mathbb{J}_d denote the $d \times d$ *all-ones matrix*. Let \circ be the Hadamard (i.e., entrywise) matrix product, and note that the all-ones matrices act as the identity for this product. Let $\vec{1}_\ell$ denote the all-ones vector of length ℓ . The Cauchy–Schwarz inequality applied to $\vec{1}_\ell$ and (v_1, \dots, v_ℓ) implies that $(\sum_{i=1}^\ell v_i)^2 \leq \ell \cdot \sum_{i=1}^\ell v_i^2$.

1.2 Semidefinite programming

The following definition is a combination of definitions from [Lov95], [Wat11], and [WSV00]. Let “SDP” be the abbreviation for “semidefinite program”.

Definition 1.1. A general *semidefinite program* (in the inequality form with multiple linear matrix inequality constraints) is an optimization problem in the form

$$\text{minimize} \quad \sum_{j=1}^m a_j x_j \tag{1.1a}$$

$$\text{subject to} \quad \sum_{j=1}^m Q_{i,j} x_j \succeq B_i \quad \text{for all } i \in [n]; \tag{1.1b}$$

$$x_j \geq 0 \quad \text{for all } j \in [m], \tag{1.1c}$$

where all B_i and $Q_{i,j}$ are Hermitian matrices of the same dimension, and all a_j are real numbers. One calls (1.1) the *primal SDP* (or simply, the *primal*, for short). Its *dual* semidefinite program is

$$\text{maximize} \quad \sum_{i=1}^n \langle B_i, Y_i \rangle \tag{1.2a}$$

$$\text{subject to} \quad \sum_{i=1}^n \langle Q_{i,j}, Y_i \rangle \leq a_j \quad \text{for all } j \in [m]; \tag{1.2b}$$

$$Y_i \succeq 0 \quad \text{for all } i \in [n]. \tag{1.2c}$$

If for any i there is equality instead of inequality in (1.1b), then Y_i is only required to be Hermitian in (1.2c). Similarly, if for any j there is equality instead of semidefinite inequality in (1.2b), then in (1.1c) it is only required that $x_j \in \mathbb{R}$. A solution of the primal (1.1) is *strictly feasible* if it strictly satisfies all inequalities in (1.1b) and (1.1c), and similarly for the dual. *Slater’s condition* implies that, if there are strictly feasible solutions for both the primal (1.1) and the dual (1.2), then the optimal values of both (1.1) and (1.2) are equivalent (a property known as *strong duality*) and they are attained by some feasible solutions of (1.1) and (1.2).

1.3 Linear group representations

The representation theory of the symmetric group is essential to the results presented in Chapters 4, 5, and 6. In this section, we present basics of the representation theory, and we later focus on the case of the symmetric group in Section 1.4. For more background on the representation theory of finite groups, the reader may refer to [Ser77], and, for group theory, to [Rot95].

1.3.1 Basics of group theory and the symmetric group

Conjugacy classes, transversals, and direct products. Let ε denote the *identity element* of a group G . Two elements g, g' of the group G are said to be *conjugate*, if there exists $h \in G$ such that $g' = hgh^{-1}$. Conjugacy is an equivalence relation, and therefore it divides G into equivalence classes, which we call *conjugacy classes* of G .

For a finite group G , let $|G|$ denote its *order*, that is, the number of elements in G . Let $H \leq G$ denote that H is a subgroup of G . Given $H \leq G$, let $G/H := \{gH : g \in G\}$ denote the set of *left cosets* of H in G .¹ Note that $|G/H| = |G|/|H|$. For every pair of distinct left cosets $g_1H \neq g_2H$ and every $g' \in G$, we have $g'g_1H \neq g'g_2H$.

An element $g' \in gH$ is called a *representative* of the coset gH , and we have $g'H = gH$. A set containing exactly one representative of each coset in G/H is called a *transversal* of the left cosets of H in G , and we denote it by $\text{Rep}(G/H)$ (note: this set is not unique; we have $|H|^{|G/H|}$ distinct transversals).

The *direct product* $G \times H$ of two groups G and H is a group itself. Let ε_H be the identity element of H , and note that the groups G and $G \times \{\varepsilon_H\} \leq G \times H$ are isomorphic. Thus, by abuse of notation, we often write G instead of $G \times \{\varepsilon_H\}$.

Symmetric group. Let \mathbb{S}_L denote the symmetric group of a finite set L , that is, the group with the permutations of L as elements, and composition as the operation. If m is a positive integer, \mathbb{S}_m denotes the isomorphism class of all symmetric groups \mathbb{S}_L with $|L| = m$.

A permutation $\pi \in \mathbb{S}_L$ is called a *cycle* if there exists a disjoint decomposition $L = L_1 \sqcup L_2$ satisfying the following: for all $\ell, \ell' \in L_1$, there exists $k \in \mathbb{N}$ such that $\pi^k(\ell) = \ell'$ and, for all $\ell \in L_2$, $\pi(\ell) = \ell$. We call L_1 the *elements* of the cycle, $|L_1|$ the *length* of the cycle, and we denote this cycle by a tuple $(\ell, \pi(\ell), \pi^2(\ell), \dots, \pi^{|L_1|-1}(\ell))$, where $\ell \in L_1$; we may omit commas in the tuple when convenient. We call two cycles *non-overlapping* if they have no elements in common, and such cycles commute.

¹If H is a *normal* subgroup of G (see [Rot95]), then G/H is a group.

We can write each permutation a product of non-overlapping cycles, which we call its *cycle factors*. And we can choose whether to write cycles of length one or not. For example, define $\pi \in \mathbb{S}_6$ as

$$1 \mapsto 5, \quad 2 \mapsto 6, \quad 3 \mapsto 3, \quad 4 \mapsto 1, \quad 5 \mapsto 4, \quad 6 \mapsto 2.$$

We can write $\pi = (154)(26)$ or $\pi = (26)(3)(415)$.

The study of symmetric group is closely related to the concept of a partition. A *partition* λ of an integer $m \in \mathbb{N}$ is a non-increasing list $(\lambda_1, \dots, \lambda_k)$ of positive integers satisfying $\lambda_1 + \dots + \lambda_k = m$ (we may also write $\lambda(i)$ instead of λ_i). We denote this relation by $\lambda \vdash m$, or write $m = |\lambda|$. 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 a permutation in \mathbb{S}_m , its *cycle partition* is the list of lengths of its cycle factors, from longest to shortest, including all the cycle factors of length one. Two permutations in \mathbb{S}_m belong to the same conjugacy class if and only if they have the same cycle partition. Hence, the number of conjugacy classes of \mathbb{S}_m equals the number of distinct partitions of m .

1.3.2 Basic definitions

Let \mathcal{X} be a Hilbert space. A *linear representation* of a group G on \mathcal{X} is a group homomorphism ρ from G to the general linear group $\text{GL}(\mathcal{X})$. We often write ρ_g instead of $\rho(g)$, call it a *representation operator*, and, by definition, ρ satisfies $\rho_{gg'} = \rho_g \rho_{g'}$ for all $g, g' \in G$. While, technically, ρ is a (linear) representation and ρ_g is an operator, we may also refer to ρ as an operator and ρ_g as a representation; their use should be clear from the context. When the map ρ is given, we also call \mathcal{X} a representation of G by abuse of terminology. We also refer to a representation $\rho: G \rightarrow \text{GL}(\mathcal{X})$ as an *action* of G on \mathcal{X} (which is also an abuse of terminology).

For finite groups, we extend the concept of representations to group algebras by linearity. The *group algebra* $\mathbb{C}G$ is the vector space \mathbb{C}^G with the multiplication law in G extended to $\mathbb{C}G$ by linearity. For example, the following is a multiplication of two elements in $\mathbb{C}\mathbb{S}_4$:

$$\begin{aligned} [6\varepsilon - 3(\mathbf{243})] \cdot [2(\mathbf{13})(\mathbf{24}) + 4(\mathbf{143})] &= 12(\mathbf{13})(\mathbf{24}) + 24(\mathbf{143}) - 6(\mathbf{123}) - 12(\mathbf{13})(\mathbf{24}) \\ &= 24(\mathbf{143}) - 6(\mathbf{123}), \end{aligned}$$

where, for clarity, we have displayed the elements of \mathbb{C} in italic.

The *dimension* of a representation $\rho: G \rightarrow \text{GL}(\mathcal{X})$ is the dimension of the space \mathcal{X} , and we denote it $\dim \rho := \dim \mathcal{X}$. Two representations $\rho: G \rightarrow \text{GL}(\mathcal{X})$ and $\rho': G \rightarrow \text{GL}(\mathcal{X}')$ are *isomorphic* if there exists a linear isomorphism $\Xi \in \text{L}(\mathcal{X}', \mathcal{X})$ between \mathcal{X}' and \mathcal{X} that satisfies $\rho_g \Xi = \Xi \rho'_g$ for all $g \in G$, and we write $\rho \cong \rho'$ and $\mathcal{X} \cong \mathcal{X}'$. An *isomorphism class* is the set of all

representations that are isomorphic to some given representation, and, clearly, all representations in such a class have the same dimension.

A representation $G \rightarrow \mathbf{U}(\mathcal{X})$ is called *unitary*. Every representation of a finite group G is isomorphic to some unitary representation, and, from now on, let us only consider unitary representations. Aside from finite groups, we also consider certain unitary representations of the unitary group $\mathbf{U}(\mathcal{X})$ itself.

1.3.3 Subrepresentations

Given a group G and a representation ρ of G on \mathcal{X} , a subspace \mathcal{Y} of \mathcal{X} is called *stable* under ρ (or under G) if \mathcal{Y} is stable under ρ_g for all $g \in G$. Suppose \mathcal{Y} is stable. Then \mathcal{Y} is also a representation of G , that is, $\rho_g|_{\mathcal{Y}} := \rho_g \Pi_{\mathcal{Y}}$ is a representation of G . Note that ρ_g and $\Pi_{\mathcal{Y}}$ commute because ρ_g is unitary (and, thus, normal). The representation $\rho|_{\mathcal{Y}}$ is called the *reduction* of ρ to \mathcal{Y} , and one also says that $\rho|_{\mathcal{Y}}$ is a *subrepresentation* of ρ .

Suppose \mathcal{X} is a representation of G and $\mathcal{Y} \subset \mathcal{X}$ is stable under G . Then $\mathcal{Y}^{\perp} := \mathcal{X} \ominus \mathcal{Y}$ is also stable under G . Hence, one can *decompose* a representation \mathcal{X} as a direct sum of two representations: $\mathcal{X} = \mathcal{Y} \oplus \mathcal{Y}^{\perp}$. Following the same procedure, if \mathcal{Y} or \mathcal{Y}^{\perp} contain *proper* stable subspaces, one can recursively further decompose \mathcal{X} as a direct sum of more than two components. This decomposition procedure can go on until all spaces in the direct sum are irreducible.

A representation \mathcal{Y} is called *irreducible* (or just *irrep*, for short) if it is not 0 and it does not contain stable subspaces other than \mathcal{Y} and 0, that is, if it does not contain proper subrepresentations. A representation that is not irreducible is called *reducible*. An essential basic result in the representation theory is the following

Lemma 1.2 (Schur's Lemma). *Suppose that ρ and ρ' are two irreducible representations of a group G on \mathcal{X} and \mathcal{X}' , respectively, and that a linear map $M \in \mathbf{L}(\mathcal{X}', \mathcal{X})$ satisfies $\rho_g M = M \rho'_g$ for all $g \in G$ (that is, M is a homomorphism). If $\mathcal{X} \not\cong \mathcal{X}'$, then $M = 0$. And, if $\mathcal{X} \cong \mathcal{X}'$, then M is an isomorphism and it is unique up to a scalar multiplier.*

Let us present some very useful consequences of Schur's lemma.

Corollary 1.3. *Suppose we are given two representations $\sigma: G \mapsto \mathbf{U}(\mathcal{X})$ and $\sigma': G \mapsto \mathbf{U}(\mathcal{X}')$ and a linear map $A \in \mathbf{L}(\mathcal{X}', \mathcal{X})$ that satisfies $\sigma_g A = A \sigma'_g$ for all $g \in G$. Also, let $\rho: G \mapsto \mathbf{U}(\mathcal{Y})$ and $\rho': G \mapsto \mathbf{U}(\mathcal{Y}')$ be irreducible and non-isomorphic subrepresentations of σ and σ' , respectively. Then $\Pi_{\mathcal{Y}} A \Pi_{\mathcal{Y}'} = 0$.*

Proof. Recall that, for all $g \in G$, $\rho_g = \sigma_g \Pi_{\mathcal{Y}} = \Pi_{\mathcal{Y}} \sigma_g$ and $\rho'_g = \sigma'_g \Pi_{\mathcal{Y}'} = \Pi_{\mathcal{Y}'} \sigma'_g$. Hence

$$\rho_g \Pi_{\mathcal{Y}} A \Pi_{\mathcal{Y}'} = \Pi_{\mathcal{Y}} \sigma_g A \Pi_{\mathcal{Y}'} = \Pi_{\mathcal{Y}} A \sigma'_g \Pi_{\mathcal{Y}'} = \Pi_{\mathcal{Y}} A \Pi_{\mathcal{Y}'} \rho'_g$$

for all $g \in G$, and Schur's lemma implies the result. \square

If we choose $\mathcal{X} = \mathcal{X}'$ and $A = \mathbb{I}_{\mathcal{X}}$ in Corollary 1.3, we get

Corollary 1.4. *Given a representation \mathcal{X} and two non-isomorphic irreducible subrepresentations $\mathcal{Y}, \mathcal{Y}' \subset \mathcal{X}$, the spaces \mathcal{Y} and \mathcal{Y}' are orthogonal.*

This result is very useful because of the following. Suppose we have decomposed a representation \mathcal{X} of a group G as a direct sum of irreducible representations (for example, by using the recursive procedure described above). Then we can write

$$\mathcal{X} = \bigoplus_{\rho} \bigoplus_{j=1}^{m_{\rho}} \mathcal{Y}_{\rho,j}, \quad (1.3)$$

where ρ runs through representatives of all isomorphism classes of irreps of G (we allow $m_{\rho} = 0$) and $\mathcal{Y}_{\rho,j}$ is an irrep isomorphic to ρ . In general, this decomposition is not necessarily unique. However, Corollary 1.4 implies that, for all ρ , the space $\mathcal{Y}_{\rho} := \bigoplus_{j=1}^{m_{\rho}} \mathcal{Y}_{\rho,j}$ is unique, as is $m_{\rho} = \dim \mathcal{Y}_{\rho} / \dim \rho$.

One calls $\mathcal{X} = \bigoplus_{\rho} \mathcal{Y}_{\rho}$ the *canonical decomposition* of \mathcal{X} . The space \mathcal{Y}_{ρ} is called the *isotypical subspace* of \mathcal{X} corresponding to ρ , or simply ρ -isotypical subspace of \mathcal{X} . We provide a method for computing the projector on this subspace in Theorem 1.9. One says that \mathcal{X} *contains m_{ρ} instances* (or *copies*) of ρ or that ρ *appears* (or *occurs*, or *is present*) in \mathcal{X} with *multiplicity* m_{ρ} (or, simply, m_{ρ} times). The representation is called *multiplicity-free*, if it contains each irrep at most once (i.e., $m_{\rho} \in \{0, 1\}$ for all ρ). The decomposition (1.3) is unique if and only if \mathcal{X} is multiplicity-free.

1.3.4 Transporter bases of isotypical subspaces

Now let us explore how many degrees of freedom one has in decomposing an isotypical subspace as a direct sum of irreps. First note that, given an irrep $\rho: G \mapsto \mathbf{U}(\mathcal{Y})$, $\Pi_{\mathcal{Y}}$ is an automorphism from ρ to itself. Thus, Schur's lemma implies that every such automorphism is proportional to $\Pi_{\mathcal{Y}}$.

Claim 1.5. *Suppose \mathcal{Y} and \mathcal{Y}' are two isomorphic irreps, and $\Xi \in \mathbf{L}(\mathcal{Y}', \mathcal{Y})$ is a non-zero isomorphism between them. Then all the $\dim \mathcal{Y}$ singular values of Ξ are equal, namely, $\Xi \Xi^* \propto \Pi_{\mathcal{Y}}$ and $\Xi^* \Xi \propto \Pi_{\mathcal{Y}'}$.*

Proof. Let $\rho: G \rightarrow \mathbf{U}(\mathcal{Y})$ and $\rho': G \rightarrow \mathbf{U}(\mathcal{Y}')$ be the irreps in question. For every $g \in G$, we have $\rho_g^* = \rho_g^{-1} = \rho_{g^{-1}}$ and the same for ρ'_g . Hence, $\rho_g \Xi = \Xi \rho'_g$ implies $\Xi^* \rho_{g^{-1}} = \rho'_{g^{-1}} \Xi^*$. By multiplying the corresponding sides (left and right) of the two equalities, we get $\rho_g \Xi \Xi^* \rho_{g^{-1}} = \Xi \Xi^*$. Since this holds for all $g \in G$, $\Xi \Xi^*$ is an automorphism of ρ , and hence it is proportional to $\Pi_{\mathcal{Y}}$. The proof for $\Xi^* \Xi$ is equivalent. \square

The following claim is from [AMRR11], where it was proved differently.

Claim 1.6. *Suppose \mathcal{Y} , \mathcal{Y}' , and \mathcal{Y}'' are isomorphic irreps. Then*

$$\mathrm{Tr}(\Pi_{\mathcal{Y}}\Pi_{\mathcal{Y}'}\Pi_{\mathcal{Y}}\Pi_{\mathcal{Y}''}) = \mathrm{Tr}(\Pi_{\mathcal{Y}}\Pi_{\mathcal{Y}'})\mathrm{Tr}(\Pi_{\mathcal{Y}}\Pi_{\mathcal{Y}''})/\dim \mathcal{Y}.$$

Proof. Let $\mathcal{Y}''' \in \{\mathcal{Y}', \mathcal{Y}''\}$. We use $\Pi_{\mathcal{Y}} = \Pi_{\mathcal{Y}}^2$ and the cyclic property of the trace:

$$\mathrm{Tr}(\Pi_{\mathcal{Y}}\Pi_{\mathcal{Y}'}\Pi_{\mathcal{Y}}\Pi_{\mathcal{Y}''}) = \mathrm{Tr}(\Pi_{\mathcal{Y}}\Pi_{\mathcal{Y}'}\Pi_{\mathcal{Y}} \cdot \Pi_{\mathcal{Y}}\Pi_{\mathcal{Y}''}\Pi_{\mathcal{Y}}), \quad \mathrm{Tr}(\Pi_{\mathcal{Y}}\Pi_{\mathcal{Y}''}) = \mathrm{Tr}(\Pi_{\mathcal{Y}}\Pi_{\mathcal{Y}''}\Pi_{\mathcal{Y}}).$$

Notice that $\Pi_{\mathcal{Y}}\Pi_{\mathcal{Y}''}\Pi_{\mathcal{Y}}$ is an automorphism on \mathcal{Y} , and therefore it is proportional to $\Pi_{\mathcal{Y}}$. The coefficient of this proportionality is $\mathrm{Tr}(\Pi_{\mathcal{Y}}\Pi_{\mathcal{Y}''})/\dim \mathcal{Y}$. \square

Now suppose \mathcal{X} is an isotypical subspace corresponding to some irrep ρ of G . For two irreps $\mathcal{Y}, \mathcal{Y}' \subseteq \mathcal{X}$, we define their *overlap* as $\mathrm{Tr}(\Pi_{\mathcal{Y}}\Pi_{\mathcal{Y}'})/\dim \rho \in [0, 1]$. Let \mathcal{X} contain m instances of the irrep ρ , and suppose we are given a fixed decomposition $\mathcal{X} = \bigoplus_{i=1}^m \mathcal{Y}_i$ of \mathcal{X} into irreps (each \mathcal{Y}_i is isomorphic to ρ). Let $\rho_i(g) \in \mathbf{U}(\mathcal{Y}_i)$ denote the representation operator corresponding to \mathcal{Y}_i and $g \in G$, and let $\sigma(g) := \bigoplus_{i=1}^m \rho_i(g)$ denote the representation operator corresponding to \mathcal{X} and $g \in G$.

For $i, j \in [m]$, let us call an isomorphism $\Xi_{j \leftarrow i} \in \mathbf{L}(\mathcal{Y}_i, \mathcal{Y}_j)$ a *transporter* from the irrep \mathcal{Y}_i to the irrep \mathcal{Y}_j if $\|\Xi_{j \leftarrow i}\| = 1$, that is, if all its singular values are 1 (they are the same by Claim 1.5). This transporter is unique up to a scalar multiple on the unit circle in the complex plane; we call a scalar on the unit circle a *global phase*. When we consider real Hilbert spaces, the transporters are unique up to a global phase ± 1 . We use the term ‘transporter’ only to refer to norm-one isomorphisms between *orthogonal* irreps.

We call $\{\Xi_{j \leftarrow i} : i, j \in [m]\}$ a *basis of transporters*, where $\Xi_{j \leftarrow i}$ is a transporter from \mathcal{Y}_i to \mathcal{Y}_j , if it satisfies the composition $\Xi_{k \leftarrow j}\Xi_{j \leftarrow i} = \Xi_{k \leftarrow i}$ and the inversion $(\Xi_{j \leftarrow i})^* = \Xi_{i \leftarrow j}$ for all $i, j, k \in [m]$. These two conditions together imply that $\Xi_{i \leftarrow i} = \Pi_{\mathcal{Y}_i}$, and note that $\Xi_{k \leftarrow j}\Xi_{j' \leftarrow i} = 0$ whenever $j \neq j'$. Also, for every transporter $\Xi_{i \leftarrow j}$, let us denote

$$\Xi_{i \leftrightarrow j} := \Xi_{i \leftarrow j} + (\Xi_{i \leftarrow j})^* = \Xi_{i \leftarrow j} + \Xi_{j \leftarrow i},$$

which is a unitary operator on $\mathcal{Y}_i \oplus \mathcal{Y}_j$ due to the orthogonality of \mathcal{Y}_i and \mathcal{Y}_j . A basis of transporters always exists. It can be obtained, for example, by first arbitrarily choosing global phases of transporters $\Xi_{i+1 \leftarrow i}$ for $i \in [m-1]$, and then using the composition and the inversion to obtain other transporters.

Claim 1.7. *For any operator $A \in \mathbf{L}(\mathcal{X})$ that satisfies $\sigma(g)A = A\sigma(g)$ for all $g \in G$, we have $A = \sum_{i,j=1}^m a_{i,j}\Xi_{i \leftarrow j}$ for some scalars $a_{i,j}$, where the Ξ 's are a basis of transporters. And we have $a_{i,j} = \mathrm{Tr}(A\Xi_{j \leftarrow i})/\dim \rho$.*

Proof. The proof of the first statement is equivalent to the proof of Corollary 1.3. The second statement then follows from $A\Xi_{j \leftarrow i} = \sum_{k=1}^m a_{k,j}\Xi_{k \leftarrow i}$ (due to the composition property) and $\mathrm{Tr}(\Xi_{k \leftarrow i}) = \delta_{i,k} \dim \rho$. \square

For any vector $\gamma = (\gamma_1, \dots, \gamma_m)$, let

$$\Pi_{(\gamma)} := \sum_{i,j=1}^m \gamma_i \bar{\gamma}_j \Xi_{i \leftarrow j}. \quad (1.4)$$

We have

$$\Pi_{(\gamma)}^* = \Pi_{(\gamma)}, \quad \Pi_{(\gamma)}^2 = \|\gamma\|^2 \Pi_{(\gamma)}, \quad \text{Tr}(\Pi_{(\gamma)}) = \|\gamma\|^2 \dim \rho, \quad \text{and} \quad \text{rank} \Pi_{(\gamma)} = \dim \rho$$

whenever $\gamma \neq 0$. We also have $\sigma(g)\Pi_{(\gamma)} = \Pi_{(\gamma)}\sigma(g)$ for all $g \in G$. Hence we have the following.

Claim 1.8. *Let $\mathcal{Y}' \subseteq \mathcal{X}$ be an irrep isomorphic to ρ . There exists a vector γ such that $\Pi_{(\gamma)} = \Pi_{\mathcal{Y}'}$. The vector γ has unit norm and it is unique up to a global phase. The converse also holds: for any unit vector γ , $\Pi_{(\gamma)}$ is a projector on an irrep isomorphic to ρ .*

1.3.5 Character theory

Let us introduce the basics of character theory. The *character* of a linear representation $\rho: G \rightarrow \text{U}(\mathcal{X})$ is the function

$$\chi_\rho: G \rightarrow \mathbb{C}: g \mapsto \text{Tr}(\rho(g)).$$

The characters of two isomorphic representations are the same. Indeed, for an isomorphism Ξ , we have $\text{Tr}(\Xi^{-1}\rho(g)\Xi) = \text{Tr}(\rho(g))$ due to the cyclic property of the trace. It is known that the converse also holds: any two representations that have the same character are isomorphic.

Given a representation ρ , its character takes the same value on every element of a given conjugacy class:

$$\chi_\rho(hgh^{-1}) = \text{Tr}(\rho(h)\rho(g)\rho(h)^{-1}) = \text{Tr}(\rho(g)) = \chi_\rho(g)$$

for all $g, h \in G$. Note that $\chi_\rho(g^{-1}) = \text{Tr}(\rho(g)^*) = \bar{\chi}_\rho(g)$. Therefore, for the symmetric group (our main group of interest), $\pi \in \mathbb{S}_n$ and π^{-1} belong to the same conjugacy class, and therefore $\chi_\rho(\pi) = \bar{\chi}_\rho(\pi)$.

Suppose G is finite. The *character orthogonality relations*, which we do not state in this thesis, imply that the number of irreducible representations of G (up to isomorphism) equals the number of conjugacy classes of G . We also have the following

Theorem 1.9. *[Ser77, Sec. 2.6] Let $\sigma: G \rightarrow \text{U}(\mathcal{X})$ be a representation of G and let ρ be an irrep of G . The projector on the ρ -isotypical subspace of \mathcal{X} is*

$$\frac{\dim \rho}{|G|} \sum_{g \in G} \bar{\chi}_\rho(g) \sigma(g). \quad (1.5)$$

1.3.6 Group action and regular representations

Suppose we are given a finite set A and a group G . The *left group action* of G on A is a function $\phi: G \times A \rightarrow A$ that satisfies $\phi(\varepsilon, a) = a$ and $\phi(gh, a) = \phi(g, \phi(h, a))$ for all $g, h \in G$ and $a \in A$. The *right group action* of G on A is a function $\psi: A \times G \rightarrow A$ that satisfies $\psi(a, \varepsilon) = a$ and $\psi(a, gh) = \psi(\psi(a, g), h)$ for all $g, h \in G$ and $a \in A$.

Suppose ϕ and ψ are, respectively, a left and a right group action of G on A . For all $g \in G$, let Φ_g and Ψ_g be linear operators on \mathbb{C}^A defined via their action on the standard basis of \mathbb{C}^A : for all $a \in A$, let $\Phi_g \mathbf{a} := \phi(g, \mathbf{a})$ and $\Psi_g \mathbf{a} := \psi(\mathbf{a}, g^{-1})$. The maps $\Phi: g \mapsto \Phi_g$ and $\Psi: g \mapsto \Psi_g$ are called the *permutation representations* corresponding to ϕ and ψ , respectively, as, in the standard basis of \mathbb{C}^A , unitary operators Φ_g and Ψ_g are permutation matrices. They are indeed valid representations as we have

$$\begin{aligned}\Phi_g \Phi_h \mathbf{a} &= \Phi_g \phi(h, \mathbf{a}) = \phi(g, \phi(h, \mathbf{a})) = \phi(gh, \mathbf{a}) = \Phi_{gh} \mathbf{a}, \\ \Psi_g \Psi_h \mathbf{a} &= \Psi_g \psi(\mathbf{a}, h^{-1}) = \psi(\psi(\mathbf{a}, h^{-1}), g^{-1}) = \psi(\mathbf{a}, h^{-1} g^{-1}) = \psi(\mathbf{a}, (gh)^{-1}) = \Psi_{gh} \mathbf{a}\end{aligned}$$

for all $g, h \in G$ and $a \in A$.

The one-dimensional irrep that maps every group element to the multiplicative identity $1 \in \mathbb{C}$ is called the *trivial representation*. One can see that the one-dimensional space spanned by the all-ones vector in the standard basis (i.e., $\sum_{a \in A} \mathbf{a}$) is stable under Φ and Ψ , and they act on it as the identity. Thus, this subspace corresponds to the trivial representation.

A case of particular interest is when the group G acts on itself, and the action is given by the group operation. The representation corresponding to this action is called the *regular representation*. Namely, the *left regular representation* $R_\ell: G \rightarrow \mathcal{U}(\mathbb{C}^G)$ and the *right regular representation* $R_r: G \rightarrow \mathcal{U}(\mathbb{C}^G)$ are defined as follows: $R_\ell(g)\mathbf{h} = \mathbf{gh}$ and $R_r(g)\mathbf{h} = \mathbf{hg}^{-1}$ for all $g, h \in G$. Both regular representations are isomorphic. Every irrep ρ of G appears in the regular representation with multiplicity $\dim \rho$, thus the dimension of the isotypical subspace corresponding to ρ is $(\dim \rho)^2$.

From now on, by “group action” we mean “left group action”, except when we talk about the right regular representation. Let

$$\phi_g: A \rightarrow A: a \mapsto \phi(g, a).$$

The *orbit* of an element $a \in X$ under the action of G (or, simply, under G) is defined as

$$G_\phi(a) := \{\phi_g(a) : g \in G\},$$

where we may choose to omit the subscript ϕ when it is clear which action we are referring to. Note that $a \in G_\phi(a')$ is an equivalence relation.

We extend the function ϕ_g to subsets of A , namely, $\phi_g(\{a_1, \dots, a_k\}) := \{\phi_g(a_1), \dots, \phi_g(a_k)\}$. This way, in effect, we have defined a group action of G on the set of all subsets of A of a given size k . And, of course, this group action also gives a rise to a representation of G .

1.3.7 Restricted and induced representations

Suppose G is a group, $H \leq G$ is a subgroup of G , and $\rho: G \rightarrow \mathbf{U}(\mathcal{X})$ is a representation of G . The representation $H \rightarrow \mathbf{U}(\mathcal{X}): h \mapsto \rho_h$ of H is called the *restriction* of ρ from G to H and denoted by $\rho \downarrow H$. By abuse of terminology, we may also say that ρ is a representation of H . Note that, if ρ is irreducible, $\rho \downarrow H$ may be reducible.

For example, the symmetric group $\mathbb{S}_{[3]}$ is generated by (12) and (23), thus a representation ρ of $\mathbb{S}_{[3]}$ can be completely specified by $\rho_{(12)}$ and $\rho_{(23)}$. In Section 1.4.4 we state that ρ defined via

$$\rho: (12) \mapsto \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \rho: (23) \mapsto \begin{pmatrix} \frac{1}{2} & \frac{\sqrt{3}}{2} \\ \frac{\sqrt{3}}{2} & -\frac{1}{2} \end{pmatrix} \quad (1.6)$$

is indeed a representation of $\mathbb{S}_{[3]}$ and that it is irreducible. However, $\rho \downarrow (\mathbb{S}_{\{1,2\}} \times \mathbb{S}_{\{3\}})$ is not irreducible any more, as $\mathcal{Y} = \text{span}\{(0, 1)\}$ is stable under both ρ_ε and $\rho_{(12)}$ and it corresponds to the trivial representation of $\mathbb{S}_{\{1,2\}} \times \mathbb{S}_{\{3\}}$.

Let us introduce a concept “opposite” to the restriction. Suppose G and $H \leq G$ are finite groups, $\sigma: G \rightarrow \mathbf{U}(\mathcal{X})$ is a representation of G , and $\mathcal{Y} \subseteq \mathcal{X}$ is stable under H . Let $\rho := (\sigma \downarrow H)|_{\mathcal{Y}}$ be the reduction to \mathcal{Y} of the restriction of σ to H and let $\text{Rep}(G/H)$ be a transversal of the left cosets of H in G . One says that the representation σ of G is *induced* by the representation ρ of H if

$$\mathcal{X} = \bigoplus_{g \in \text{Rep}(G/H)} \sigma_g \mathcal{Y}. \quad (1.7)$$

Note that it is irrelevant which transversal we consider, as, for every coset gH and every $g' \in gH$, the subspace $\sigma_{g'} \mathcal{Y}$ is the same. Also note that (1.7) implies $\dim \mathcal{X} / \dim \mathcal{Y} = |G|/|H|$. Let $\rho \uparrow G$ denote a representation of G induced by a representation ρ of H . It is known that all such representations are isomorphic.

Suppose $\sigma: G \rightarrow \mathbf{U}(\mathcal{X})$ is induced by $\rho: H \rightarrow \mathbf{U}(\mathcal{Y})$, and $\mathcal{Y}' \subset \mathcal{Y}$ is stable under H . Then, one can see that $\mathcal{X}' := \bigoplus_{g \in \text{Rep}(G/H)} \sigma_g \mathcal{Y}'$ is stable under G . Namely, $G \rightarrow \mathbf{U}(\mathcal{X}'): g \mapsto \sigma_g \Pi_{\mathcal{X}'}$ is a representation of G induced by a representation $H \rightarrow \mathbf{U}(\mathcal{Y}'): h \mapsto \rho_h \Pi_{\mathcal{Y}'}$ of H .

The following theorem establishes a useful connection between the restriction and the induction.

Theorem 1.10 (Frobenius reciprocity). *Let G and $H \leq G$ be finite groups, and let σ and ρ be irreps of G and H , respectively. The number of times ρ appears in $\sigma \downarrow H$ equals the number of times σ appears in $\rho \uparrow G$.*

1.3.8 Composition of representations

Let us first consider two ways of composing representations of the same group. Let $\rho: G \rightarrow \mathbf{U}(\mathcal{X})$ and $\sigma: G \rightarrow \mathbf{U}(\mathcal{Y})$. Their *direct sum* is the representation

$$\rho \oplus \sigma: G \rightarrow \mathbf{U}(\mathcal{X} \oplus \mathcal{Y}): g \mapsto \rho_g \oplus \sigma_g$$

and their (*inner*) *tensor product* is the representation

$$\rho \otimes \sigma: G \rightarrow \mathbf{U}(\mathcal{X} \otimes \mathcal{Y}): g \mapsto \rho_g \otimes \sigma_g.$$

If we are given a decomposition of each ρ and σ into irreps, it is trivial to decompose into irreps the direct sum $\rho \oplus \sigma$. This task is not as simple for the tensor product $\rho \otimes \sigma$, and we will later present theorems concerning particular cases that we are interested in. Note that, if σ_{id} is the trivial representation, then $\rho \otimes \sigma_{\text{id}} \cong \rho$.

Now let us consider representations of the direct product $G \times H$. Suppose $\rho: G \rightarrow \mathbf{U}(\mathcal{X})$ and $\sigma: H \rightarrow \mathbf{U}(\mathcal{Y})$ are representations of G and H respectively. Then, their (*outer*) *tensor product* $\rho \times \sigma$ is the representation

$$\rho \times \sigma: G \times H \rightarrow \mathbf{U}(\mathcal{X} \otimes \mathcal{Y}): (g, h) \mapsto \rho_g \otimes \sigma_h.$$

If ρ and σ are irreps of G and H , respectively, then $\rho \times \sigma$ is an irrep of $G \times H$. Conversely, every irrep of $G \times H$ can be written as $\rho \times \sigma$, where ρ and σ are irreps of G and H , respectively. Note: given an irrep ρ of G , the outer tensor product $\rho \times \rho$ is an irrep of $G \times G$, but the inner tensor product $\rho \otimes \rho$ is not necessarily an irrep of G .

Using group isomorphism, let us also think of G and H as subgroups of $G \times H$. Suppose ρ and σ are irreps of G and H , respectively, and \mathcal{X} is a representation of $G \times H$. Then the $(\rho \times \sigma)$ -isotypical subspace of \mathcal{X} equals the intersection of the ρ -isotypical subspace of $\mathcal{X} \downarrow G$ and the σ -isotypical subspace of $\mathcal{X} \downarrow H$. Also note that $(\rho \times \sigma) \downarrow G$ consists of $\dim \sigma$ instances of the irrep ρ of G .

1.4 Representation theory of the symmetric group

In this section we present basics of the representation theory of the symmetric and unitary groups. The material presented here is based on various textbooks [JK81, Sag01, Boe63] and notes [Aud06] on the subject. The proof of Lemma 1.11 is from Ref. [BR14], yet we do not claim its originality.

In Section 1.3.1 we already established a bijection between partitions of m and conjugacy classes of \mathbb{S}_m . We know that the number of conjugacy classes of a group equals the number of irreducible representations, so now let us establish a bijection between partitions of m and irreducible representations of \mathbb{S}_m . And then, let us see how these partitions allow us to reason about the representation theory of \mathbb{S}_m .

1.4.1 Young diagrams and Young tableaux

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 example, the Young diagram representing $(5, 3, 3, 2) \vdash 13$ is shown in Figure 1.1. We often use the terms ‘partition of m ’ and ‘ m -box Young diagram’ interchangeably.

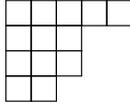


Figure 1.1: The Young diagram corresponding to the partition $(5, 3, 3, 2)$.

For a Young diagram λ , let λ^\top denote the transposed diagram of λ , i.e., the number of boxes in the i -th row of λ^\top equals to the number of boxes in the i -th column of λ . Note that λ_1^\top is the number of rows in λ . We say that a box (i, j) is *present* in λ and write $(i, j) \in \lambda$ if $\lambda_i \geq j$ (equivalently, $\lambda_j^\top \geq i$). The *hook-length* $h_\lambda(\mathbf{b})$ of a box $\mathbf{b} := (i, j) \in \lambda$ is the sum of the number of boxes on the right from \mathbf{b} in the same row (i.e., $\lambda_i - j$) and the number of boxes below \mathbf{b} in the same column (i.e., $\lambda_j^\top - i$) plus one (i.e., the box \mathbf{b} itself). For example, the hook lengths of the boxes of the Young diagrams $(5, 3, 3, 2) \vdash 13$ and $(m-3, 3) \vdash m$ are given in Figure 1.2.

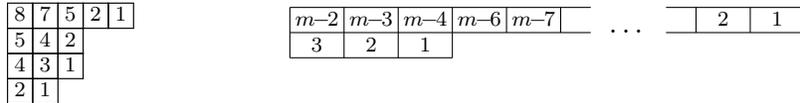


Figure 1.2: Hook lengths of the boxes of the Young diagrams $(5, 3, 3, 2)$ (left) and $(m-3, 3)$ (right).

Given an m -box Young diagram λ , let us say that a box is an *inner corner* of λ if its hook length is one, in other words, if the box is last in its row and last in its column. We can *remove* such a box from λ , thereby obtaining an $(m-1)$ -box Young diagram. By employing this procedure multiple times, we can also remove multiple boxes (note: some hook lengths will change after each removal of a box). Suppose m and $n \leq m$ are non-negative integers, and $\lambda \vdash m$ and $\mu \vdash n$ are two Young diagrams. We say that μ is *contained* in λ if μ can be obtained from λ by removing $m-n$ boxes, that is, if $\mu_i \leq \lambda_i$ for all i . In particular, let $\mu \subset \lambda$ and $\mu \subset\subset \lambda$ denote that a Young diagram μ is obtained from λ by removing exactly one box and exactly two boxes, respectively. Given $\mu \subset\subset \lambda$, let us write $\mu \subset_r \lambda$ or $\mu \subset_c \lambda$ if the two boxes removed from λ to obtain μ are, respectively, in different rows or different columns. Let $\mu \subset_{rc} \lambda$ be a shorthand for $(\mu \subset_r \lambda) \& (\mu \subset_c \lambda)$.

For every m , we use “ $<$ ” to denote the *lexicographical order* on the set of partitions of m . For

example, the lexicographical order of the partitions of $m = 5$ is

$$(1, 1, 1, 1, 1) < (2, 1, 1, 1) < (2, 2, 1) < (3, 2) < (3, 1, 1) < (4, 1) < (5).$$

Given a partition $\lambda \vdash m$, a *Young tableau* (or just *tableau*, for short) of *shape* λ is a bijection that assigns to every box of the Young diagram λ a number in $[m]$. For a Young tableau \mathfrak{T} of shape λ and a box $\mathfrak{b} \in \lambda$, we call $\mathfrak{T}(\mathfrak{b})$ the *entry* of the box \mathfrak{b} . A Young tableau is called *standard* if the numbers in every row and every column are strictly increasing. Figure 1.3 shows an example of a standard and a non-standard tableau.

1	2	5	8	13
3	6	7		
4	9	11		
10	12			

8	6	2	12	5
3	11	10		
13	1	4		
9	7			

Figure 1.3: A standard (left) and a non-standard (right) Young tableau of shape $(5, 3, 3, 2)$.

We will associate every irreducible representations of \mathbb{S}_m with a partition $\lambda \vdash m$, and the dimension of this irrep will be the total number of standard tableaux of shape λ . Therefore, let us denote this number by $\dim \lambda$, and it is given by the *hook-length formula*:

$$\dim \lambda = |\lambda|! / h(\lambda), \quad \text{where} \quad h(\lambda) = \prod_{\mathfrak{b} \in \lambda} h_{\lambda}(\mathfrak{b}). \quad (1.8)$$

For example, the hook-length formula and Figure 1.2 tells us that there are

$$13! / (8 \cdot 7 \cdot 5 \cdot 2 \cdot 1 \cdot 5 \cdot 4 \cdot 2 \cdot 4 \cdot 3 \cdot 1 \cdot 2 \cdot 1) = 11\,583$$

standard tableaux of shape $(5, 3, 3, 2)$ and

$$m! / ((m-2) \cdot (m-3) \cdot (m-4) \cdot (m-6)! \cdot 3 \cdot 2 \cdot 1) = m(m-1)(m-5)/6$$

standard tableaux of shape $(m-3, 3)$.

Suppose we fix the part of λ below the the first row and allow m to vary, namely, let $\lambda := (m - |\mu|, \mu)$, where μ is fixed. Then $\dim(m - |\mu|, \mu)$ is a polynomial in m of degree $|\mu|$ and its leading coefficient is $1/h(\mu)$. For example, see Table 1.1 for “dimensions” of the last seven partitions in the lexicographical order.

The group \mathbb{S}_m acts on the set of tableaux of shape $\lambda \vdash m$ in the following natural way. For all $\pi \in \mathbb{S}_m$, all tableaux \mathfrak{T} of shape λ , and all boxes $\mathfrak{b} \in \lambda$, if $\mathfrak{T}(\mathfrak{b}) = i$, then $\pi(\mathfrak{T})(\mathfrak{b}) = \pi(i)$. (This action corresponds to the regular representation of \mathbb{S}_m .)

λ	$\dim \lambda$
(m)	1
$(m-1, 1)$	$m-1$
$(m-2, 2)$	$m(m-3)/2$
$(m-2, 1, 1)$	$(m-1)(m-2)/2$
$(m-3, 3)$	$m(m-1)(m-5)/6$
$(m-3, 2, 1)$	$m(m-2)(m-4)/3$
$(m-3, 1, 1, 1)$	$(m-1)(m-2)(m-3)/6$

Table 1.1: The number of standard tableaux of shape λ for the last seven λ in the lexicographical order.

1.4.2 Specht modules

Suppose $\lambda = (\lambda_1, \dots, \lambda_k) \vdash m$. Given a tableau \mathfrak{T} of shape λ , let $R_i(\mathfrak{T})$ and $C_j(\mathfrak{T})$ be the set of entries in the i -th row and j -th column of \mathfrak{T} , respectively. For brevity, let

$$\mathbb{S}_{R(\mathfrak{T})} := \times_{i=1}^k \mathbb{S}_{R_i(\mathfrak{T})} \quad \text{and} \quad \mathbb{S}_{C(\mathfrak{T})} := \times_{j=1}^{\lambda_1} \mathbb{S}_{C_j(\mathfrak{T})},$$

which are subgroups of \mathbb{S}_m . A *tabloid* of shape λ is an equivalence class of all tableaux \mathfrak{T} that have the same $R_i(\mathfrak{T})$ for all i . Let $\{\mathfrak{T}\}$ denote the tabloid associated to \mathfrak{T} . In other words,

$$\{\mathfrak{T}\} = \mathbb{S}_{R(\mathfrak{T})} \mathfrak{T} := \{\pi(\mathfrak{T}) : \pi \in \mathbb{S}_{R(\mathfrak{T})}\}.$$

Essentially, a tabloid is a tableau for which one ignores the order of entries in each row. We represent tabloids like tableaux except omitting vertical boundaries of all boxes (see Figure 1.4).

$$\frac{\overline{2 \ 5}}{\overline{1 \ 3}} \overline{4} = \left\{ \begin{array}{|c|c|} \hline 2 & 5 \\ \hline 3 & 1 \\ \hline 4 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 2 & 5 \\ \hline 1 & 3 \\ \hline 4 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 5 & 2 \\ \hline 3 & 1 \\ \hline 4 & \\ \hline \end{array}, \begin{array}{|c|c|} \hline 5 & 2 \\ \hline 1 & 3 \\ \hline 4 & \\ \hline \end{array} \right\}$$

Figure 1.4: A tabloid of shape $(2, 2, 1)$.

For a tabloid $\{\mathfrak{T}\}$ of shape $\lambda \vdash m$ and $\pi \in \mathbb{S}_m$, let $\pi(\{\mathfrak{T}\}) := \{\pi(\mathfrak{T})\}$, which defines a left group action of \mathbb{S}_m on the set of tabloids of shape λ . (Note: this action is well defined because, for all $\mathfrak{T}' \in \{\mathfrak{T}\}$, we have $\{\pi(\mathfrak{T}')\} = \{\pi(\mathfrak{T})\}$.) For now, let \mathcal{X}^λ denote the space corresponding to the set of tabloids of shape λ . The space \mathcal{X}^λ together with the action of \mathbb{S}_m on it is known as the *permutation module* corresponding to λ , and it is well studied how \mathcal{X}^λ decomposes into irreps. We will be interested in one particular irrep appearing in \mathcal{X}^λ .

Recall the group algebra $\mathbb{C}\mathbb{S}_m$. Given a subgroup H of \mathbb{S}_m , let

$$\kappa(H) := \frac{1}{|H|} \sum_{\pi \in H} \text{sgn}(\pi) \pi \in \mathbb{C}\mathbb{S}_m,$$

where $\text{sgn}(\pi)$ is the sign of the permutation π . Note that,

$$\forall \pi \in H: \text{sgn}(\pi) \pi \cdot \kappa(H) = \kappa(H),$$

thus $\kappa(H)\kappa(H) = \kappa(H)$. Given a tableau \mathfrak{T} of shape λ , let $\kappa_{\mathfrak{T}} := \kappa(\mathbb{S}_{C(\mathfrak{T})})$, and let

$$e_{\mathfrak{T}} := \kappa_{\mathfrak{T}}\{\mathfrak{T}\} \in \mathcal{X}^{\lambda}.$$

For example, for

$$\mathfrak{T} = \begin{array}{|c|c|c|} \hline 1 & 4 & 3 \\ \hline 5 & 2 & \\ \hline \end{array},$$

we have $\mathbb{S}_{C(\mathfrak{T})} = \mathbb{S}_{\{1,5\}} \times \mathbb{S}_{\{2,4\}}$ and

$$\{\mathfrak{T}\} = \frac{\overline{1\ 3\ 4}}{\underline{2\ 5}},$$

therefore

$$e_{\mathfrak{T}} \propto (\varepsilon - (\mathbf{24}) - (\mathbf{15}) + (\mathbf{15})(\mathbf{24})) \frac{\overline{1\ 3\ 4}}{\underline{2\ 5}} = \frac{\overline{1\ 3\ 4}}{\underline{2\ 5}} - \frac{\overline{1\ 2\ 3}}{\underline{4\ 5}} - \frac{\overline{3\ 4\ 5}}{\underline{1\ 2}} + \frac{\overline{2\ 3\ 5}}{\underline{1\ 4}},$$

where we have omitted the scalar $\frac{1}{4}$ for clarity. Let \mathcal{S}^{λ} be the space spanned by all $e_{\mathfrak{T}}$, where \mathfrak{T} is a tableau of shape λ . \mathcal{S}^{λ} is an irreducible representation of \mathbb{S}_m known as the *Specht module*.

The only two one-dimensional representations of \mathbb{S}_m are the following:

- *The trivial representation $\mathcal{S}^{(m)}$* . There is only one tabloid of shape (m) and $e_{\mathfrak{T}}$ equals it for all tableaux \mathfrak{T} of shape (m) because $\kappa_{\mathfrak{T}} = \varepsilon$. All $\pi \in \mathbb{S}_m$ map this tabloid to itself.
- *The sign representation $\mathcal{S}^{(1^m)}$* , where $(1^m) := (1, 1, \dots, 1)$. For a tableau \mathfrak{T} of shape (1^m) , let $\text{sgn}(\mathfrak{T})$ be the sign of the unique permutation in \mathbb{S}_m that maps $i \in [m]$ to the sole entry in i -th row of \mathfrak{T} . Then we have $e_{\mathfrak{T}} = \text{sgn}(\mathfrak{T}) \sum_{\mathfrak{T}'} \text{sgn}(\mathfrak{T}') \{\mathfrak{T}'\} / m!$ and $\pi: e_{\mathfrak{T}} \mapsto \text{sgn}(\pi) e_{\mathfrak{T}}$ for all $\pi \in \mathbb{S}_m$.

The set of vectors $e_{\mathfrak{T}}$ such that \mathfrak{T} is a standard tableau of shape λ forms a basis for \mathcal{S}^{λ} . Therefore, $\dim \mathcal{S}^{\lambda} = \dim \lambda$, justifying the notation introduced earlier. For two distinct partitions $\lambda, \lambda' \vdash m$, the irreps \mathcal{S}^{λ} and $\mathcal{S}^{\lambda'}$ are non-isomorphic. Recall that the number of irreducible representations (up to isomorphism) and the number of conjugacy classes are equal, and there is a one-to-one correspondence between conjugacy classes of \mathbb{S}_m and partitions of m . Hence, as λ runs over all partitions of m , \mathcal{S}^{λ} runs over all (up to isomorphism) irreps of \mathbb{S}_m .

Lemma 1.11. *Suppose m and $k \leq m/2$ are positive integers and $a_1, b_1, \dots, a_k, b_k$ are some distinct fixed elements of $[m]$. Let*

$$\kappa := \frac{1}{2^k} (\varepsilon - (\mathbf{a}_1, \mathbf{b}_1))(\varepsilon - (\mathbf{a}_2, \mathbf{b}_2)) \cdots (\varepsilon - (\mathbf{a}_k, \mathbf{b}_k)) \in \mathbb{C}\mathbb{S}_m.$$

For any $\mu \vdash k$, the irrep $\mathcal{S}^{(m-k, \mu)}$ contains a non-zero vector v satisfying $\kappa v = v$.

Proof. Let $\ell = \mu_1$, and let \mathfrak{T} be a tableau of the shape $(m - k, \mu)$ with a_1, \dots, a_k being the first k elements of the first row, and b_1, \dots, b_k being the elements of the remaining rows, so that b_1, \dots, b_ℓ form the second row. The vector $e_{\mathfrak{T}} \in \mathcal{S}^{(m-k, \mu)}$ is not zero and it satisfies $\kappa_{\mathfrak{T}} e_{\mathfrak{T}} = e_{\mathfrak{T}}$.

Take $v := 2^{k-\ell} \kappa e_{\mathfrak{T}}$, which clearly satisfies $\kappa v = v$ as $\kappa \kappa = \kappa$. Since $(a_i, b_i) \in \mathbb{S}_{C(\mathfrak{T})}$ implies $\frac{1}{2}(\varepsilon - (\mathbf{a}_i, \mathbf{b}_i)) \kappa_{\mathfrak{T}} = \kappa_{\mathfrak{T}}$ for all $i \leq \ell$, we have $v = (\varepsilon - (\mathbf{a}_{\ell+1}, \mathbf{b}_{\ell+1})) \cdots (\varepsilon - (\mathbf{a}_k, \mathbf{b}_k)) e_{\mathfrak{T}}$. And $v \neq 0$ because no tabloid present in $e_{\mathfrak{T}}$ can be cancelled by other terms of v , because they have different content of the first row. \square

From now on, let \mathcal{S}^λ denote equivalence class of all irreps of $\mathbb{S}_{|\lambda|}$ isomorphic to the Specht module \mathcal{S}^λ . Note that Lemma 1.11 holds for all irreps in the isomorphism class.

1.4.3 Induction and restriction of representations

Given an irrep \mathcal{S}^μ of \mathbb{S}_m , where $\mu \vdash m$, the *branching rule* states that

$$\mathcal{S}^\mu \downarrow \mathbb{S}_{m-1} \cong \bigoplus_{\nu \subset \mu} \mathcal{S}^\nu$$

and, by Frobenius reciprocity,

$$\mathcal{S}^\mu \uparrow \mathbb{S}_{m+1} \cong \bigoplus_{\lambda \supset \mu} \mathcal{S}^\lambda. \quad (1.9)$$

Note that, since $|\mathbb{S}_{m+1}|/|\mathbb{S}_m| = m + 1$, the branching rule implies that

$$\dim \mu = \sum_{\nu \subset \mu} \dim \nu \quad \text{and} \quad (m + 1) \dim \mu = \sum_{\lambda \supset \mu} \dim \lambda.$$

The branching rule states what happens when one induces from $\mathbb{S}_{m-1} \times \mathbb{S}_1$ to \mathbb{S}_m (or restricts in the opposite direction). The more general *Littlewood–Richardson rule* describes, for every $k \in [0..m]$, what happens when one induces from $\mathbb{S}_{m-k} \times \mathbb{S}_k$ to \mathbb{S}_m . To state the Littlewood–Richardson rule fully, we would have to introduce concepts such as *skew shapes*, *semistandard tableaux*, the *weight* of such tableaux, and others. We choose not to do that because we will employ only two special cases of the rule:

1. Let $\lambda \vdash m$, so that \mathcal{S}^λ is an irrep of \mathbb{S}_m . We have

$$\mathcal{S}^\lambda \downarrow (\mathbb{S}_{m-2} \times \mathbb{S}_2) \cong \bigoplus_{\nu \subset \varepsilon \lambda} (\mathcal{S}^\nu \times \mathcal{S}^{(2)}) \oplus \bigoplus_{\nu \subset \varepsilon_r \lambda} (\mathcal{S}^\nu \times \mathcal{S}^{(1,1)}), \quad (1.10)$$

and the Frobenius reciprocity then tells us what happens when one induces from $\mathbb{S}_{m-2} \times \mathbb{S}_2$ to \mathbb{S}_m : for $\nu \vdash m - 2$, we have

$$(\mathcal{S}^\nu \times \mathcal{S}^{(2)}) \uparrow \mathbb{S}_m \cong \bigoplus_{\lambda \supset \varepsilon \nu} \mathcal{S}^\lambda \quad \text{and} \quad (\mathcal{S}^\nu \times \mathcal{S}^{(1,1)}) \uparrow \mathbb{S}_m \cong \bigoplus_{\lambda \supset \varepsilon_r \nu} \mathcal{S}^\lambda. \quad (1.11)$$

2. Let $\nu = (\nu_1, \nu_2, \dots, \nu_\ell) \vdash k$ and let $\Lambda(\nu)$ be the set of all Young diagrams μ that can be obtained from ν by removing at most one box per column (in other words, μ such that (ν_2, \dots, ν_ℓ) is contained in μ and μ is contained in ν). We have

$$(\mathcal{S}^{(m-k)} \times \mathcal{S}^\nu) \uparrow \mathbb{S}_m \cong \bigoplus_{\substack{\mu \in \Lambda(\nu) \\ m-|\mu| \geq \mu_1}} \mathcal{S}^{(m-|\mu|, \mu)}. \quad (1.12)$$

(Recall that $\mathcal{S}^{(m-k)}$ is the trivial representation of \mathbb{S}_{m-k} .)

1.4.4 The orthogonal form of \mathcal{S}^λ

In Section 1.4.2, we presented the irrep \mathcal{S}^λ by considering the action of the group algebra $\mathbb{C}\mathbb{S}_{|\lambda|}$ on tabloids of shape λ . Here we present \mathcal{S}^λ in another useful form: we define $\omega^\lambda \cong \mathcal{S}^\lambda$ known as the *orthogonal form* of \mathcal{S}^λ . The irrep ω^λ acts on the space corresponding to the set of standard tableaux of shape λ , which we here denote by \mathcal{X}^λ (do not mistake it for the permutation module discussed in and only in Section 1.4.2). We will express the operators ω^λ as unitary matrices in an orthonormal basis, whose basis vectors are labeled by the standard tableaux of shape λ ordered according to the *last letter order*.

Given a partition $\lambda \vdash m$, the last letter order ' $<$ ' on the set of standard tableaux of shape λ is a total ordering defined as follows. For \mathfrak{T} and \mathfrak{T}' being two standard tableaux of shape λ , one defines $\mathfrak{T} < \mathfrak{T}'$ if there exists $i \in [m]$ such that both

1. for all $j > i$, j appears in the same row (and the same column) in both \mathfrak{T} and \mathfrak{T}' ;
2. i appears in \mathfrak{T} in a higher row than in \mathfrak{T}' .

Namely, for a standard tableau \mathfrak{T} of shape $\lambda \vdash m$ and $j \in [m]$, let \mathfrak{T}_j be the standard tableau obtained from \mathfrak{T} by removing boxes corresponding to $j, j+1, \dots, m$ and let $\mu(\mathfrak{T}_j) \vdash j-1$ be its shape. We have $\mathfrak{T} < \mathfrak{T}'$, if the exists $i \in [m]$ such that

$$\forall j > i: \mu(\mathfrak{T}_j) = \mu(\mathfrak{T}'_j) \quad \text{and} \quad \mu(\mathfrak{T}_i) < \mu(\mathfrak{T}'_i),$$

where, for the comparison of shapes, the lexicographical order is used. For example, Figure 1.5 illustrates the last letter order of all five standard tableaux of shape $(3, 2)$.

The *axial distance* between two boxes $\mathfrak{b} = (i, j)$ and $\mathfrak{b}' = (i', j')$ is defined as

$$d(\mathfrak{b}, \mathfrak{b}') := (i' - i) + (j - j').$$

Similarly, we call $|i - i'| + |j - j'|$ the *distance* between \mathfrak{b} and \mathfrak{b}' . Note that, if \mathfrak{b}' is below \mathfrak{b} (i.e., $i' > i$), to the left from \mathfrak{b} (i.e., $j' < j$), or both, then $d(\mathfrak{b}, \mathfrak{b}') > 0$. On the other hand, if \mathfrak{b}' is

$$\begin{array}{|c|c|c|} \hline 1 & 3 & 5 \\ \hline 2 & 4 & \\ \hline \end{array} < \begin{array}{|c|c|c|} \hline 1 & 2 & 5 \\ \hline 3 & 4 & \\ \hline \end{array} < \begin{array}{|c|c|c|} \hline 1 & 3 & 4 \\ \hline 2 & 5 & \\ \hline \end{array} < \begin{array}{|c|c|c|} \hline 1 & 2 & 4 \\ \hline 3 & 5 & \\ \hline \end{array} < \begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline 4 & 5 & \\ \hline \end{array}$$

Figure 1.5: The last letter order of the standard tableaux of shape $(3, 2)$.

$$\begin{aligned} \omega^{(3,2)} : (12) &\mapsto \begin{pmatrix} -1 & & & & \\ & 1 & & & \\ & & -1 & & \\ & & & 1 & \\ & & & & 1 \end{pmatrix}, \\ \omega^{(3,2)} : (23) &\mapsto \begin{pmatrix} \frac{1}{2} & \frac{\sqrt{3}}{2} & & & \\ \frac{\sqrt{3}}{2} & -\frac{1}{2} & & & \\ & & \frac{1}{2} & \frac{\sqrt{3}}{2} & \\ & & \frac{\sqrt{3}}{2} & -\frac{1}{2} & \\ & & & & 1 \end{pmatrix}, \\ \omega^{(3,2)} : (34) &\mapsto \begin{pmatrix} -1 & & & & \\ & 1 & & & \\ & & 1 & & \\ & & & \frac{1}{3} & \frac{2\sqrt{2}}{3} \\ & & & \frac{2\sqrt{2}}{3} & -\frac{1}{3} \end{pmatrix}, \\ \omega^{(3,2)} : (45) &\mapsto \begin{pmatrix} \frac{1}{2} & & \frac{\sqrt{3}}{2} & & \\ & \frac{1}{2} & & \frac{\sqrt{3}}{2} & \\ \frac{\sqrt{3}}{2} & & -\frac{1}{2} & & \\ & \frac{\sqrt{3}}{2} & & -\frac{1}{2} & \\ & & & & 1 \end{pmatrix}. \end{aligned}$$

Figure 1.6: The generating matrices of the orthogonal form $\omega^{(3,2)}$. (Here we have omitted entries 0.)

above \mathfrak{b} (i.e., $i' < i$), to the right from \mathfrak{b} (i.e., $j' > j$), or both, then $d(\mathfrak{b}, \mathfrak{b}') < 0$. Also note that $d(\mathfrak{b}, \mathfrak{b}') = -d(\mathfrak{b}', \mathfrak{b})$.

Given a tableau \mathfrak{T} and $i \in [m]$, let $\mathfrak{b}_{\mathfrak{T}}(i)$ be the box in \mathfrak{T} containing i . For a transposition $(i, i+1)$, where $i \in [m-1]$, the operator $\omega_{(i, i+1)}^{\lambda}$ acts on the standard basis of \mathcal{X}^{λ} as follows: for

a standard tableau \mathfrak{T} of shape λ ,

$$\omega_{(i,i+1)}^\lambda: \mathfrak{T} \mapsto \frac{1}{d(\mathfrak{b}_{\mathfrak{T}}(i+1), \mathfrak{b}_{\mathfrak{T}}(i))} \mathfrak{T} + \sqrt{1 - \frac{1}{|d(\mathfrak{b}_{\mathfrak{T}}(i+1), \mathfrak{b}_{\mathfrak{T}}(i))|^2}} (i, i+1)\mathfrak{T}.$$

Note that, since \mathfrak{T} is a standard tableau, we have the following, where we suppose that $\mathfrak{T}' := (i, i+1)\mathfrak{T}$.

- If i and $i+1$ are in the same row or the same column of \mathfrak{T} , then they are adjacent and we have $\omega_{(i,i+1)}^\lambda: \mathfrak{T} \mapsto \mathfrak{T}$ or $\omega_{(i,i+1)}^\lambda: \mathfrak{T} \mapsto -\mathfrak{T}$, respectively. That is, one does not have to worry about the fact that \mathfrak{T}' is a non-standard tableaux in this case.
- If i and $i+1$ are in different rows and different columns, then \mathfrak{T}' is also a standard tableau. For simplicity, denote

$$d := d(\mathfrak{b}_{\mathfrak{T}}(i+1), \mathfrak{b}_{\mathfrak{T}}(i)) = -d(\mathfrak{b}_{\mathfrak{T}'}(i+1), \mathfrak{b}_{\mathfrak{T}'}(i)).$$

We have $d \geq 2$ if $\mathfrak{T} < \mathfrak{T}'$ and $d \leq -2$ if $\mathfrak{T} > \mathfrak{T}'$. Assuming the former case (i.e., d is positive), the submatrix of $\omega_{(i,i+1)}^\lambda$ corresponding to the coordinates $(\mathfrak{T}, \mathfrak{T}')$ is

$$\begin{pmatrix} \frac{1}{d} & \frac{\sqrt{d^2-1}}{d} \\ \frac{\sqrt{d^2-1}}{d} & -\frac{1}{d} \end{pmatrix}.$$

The set of transpositions $\{(1, 2), (2, 3), \dots, (m-1, 1)\}$ generates \mathbb{S}_m , therefore we have specified the action of ω_π^λ on \mathcal{X}^λ for all $\pi \in \mathbb{S}_m$. For example, the four generating matrices of $\omega^{(3,2)}$ are given in Figure 1.6 (this example is from [JK81, Sec. 3.4], and, for clarity, we present it below Figure 1.5). Another example, the orthogonal form $\omega^{(2,1)}$, was already given in (1.6).

First, in order to understand the orthogonal form better, let us show how it yields the branching rule. Afterwards, we use the orthogonal form to prove an equality that we will require later in the thesis.

In every standard tableau of shape $\lambda \vdash m$, the box containing the entry m is an inner corner. The box containing $m-1$ is either adjacent to the box containing m or it is an inner corner as well. Suppose \mathfrak{b} and \mathfrak{b}' are two distinct inner corners of λ (assume $\lambda \neq (m)$ and $\lambda \neq (1^m)$). Let $\lambda \setminus \mathfrak{b} \vdash m-1$, $\lambda \setminus \mathfrak{b}' \vdash m-1$, and $\lambda \setminus \mathfrak{b}\mathfrak{b}' \vdash m-2$ denote, respectively, the Young diagrams obtained from λ by removing \mathfrak{b} , \mathfrak{b}' , and both \mathfrak{b} and \mathfrak{b}' . Let $\mathcal{X}_\mathfrak{b}^\lambda \subseteq \mathcal{X}^\lambda$ be the subspace spanned by all basis vectors \mathfrak{T} such that $\mathfrak{T}(\mathfrak{b}) = m$, let $\mathcal{X}_{\mathfrak{b}\mathfrak{b}', \mathfrak{b}}^\lambda \subseteq \mathcal{X}_\mathfrak{b}^\lambda$ spanned by all \mathfrak{T} such that $\mathfrak{T}(\mathfrak{b}) = m$ and $\mathfrak{T}(\mathfrak{b}') = m-1$, and let $\mathcal{X}_{\mathfrak{b}'}^\lambda$ and $\mathcal{X}_{\mathfrak{b}\mathfrak{b}', \mathfrak{b}'}$ be defined analogously.

The space $\mathcal{X}_\mathfrak{b}^\lambda$ is stable under $\omega_{(i,i+1)}^\lambda$ for all $i \in [m-2]$ as these operators do not “move” the entry m , and therefore it is stable under ω_π^λ for all $\pi \in \mathbb{S}_{m-1}$. We can ignore the “fixed” box

\mathfrak{b} , and we have $(\omega^\lambda \downarrow \mathbb{S}_{m-1})|_{\mathcal{X}_\mathfrak{b}^\lambda} = \omega^{\lambda \setminus \mathfrak{b}}$, where we use “=” instead of “ \cong ” to stress that these matrices are identical (for all $\pi \in \mathbb{S}_{m-1}$). Since every $\mu \subset \lambda$ equals $\lambda \setminus \mathfrak{b}$ for a unique inner corner $\mathfrak{b} \in \lambda$, this is the branching rule.

Similarly, the restriction of ω^λ to \mathbb{S}_{m-2} contains two instances of the irrep $\mathcal{S}^{\lambda/bb'}$, and $\mathcal{X}_{bb'}^\lambda := \mathcal{X}_{bb',b}^\lambda \oplus \mathcal{X}_{bb',b'}^\lambda$ is the isotypical subspace of \mathcal{X}^λ corresponding to $\mathcal{S}^{\lambda/bb'}$. We have

$$(\omega^\lambda \downarrow \mathbb{S}_{m-2})|_{\mathcal{X}_{bb',b}^\lambda} = (\omega^\lambda \downarrow \mathbb{S}_{m-2})|_{\mathcal{X}_{bb',b'}^\lambda} = \omega^{\lambda \setminus bb'}.$$

Note that we can write the projectors $\Pi_{bb',b}^\lambda$ and $\Pi_{bb',b'}^\lambda$ on $\mathcal{X}_{bb',b}^\lambda$ and $\mathcal{X}_{bb',b'}^\lambda$, respectively, as

$$\Pi_{bb',b}^\lambda = \sum_{\mathfrak{T}: \mathfrak{T}(\mathfrak{b})=m, \mathfrak{T}(\mathfrak{b}')=m-1} \mathfrak{T}\mathfrak{T}^* \quad \text{and} \quad \Pi_{bb',b'}^\lambda = \sum_{\mathfrak{T}: \mathfrak{T}(\mathfrak{b}')=m, \mathfrak{T}(\mathfrak{b})=m-1} \mathfrak{T}\mathfrak{T}^*,$$

and note that

$$\Xi_{bb',b' \leftrightarrow bb',b}^\lambda := \sum_{\mathfrak{T}: \mathfrak{T}(\mathfrak{b})=m, \mathfrak{T}(\mathfrak{b}')=m-1} ((m-1, m)\mathfrak{T})\mathfrak{T}^*$$

is a transporter from $\mathcal{X}_{bb',b}^\lambda$ to $\mathcal{X}_{bb',b'}^\lambda$. Suppose \mathfrak{b} is above and to the right of \mathfrak{b}' , implying that $(\lambda \setminus \mathfrak{b}) < (\lambda \setminus \mathfrak{b}')$ in the lexicographical order and $d(\mathfrak{b}, \mathfrak{b}') \geq 2$. Then, all basis vectors \mathfrak{T} spanning $\mathcal{X}_{bb',b}^\lambda$ appear in the ordered basis of \mathcal{X}^λ before all basis vectors \mathfrak{T}' spanning $\mathcal{X}_{bb',b'}^\lambda$ (i.e., $\mathfrak{T} < \mathfrak{T}'$ in the last letter order). The $\mathcal{S}^{\lambda/bb'}$ -isotypical subspace $\mathcal{X}_{bb'}^\lambda$ is invariant under $\omega_{(m-1,m)}^\lambda$, and we have

$$\omega_{(m-1,m)}^\lambda|_{\mathcal{X}_{bb'}^\lambda} = \frac{1}{d(\mathfrak{b}, \mathfrak{b}')} \Pi_{bb',b}^\lambda - \frac{1}{d(\mathfrak{b}, \mathfrak{b}')} \Pi_{bb',b'}^\lambda + \frac{\sqrt{d(\mathfrak{b}, \mathfrak{b}')^2 - 1}}{d(\mathfrak{b}, \mathfrak{b}')} \Xi_{bb',b' \leftrightarrow bb',b}^\lambda. \quad (1.13)$$

The equality (1.13) plays a mayor role in both the adversary lower bound for the COLLISION and SET EQUALITY problems (Chapter 4) and the adversary lower bound for the ELEMENT DISTINCTNESS problem with small range (Chapter 5). In Chapter 5 we will also have to consider the restriction of an irrep of \mathbb{S}_m to \mathbb{S}_{m-3} , and, using the orthogonal form, we handle that case similarly.

Note that, due to isomorphism, an equality analogous to (1.13) holds for every irrep in the isomorphism class \mathcal{S}^λ . Also, here we essentially reasoned about $\mathbb{S}_m = \mathbb{S}_{[m]}$ and we considered the restrictions of $\mathbb{S}_{[m]}$ to $\mathbb{S}_{[m-1]}$ and then further to $\mathbb{S}_{[m-2]}$. Due to symmetry, the equality still holds when we consider instead the restrictions of $\mathbb{S}_{[m]}$ to $\mathbb{S}_{[m] \setminus \{i\}}$ and then further to $\mathbb{S}_{[m] \setminus \{i,j\}}$ for any $i, j \in [m]$ (i.e., no element in $[m]$ is “more special” than any other element in $[m]$).

1.4.5 Decomposition of inner tensor products

Suppose $\lambda, \mu \vdash m$, and we interested how a representation $\mathcal{S}^\lambda \otimes \mathcal{S}^\mu$ of \mathbb{S}_m decomposes into irreps. For general λ and μ , this can be done using the *determinantal form* of irreps of \mathbb{S}_m and the general

version of the Littlewood–Richardson rule. In this thesis we choose to introduce neither, because we will be interested only in the special case given by the following lemma.

Lemma 1.12. *Consider \mathbb{S}_m and suppose $1 < i, j < m/2$. The multiplicity of the irrep \mathcal{S}^ν in $\mathcal{S}^{(m-i,i)} \otimes \mathcal{S}^{(m-j,j)}$ equals the multiplicity of \mathcal{S}^ν in $\mathcal{S}^{(m-j,j)} \downarrow (\mathbb{S}_{m-i} \times \mathbb{S}_i) \uparrow \mathbb{S}_m$ minus the multiplicity of \mathcal{S}^ν in $\mathcal{S}^{(m-j,j)} \downarrow (\mathbb{S}_{m-i+1} \times \mathbb{S}_{i-1}) \uparrow \mathbb{S}_m$.*

In particular, we will care only about the case when $i = 1$ (and the trivial $i = 0$ or $j = 0$).

Corollary 1.13. *For $1 \leq j \leq m/2 - 1$, we have*

$$\mathcal{S}^{(m-1,1)} \otimes \mathcal{S}^{(m-j,j)} \cong \mathcal{S}^{(m-j+1,j-1)} \oplus \mathcal{S}^{(m-j,j)} \oplus (\mathcal{S}^{(m-j,j-1,1)}) \oplus \mathcal{S}^{(m-j-1,j+1)} \oplus \mathcal{S}^{(m-j-1,j,1)},$$

where we omit the term $\mathcal{S}^{(m-j,j-1,1)}$ when $j = 1$ as $(m-1, 0, 1)$ is not a partition.

Proof. For $i = 1$, $\mathbb{S}_{m-1} \times \mathbb{S}_1 \cong \mathbb{S}_{m-1}$ and $\mathbb{S}_{m-0} \times \mathbb{S}_0 \cong \mathbb{S}_m$. By the branching rule:

$$\mathcal{S}^{(m-j,j)} \downarrow \mathbb{S}_{m-1} \cong \mathcal{S}^{(m-j,j-1)} \oplus \mathcal{S}^{(m-j-1,j)},$$

and

$$\begin{aligned} \mathcal{S}^{(m-j,j-1)} \uparrow \mathbb{S}_m &\cong \mathcal{S}^{(m-j+1,j-1)} \oplus \mathcal{S}^{(m-j,j)} \oplus (\mathcal{S}^{(m-j,j-1,1)}), \\ \mathcal{S}^{(m-j-1,j)} \uparrow \mathbb{S}_m &\cong \mathcal{S}^{(m-j,j)} \oplus \mathcal{S}^{(m-j-1,j+1)} \oplus \mathcal{S}^{(m-j-1,j,1)}. \end{aligned}$$

From the direct sum of these, we have to “subtract” $\mathcal{S}^{(m-j,j)} \downarrow \mathbb{S}_m \uparrow \mathbb{S}_m = \mathcal{S}^{(m-j,j)}$. □

1.4.6 Representation theory of the unitary group

Suppose \mathcal{Y}_r is a Hilbert space of dimension r , and consider the unitary group $\mathbf{U}(\mathcal{Y}_r)$. A representation of $\rho: \mathbf{U}(\mathcal{Y}_r) \rightarrow \mathbf{U}(\mathcal{X})$ is called *polynomial* (or, by some authors, *integral*) if the matrix elements of $\rho(U)$ are polynomials in the elements of the represented matrix $U \in \mathbf{U}(\mathcal{Y}_r)$ (in some fixed bases of \mathcal{Y}_r and \mathcal{X}). There is one-to-one correspondence between polynomial irreps of $\mathbf{U}(\mathcal{Y})$ and partitions $\lambda = (\lambda_1, \dots, \lambda_k)$ having $k \in [r]$ (i.e., Young diagrams having at most r rows). For such a partition λ , let \mathcal{W}_r^λ denote the corresponding irrep of $\mathbf{U}(\mathcal{Y})$. It is known as the *Weyl module*.

Let m be a positive integer, and consider the space $\mathcal{Y}_r^{\otimes m}$. Then $U \in \mathbf{U}(\mathcal{Y}_r)$ acts on this space by simultaneous matrix multiplication, that is,

$$U: v_1 \otimes \dots \otimes v_m \mapsto Uv_1 \otimes \dots \otimes Uv_m = U^{\otimes m}(v_1 \otimes \dots \otimes v_m). \quad (1.14)$$

The symmetric group \mathbb{S}_m acts on $\mathcal{Y}_r^{\otimes m}$ by permuting the tensor factors, that is, for $\pi \in \mathbb{S}_m$, we have

$$\pi: v_1 \otimes \dots \otimes v_m \mapsto v_{\pi^{-1}(1)} \otimes \dots \otimes v_{\pi^{-1}(m)}. \quad (1.15)$$

These actions of $\mathbf{U}(\mathcal{Y}_r)$ and \mathbb{S}_m commute, thus, they define a representation of their direct product, $\mathbf{U}(\mathcal{Y}_r) \times \mathbb{S}_m$.

Theorem 1.14 (Schur–Weyl duality). *The above representation of $\mathbf{U}(\mathcal{Y}_r) \times \mathbb{S}_m$ on $\mathcal{Y}_r^{\otimes m}$ can be decomposed as a direct sum of irreps $\mathcal{W}_r^\lambda \times \mathcal{S}^\lambda$ taken over all $\lambda \vdash m$ such that the Young diagram λ has at most r rows.*

1.5 Association schemes

Definition 1.15. An *association scheme* is a set $\{A_0, A_1, \dots, A_k\}$ of symmetric $(0, 1)$ -matrices of the same dimensions $d \times d$ such that

1. A_0 is the identity matrix \mathbb{I}_d ,
2. $A_i A_j = A_j A_i$ for all i and j ,
3. $\sum_{i=0}^k A_i$ is the all-ones matrix \mathbb{J}_d .

The second condition ensures that all these matrices share the same eigenspaces, and we call them the eigenspaces of the association scheme. Let us introduce here the basic theory of two widely used association schemes: the Hamming scheme and the Johnson scheme. For more background on association schemes, refer to [God05].

1.5.1 Hamming scheme

Suppose Σ is a finite alphabet. For two strings x and y over Σ of the same length, their *Hamming distance* is the number of positions at which the corresponding entries of x and y differ, and it is denoted by $|x - y|$. Assuming Σ contains the *zero-symbol*, the *Hamming weight* of a string x over Σ , denoted by $|x|$, is the number of positions in x containing a non-zero symbol.

Definition 1.16. Let Σ be a finite alphabet of size q and let n be a positive integer. The *Hamming (association) scheme* is the set of $q^n \times q^n$ matrices $\{A_0^{(n)}, A_1^{(n)}, \dots, A_n^{(n)}\}$ whose rows and columns are labeled by all strings $x \in \Sigma^n$ and $A_i^{(n)} \llbracket x, y \rrbracket = 1$ if and only if $|x - y| = i$.

Suppose Σ is ordered, which provides the lexicographical order of Σ^n . And suppose the rows and the columns of all $A_i^{(n)}$ are ordered according to this lexicographical order. Then, for all i , we have

$$A_i^{(n)} = \sum_{b \in \{0,1\}^n: |b|=i} \bigotimes_{j=1}^n A_{b_j}$$

where are $A_0 := A_0^{(1)} = \mathbb{I}_q$ and $A_1 := A_1^{(1)} = \mathbb{J}_q - \mathbb{I}_q$. Since $\Pi_0 := \mathbb{J}_q/q$ and $\Pi_1 := \mathbb{I}_q - \Pi_0$ are projectors on the eigenspaces of A_0 and A_1 , one can see that, for all $i \in [0..n]$,

$$\Pi_i^{(n)} := \sum_{b \in \{0,1\}^n: |b|=i} \bigotimes_{j=1}^n \Pi_{b_j} \quad (1.16)$$

projects on an eigenspace of the Hamming scheme.

1.5.2 Johnson scheme

Definition 1.17. Let n and $k \leq n/2$ be positive integers. The *Johnson (association) scheme* is the set of $\binom{n}{k} \times \binom{n}{k}$ matrices $\{A_0, A_1, \dots, A_k\}$ whose rows and columns are labeled by all subsets x of $[n]$ of size k and $A_i[x, y] = 1$ if and only if $|x \setminus y| = i$.

Let L_j denote the set of all subsets of $[n]$ of size j . Hence, the matrices A_i can be thought to act on the space $\mathcal{X} := \mathbb{C}^{L_k}$, and we can write

$$A_i = \sum_{x, y \in L_k: |x \setminus y|=i} \mathbf{x} \mathbf{y}^*. \quad (1.17)$$

The symmetric group $\mathbb{S}_{[n]}$ acts on L_k in a natural way as follows:

$$\pi \in \mathbb{S}_{[n]}: \{x_1, \dots, x_k\} \mapsto \{\pi(x_1), \dots, \pi(x_k)\}.$$

This action defines a permutation representation $P: \mathbb{S}_{[n]} \rightarrow \mathbf{U}(\mathcal{X})$. In order to see how the representation P decomposes into irreps, let us use induction.

Fix $x \in L_k$ and consider the one dimensional space $\mathcal{X}_x := \text{span}\{\mathbf{x}\}$. The subspace \mathcal{X}_x is stable under the subgroup $\mathbb{S}_x \times \mathbb{S}_{[n] \setminus x}$ of $\mathbb{S}_{[n]}$, and the action of this group corresponds to the trivial representation, that is, $\mathcal{X}_x \cong \mathcal{S}^{(k)} \times \mathcal{S}^{(n-k)}$. For every $x' \in L_k$, there exists $\pi \in \mathbb{S}_{[n]}$ such that $\pi(x) = x'$, and $\dim \mathcal{X} / \dim \mathcal{X}_x = |\mathbb{S}_{[n]}| / |\mathbb{S}_x \times \mathbb{S}_{[n] \setminus x}|$. Hence,

$$\mathcal{X} = \mathcal{X}_x \uparrow \mathbb{S}_{[n]} \cong (\mathcal{S}^{(k)} \times \mathcal{S}^{(n-k)}) \uparrow \mathbb{S}_n \cong \bigoplus_{h=0}^k \mathcal{S}^{(n-h, h)}, \quad (1.18)$$

where the last isomorphism is due to the Littlewood–Richardson rule (1.12). Let Π_h be the projector on the subspace of \mathcal{X} isomorphic to $\mathcal{S}_{(n-h, h)}$.

Note that $P_\pi A_i P_\pi^{-1} = A_i$ for all $\pi \in \mathbb{S}_{[n]}$. Since, according to (1.18), P is multiplicity-free, Schur's lemma implies that we can express A_i as a linear combination of Π_h . To present the coefficients in this linear combination, it helps to introduce the matrices

$$C_j := \sum_{z \in L_j} \zeta_z \zeta_z^* \quad (1.19)$$

for all $j \in [0..k]$, where

$$\zeta_z := \sum_{x \in L_k: z \subseteq x} \mathbf{x}. \quad (1.20)$$

It is known that

$$\forall i: A_i = \sum_{j=k-i}^k (-1)^{j-k+i} \binom{j}{k-i} C_j \quad \text{and} \quad \forall j: C_j = \sum_{i=0}^{k-j} \binom{k-i}{j} A_i, \quad (1.21)$$

and also

$$\forall j: C_j = \sum_{h=0}^j \binom{n-j-h}{n-k-h} \binom{k-h}{j-h} \Pi_h \quad (1.22)$$

[God05, Chp. 7]. Hence, we can express A_i uniquely as a linear combination of orthogonal projectors Π_h , and the coefficients corresponding to these projectors are the eigenvalues of A_i .

In Chapter 6, we are interested in the converse: expressing Π_1 and Π_2 as linear combinations of A_i . From (1.22) one can see that

$$\Pi_h = (n - 2h + 1) \sum_{j=0}^h (-1)^{j-h} \frac{\binom{k-j}{h-j}}{(k-j+1) \binom{n-j-h+1}{n-k-h}} C_j \quad (1.23)$$

for $h = 0, 1, 2$ (a symbolic calculation for h up to 6 suggests that (1.23) might hold for all h). From (1.23) and the right hand side of (1.21) we get

$$\Pi_1 = \frac{1}{\binom{n-2}{k-1}} \sum_{i=0}^k \left((k-i) - \frac{k^2}{n} \right) A_i, \quad (1.24)$$

$$\Pi_2 = \frac{1}{\binom{n-4}{k-2}} \sum_{i=0}^k \left(\binom{k-i}{2} - \frac{(k-1)^2}{n-2} (k-i) + \frac{k^2(k-1)^2}{2(n-1)(n-2)} \right) A_i. \quad (1.25)$$

Chapter 2

Quantum query complexity

In this chapter we introduce quantum query complexity and concepts related to it. It is recommended (but not necessary) that the reader has at least introductory knowledge on quantum computing. For more background on quantum computing, one may refer to [NC00].

This chapter is organized as follows. In Section 2.1 we introduce the main computational problems considered in this thesis and concepts related to certificates. Then, in Section 2.2, we define the quantum query algorithm and we describe how one can “symmetrize” it and run it on a superposition of inputs. In Section 2.3 we define the adversary bound, sketch the idea behind its proof, and present some tools that simplify its application. We also give the basic intuition behind adversary bounds presented in later chapters. Finally, in Section 2.4, we introduce a computational model of learning graph.

2.1 Computational problems

Definition 2.1. We define a *computational problem* to be a binary relation $\mathcal{P} \subseteq \Sigma^n \times R$, where $n \in \mathbb{N}$ is called the *input length*, Σ is a finite set called the *input alphabet* (or, simply, the *alphabet*), elements in Σ are called *symbols* or *characters*, and R is a finite set called the *codomain*. A *family of computational problems* is a function that maps (not necessarily every) $n \in \mathbb{N}$ to a computational problem of input length n (we allow the input alphabet and the codomain to depend on n .)

We call $x \in \Sigma^n$ an *input* or an *input string*. For $i \in [n]$, we call x_i an *input variable* when we do not have a specific value (i.e., a symbol in Σ) of x_i in mind. We call $r \in R$ an *output*. We interpret $(x, r) \in \mathcal{P}$ as r being a *correct solution* to a problem \mathcal{P} on an input x , and we do not ask this solution to be unique. Let us also use the notation

$$\mathcal{P}(x) := \{r \in R : (x, r) \in \mathcal{P}\}.$$

The *domain* of the problem \mathcal{P} is

$$\mathcal{D} := \{x \in \Sigma^n : \mathcal{P}(x) \neq \emptyset\},$$

and, given an input $x \in \mathcal{D}$, the task of an algorithm for \mathcal{P} is to output a correct solution $r \in \mathcal{P}(x)$.

Without loss of generality, we assume that Σ is a group in order to later define an oracle that will allow us to interact with the input (see (2.1)). Some problems, like k -SUM and TRANGLE-SUM, require Σ to be a group regardless of how one chooses to access the input.

A problem \mathcal{P} is a *function* if, for every $x \in \mathcal{D}$, there exists a unique $r \in R$ such that $(x, r) \in \mathcal{P}$, and we write $\mathcal{P}(x) = r$ instead of $\mathcal{P}(x) = \{r\}$ and $\mathcal{P} : \Sigma^n \rightarrow R$ (or $\mathcal{P} : \mathcal{D} \rightarrow R$) instead of $\mathcal{P} \subseteq \Sigma^n \times R$. For a function \mathcal{P} , let

$$\mathcal{P}^{-1}(r) := \{x \in \mathcal{D} : \mathcal{P}(x) = r\},$$

and note that, for two distinct $r, r' \in R$, $\mathcal{P}^{-1}(r) \cap \mathcal{P}^{-1}(r') = \emptyset$. We also commonly use the notation \mathcal{D}_r instead of $\mathcal{P}^{-1}(r)$, so that $\mathcal{D} = \bigsqcup_{r \in R} \mathcal{D}_r$. Boolean-valued functions, that is, functions for which $R = \{0, 1\}$, are also referred to as *decision problems*. We call $x \in \mathcal{P}^{-1}(1)$ a *yes-input*, a *positive input*, or a *1-input* and $y \in \mathcal{P}^{-1}(0)$ a *no-input*, a *negative input*, or a *0-input*. For decision problems, we typically use x and y to refer to a yes-input and a no-input, respectively. We may also use x to refer to a general input, or use z for this purpose.

Problems for which $\mathcal{D} \neq \Sigma^n$ are called *promise* problems. Functions that are promise problems are called *partial* functions, and functions for which $\mathcal{D} = \Sigma^n$ are called *total* functions.

2.1.1 Common computational problems

Let us introduce the main computational problems that we consider in this thesis. For some computational problems, the custom is to call the *size* of the problem a quantity other than the input length. In such cases, let n denote the size of the problem, and we will explain how it relates to the input length.

In this thesis, we consider in detail only one computational problem that is not a function: the FIND-TWO problem. This problem is closely related to the (unstructured) SEARCH, so let us introduce both of them. For SEARCH and FIND-TWO, the input alphabet is *binary* (i.e., $\Sigma := \{0, 1\}$), and we call $i \in [n]$ *marked* if $x_i = 1$.

1. The SEARCH problem is to find a marked index i (if such index exists).
2. The FIND-TWO problem is to find a pair of marked indices i, j , where $i \neq j$ (if such a pair exists).

All the other problems that we consider in detail are functions; in particular, they are decision problems. Let us now introduce them.

3. The THRESHOLD- k problem is to decide whether the Hamming weight of a binary input string is at least k (answer: ‘yes’) or strictly less than k (‘no’). The OR function is equal to THRESHOLD-1 and the AND function is equal to THRESHOLD- n , where n is the length of the input.
4. The ELEMENT DISTINCTNESS problem is to decide whether there is a symbol in Σ that appears in the input string at least twice (‘yes’) or each symbol in the input string is unique (‘no’).
5. The k -DISTINCTNESS problem is a generalization of ELEMENT DISTINCTNESS that asks whether there is a symbol that appears in the input string at least k -times (‘yes’) or each symbol appears at most $k - 1$ times (‘no’).
6. For the k -SUM problem, we require that Σ is an additive group, and the problem is to decide whether there exists k distinct indices $i_1, i_2, \dots, i_k \in [n]$ such that $x_{i_1} + x_{i_2} + \dots + x_{i_k} = 0$ (‘yes’) or not (‘no’).

For the COLLISION, SET EQUALITY, and HIDDEN SHIFT problems, let the length of the input be $2n$.

7. The COLLISION problem is to decide whether each symbol present in the input string is unique (‘no’) or appears in it exactly twice (‘yes’), given a promise that either case holds.
8. The SET EQUALITY problem is a special case of COLLISION given an additional promise that each symbol of the first half of the input string is unique (thus, each symbol of the second half is unique too).
9. The HIDDEN SHIFT (decision) problem is a special case of SET EQUALITY given an additional promise that, for all yes-inputs x , there exists a unique $s \in [n]$ such that, for all $i \in [1..n]$ and all $j \in [n + 1..2n]$, $x_i = x_j$ if and only if $j \equiv i + s \pmod n$. (The “non-decision” version of the problem is to find the *hidden shift* s).

For the TRIANGLE and TRANGLE-SUM problems, let the length of the input be $\binom{n}{2}$ and let the input variables be labeled as x_{ij} where $1 \leq i < j \leq n$. Here n represents an order of a graph and $i \in [n]$ represents a vertex in the graph. A pair (i, j) with $1 \leq i < j \leq n$ represents an undirected edge, and an input $x \in \Sigma^{\binom{n}{2}}$ effectively assigns a symbol to every edge.

10. For the TRIANGLE problem, $\Sigma = \{0,1\}$, and we interpret x_{ij} as the indicator function whether there is ($x_{ij} = 1$) or is not ($x_{ij} = 0$) an edge between vertices i and j . The TRIANGLE problem is to decide whether the graph given by x contains a triangle ('yes') or not ('no'), that is, decide whether there is a triple $1 \leq a < b < c \leq n$ such that $x_{ab} = x_{ac} = x_{bc} = 1$.
11. For the TRANGLE-SUM problem, similarly to k -SUM, we require that Σ is an additive group, and the problem is to decide whether there is a triple $1 \leq a < b < c \leq n$ such that $x_{ab} + x_{ac} + x_{bc} = 0$ ('yes') or not ('no').

When we think of k -DISTINCTNESS and k -SUM as families of problems, we think of k as a constant independent from n . For ELEMENT DISTINCTNESS and k -DISTINCTNESS, we assume that $|\Sigma| \geq n$ and $|\Sigma| \geq n/(k-1)$, respectively, as these problems become trivial for smaller input alphabets. We later generalize k -SUM and the TRANGLE-SUM problems as CERTIFICATE-SUM problems and further as ORTHOGONAL ARRAY problems.

2.1.2 Certificates for decision problems

We can think of every input string $x \in \Sigma^n$ as a function $[n] \rightarrow \Sigma$ that assigns to every index i in $[n]$ the symbol x_i . For this reason, Σ is also commonly called the *range*. Given a subset of indices $S \subseteq [n]$, let $a \in \Sigma^S$ denote a function $S \rightarrow \Sigma$ that assigns a character to every index in S ; we call such a function a (*partial*) *assignment*. Recall that, given $x \in \Sigma^n$, $x_S \in \Sigma^S$ is the *restriction* of x to S . (Note that we can also further restrict partial assignments.) We say that $x \in \Sigma^n$ is *compatible* with $a \in \Sigma^S$ if $x_S = a$.

Suppose $\mathcal{P}: \Sigma^n \rightarrow \{0,1\}$ is a decision problem with domain $\mathcal{D} = \mathcal{D}_0 \sqcup \mathcal{D}_1$. For a value $b \in \{0,1\}$ and a subset $S \subseteq [n]$, the partial assignment $a \in \Sigma^S$ is called a *b-certificate* if $\mathcal{P}(x) = b$ for all $x \in \mathcal{D}$ such that $x_S = a$; one also says that a is a certificate of x . Given $x \in \mathcal{D}_b$, let $M_x \subseteq 2^{[n]}$ be the set of all subsets $S \subseteq [n]$ such that x_S is a b -certificate. Note that M_x is closed under taking supersets: if $S \in M_x$, then $S' \in M_x$ for all $S' \supseteq S$. We call a subset $M \subseteq 2^{[n]}$ that is closed under taking supersets a *certificate placement*.

Example 2.2. Consider the 4-bit AND-OF-OR's function: $\mathcal{P}(x) := (x_1 \vee x_2) \wedge (x_3 \vee x_4)$. For two 1-inputs $x := (1, 1, 0, 1)$ and $x := (0, 1, 0, 1)$, their corresponding certificate placements are

$$M_{(1,1,0,1)} = \{\{1, 4\}, \{2, 4\}, \{1, 2, 4\}, \{1, 3, 4\}, \{2, 3, 4\}, \{1, 2, 3, 4\}\},$$

$$M_{(0,1,0,1)} = \{\{2, 4\}, \{1, 2, 4\}, \{2, 3, 4\}, \{1, 2, 3, 4\}\}.$$

The *certificate complexity* $C(x)$ of an input $x \in \mathcal{D}_b$ is the minimum among the sizes of certificates of x , namely, $\min_{S \in M_x} |S|$. For $b \in \{0,1\}$, the *b-certificate complexity* of the problem

\mathcal{P} is defined as

$$C_b := \max_{x \in \mathcal{D}_b} C(x).$$

We will be particularly interested in problems having small 1-certificates. Note that it is the case for all decision problems we introduced above in Section 2.1.1. Because of this, most common algorithms for such problems try to look for a 1-certificate, and the result they output is determined depending on whether they succeed to find one.

Definition 2.3 (Certificate Structure). A *certificate structure* \mathcal{C} on n variables is a collection of non-empty certificate placements $M \subseteq 2^{[n]}$. We say that a decision problem $\mathcal{P}: \mathcal{D} \rightarrow \{0, 1\}$ has certificate structure \mathcal{C} if, for every $x \in \mathcal{D}_1$, one can find $M' \in \mathcal{C}$ such that $M' \subseteq M_x$, that is, x_S is a 1-certificate for all $S \in M'$.

A decision problem can have multiple certificate structures. For example, all decision problems on n variables have the trivial certificate structure $\{\{[n]\}\}$. Every problem \mathcal{P} also has the certificate structure $\{M_x: x \in \mathcal{D}_1\}$, which we call its *full certificate structure*. If we take the inclusion-wise minimal elements of the full certificate structure, we obtain the *minimal certificate structure* of the problem, which we denote by $\mathcal{C}_{\mathcal{P}}$. The minimal certificate structure, in some sense, corresponds to 1-inputs that are hardest to distinguish from 0-inputs.

Example 2.4. Consider again the function $\mathcal{P}(x) = (x_1 \vee x_2) \wedge (x_3 \vee x_4)$. Its minimal certificate structure is

$$\begin{aligned} \mathcal{C} = & \left\{ \left\{ \{1, 3\}, \{1, 2, 3\}, \{1, 3, 4\}, \{1, 2, 3, 4\} \right\}, \right. \\ & \left\{ \{1, 4\}, \{1, 2, 4\}, \{1, 3, 4\}, \{1, 2, 3, 4\} \right\}, \\ & \left\{ \{2, 3\}, \{1, 2, 3\}, \{2, 3, 4\}, \{1, 2, 3, 4\} \right\}, \\ & \left. \left\{ \{2, 4\}, \{1, 2, 4\}, \{2, 3, 4\}, \{1, 2, 3, 4\} \right\} \right\}. \end{aligned}$$

Each row in this expression is a certificate placement.

Assuming that the size of the input alphabet Σ is sufficiently large, both k -DISTINCTNESS and k -SUM share the same minimal certificate structure (which is also the minimal certificate structure of THRESHOLD- k):

Definition 2.5. The *k -subset certificate structure* \mathcal{C} on n elements with $k = O(1)$ is defined as follows. It has $\binom{n}{k}$ elements and, for each subset $A \subseteq [n]$ of size k , there exists a unique certificate placement $M \in \mathcal{C}$ such that $S \in M$ if and only if $A \subseteq S \subseteq [n]$.

One could define the unique certificate structure of a function to be its minimal certificate structure. All the results presented in this thesis would still hold in this case, because they all

consider functions of large input alphabets Σ . Nonetheless, the proofs of these results would become slightly more complex, because, when we design a function to have some given certificate structure, for smaller input alphabets, the function might become “simpler” and its minimal certificate structure might change. For example, when $|\Sigma| < n/k$, the minimal certificate structure of k -DISTINCTNESS becomes $\{2^{[n]}\}$.

For each COLLISION, SET EQUALITY, and HIDDEN SHIFT, its full and minimal certificate structure is the same and defined as follows. Note that every positive input x of these problems decomposes the set of indices $[2n]$ as a disjoint union of pairs $\{\mu_{i,1}, \mu_{i,2}\}$ such that $x[\mu_{i,1}] = x[\mu_{i,2}]$.

Definition 2.6. Each of the following certificate structures is defined on $2n$ input variables. In the *collision certificate structure*, there is a unique certificate placement M for each decomposition

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

and $S \in M$ if and only if $S \supseteq \{\mu_{i,1}, \mu_{i,2}\}$ for some $i \in [n]$. The *set equality certificate structure* contains only those M from the collision certificate structure that correspond to decompositions with $\mu_{i,1} \in [1..n]$ and $\mu_{i,2} \in [n+1..2n]$ for all i . Finally, the *hidden shift certificate structure* contains only those M from the set equality certificate structure that correspond to decompositions such that $s \in [n]$ exists with the property $\mu_{i,2} \equiv \mu_{i,1} + s \pmod{n}$ for all $i \in [n]$.

Example 2.7. The collision certificate structure for $n = 2$ is

$$\begin{aligned} \mathcal{C} = \{ & \{ \{1, 2\}, \{3, 4\}, \{1, 2, 3\}, \{1, 2, 4\}, \{1, 3, 4\}, \{2, 3, 4\}, \{1, 2, 3, 4\} \}, \\ & \{ \{1, 3\}, \{2, 4\}, \{1, 2, 3\}, \{1, 2, 4\}, \{1, 3, 4\}, \{2, 3, 4\}, \{1, 2, 3, 4\} \}, \\ & \{ \{1, 4\}, \{2, 3\}, \{1, 2, 3\}, \{1, 2, 4\}, \{1, 3, 4\}, \{2, 3, 4\}, \{1, 2, 3, 4\} \} \}. \end{aligned}$$

To further illustrate that a problem can have multiple certificate structures, notice that both 2-DISTINCTNESS (i.e., ELEMENT DISTINCTNESS) and COLLISION have the 2-subset certificate structure, but, between them, only COLLISION has the collision certificate structure.

Recall that, for TRIANGLE and TRANGLE-SUM, the input variables correspond to edges of an undirected graph. The minimal certificate structure for TRIANGLE and (assuming sufficiently large Σ) TRANGLE-SUM is as follows.

Definition 2.8. The *triangle certificate structure* \mathcal{C} on n vertices is a certificate structure on $\binom{n}{2}$ variables defined as follows. Assume that the variables are labelled as x_{ij} where $1 \leq i < j \leq n$. The certificate structure has $\binom{n}{3}$ elements, and, for every triple $1 \leq a < b < c \leq n$, there exists a unique certificate placement $M \in \mathcal{C}$ such that $S \in M$ if and only if $S \supseteq \{ab, bc, ac\}$.

Let us generalize the k -SUM and TRANGLE-SUM problems. For that reason, we consider a certain type of certificate structures.

Definition 2.9. A certificate structure \mathcal{C} is *boundedly generated* if, for every $M \in \mathcal{C}$, there is a (unique) subset $A_M \subseteq [n]$ such that $|A_M| = O(1)$, and $S \in M$ if and only if $S \supseteq A_M$.

Given a boundedly generated certificate structure \mathcal{C} , the \mathcal{C} -SUM problem is: given $x \in \Sigma^n$, decide whether there exists $M \in \mathcal{C}$ such that $\sum_{j \in A_M} x_j = 0$ (‘yes’) or not (‘no’). When we do not have a specific boundedly generated \mathcal{C} in mind, we call \mathcal{C} -SUM a CERTIFICATE-SUM problem. Note that k -SUM and TRANGLE-SUM correspond to \mathcal{C} being the k -subset and the triangle certificate structure, respectively. We further generalize CERTIFICATE-SUM problems as ORTHOGONAL ARRAY problems in Chapter 4.

2.2 Quantum query algorithm

Let us now define quantum query algorithms. Since in this thesis we only consider query complexity, we ignore aspects of quantum query algorithms that are only of importance when quantum circuit complexity is considered. For example, we do not consider qubits, which are basic units of quantum information, and, instead of defining a register of a quantum algorithm to be a collection of qubits, we will simply define it as a finite-dimensional space. We also do not worry about the *implementation costs* of unitary transformations.

For the illustrative examples of *quantum circuits* (Figures 2.1, 2.2, 2.3, 2.4, and 2.5), we assume that the reader is familiar with the basic notation of *quantum circuit diagrams* (see, e.g., [NC00]). In these diagrams, the computation proceeds from the left to the right and each *wire* (i.e., horizontal line) corresponds to one quantum register.

2.2.1 Registers and states of the computation

An algorithm will access the input $x \in \Sigma^{[n]}$ via an oracle, and we call each such access a *query*. Informally, an oracle receives an index $i \in [n]$ and returns the i -th entry of the input, i.e., $x_i \in \Sigma$. In order to make this operation reversible, we require the alphabet Σ to be a group.

Given a computational problem $\mathcal{P} \subseteq \Sigma^n \times R$, let $\mathcal{X}_{Q'} := \mathbb{C}^{[n]}$, $\mathcal{X}_{Q''} := \mathbb{C}^\Sigma$, and $\mathcal{X}_Q := \mathcal{X}_{Q'} \otimes \mathcal{X}_{Q''}$. We call these spaces—as well as similar spaces \mathcal{X}_W , \mathcal{X}_R , \mathcal{X}_A , \mathcal{X}_I , \mathcal{X}_S introduced later—*registers*. We call \mathcal{X}_Q the *query* register, $\mathcal{X}_{Q'}$ the *query index* (or, simply, *index*) register, and $\mathcal{X}_{Q''}$ the *query symbol* register. The (*standard*) *oracle* is a function \mathcal{O} that for every $x \in \Sigma^n$ assigns the unitary $\mathcal{O}(x) \in U(\mathcal{X}_Q)$ defined as

$$\mathcal{O}(x): \mathbf{i} \otimes \mathbf{a} \mapsto \mathbf{i} \otimes (\mathbf{a} + \mathbf{x}_i) \tag{2.1}$$

for all $i \in [n]$ and $\mathbf{a} \in \Sigma$ (see Figure 2.1), where, according to our notation, $(\mathbf{a} + \mathbf{x}_i)$ is the vector in the standard basis of \mathbb{C}^Σ that corresponds to $(\mathbf{a} + \mathbf{x}_i) \in \Sigma$. We also use the term ‘oracle’ to refer to the operator $\mathcal{O}(x)$.

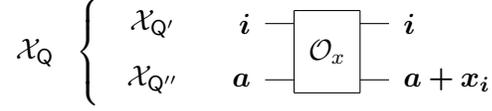


Figure 2.1: The circuit diagram of the standard quantum oracle. For the sake of conciseness, in circuit diagrams we write \mathcal{O}_x instead of $\mathcal{O}(x)$.

For now, let us simply refer to computational problems as ‘problems’. Later, in Chapter 6, we consider the ENHANCED FIND-TWO problem, which is technically not a computational problem. There we will also consider other (i.e., non-standard) oracles, as well as, provide an algorithm with additional information on the input.

Recall that R denotes the codomain of \mathcal{P} . Similarly to the query register, let $\mathcal{X}_R := \mathbb{C}^R$, which is called the *output* register. A vector in a register is called a *state* of this register if it has unit norm, and similarly for vectors in tensor products of registers. In order to clarify that a state belongs to (or an operator acts on) a register \mathcal{X}_{reg} , where *reg* is a subscript identifying the register, we frequently use the same subscript for the state (or the operator) itself. This also allows us to change the order of registers in expressions, when convenient. For example, given $u \in \mathcal{X}_Q$ and $v \in \mathcal{X}_R$, we think of $u_Q \otimes v_R$ and $v_R \otimes u_Q$ as the same state in $\mathcal{X}_Q \otimes \mathcal{X}_R$. We may also concatenate subscripts when we address multiple registers at once, for example, we may write \mathbb{I}_{QR} instead of $\mathbb{I}_Q \otimes \mathbb{I}_R$.

For simplicity, let Tr_{reg} denote a partial trace over \mathcal{X}_{reg} (instead of $\text{Tr}_{\mathcal{X}_{\text{reg}}}$). Given a state $w \in \mathcal{X}_{\text{reg}_1} \otimes \mathcal{X}_{\text{reg}_2}$, we refer to the density operator $\text{Tr}_{\text{reg}_2}(ww^*)$ as the state of the register $\mathcal{X}_{\text{reg}_1}$ or as the $\mathcal{X}_{\text{reg}_1}$ -*part* of the state w . We say that $\mathcal{X}_{\text{reg}_1}$ and $\mathcal{X}_{\text{reg}_2}$ are *entangled* if $\text{rank}(\text{Tr}_{\text{reg}_2}(ww^*)) > 1$.

Definition 2.10. A *quantum query algorithm* \mathcal{A} for a problem $\mathcal{P} \subseteq \Sigma^n \times R$ is a quadruple

$$\mathcal{A} := (\mathcal{X}_W, \phi_0 \in \mathcal{X}_A, T \in \mathbb{N}, \{V_t \in \text{U}(\mathcal{X}_A) : t \in [0..T]\}),$$

where \mathcal{X}_W is a finite, non-zero Hilbert space, $\mathcal{X}_A := \mathcal{X}_Q \otimes \mathcal{X}_W \otimes \mathcal{X}_R$, and $\|\phi_0\| = 1$. We call \mathcal{X}_W the *work* register and \mathcal{X}_A the *algorithm* registers. The state ϕ_0 is called the *initial state* of the algorithm, T is the *number of queries*. V_0 and V_T are called, respectively, the *initial* and the *final* unitary transformation, and, for $t \in [T - 1]$, V_t is called the unitary transformation between the queries (number) t and $t + 1$. We also refer to T as the *quantum query complexity* of \mathcal{A} .

The circuit diagram of a generic quantum query algorithm \mathcal{A} is given in Figure 2.2. Let $\phi_t(x)$ be the state of \mathcal{A} *just after* the query t and $\psi_t(x)$ the state of \mathcal{A} *just before* the query $t + 1$. And let $\psi_T(x)$ be the *final state* of the algorithm \mathcal{A} .

More formally, for $t \in [0..T]$, we define the states $\phi_t(x)$ and $\psi_t(x)$ in \mathcal{X}_A recursively as follows. As the base case, we use the initial state $\phi_0(x) = \phi_0$. (Note: for the ENHANCED FIND-TWO

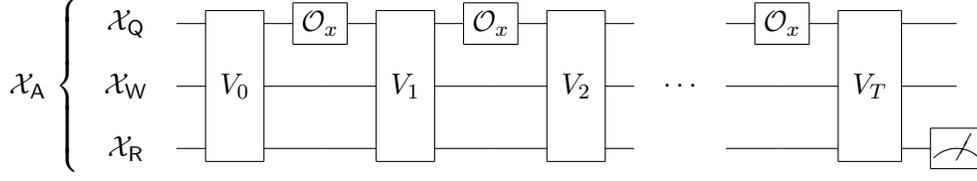


Figure 2.2: The circuit diagram of a generic quantum query algorithm.

problem considered in Chapter 6, the initial state $\phi_0(x)$ will depend on x . It is also the case in state conversion problems considered in [LMR⁺11].) Then, let

$$\forall t \in [0..T]: \psi_t(x) := V_t \phi_t(x) \quad \text{and} \quad \forall t \in [1..T]: \phi_t(x) := (\mathcal{O}(x) \otimes \mathbb{I}_{\mathcal{W}\mathcal{R}}) \psi_{t-1}(x). \quad (2.2)$$

Note that

$$\langle \psi_t(x), \psi_t(y) \rangle = \langle V_t \phi_t(x), V_t \phi_t(y) \rangle = \langle \phi_t(x), \phi_t(y) \rangle. \quad (2.3)$$

At the end of the computation, the algorithm \mathcal{A} returns r with the probability

$$\|(\mathbb{I}_{\mathcal{Q}\mathcal{W}} \otimes \mathbf{r}_{\mathcal{R}})^* \psi_T(x)\|^2,$$

which, in the standard terminology of quantum computing, is the probability that, upon a *measurement* of $\psi_T(x)$ in the standard basis of $\mathcal{X}_{\mathcal{R}}$, the outcome of the measurement is r (see [NC00]). The algorithm is successful if the returned value r is a correct solution on the input x . Hence, the *success probability* of \mathcal{A} on an input x is

$$p_{\mathcal{A}}(x) := \sum_{r \in \mathcal{P}(x)} \|(\mathbb{I}_{\mathcal{Q}\mathcal{W}} \otimes \mathbf{r}_{\mathcal{R}})^* \psi_T(x)\|^2, \quad (2.4)$$

and we call $1 - p_{\mathcal{A}}(x)$ the *error probability*. We study the *worst case* success probability of \mathcal{A} , which is defined as $p_{\mathcal{A}} := \min\{p_{\mathcal{A}}(x) : x \in \mathcal{D}\}$. We say that \mathcal{A} *solves* the problem \mathcal{P} with the error probability $1 - p_{\mathcal{A}}$.

Given $0 < \varepsilon < 1/2$, the ε -*error quantum query complexity* of a problem \mathcal{P} is the minimum number of queries (i.e., T) required by any algorithm that solves \mathcal{P} with the error probability at most ε , and we denote this complexity by $Q_{\varepsilon}(\mathcal{P})$. We define the *bounded-error quantum query complexity* (or, simply, the *quantum query complexity*) of \mathcal{P} to be $Q_{1/3}(\mathcal{P})$.

If \mathcal{P} is a function, the choice for the constant $\varepsilon = 1/3$ is arbitrary in the following sense. We are typically interested in a family of problems $\{\mathcal{P}_n : n \in \mathbb{N}\}$, and we only care about the asymptotic query complexity. Given any constant $\varepsilon > 0$ independent from n , we can run an algorithm of error probability $1/3$ multiple times and take the majority of answers, thereby reducing the error probability below ε . (The number of times we have to repeat the algorithm depends only on ε , so this number is a constant with respect to n .)

We define the *quantum query complexity of a certificate structure* as the maximum bounded-error quantum query complexity over all decision problems possessing this certificate structure.

Remark 2.11. We could have also allowed quantum query algorithms to have a “control” register $\mathcal{X}_C := \mathbb{C}^{\{0,1\}}$ that determines whether to query the oracle: query the oracle if \mathcal{X}_C is in the state $\mathbf{1}$ and do not query the oracle if \mathcal{X}_C is in the state $\mathbf{0}$. Nevertheless, such a register would not increase the power of quantum query algorithms because we can simulate this controlled oracle using the fact that

$$\mathcal{O}(x)(\mathbf{i} \otimes \xi) = \mathbf{i} \otimes \xi \quad (2.5)$$

for all $x \in \Sigma^n$ and $i \in [n]$, where

$$\xi := \frac{1}{\sqrt{|\Sigma|}} \sum_{a \in \Sigma} \mathbf{a}.$$

To perform this simulation, we introduce an ancillary register $\mathcal{X}'_{Q''}$ isomorphic to $\mathcal{X}_{Q''}$ that is initialized to ξ . Then, right before and right after each query, we swap the registers $\mathcal{X}_{Q''}$ and $\mathcal{X}'_{Q''}$ if and only if \mathcal{X}_C is in the state $\mathbf{0}$. This simulation works because, as shown by (2.5), applying the oracle to $\mathbf{i} \otimes \xi$ is equivalent to not querying the oracle at all.

2.2.2 Automorphisms of problems and symmetrization

Many problems that we are interested in, including all the problems we defined above, possess many symmetries. We describe this symmetry using automorphism groups. The *automorphism principle*, which we present in Section 2.3.2, states that, without loss of generality, one can assume that adversary matrices are symmetric under automorphism groups [HLŠ07]. This symmetry was heavily utilized for the $\Omega(n^{1/2})$ lower bound for the INDEX ERASURE problem [AMRR11]. Automorphism groups also play a role in the *symmetrization* of algorithms, which we present later in this section.

In order to define automorphism groups, let us consider the (left) group actions of $\mathbb{S}_{[n]}$, \mathbb{S}_Σ , and $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$ on the set of inputs Σ^n defined as follows. Given a permutation of indices $\pi \in \mathbb{S}_{[n]}$ and a permutation of symbols $\tau \in \mathbb{S}_\Sigma$, we define the group actions of $\mathbb{S}_{[n]}$ and \mathbb{S}_Σ as, respectively,

$$\pi: (x_1, \dots, x_n) \mapsto (x_{\pi^{-1}(1)}, \dots, x_{\pi^{-1}(n)}), \quad (2.6)$$

$$\tau: (x_1, \dots, x_n) \mapsto (\tau(x_1), \dots, \tau(x_n)). \quad (2.7)$$

These two actions commute, $\pi(\tau(x)) = \tau(\pi(x))$, so we define the group action of $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$ as

$$(\pi, \tau): (x_1, \dots, x_n) \mapsto (\tau(x_{\pi^{-1}(1)}), \dots, \tau(x_{\pi^{-1}(n)})). \quad (2.8)$$

Let $U_{Q', \pi} \in \mathbf{U}(\mathcal{X}_{Q'})$ and $U_{Q'', \tau} \in \mathbf{U}(\mathcal{X}_{Q''})$ be the permutation representations corresponding to the group actions $\pi: i \mapsto \pi(i)$ and $\tau: a \mapsto \tau(a)$, respectively, where $i \in [n]$ and $a \in \Sigma$. They

are known as the *natural* representations of the symmetric group. Let $U_{\mathbb{Q},(\pi,\tau)} := U_{\mathbb{Q}',\pi} \otimes U_{\mathbb{Q}'',\tau}$, which is a representation of $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$.

Definition 2.12. An *automorphism* of a problem $\mathcal{P} \subseteq \Sigma^n \times R$ is a subgroup $G \leq \mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$ such that there exists a group action $\omega: G \times R \rightarrow R$ satisfying

$$\forall \gamma \in G, \forall x \in \Sigma^n: \omega_\gamma(\mathcal{P}(x)) = \mathcal{P}(\gamma(x)), \quad (2.9)$$

where $\gamma(x)$ is defined in (2.8). We call the automorphism G of \mathcal{P} an *oracle automorphism* if

$$U_{\mathbb{Q},\gamma} \cdot \mathcal{O}(x) \cdot U_{\mathbb{Q},\gamma}^{-1} = \mathcal{O}(\gamma(x)) \quad (2.10)$$

for all $x \in \mathcal{D}$ and all $\gamma \in G$.

Example 2.13. Consider the ELEMENT DISTINCTNESS problem. The group $G := \mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$ together with the trivial group action (namely, $\omega_\gamma(r) = r$ for all $\gamma \in G$ and $r \in \{0, 1\}$) is an automorphism for ELEMENT DISTINCTNESS. However, it is not an oracle automorphism. For $\tau \in \mathbb{S}_\Sigma$, the action of $U_{\mathbb{Q},(\varepsilon,\tau)} \mathcal{O}(x) U_{\mathbb{Q},(\varepsilon,\tau)}^{-1}$ on the standard basis is

$$\mathbf{i} \otimes \mathbf{a} \xrightarrow{\tau^{-1}} \mathbf{i} \otimes \tau^{-1}(\mathbf{a}) \xrightarrow{\mathcal{O}(x)} \mathbf{i} \otimes (\tau^{-1}(\mathbf{a}) + \mathbf{x}_i) \xrightarrow{\tau} \mathbf{i} \otimes \tau(\tau^{-1}(\mathbf{a}) + \mathbf{x}_i),$$

while, for $\mathcal{O}((\varepsilon,\tau)(x))$, it is

$$\mathbf{i} \otimes \mathbf{a} \xrightarrow{\mathcal{O}(\tau(x))} \mathbf{i} \otimes (\mathbf{a} + \tau(\mathbf{x}_i)),$$

where ε is the identity element of $\mathbb{S}_{[n]}$. These actions are not necessarily equal. For example, consider $\Sigma = \{0, 1, 2, \dots, |\Sigma| - 1\}$ with the addition modulo $|\Sigma|$ as the group operation and the transposition $\tau := (0, 1)$. Then, for $a = 1$ and $x_i = 0$, we have

$$\tau(\tau^{-1}(a) + x_i) = 1 \quad \text{and} \quad a + \tau(x_i) = 2,$$

therefore (2.10) does not hold.

While automorphisms play a role in the adversary bound, oracle automorphisms, which are more restricted, play a role in the symmetrization of algorithms. The condition (2.10) does not hold for the standard oracle \mathcal{O} , unless $\tau(a) + \tau(a') = \tau(a + a')$ for all $a, a' \in \Sigma$ and all $(\pi, \tau) \in G$. On the other hand, consider automorphisms in the form $G^* \times \{\varepsilon\}$, where G^* is a subgroup of $\mathbb{S}_{[n]}$ and ε is the identity element of \mathbb{S}_Σ . For $\pi \in G^*$, the action of $U_{\mathbb{Q},(\pi,\varepsilon)} \mathcal{O}(x) U_{\mathbb{Q},(\pi,\varepsilon)}^{-1}$ on the standard basis is

$$\begin{aligned} \mathbf{i} \otimes \mathbf{a} \xrightarrow{\pi^{-1}} \pi^{-1}(\mathbf{i}) \otimes \mathbf{a} \xrightarrow{\mathcal{O}(x)} \pi^{-1}(\mathbf{i}) \otimes (\mathbf{a} + \mathbf{x}_{\pi^{-1}(\mathbf{i})}) &= \pi^{-1}(\mathbf{i}) \otimes (\mathbf{a} + \pi(\mathbf{x})_i) \\ &\xrightarrow{\pi} \mathbf{i} \otimes (\mathbf{a} + \pi(\mathbf{x})_i), \end{aligned}$$

which equals the action of $\mathcal{O}((\pi, \varepsilon)(x))$. Hence, the condition (2.10) holds.¹

Suppose G is an oracle automorphism of a problem \mathcal{P} with a group action ω_γ on R , and let $U_{R,\gamma}$ be the permutation representation of G corresponding to ω_γ . Given a quantum query algorithm \mathcal{A} for \mathcal{P} , one can “symmetrize” it so that the *symmetrized algorithm* $\bar{\mathcal{A}}$ performs equally well within each orbit $G(x)$, no worse than \mathcal{A} in the worst case, and uses the same number of queries as \mathcal{A} . To use symmetrization as a tool for proving lower bounds was first considered by Ambainis in [Amb10], and later used in [AŠdW09]. We “symmetrize” \mathcal{A} as follows.

Without loss of generality, we assume that the initial state satisfies

$$(U_{Q,\gamma} \otimes \mathbb{I}_{WR})\phi_0 = \phi_0 \quad (2.11)$$

for all $\gamma \in G$. (In Chapter 6, where $\phi_0(x)$ depends on x , we will require

$$(U_{Q,\gamma} \otimes \mathbb{I}_{WR})\phi_0(x) = \phi_0(\gamma(x))$$

for all $x \in \mathcal{D}$ and all $\gamma \in G$.) Consider $\gamma \in G$. Let the algorithm $\mathcal{A}^{(\gamma)}$ be obtained from \mathcal{A} by replacing V_t with

$$V_t^{(\gamma)} := (U_{Q,\gamma} \otimes \mathbb{I}_{WR})^{-1}V_t(U_{Q,\gamma} \otimes \mathbb{I}_{WR})$$

for all $t \in [0..T-1]$ and V_T with

$$V_T^{(\gamma)} := (\mathbb{I}_{QW} \otimes U_{R,\gamma})^{-1}V_T(U_{Q,\gamma} \otimes \mathbb{I}_{WR}).$$

Hence, because of the conditions (2.10) and (2.11), the state of $\mathcal{A}^{(\gamma)}$ running on an input x just before the query $t+1$ is

$$(U_{Q,\gamma} \otimes \mathbb{I}_{WR})^{-1}\psi_t(\gamma(x)), \quad (2.12)$$

where $\psi_t(\gamma(x))$ is the state of \mathcal{A} on $\gamma(x)$ just before the query $t+1$. Similarly, the final state of $\mathcal{A}^{(\gamma)}$ on x is

$$(\mathbb{I}_{QW} \otimes U_{R,\gamma})^{-1}\psi_T(\gamma(x)).$$

Because of the condition (2.9) on ω_γ , the success probability of $\mathcal{A}^{(\gamma)}$ on x equals the success probability of \mathcal{A} on $\gamma(x)$.

To average out success probabilities within each orbit $G(x)$, let us effectively run algorithms $\mathcal{A}^{(\gamma)}$ in *superposition* over all $\gamma \in G$. Namely, we add the *symmetrization* register $\mathcal{X}_S := \mathbb{C}^G$ to the registers of the algorithm \mathcal{A} . (To see that this fits with Definition 2.10, one can think of $\mathcal{X}_S \otimes \mathcal{X}_W$ as the new work register.) Initially, let \mathcal{X}_S hold the uniform superposition

$$\frac{1}{\sqrt{|G|}} \sum_{\gamma \in G} \gamma$$

¹Suppose, aside from the standard oracle \mathcal{O} , we were also given the oracle \mathcal{O}^{-1} . Then, a pair of one query to $\mathcal{O}(x)$ and one query to $\mathcal{O}^{-1}(x)$ can simulate a query to $\mathcal{O}(\gamma(x))$ for all $\gamma \in \mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$ (see, for example, [AMRR11]).

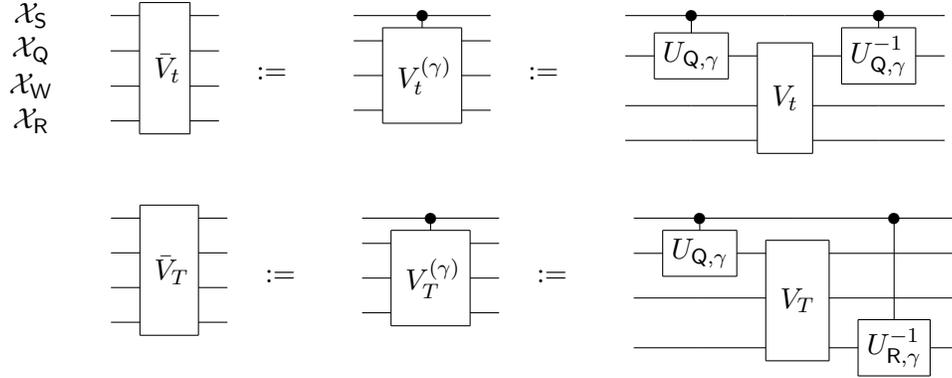


Figure 2.3: Unitary transformations of the symmetrized algorithm $\bar{\mathcal{A}}$: (top) the transformation between the queries t and $t+1$, if $t \in [T-1]$, or the initial transformation, if $t = 0$; (bottom) the final transformation (i.e., $t = T$). The controlled unitary transformations are controlled by $\gamma \in G$.

over all “permutations” $\gamma \in G$. And, for $\bar{\mathcal{A}}$, we substitute each unitary $V_t \in \mathbf{U}(\mathcal{X}_A)$ of \mathcal{A} with

$$\bar{V}_t := \sum_{\gamma \in G} \gamma \gamma^* \otimes V_t^{(\gamma)} \in \mathbf{U}(\mathcal{X}_S \otimes \mathcal{X}_A)$$

(see Figure 2.3). Thereby we ensure that, for all $x \in \Sigma^n$,

$$p_{\bar{\mathcal{A}}}(x) = \frac{1}{|G(x)|} \sum_{x' \in G(x)} p_{\mathcal{A}}(x') \geq \min_{x' \in G(x)} p_{\mathcal{A}}(x') \geq p_{\mathcal{A}}.$$

2.2.3 Algorithms with an input register

We just described how, using the symmetrization register, we can run multiple algorithms in superposition on one given input x . Analogously, we now consider how to run one given algorithm on a superposition of inputs. This argument is the basis of the adversary lower bound method as well as the lower bound we will prove for the ENHANCED FIND-TWO problem in Chapter 6.

Given a problem $\mathcal{P} : \mathcal{D} \rightarrow R$ and a unit vector $(\delta_x \in \mathbb{C} : x \in \mathcal{D})$, let us recast \mathcal{A} into a different form \mathcal{A}^+ , introducing the *input* register $\mathcal{X}_I := \mathbb{C}^{\mathcal{D}}$ that stores the input. The initial state of the algorithm \mathcal{A}^+ is

$$\phi_0^+ := \sum_{x \in \mathcal{D}} \delta_x \mathbf{x}_I \otimes \phi_0(x)_A \in \mathcal{X}_{IA}. \quad (2.13)$$

For \mathcal{A}^+ , each oracle \mathcal{O} is replaced by

$$\mathcal{O}^+ := \sum_{x \in \mathcal{D}} (\mathbf{x} \mathbf{x}^*)_I \otimes \mathcal{O}(x)_Q, \quad (2.14)$$

which we also call an *oracle* (see Figure 2.4), and each unitary transformation $V_t \in \mathbf{U}(\mathcal{X}_A)$ by $\mathbb{I}_1 \otimes V_t \in \mathbf{U}(\mathcal{X}_{1A})$. Similarly to (2.2), let $\phi_t^+ := (\mathbb{I}_1 \otimes V_t)\phi_t^+$ and $\phi_t^+ := (\mathcal{O}^+ \otimes \mathbb{I}_{WR})\psi_{t-1}^+$. Since $\mathbb{I}_1 \otimes V_t$ does not interact with the \mathcal{X}_1 register and the state of \mathcal{X}_1 only controls which input to consider, we have

$$\phi_t^+ = \sum_{x \in \mathcal{D}} \delta_x \mathbf{x}_1 \otimes \phi_t(x)_A \quad \text{and} \quad \psi_t^+ = \sum_{x \in \mathcal{D}} \delta_x \mathbf{x}_1 \otimes \psi_t(x)_A$$

for all $t \in [0..T]$.

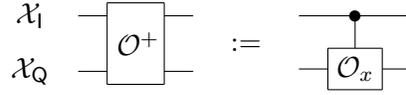


Figure 2.4: The oracle controlled by the input.

Let

$$\rho_t := \text{Tr}_A(\phi_t^+(\phi_t^+)^*) = \sum_{x,y \in \mathcal{D}} \delta_x \bar{\delta}_y \mathbf{x} \mathbf{y}^* \langle \phi_t(y), \phi_t(x) \rangle, \quad (2.15)$$

and, note that, due to (2.3), we also have

$$\rho_t = \text{Tr}_A(\psi_t^+(\psi_t^+)^*) = \sum_{x,y \in \mathcal{D}} \delta_x \bar{\delta}_y \mathbf{x} \mathbf{y}^* \langle \psi_t(y), \psi_t(x) \rangle. \quad (2.16)$$

That is, the state ρ_t of the register \mathcal{X}_1 stays the same between the queries t and $t+1$ (interpreting the queries -1 and $T+1$ as the beginning and the end of the algorithm, respectively).

The final state of the algorithm \mathcal{A}^+ is ψ_T^+ , and we define the success probability of \mathcal{A}^+ as

$$p_{\mathcal{A}^+} := \sum_{(x,r) \in \mathcal{P}} \|(\mathbf{x}_1 \otimes \mathbb{I}_{QW} \otimes \mathbf{r}_R)^* \psi_T^+\|^2. \quad (2.17)$$

From the definition (2.4) of the success probability of \mathcal{A} on x , one can see that

$$p_{\mathcal{A}^+} = \sum_{x \in \mathcal{D}} |\delta_x|^2 p_{\mathcal{A}}(x) \geq p_{\mathcal{A}}.$$

Therefore, if one lower bounds the number of queries to \mathcal{O}^+ required by any \mathcal{A}^+ in order to err with probability at most ε , this is automatically a lower bound on $Q_\varepsilon(\mathcal{P})$.

2.2.4 Symmetries of the input register

Now suppose we add both registers \mathcal{X}_1 and \mathcal{X}_S to the algorithm \mathcal{A} , that is, we both symmetrize the algorithm and run it in superposition over multiple inputs. Let $\bar{\mathcal{A}}^+$ denote the resulting

algorithm, whose circuit diagram is given in Figure 2.5. Again suppose that G is an oracle automorphism of a problem \mathcal{P} with U_R being the permutation representation corresponding to the group action of G on R , and suppose (δ_x) is the unit vector determining the superposition over inputs. Let us assume that δ_x is the same within each orbit of G , namely, $\delta_x = \delta_{\gamma(x)}$ for all $x \in \mathcal{D}$ and $\gamma \in G$.

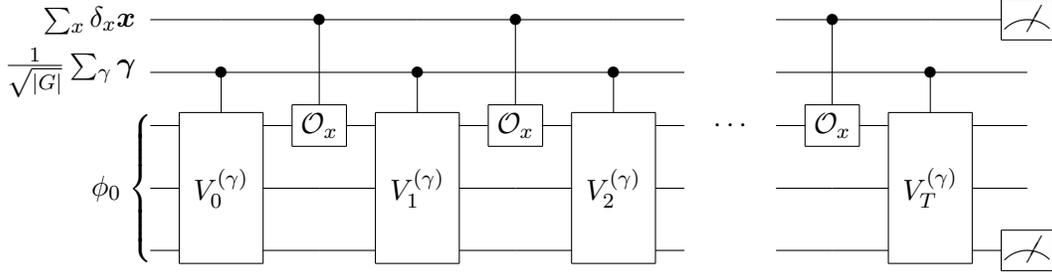


Figure 2.5: The circuit diagram of a generic quantum query algorithm with the input and symmetrization registers.

Note that, for $\gamma = (\pi, \tau) \in \mathbb{S}_{[n]} \times \mathbb{S}_{\Sigma}$ and $\ell = (i, a) \in [n] \times \Sigma$, the unitary transformation $U_{Q,\gamma}$ corresponds to the group action $\gamma: \ell \mapsto (\pi(i), \tau(a))$. Let us write the state of the algorithm \mathcal{A} on x just before the query $t + 1$ as

$$\psi_t(x) = \sum_{\ell \in [n] \times \Sigma} \ell_Q \otimes \psi_{t,\ell}(x)_{WR},$$

where $\psi_{t,\ell}(x)$ may be unnormalized. Therefore, the state of the algorithm $\mathcal{A}^{(\gamma)}$ on x just before the query $t + 1$, according to (2.12), is

$$\sum_{\ell \in [n] \times \Sigma} \gamma^{-1}(\ell)_Q \otimes \psi_{t,\ell}(\gamma(x))_{WR},$$

and the state of the algorithm $\bar{\mathcal{A}}^+$ just before the query $t + 1$ is

$$\bar{\psi}_t^+ = \sum_{x \in \mathcal{D}} \delta_x \mathbf{x}_I \otimes \frac{1}{\sqrt{|G|}} \sum_{\gamma \in G} \gamma_S \otimes \sum_{\ell \in [n] \times \Sigma} \gamma^{-1}(\ell)_Q \otimes \psi_{t,\ell}(\gamma(x))_{WR}. \quad (2.18)$$

Let $U_{I,\gamma} \in \mathbf{U}(\mathcal{X}_I)$ be the representation corresponding to the group action (2.8), and let $U_{S,\gamma} \in \mathbf{U}(\mathcal{X}_S)$ be the right regular representation of G . Similarly as before, for any two representations

$U_{\text{reg}_1, \gamma}$ and $U_{\text{reg}_2, \gamma}$ of G , let $U_{\text{reg}_1, \text{reg}_2, \gamma} := U_{\text{reg}_1, \gamma} \otimes U_{\text{reg}_2, \gamma}$. For all $\kappa \in G$, we have

$$\begin{aligned} (U_{\text{ISQ}, \kappa} \otimes \mathbb{I}_{\text{WR}}) \bar{\psi}_t^+ &= \sum_{x \in \mathcal{D}} \delta_x \kappa(x)_1 \otimes \frac{1}{\sqrt{|G|}} \sum_{\gamma \in G} \gamma \kappa_S^{-1} \otimes \sum_{\ell \in [n] \times \Sigma} \kappa \gamma^{-1}(\ell)_Q \otimes \psi_{t, \ell}(\gamma(x))_{\text{WR}} \\ &= \sum_{x \in \mathcal{D}} \delta_{\kappa(x)} \kappa(x)_1 \otimes \frac{1}{\sqrt{|G|}} \sum_{\gamma \in G} \gamma \kappa_S^{-1} \otimes \sum_{\ell \in [n] \times \Sigma} (\gamma \kappa^{-1})^{-1}(\ell)_Q \otimes \psi_{t, \ell}((\gamma \kappa^{-1})(\kappa(x)))_{\text{WR}} = \bar{\psi}_t^+. \end{aligned} \quad (2.19)$$

Let $\rho'_t := \text{Tr}_{\text{SWR}}(\bar{\psi}_t^+ (\bar{\psi}_t^+)^*)$ for $t \in \{0, 1, \dots, T-1\}$ and let $\rho'_T := \text{Tr}_{\text{SQW}}(\bar{\psi}_T^+ (\bar{\psi}_T^+)^*)$, so that $\text{Tr}_{\text{Q}}(\rho'_t) = \rho_t$ and $\text{Tr}_{\text{R}}(\rho'_T) = \rho_T$. Due to (2.19), for all $t \in [0..T-1]$, we have

$$U_{\text{IQ}, \gamma} \rho'_t U_{\text{IQ}, \gamma}^{-1} = \rho'_t \quad \text{and} \quad U_{1, \gamma} \rho_t U_{1, \gamma}^{-1} = \rho_t \quad \text{for all } \gamma \in G. \quad (2.20)$$

Similarly to (2.19), one can show that

$$(U_{\text{ISR}, \kappa} \otimes \mathbb{I}_{\text{QW}}) \bar{\psi}_T^+ = \bar{\psi}_T^+ \quad \text{for all } \gamma \in G, \quad (2.21)$$

and, thus,

$$U_{\text{IR}, \gamma} \rho'_T U_{\text{IR}, \gamma}^{-1} = \rho'_T \quad \text{and} \quad U_{1, \gamma} \rho_T U_{1, \gamma}^{-1} = \rho_T \quad \text{for all } \gamma \in G. \quad (2.22)$$

We will use the symmetries (2.19), (2.20), (2.21), (2.22) in Chapter 6.

2.3 Adversary bound

Let us now introduce the main lower bound technique studied in this thesis, the adversary bound. The adversary bound addresses the bounded-error quantum query complexity of function evaluation. Consider a function $\mathcal{P}: \Sigma^n \rightarrow R$, and let \mathcal{D} be the domain of \mathcal{P} .

Definition 2.14. An *adversary matrix* for \mathcal{P} is a non-zero, real, symmetric $|\mathcal{D}| \times |\mathcal{D}|$ -matrix Γ whose rows and columns are symmetrically labeled by inputs $x \in \mathcal{D}$ and which satisfies

$$\Gamma[[x, y]] = 0 \quad \text{whenever} \quad \mathcal{P}(x) = \mathcal{P}(y). \quad (2.23)$$

For $i \in [n]$, the *difference matrices* Δ_i and $\bar{\Delta}_i$ are the matrices of the same dimensions and the same row and column labeling as Γ that are defined by

$$\Delta_i[[x, y]] := \begin{cases} 0, & \text{if } x_i = y_i, \\ 1, & \text{if } x_i \neq y_i, \end{cases} \quad \text{and} \quad \bar{\Delta}_i[[x, y]] := \begin{cases} 1, & \text{if } x_i = y_i, \\ 0, & \text{if } x_i \neq y_i. \end{cases} \quad (2.24)$$

Theorem 2.15 (Adversary bound, [HLŠ07]). *In the notation of Definition 2.14, the ε -error quantum query complexity of \mathcal{P} satisfies*

$$Q_\varepsilon(\mathcal{P}) \geq \frac{1 - 2\sqrt{\varepsilon}}{2} \text{Adv}(\mathcal{P}), \quad (2.25)$$

where

$$\text{Adv}(\mathcal{P}) := \max_{\Gamma} \frac{\|\Gamma\|}{\max_i \|\Delta_i \circ \Gamma\|}. \quad (2.26)$$

Hence, the bounded-error quantum query complexity of \mathcal{P} is $Q(\mathcal{P}) = \Omega(\text{Adv}(\mathcal{P}))$. If $R = \{0, 1\}$, one can use $2\sqrt{\varepsilon(1-\varepsilon)}$ instead of $2\sqrt{\varepsilon}$ in (2.25).²

We can always scale the adversary matrix Γ so that the denominator in (2.26) is at most 1, therefore, $\text{Adv}(\mathcal{P})$ is equivalent to the optimal value of the optimization problem

$$\text{maximize} \quad \|\Gamma\| \quad (2.27a)$$

$$\text{subject to} \quad \|\Delta_i \circ \Gamma\| \leq 1 \quad \text{for all } i \in [n], \quad (2.27b)$$

where the maximization is over all adversary matrices Γ for \mathcal{P} . In fact (2.27) is a semidefinite program (see [Rei09]), and, for decision problems, we consider its dual in Section 2.4.

We prefer to use (2.27) over (2.26). Note that every feasible solution to the semidefinite program (2.27) yields a lower bound on the quantum query complexity of \mathcal{P} . In practice, we typically care only about the asymptotic behaviour of the adversary bound and we use the condition $\|\Delta_i \circ \Gamma\| = O(1)$ instead of $\|\Delta_i \circ \Gamma\| \leq 1$. Also note that $\Delta_i \circ \Gamma = \Gamma - \bar{\Delta}_i \circ \Gamma$.

The adversary bound was first introduced by Ambainis in [Amb02], essentially considering the case when each entry of the adversary matrix Γ is either 0 or 1. Many generalizations of the bound were subsequently proposed, which were later all shown to be equivalent to the case when all entries (also called *weights*) of Γ are non-negative [ŠS06]. This version of the bound is known as the *positive-weights adversary bound*, and it suffers certain limitations we describe in Section 2.3.4. The adversary bound was further generalized in [HLŠ07] by allowing both positive and negative weights in Γ . Reichardt et al. showed that the (general) adversary bound is optimal for every function (up to a constant depending only on ε) [Rei11, LMR⁺11].

We do not repeat the proof of the adversary bound in this thesis. Instead, we provide the basic intuition behind it, which partially also applies for the lower bound for the ENHANCED FIND-TWO problem considered in Chapter 6.

²The adversary bound in [HLŠ07] had $2\sqrt{\varepsilon(1-\varepsilon)} + 2\varepsilon$ instead of $2\sqrt{\varepsilon}$ in (2.25). The proof having $2\sqrt{\varepsilon}$ can be found in [Bell3].

2.3.1 Intuition behind the bound

Consider a quantum query algorithm \mathcal{A} for a problem \mathcal{P} . As in Section 2.2.3, let \mathcal{A}^+ be the query algorithm running \mathcal{A} on a superposition of inputs determined by a unit vector $(\delta_x : x \in \mathcal{D})$. All quantum query lower bounds that we prove in this thesis are based on the following three observations regarding the entanglement between the input register \mathcal{X}_I and the algorithm registers \mathcal{X}_A . Informally,

1. At the beginning of the algorithm, \mathcal{X}_I and \mathcal{X}_A are not very entangled (for the problems that we consider, they are not entangled at all, except for the ENHANCED FIND-TWO problem in Chapter 6).
2. In order for the algorithm to err with small probability, at the end of the algorithm, \mathcal{X}_I and \mathcal{X}_A (in particular, its output “subregister” \mathcal{X}_R) have to be highly entangled.
3. A single query to the oracle cannot increase the entanglement between \mathcal{X}_I and \mathcal{X}_A by more than a certain amount, and unitary transformations V_t do not affect this entanglement whatsoever.

Of course, we need quantitative means for measuring this entanglement, and we will consider somewhat different measures for function evaluation and the ENHANCED FIND-TWO problem. Once we lower bound the total difference in the entanglement required between the beginning of the algorithm and the end of the algorithm and upper bound the possible change of the entanglement per one query, the fraction of these two bounds yields a lower bound on the quantum query complexity of \mathcal{P} . We will discuss the ENHANCED FIND-TWO problem in Chapter 6, and now let us focus on our other interest: function evaluation.

Let $\mathcal{P} : \Sigma^n \rightarrow R$ be a function, and let \mathcal{D} be the domain of \mathcal{P} . Suppose that \mathcal{A} is a quantum query algorithm for \mathcal{P} with worst case error probability at most $\varepsilon < 1/2$, and suppose that \mathcal{A}^+ is the query algorithm running \mathcal{A} on a superposition of inputs determined by a unit vector (δ_x) . Without loss of generality, let all unitary transformations V_t used by \mathcal{A} have real entries and let δ_x be real for all x . In particular, given an adversary matrix Γ for \mathcal{P} , we choose (δ_x) to be a principal eigenvector of Γ .

Suppose $x, y \in \mathcal{D}$ are two inputs such that $\mathcal{P}(x) \neq \mathcal{P}(y)$, and recall that $\langle \psi_t(x), \psi_t(y) \rangle = \langle \phi_t(x), \phi_t(y) \rangle$ for all t . At the beginning of the algorithm, since χ is constant for standard query problems, we have $\langle \phi_0(x), \phi_0(y) \rangle = 1$. At the end of the algorithm, however, we need that $\langle \psi_T(x), \psi_T(y) \rangle$ is small as $\psi_T(x)$ must lay mostly within the subspace $\mathcal{X}_{\text{QW}} \otimes \mathcal{P}(\mathbf{x})_{\text{R}}$ and $\psi_T(y)$ within $\mathcal{X}_{\text{QW}} \otimes \mathcal{P}(\mathbf{y})_{\text{R}}$. More precisely, we need that $|\langle \psi_T(x), \psi_T(y) \rangle| \leq 2\sqrt{\varepsilon(1-\varepsilon)}$. The adversary lower bound method is essentially based on the observation that is hard to reduce these inner products simultaneously for all pairs of $x, y \in \mathcal{D}$ satisfying $\mathcal{P}(x) \neq \mathcal{P}(y)$. On the other hand, if

$\mathcal{P}(x) = \mathcal{P}(y)$, we can make no useful claims about $\langle \psi_T(x), \psi_T(y) \rangle$, which is the reason behind the condition (2.23) on the adversary matrix.

Recall from (2.15) and (2.16) the density matrix ρ_t describing the state of the input register \mathcal{X}_I between the queries t and $t + 1$. The need to reduce the inner products $\langle \phi_t(x), \phi_t(y) \rangle$ for all $x, y \in \mathcal{D}$ satisfying $\mathcal{P}(x) \neq \mathcal{P}(y)$ motivates the definition of the *progress function*:

$$W_t := \sum_{\substack{x, y \in \mathcal{D} \\ \mathcal{P}(x) \neq \mathcal{P}(y)}} \Gamma[x, y] \cdot \delta_x \delta_y \langle \phi_t(x), \phi_t(y) \rangle = \text{Tr}(\Gamma \rho_t),$$

where the latter equality is from (2.23). The progress function W_t , in a way, quantifies the entanglement between \mathcal{X}_I and \mathcal{X}_A . Since the unit vector (δ_x) is the principal eigenvector of Γ , we get

$$W_0 = \sum_{x, y \in \mathcal{D}} \Gamma[x, y] \cdot \delta_x \delta_y = \pm \|\Gamma\|.$$

On the other hand, one can show that in order for the algorithm to err with probability at most ε , we must have $|W_T| \leq 2\sqrt{\varepsilon}\|\Gamma\|$ (for the proof, refer to [Bel13]).

Regarding the denominator $\max_i \|\Delta_i \circ \Gamma\|$ in (2.26), reducing $\langle \phi_t(x), \phi_t(y) \rangle$ corresponds to distinguishing between inputs x and y , which can be done only by querying indices $i \in [n]$ such that $x_i \neq y_i$. Thereby one can show that $|W_{t+1} - W_t| \leq 2\|\Delta_i \circ \Gamma\|$ for all i (refer to [HLS07]). This together with $|W_T - W_0| \geq (1 - 2\sqrt{\varepsilon})\|\Gamma\|$ yields the adversary bound.

2.3.2 Simplification tools

When constructing an adversary bound for a problem, it is hard to choose a good adversary matrix Γ and, once the adversary matrix is chosen, it is often hard to estimate the norms $\|\Gamma\|$ and $\|\Delta_i \circ \Gamma\|$. Here we present some tools that simplify both of these tasks.

First of all, let us use the notations Δ_i and $\bar{\Delta}_i$ to denote any matrices whose rows and columns correspond to inputs $x \in \Sigma^n$ and that are defined according to the definition (2.24) of the difference matrices, and we may use the name “difference matrices” for them too. We call $\Delta_i \circ A$ the *action* of Δ_i on a matrix A , and the row and column labeling of A determines which Δ_i in particular we are considering.

In this thesis, we consider only adversary bounds for decision problems. Suppose we are given a decision problem $\mathcal{P}: \mathcal{D} \rightarrow \{0, 1\}$, and, as before, let us decompose its domain $\mathcal{D} \subseteq \Sigma^n$ as $\mathcal{D}_1 \sqcup \mathcal{D}_0$.

Reduction of the adversary bound to a quadrant. Suppose Γ is an adversary matrix for the decision problem \mathcal{P} , and let Γ' be its $|\mathcal{D}_1| \times |\mathcal{D}_0|$ -submatrix corresponding to rows labeled by

yes-inputs \mathcal{D}_1 and columns labeled by no-inputs \mathcal{D}_0 . Due to (2.23), we have

$$\Gamma = \begin{pmatrix} 0 & \Gamma' \\ (\Gamma')^\top & 0 \end{pmatrix} \quad \text{and} \quad \Delta_i \circ \Gamma = \begin{pmatrix} 0 & \Delta_i \circ \Gamma' \\ (\Delta_i \circ \Gamma')^\top & 0 \end{pmatrix},$$

where the first block of rows and columns correspond to \mathcal{D}_1 and the second to \mathcal{D}_0 . “Bipartiteness” of these matrices implies that $\|\Gamma\| = \|\Gamma'\|$ and $\|\Delta_i \circ \Gamma\| = \|\Delta_i \circ \Gamma'\|$. By abuse of terminology, we call Γ' an adversary matrix for \mathcal{P} , and we remove the prime symbol ($'$) from its notation for simplicity, as the value $\text{Adv}(\mathcal{P})$ given by (2.27) stays the same for this submatrix.

Restriction to hard-to-distinguish inputs. Often, when we prove lower bounds for function evaluation, it suffices to consider inputs that have different values, yet are hard to distinguish one from another (i.e., they are equal in most positions). For example, for the OR function, one typically has to compare only the input 0^n and the inputs containing a unique 1.

Remark 2.16. If one restricts a problem to its subdomain, the restricted (promise) problem can only become easier. Therefore, any lower bound for the restricted problem is also a lower bound for the original problem.

For the adversary method, this restriction is manifested by restricting the adversary matrix to rows and columns corresponding to only those inputs that we care about. (This is equivalent to placing only zeros in all the rows and columns corresponding the inputs that we ignore.)

Automorphism principle. Recall that an element $\gamma = (\pi, \tau)$ of the group $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$ acts on an input $x \in \mathcal{D}$ according to (2.8). Suppose $G \leq \mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$ is an automorphism of \mathcal{P} (we do not require G to be an oracle automorphism, namely, we do not require conditions (2.10) and (2.11) to hold). And let us assume that $\mathcal{P}(\gamma(x)) = \mathcal{P}(x)$ for all $x \in \mathcal{D}$ and $\gamma \in G$ (i.e., the group action ω in (2.9) is trivial). The *automorphism principle*, introduced in [HLŠ07], states that without loss of generality, we can assume that the adversary matrix Γ is fixed under G . More precisely, one can restrict the maximization in (2.27) to adversary matrices Γ satisfying

$$\Gamma[x, y] = \Gamma[\gamma(x), \gamma(y)] \quad \text{for all } x \in \mathcal{D}_1, y \in \mathcal{D}_0, \gamma \in G \quad (2.28)$$

without affecting the optimal value of $\text{Adv}(\mathcal{P})$.³

We also view Γ as a linear map in $\mathbb{L}(\mathbb{C}^{\mathcal{D}_0}, \mathbb{C}^{\mathcal{D}_1})$. Let $U_{0,\gamma} \in \mathbb{U}(\mathbb{C}^{\mathcal{D}_0})$ and $U_{1,\gamma} \in \mathbb{U}(\mathbb{C}^{\mathcal{D}_1})$ be the permutation representations of G corresponding to its group action (2.8) on \mathcal{D}_0 and \mathcal{D}_1 , respectively. Then (2.28) is equivalent to

$$U_{1,\gamma}\Gamma = \Gamma U_{0,\gamma} \quad (2.29)$$

³Since we are proving lower bounds, the automorphism principle is not formally needed, as we can assume the symmetry (2.28) without it.

for all $\gamma \in G$. Schur’s lemma (Lemma 1.2) and (2.29) imply that we can write Γ as a linear combination of transporters from irreps of G in $\mathbb{C}^{\mathcal{D}_0}$ to isomorphic irreps in $\mathbb{C}^{\mathcal{D}_1}$.

One says that G is \mathcal{P} -transitive if, for every x, y such that $\mathcal{P}(x) = \mathcal{P}(y)$, there is $\gamma \in G$ such that $\gamma(x) = y$. The automorphism principle also states that, if G is \mathcal{P} -transitive, without loss of generality, one can restrict the maximization in (2.27) to adversary matrices whose both right and left principal singular vectors are the all-ones vectors (see [HLŠ07]). These singular vectors correspond to the unique trivial representations in $U_{0,\gamma}$ and $U_{1,\gamma}$, respectively, and the transporter between them is the $|\mathcal{D}_1| \times |\mathcal{D}_0|$ -matrix Ξ_{id} of all entries equal to $1/\sqrt{|\mathcal{D}_0| \cdot |\mathcal{D}_1|}$.

Automorphisms G that we will consider (explicitly or implicitly) for the ELEMENT DISTINCTNESS, COLLISION, and SET EQUALITY problems will all be transitive, and the adversary matrices that we will construct for these problems will all satisfy (2.28),(2.29). While we will not show that the all-ones vectors correspond to the highest singular values of these adversary matrices, when constructing these matrices as a linear combination of transporters, the coefficient we assign to Ξ_{id} will be the target value of our desired adversary bound.

Adversary matrices with reoccurring rows labels. The adversary bound still holds if we allow multiple rows of the adversary matrix Γ to correspond to the same yes-input $x \in \mathcal{D}_1$ [BŠ13]. In this case, we label each row of the matrices Γ and Δ_i by a pair (x, a) , where $x \in \mathcal{D}_1$ and a serves to distinguish pairs with the same first element. Now, technically, $\Delta_i[(x, a), y] = 1$ if and only if $x_i \neq y_i$. (We can simultaneously do the same for columns and no-inputs, but that will not be necessary in our applications.)

When considering symmetries of this generalized adversary matrix using the automorphism principle, we now also have to consider how the group $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$ acts on labels a . This will be done in a natural way as every a will essentially correspond to a subset of the power set $2^{[n]}$.

Embedded adversary matrices. It is sometimes helpful to embed the adversary matrix Γ into a larger matrix $\tilde{\Gamma}$, as it becomes easier to argue about properties of Γ and $\Delta_i \circ \Gamma$ using $\tilde{\Gamma}$ and $\Delta_i \circ \tilde{\Gamma}$ instead. This approach was first introduced in [Bel12b] for the ELEMENT DISTINCTNESS problem, and later used for k -SUM [BŠ13] and other problems [Špa13]. We use the same approach for constructing and analyzing adversary matrices for the COLLISION, SET EQUALITY, and ORTHOGONAL ARRAY problems.

We label the columns of $\tilde{\Gamma}$ by all inputs in Σ^n —yes-inputs, no-inputs, and the inputs outside the domain \mathcal{D} —and rows by yes-inputs and certain inputs outside the domain. As discussed above, we allow multiple rows to correspond to the same input. We call rows and columns of $\tilde{\Gamma}$ that correspond to input outside \mathcal{D}_1 and \mathcal{D}_0 , respectively, *illegal*. We extract Γ from $\tilde{\Gamma}$ by deleting all the illegal rows and columns. Figure 2.6 illustrates these concepts.

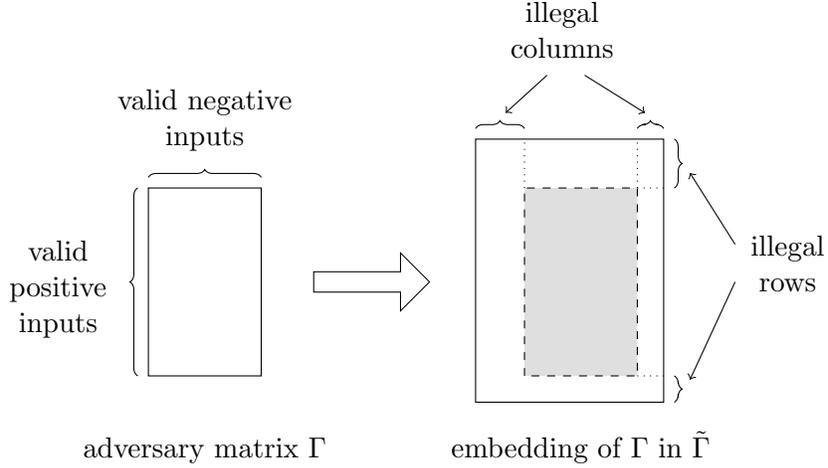


Figure 2.6: Embedding the adversary matrix Γ into a larger matrix $\tilde{\Gamma}$.

Because $\Delta_i \circ \Gamma$ is a submatrix of $\Delta_i \circ \tilde{\Gamma}$, we clearly have

$$\|\Delta_i \circ \Gamma\| \leq \|\Delta_i \circ \tilde{\Gamma}\| \quad (2.30)$$

If we could show that $\|\Gamma\|$ is not much smaller than $\|\tilde{\Gamma}\|$, that would allow us to use $\tilde{\Gamma}$ instead of Γ in the adversary bound, Theorem 2.15. (Note: we are interested only in the asymptotic value of $\text{Adv}(\mathcal{P})$.) That is not true for every choice of $\tilde{\Gamma}$, however. For instance, if $\tilde{\Gamma}$ contains non-zero entries only in illegal rows or columns, then $\Gamma = 0$. Nevertheless, in our applications, we will ensure that $\|\Gamma\| \approx \|\tilde{\Gamma}\|$, and, for that, the condition that \mathcal{D}_0 and Σ^n have approximately the same size will be essential.

We will construct “adversary matrices” $\tilde{\Gamma}$ using the projectors on the eigenspaces of the Hamming scheme (see (1.16) in Section 1.5.1) or similar projectors (namely, (4.4)). The orthogonality of these projectors is the main reason why it is much easier to evaluate the norms of $\tilde{\Gamma}$ and $\Delta_i \circ \tilde{\Gamma}$ rather than those of Γ and $\Delta_i \circ \Gamma$.

Approximation of the Δ -action. Precise calculation of $\|\Delta_i \circ \Gamma\|$ may be tedious, but we can upper bound $\|\Delta_i \circ \Gamma\|$ using the following trick first introduced in [Bel12b] and later used in [BŠ13, Špa13]. By considering a matrix norm called the γ_2 norm, [LMR⁺11] shows that

Lemma 2.17. *For any matrix A whose rows and columns correspond to inputs in Σ^n ,*

$$\|\Delta_j \circ A\| \leq 2 \|A\|.$$

For any matrix A , we call a matrix B satisfying

$$\Delta_i \circ B = \Delta_i \circ A \quad (= \Delta_i \circ \Delta_i \circ A)$$

an *approximation* of $\Delta_i \circ A$ and denote it $\Delta_i \diamond A$. Or, we write $A \overset{\Delta_i}{\rightsquigarrow} B$. From Lemma 2.17, 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 any $\Delta_i \diamond \Gamma$. (Note: the approximations that we will consider will depend on i .) That is, it suffices to show that we can change entries of Γ with $x_i = y_i$ in a way that the spectral norm of the resulting matrix is constantly bounded.

2.3.3 Structure of adversary constructions

In this section, we present the high level ideas behind adversary bounds in Chapters 4 and 5, which complement matrix embedding, approximate Δ -action, and other simplification tools from the previous section. For some adversary constructions, these ideas are followed explicitly, for some, only implicitly. Most arguments in this section are informal and only serve to establish intuition behind the constructions.

As mentioned in the introduction, negative weights allow to construct the adversary matrix as any linear combination of matrices of the same dimensions. We construct the adversary matrix as a linear combination $\Gamma := \sum_{\ell} \alpha_{\ell} W_{\ell}$, where W_{ℓ} are mutually orthogonal matrices of (spectral) norm approximately 1. For problems with many symmetries (ELEMENT DISTINCTNESS, COLLISION, SET EQUALITY), indices ℓ run over integers in $[n]$, for other problems (CERTIFICATE-SUM, ORTHOGONAL ARRAY), over subsets of $[n]$. Due to the orthogonality of the W_{ℓ} matrices, the norm of Γ is approximately $\max_{\ell} |\alpha_{\ell}|$.

The adversary bound requires that the norm of $\Delta_i \circ \Gamma$ is $O(1)$ for all $i \in [n]$. After the application of Δ_i , the matrices $\Delta_i \circ W_{\ell}$ are not orthogonal any more. For example, a “part” of $\Delta_i \circ W_{\ell}$ may also appear in $\Delta_i \circ W_{\ell'}$ for some $\ell' \neq \ell$, but with the opposite sign. Thus, if the coefficients α_{ℓ} and $\alpha_{\ell'}$ are close, this part *cancel out* in $\Delta_i \circ \Gamma = \sum_{\ell} \alpha_{\ell} (\Delta_i \circ W_{\ell})$. (Sometimes more than two W_{ℓ} ’s may be required to cancel out a certain part, as this part may appear with different “weights” in different $(\Delta_i \circ W_{\ell})$ ’s.) Thus, even though each $\Delta_i \circ W_{\ell}$ may have norm close to 1 and $\alpha_{\ell} = \omega(1)$, we can still have $\|\Delta_i \circ \Gamma\| = O(1)$ because of the cancellation.

Problems with many symmetries. We construct adversary matrices for ELEMENT DISTINCTNESS, COLLISION, and SET EQUALITY (implicitly or explicitly) as a linear combination

$$\Gamma := \sum_{k=1}^{\mathcal{T}} \alpha_k W_k, \quad (2.31)$$

where \mathcal{T} is the target bound we aim to prove, W_0 is a matrix with all entries equal, and, as before, W_k 's are mutually orthogonal and $\|W_k\| \approx 1$.

When analyzing $\Delta_i \circ \Gamma$, we decompose W_k into three parts,

$$W_k = X_k + Y_k + Z_k,$$

with the exception of $W_0 = X_0$. (This decomposition of W_k depends on i .) We choose these matrices so that, informally,

1. $\Delta_i \circ X_k \approx X_k$, $\Delta_i \circ Y_k \approx Y_k$, and $\Delta_i \circ Z_k \approx -X_{k-1}$;
2. X_k 's are mutually orthogonal and of norm approximately 1;
3. Y_k 's are mutually orthogonal, their norm is $o(1)$, but it increases as k increases.

Due to Points 1 and 2, the “part” X_k almost completely cancels out in $\Delta_i \circ \Gamma$ if W_k and W_{k+1} appear in (2.31) with similar coefficients, namely, $|\alpha_k - \alpha_{k+1}| = O(1)$ (see Figure 2.7). For this reason, in (2.31), we choose

$$\alpha_k := \mathcal{T} - k. \quad (2.32)$$

Now consider the “part” Y_k of $\Delta_i \circ W_k$ that does not cancel out with anything. These parts are the main reason why we cannot choose an arbitrarily large \mathcal{T} . Namely, since $\|Y_k\|$ grows with k and we need $\|\alpha_k Y_k\| \leq 1$, the coefficients α_k have to be small for large k . This holds only if \mathcal{T} is not too large in (2.32).

Cancellation among irreps. For COLLISION (Section 4.2) and ELEMENT DISTINCTNESS with small range (Chapter 5), one can decompose each W_k as a sum

$$W_k = \sum_{\lambda \vdash n: \lambda_1 = n-k} W_\lambda,$$

where W_λ corresponds to the irrep \mathcal{S}^λ of $\mathbb{S}_{[n]}$.⁴ (There is also a similar decomposition for SET EQUALITY.) More formally, we have $U_{1,\pi} W_\lambda = W_\lambda U_{0,\pi}$ for all $\pi \in \mathbb{S}_{[n]}$, where, for $b \in \{0, 1\}$, $U_{b,\pi} \in \mathbf{U}(\mathbb{C}^{\mathcal{D}^b})$ is the permutation representation of $\mathbb{S}_{[n]}$ as in (2.29) of the automorphism principle

⁴Note: in Section 4.2, where one considers COLLISION and SET EQUALITY, W_k in (2.31) is denoted \bar{W}_k , as there, the notation W_k corresponds to a matrix related to \bar{W}_k . And, instead of n , one uses $2n$.

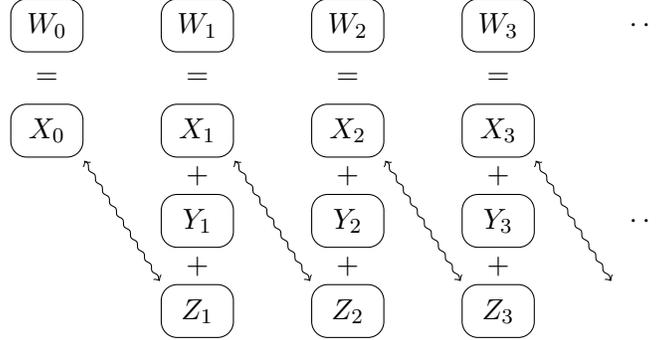


Figure 2.7: Decomposition of W_k as the sum of three matrices. Wiggly arrows connect matrices that cancel out after the application of Δ_i .

(here, by abuse of notation, we write $U_{b,\pi}$ instead of $U_{b,(\pi,\varepsilon)}$, where ε is the identity permutation of \mathbb{S}_Σ). W_λ is such that its column and row spaces are contained in the \mathcal{S}^λ -isotypical subspaces of $\mathbb{C}^{\mathcal{D}^1}$ and $\mathbb{C}^{\mathcal{D}^0}$, respectively. Thus, Schur's lemma implies that the matrices W_λ are orthogonal for different λ .

Regarding the application of Δ_i , we have

$$\Delta_i \circ W_\lambda = W_\lambda - \bar{\Delta}_i \circ W_\lambda = W_\lambda - \sum_{a \in \Sigma} \hat{\Pi}_1^{(i,a)} W_\lambda \hat{\Pi}_0^{(i,a)},$$

where, for $b \in \{0, 1\}$,

$$\hat{\Pi}_b^{(i,a)} := \sum_{z \in \mathcal{D}_b: z_i=a} z z^*.$$

Note that

$$U_{b,\pi} \hat{\Pi}_b^{(i,a)} U_{b,\pi}^{-1} = \hat{\Pi}_b^{(i,a)}$$

for all $\pi \in \mathbb{S}_{[n] \setminus \{i\}}$, which means that the space corresponding to $\hat{\Pi}_b^{(i,a)}$ can be decomposed into irreps of $\mathbb{S}_{[n] \setminus \{i\}}$. Because of this, the column and row spaces of $\Delta_i \circ W_\lambda$ are respectively contained in the direct sum of the isotypical subspaces of $\mathbb{C}^{\mathcal{D}^1}$ and $\mathbb{C}^{\mathcal{D}^0}$ corresponding to irreps \mathcal{S}^μ of $\mathbb{S}_{[n] \setminus \{i\}}$ such that $\mu \subset \lambda$ (i.e., the Young diagram corresponding to μ can be obtained from the one corresponding λ by removing a box). Hence, for $\lambda, \lambda' \vdash n$, the matrices $\Delta_i \circ W_\lambda$ and $\Delta_i \circ W_{\lambda'}$ are orthogonal and no cancellation can occur between them unless there is (a unique) $\mu \vdash n-1$ such that both $\mu \subset \lambda$ and $\mu \subset \lambda'$ (see Figure 2.8). And that can only happen when λ_1 and λ'_1 differ by at most 1.

The following claim will ensure that all entries of all matrices W_λ that we consider are real.

Claim 2.18. *Suppose \mathbb{S}_n acts on a finite set A , so that \mathbb{C}^A is a permutation representation of \mathbb{S}_n . In the standard basis of \mathbb{C}^A , for all $\lambda \vdash n$, all entries of the projector on the \mathcal{S}^λ -isotypical subspace of \mathbb{C}^A are real.*

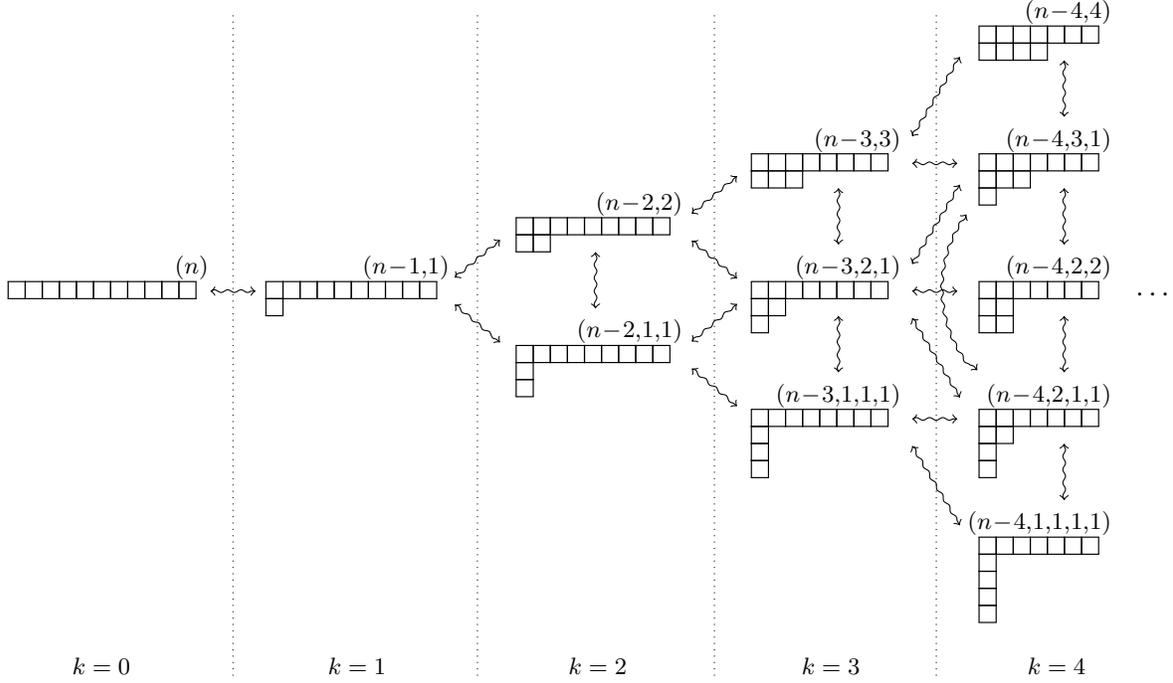


Figure 2.8: The matrix W_k corresponds to Young diagrams λ with k boxes below the first row. Wiggly arrows connect Young diagrams λ whose corresponding matrices $\Delta_i \circ W_\lambda$ have a potential for cancellation.

Proof. Let $U_\pi \in U(\mathbb{C}^A)$ be the corresponding representation operator, whose entries in the standard basis are either 0 or 1. Recall from Section 1.3.5 that all characters of the symmetric group are real. Thus, the claim holds because of (1.5), which is the projector under consideration. \square

2.3.4 Limitations of positive-weights adversary bound

Until recently, the vast majority of adversary lower bounds were obtained using the positive-weights version of the bound. However, the positive-weights adversary bound is subject to some severe constraints like the property testing barrier [HLS07] and the certificate complexity barrier [ŠS06, Zha05]:

- The property testing barrier states that, if every yes-input differs from every no-input in at least an α fraction of the input variables, no positive-weights adversary can prove a lower bound better than $\Omega(1/\alpha)$.
- The certificate complexity barrier states that no positive-weights adversary can prove a lower bound better than $\Omega(\sqrt{n \cdot \min\{C_0, C_1\}})$ and, if \mathcal{P} is a total function, $\Omega(\sqrt{C_0 \cdot C_1})$.

We note that, by the property testing barrier, no positive-weights adversary can give a non-trivial (i.e., better than $\Omega(1)$) lower bound for COLLISION or SET EQUALITY. And, by the certificate complexity barrier, no positive-weights adversary can give a lower bound for ELEMENT DISTINCTNESS, k -DISTINCTNESS, and k -SUM better than $\Omega(\sqrt{n})$.

In Ref. [HLŠ07] that originally introduced the general adversary bound, Høyer, Lee, and Špalek provide an intuition why negative weights in the adversary matrix may improve the bound:

“While it is clear that ADV^\pm is always least as large as ADV , it might at first seem surprising that ADV^\pm can achieve bounds super-linear in ADV . An intuition for why negative weights help is that it is good to give negative weight to entries with large Hamming distance, entries which are easier to distinguish by queries. Consider an entry (x, y) where x and y have large Hamming distance. This entry appears in several $\Gamma \circ D_i$ matrices but only appears in the Γ matrix once. Thus by giving this entry negative weight we can simultaneously decrease $\|\Gamma \circ D_i\|$ for several i 's, while doing relatively little damage to the large Γ matrix.”

Here they use ADV^\pm to denote the general adversary, ADV to denote the positive-weights adversary, and D_i instead of Δ_i .

2.4 Span programs and learning graphs

As we mentioned before, the maximization problem (2.27) yielding the adversary bound $\text{Adv}(\mathcal{P})$ can be expressed as an SDP, and it is called the *adversary SDP* (for the function \mathcal{P}). Recall that every feasible solution to the adversary SDP—an adversary matrix Γ —gives a lower bound on the quantum query complexity of \mathcal{P} . The dual program of the adversary SDP is simply called the *dual-adversary SDP*. Every feasible solution to the dual-adversary SDP yields a quantum algorithm for \mathcal{P} having the quantum query complexity of the objective value of the SDP [LMR⁺11]. (The feasible solutions of the dual-adversary SDP are commonly called *span programs*.) The strong duality of these SDPs implies the optimality of the adversary bound.

For illustrative purposes, let us present the dual-adversary SDP for decision problems. For a decision problem \mathcal{P} with domain $\mathcal{D} = \mathcal{D}_0 \sqcup \mathcal{D}_1$, the dual-adversary SDP is

$$\text{minimize} \quad \max_{x \in \mathcal{D}} \sum_{i \in [n]} X_i[x, x] \tag{2.33a}$$

$$\text{subject to} \quad \sum_{i: x_i \neq y_i} X_i[x, y] = 1 \quad \text{for all } x \in \mathcal{D}_1 \text{ and } y \in \mathcal{D}_0; \tag{2.33b}$$

$$X_i \succeq 0 \quad \text{for all } i \in [n], \tag{2.33c}$$

where, just like for the adversary matrix Γ (according to its original definition, Definition 2.14), rows and columns of positive semidefinite matrices X_i are labeled by inputs $x \in \mathcal{D}$.

Constructing feasible solutions to the dual-adversary SDP is already a hard task, let alone trying to minimize the objective value. For Boolean-valued functions, Belovs developed the computational model of *learning graph* [Bel12d] that, by its design, can be translated into a feasible solution of the dual-adversary SDP (i.e., a span program) [Bel12d, BL11], so one can focus on minimizing the objective value.

Informal description. Informally speaking, a learning graph for a function $\mathcal{P}: \mathcal{D} \rightarrow \{0, 1\}$ with $\mathcal{D} \subseteq \Sigma^n$ is a collection of directed probabilistic walks on an n -dimensional weighted hypercube whose vertices (or *nodes*) are subsets S of input indices (i.e., $S \subseteq [n]$) and whose edges (or *arcs*) are of the form $(S, S \cup \{j\})$. For every yes-input $x \in \mathcal{D}_1$, the walk originates from the empty set \emptyset , it learns the value of x_j when it follows an edge $(S, S \cup \{j\})$, and it culminates in a vertex S' such that $x_{S'}$ is a 1-certificate (i.e., the walk has learned that the value of the function is 1). When the walk follows $(S, S \cup \{j\})$, we say that the learning graph *loads* j . There is a complexity associated with each learning graph.

Formal definition. For $S \subset [n]$, let $j \notin S$ be short for $j \in [n] \setminus S$ and (S, j) be short for $(S, S \cup \{j\})$. Let $\mathcal{E} := \{(S, j): S \subset [n], j \notin S\}$ be the set of all edges of the hypercube described above. Formally, the learning graph for a Boolean-valued function $\mathcal{P}: \mathcal{D} \rightarrow \{0, 1\}$ is a pair of two functions—the *weight function* w and the *flow function* p —that are defined as follows.

The weight function w maps every edge $(S, j) \in \mathcal{E}$ and every assignment $a \in \Sigma^S$ of S to a non-negative *weight* $w_{S,j}(a)$. Essentially, the weight of an edge “originating” from a vertex S can depend on $x_S \in \Sigma^S$, that is, the symbols of $x \in \mathcal{D}$ that one has “learned” so far. We consider that the edges of weight 0 are *not present* in the graph.

The flow function p maps every yes-input $x \in \mathcal{D}_1$ to a unit flow $p(x)$ on the hypercube such that $S = \emptyset$ is the only source of the flow and only nodes S such that x_S is a 1-certificate can be sinks. More precisely, p maps every edge $(S, j) \in \mathcal{E}$ and every yes-input $x \in \mathcal{D}_1$ to a flow $p_{S,j}(x) \in \mathbb{R}$ that satisfies

- $\sum_{j \in [n]} p_{\emptyset,j}(x) = 1$ for all $x \in \mathcal{D}_1$;
- $\sum_{j \in S} p_{S \setminus \{j\},j}(x) = \sum_{j \notin S} p_{S,j}(x)$ for all $x \in \mathcal{D}_1$ and $S \neq \emptyset$ such that x_S is not a 1-certificate;
- $p_{S,j}(x) = 0$ whenever $w_{S,j}(x_S) = 0$ (that is, we do not allow any flow on absent edges).

Learning graph complexity. Suppose we are given a learning graph \mathcal{G} of a decision problem \mathcal{P} . For every yes-input $x \in \mathcal{D}_1$ and every no-input $y \in \mathcal{D}_0$, their respective *complexities* in the learning graph are

$$LG_1(\mathcal{G}, x) := \sum_{(S,j) \in \mathcal{E}} \frac{p_{S,j}(x)^2}{w_{S,j}(x_S)} \quad \text{and} \quad LG_0(\mathcal{G}, y) := \sum_{(S,j) \in \mathcal{E}} w_{S,j}(y_S), \quad (2.34)$$

where we assume $0/0 = 0$. The 1-*complexity* and the 0-*complexity* of the learning graph are

$$LG_1(\mathcal{G}) := \max_{x \in \mathcal{D}_1} LG_1(\mathcal{G}, x) \quad \text{and} \quad LG_0(\mathcal{G}) := \max_{y \in \mathcal{D}_0} LG_0(\mathcal{G}, y), \quad (2.35)$$

respectively. The complexity of the learning graph is their geometric mean:

$$LG(\mathcal{G}) := \sqrt{LG_1(\mathcal{G}) \cdot LG_0(\mathcal{G})}. \quad (2.36)$$

Note that, if we multiply all weights by the same scalar $c > 1$, $LG_1(\mathcal{G})$ decreases c times, $LG_0(\mathcal{G})$ increases c times, and $LG(\mathcal{G})$ remains unchanged.

A learning graph is said to be *non-adaptive* if its the weight function $w_{S,j}(a)$ is independent from a . If this restriction is not imposed, the learning graph is said to be *adaptive*. Adaptive learning graphs can be more powerful, meaning that their complexity can be smaller. The *non-adaptive (adaptive) learning graph complexity* of the function \mathcal{P} is the minimum complexity among all non-adaptive (adaptive) learning graphs for \mathcal{P} . The translation from learning graphs to span programs implies

Theorem 2.19 ([Bel12d, BL11]). *The quantum query complexity of a Boolean-valued function \mathcal{P} is at most the adaptive learning graph complexity of \mathcal{P} . (Note: the non-adaptive learning graph complexity is least as big as the adaptive learning graph complexity.)*

Remark 2.20. For functions with binary input alphabet, adaptive learning graphs can be generalized so that the weight of an edge $(S, j) \in \mathcal{E}$, in addition to x_S , can also depend on $x_j \oplus \mathcal{P}(x)$, where \oplus stands for the exclusive or (see [Bel12c]). However, it is not clear how to extend this generalization to functions with non-binary input alphabets.

Examples of non-adaptive and adaptive learning graphs. Let us present a non-adaptive learning graph for the OR function and an adaptive learning graph for the THRESHOLD-2 function, therefore proving the following two upper bounds.

Proposition 2.21. *The non-adaptive learning graph complexity of OR is $O(\sqrt{n})$.*

Proof. Let us construct a non-adaptive learning graph \mathcal{G} for OR by specifying the weight and the flow functions. We set $w_{\emptyset,j} := 1$ for all $j \in [n]$ and $w_{S,j} := 0$ whenever $S \neq \emptyset$ (here we omit

$a \in \Sigma^S$ from the notation $w_{S,j}(a)$ as the weights do not depend on a). And, for all $x \in \mathcal{D}_1$, we set

$$p_{\emptyset,j}(x) := \begin{cases} \frac{1}{|x|}, & \text{if } x_j = 1 \\ 0, & \text{if } x_j = 0 \end{cases} \quad \text{for all } j \in [n]$$

and $p_{S,j}(x) := 0$ whenever $S \neq \emptyset$, where $|x|$ is the Hamming weight of x . One can see that $p_{S,j}(x)$ satisfies all the necessary conditions of a flow. According to (2.34), we have

$$LG_1(\mathcal{G}, x) = \sum_{j \in [n]} \frac{p_{\emptyset,j}(x)^2}{w_{\emptyset,j}} = \sum_{j: x_j=1} \frac{1}{|x|^2} = \frac{1}{|x|} \leq 1$$

for all $x \in \mathcal{D}_1$ and

$$LG_0(\mathcal{G}, 0^n) = \sum_{j \in [n]} w_{\emptyset,j} = n,$$

where the all-zeros input 0^n is the unique negative input of OR. Thus from (2.35) and (2.36) one can see that the complexity of this learning graph is at most \sqrt{n} . \square

Proposition 2.22. *The adaptive learning graph complexity of THRESHOLD-2 is $O(\sqrt{n})$.*

Proof. Let us construct an adaptive learning graph \mathcal{G} for THRESHOLD-2. We set $w_{\emptyset,j}(\varepsilon) := 1$ for all $j \in [n]$, where ε denotes the empty string, $w_{\{i\},j}(1) := 1$ and $w_{\{i\},j}(0) := 0$ for all $i, j \in [n]$ (with $i \neq j$), and $w_{S,j}(a) := 0$ whenever $|S| \geq 2$. And, for all positive inputs $x \in \mathcal{D}_1$ (i.e., inputs with the Hamming weight $|x| \geq 2$), we set

$$p_{\emptyset,j}(x) := \begin{cases} \frac{1}{|x|}, & \text{if } x_j = 1 \\ 0, & \text{if } x_j = 0 \end{cases} \quad \text{and} \quad p_{\{i\},j}(x) := \begin{cases} \frac{1}{|x|^2 - |x|}, & \text{if } x_i = x_j = 1 \\ 0, & \text{otherwise} \end{cases} \quad \text{for all } i, j \in [n],$$

and $p_{S,j}(x) := 0$ whenever $|S| \geq 2$. According to (2.34), for $x \in \mathcal{D}_1$, we have

$$\begin{aligned} LG_1(\mathcal{G}, x) &= \sum_{j \in [n]} \left(\frac{p_{\emptyset,j}(x)^2}{w_{\emptyset,j}(\varepsilon)} + \sum_{i: i \neq j} \frac{p_{\{j\},i}(x)^2}{w_{\{j\},i}(x_j)} \right) \\ &= \sum_{j: x_j=1} \frac{1}{|x|^2} + \sum_{\substack{i,j: i \neq j, \\ x_i=x_j=1}} \frac{1}{(|x|^2 - |x|)^2} = \frac{1}{|x|} + \frac{1}{|x|^2 - |x|} \leq 1. \end{aligned}$$

On the other hand, for the all-zeros input $0^n \in \mathcal{D}_0$, we have

$$LG_0(\mathcal{G}, 0^n) = \sum_{j \in [n]} w_{\emptyset,j} = n,$$

and, for $y^{(i)} \in \mathcal{D}_0$ being the negative input having the unique one in the position i , we have

$$LG_0(\mathcal{G}, y^{(i)}) = \sum_{j \in [n]} w_{\emptyset, j}(\varepsilon) + \sum_{j: j \neq i} w_{\{i\}, j}(1) = 2n - 1.$$

Hence, the complexity of this learning graph \mathcal{G} is at most $\sqrt{2n}$. □

The minimal certificate structure of the THRESHOLD-2 function is the 2-subset certificate structure (see Definition 2.5). Thus, Proposition 3.4 shows that the non-adaptive learning graph complexity of THRESHOLD-2 is $\Omega(n^{2/3})$, in particular, it has the same non-adaptive learning graph complexity as the ELEMENT DISTINCTNESS problem. On the other hand, Proposition 2.22 shows that the adaptive learning graph complexity of THRESHOLD-2 is $O(\sqrt{n})$. This shows that adaptive learning graphs are strictly stronger than non-adaptive learning graphs.

In turn, quantum query algorithms are strictly stronger than adaptive learning graphs. To see that, consider the AND function. In Section 3.3.1 we show that the adaptive learning graph complexity of AND is $\Omega(n)$, yet Grover’s search algorithm (see [Gro96, BBHT98]) can solve this function using only $O(\sqrt{n})$ queries.

Notice that the learning graph complexities (both, adaptive and non-adaptive) of OR and AND differ. That is a consequence of the fact that the definition of the learning graph do not “treat” positive inputs and negative inputs the same way.

Part II

Results

Chapter 3

Lower bounds on learning graph complexity

In this chapter we obtain general results on both non-adaptive and adaptive learning graph complexity, as well as problem-specific results. We start by expressing the learning graph complexity as a semidefinite program, which allows us to obtain its dual program. *Dual learning graphs* are lower bounds on the learning graph complexity, and in Chapter 4 they will help us to obtain general results on the quantum query complexity of certificate structures.

3.1 Learning graph complexity as a semidefinite program

Let us first consider the general case, adaptive learning graphs. The semidefinite program for non-adaptive learning graphs will be very similar.

3.1.1 SDPs for adaptive learning graph complexity

Recall the definitions of the learning graph and the learning graph complexity from Section 2.4. Consider a decision problem $\mathcal{P}: \Sigma^n \rightarrow \{0, 1\}$ with a domain $\mathcal{D} = \mathcal{D}_0 \sqcup \mathcal{D}_1$. Suppose we are given a learning graph \mathcal{G} for \mathcal{P} , and its complexity is $\sqrt{K} := \sqrt{LG_1(\mathcal{G}) \cdot LG_0(\mathcal{G})}$, where $LG_1(\mathcal{G})$ and $LG_0(\mathcal{G})$ are defined via (2.34) and (2.35). Let \mathcal{G}' be the learning graph of the same complexity as \mathcal{G} that is obtained from \mathcal{G} by multiplying all its weights by $LG_1(\mathcal{G})$. Therefore $LG_1(\mathcal{G}', x) \leq 1$ and $LG_0(\mathcal{G}', y) \leq K$ for all $x \in \mathcal{D}_1$ and $y \in \mathcal{D}_0$. Since we can do this for any learning graph \mathcal{G} ,

the adaptive learning graph complexity of \mathcal{P} is the optimal value of the optimization problem

$$\text{minimize } \sqrt{K} \tag{3.1a}$$

$$\text{subject to } \sum_{(S,j) \in \mathcal{E}} \frac{p_{S,j}(x)^2}{w_{S,j}(x_S)} \leq 1 \quad \text{for all } x \in \mathcal{D}_1; \tag{3.1b}$$

$$\sum_{(S,j) \in \mathcal{E}} w_{S,j}(y_S) \leq K \quad \text{for all } y \in \mathcal{D}_0; \tag{3.1c}$$

$$\sum_{j \in S} p_{S \setminus \{j\},j}(x) = \sum_{j \notin S} p_{S,j}(x) \quad \text{for all } x \in \mathcal{D}_1 \text{ and all } S \neq \emptyset \text{ such that } x_S \text{ is not a 1-certificate}; \tag{3.1d}$$

$$\sum_{j \in [n]} p_{\emptyset,j}(x) = 1 \quad \text{for all } x \in \mathcal{D}_1; \tag{3.1e}$$

$$p_{S,j}(x) \in \mathbb{R}, w_{S,j}(a) \geq 0 \quad \text{for all } (S,j) \in \mathcal{E}, x \in \mathcal{D}_1, \text{ and } a \in \Sigma^S, \tag{3.1f}$$

where $0/0$ in (3.1b) is defined to be 0 and, as before, $(S,j) \in \mathcal{E}$ is short for $S \subset [n]$ and $j \in [n] \setminus S$.

This optimization problem can be expressed as an SDP. To express it in the form (1.1), let us consider this optimization problem with its objective value squared, namely, K instead of \sqrt{K} . The form (1.1) will allow us to obtain the dual of (3.1).

Note that we can rewrite (3.1b) as two conditions

$$\begin{aligned} \sum_{(S,j) \in \mathcal{E}} r_{S,j}(x) &\leq 1 && \text{for all } x \in \mathcal{D}_1; \\ \begin{pmatrix} r_{S,j}(x) & p_{S,j}(x) \\ p_{S,j}(x) & w_{S,j}(x_S) \end{pmatrix} &\succeq 0 && \text{for all } (S,j) \in \mathcal{E} \text{ and } x \in \mathcal{D}_1. \end{aligned}$$

Note that the latter condition already implies that $r_{S,j}(x) \geq 0$, thus it suffices to require that the variable $r_{S,j}(x)$ is real¹. Also note that this respects the $0/0 = 0$ condition. That is, if $w_{S,j}(x_S) = 0$, we have to choose $p_{S,j}(x) = 0$, and then there is no benefit in setting $r_{S,j}(x)$ to anything but 0. Technically, to comply with the form (1.1), one would have to use

$$\begin{pmatrix} r_{S,j}(x) & p_{S,j}(x) \\ p_{S,j}(x) & w_{S,j}(x_S) \end{pmatrix} = r_{S,j}(x) \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} + w_{S,j}(x_S) \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} + p_{S,j}(x) \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

since variables in (1.1) are numbers, not matrices. Noting this equality, the SDP (3.1) with its

¹Similarly we could remove the requirement $w_{S,j}(x_S) \geq 0$, but we do not because it would place stricter conditions on the dual. In particular, it would result in having equality in (3.7c)

objective value squared in the form (1.1) is

$$\text{minimize } K \tag{3.3a}$$

$$\text{subject to } \begin{pmatrix} r_{S,j}(x) & p_{S,j}(x) \\ p_{S,j}(x) & w_{S,j}(x_S) \end{pmatrix} \succeq 0 \quad \text{for all } (S,j) \in \mathcal{E} \text{ and } x \in \mathcal{D}_1; \tag{3.3b}$$

$$- \sum_{(S,j) \in \mathcal{E}} r_{S,j}(x) \geq -1 \quad \text{for all } x \in \mathcal{D}_1; \tag{3.3c}$$

$$K - \sum_{(S,j) \in \mathcal{E}} w_{S,j}(y_S) \geq 0 \quad \text{for all } y \in \mathcal{D}_0; \tag{3.3d}$$

$$\sum_{j \in S} p_{S \setminus \{j\},j}(x) - \sum_{j \notin S} p_{S,j}(x) = 0 \quad \text{for all } x \in \mathcal{D}_1 \text{ and all } S \neq \emptyset \text{ such that } x_S \text{ is not a 1-certificate}; \tag{3.3e}$$

$$\sum_{j \in [n]} p_{\emptyset,j}(x) = 1 \quad \text{for all } x \in \mathcal{D}_1; \tag{3.3f}$$

$$p_{S,j}(x) \in \mathbb{R}, r_{S,j}(x) \in \mathbb{R} \quad \text{for all } (S,j) \in \mathcal{E} \text{ and } x \in \mathcal{D}_1; \tag{3.3g}$$

$$K \geq 0, w_{S,j}(a) \geq 0 \quad \text{for all } (S,j) \in \mathcal{E} \text{ and } a \in \Sigma^S. \tag{3.3h}$$

Here (and in its dual (3.4) below) we do not distinguish scalars from 1×1 matrices, and, for them, we can use ‘ \geq ’ and ‘ \succeq ’ interchangeably.

To construct a strictly feasible solution of (3.3), first choose any $p_{S,j}(x) \in \mathbb{R}$ satisfying (3.3e) and (3.3f) and any $r_{S,j} > 0$ strictly satisfying (3.3c). Then choose $w_{S,j}(a) > 0$ large enough to strictly satisfy (3.3b). And, finally, choose K large enough to strictly satisfy (3.3d).

Using the duality of (1.1) and (1.2), we obtain the dual of (3.3) as follows. We introduce the following variables for the dual: let variables

$$\begin{pmatrix} \phi_{S,j}(x) & \rho_{S,j}(x) \\ \rho_{S,j}(x) & \omega_{S,j}(x) \end{pmatrix}, \quad \mu(x), \quad \kappa_y, \quad \nu_S(x), \quad \text{and} \quad \nu_\emptyset(x)$$

correspond to the conditions (3.3b), (3.3c), (3.3d), (3.3e), and (3.3f) of the primal, respectively. On the other hand, let conditions (3.4b), (3.4c), (3.4d), (3.4e) of the dual correspond to the variables

$$K, \quad r_{S,j}(x), \quad w_{S,j}(a), \quad p_{S,j}(x)$$

of the primal, respectively. Finally, suppose that x_S is not a 1-certificate, but $x_{S \cup \{j\}}$ is. That means that $p_{S,j}(x)$ is contained only in two conditions of the primal, corresponding to the variables

$$\begin{pmatrix} \phi_{S,j}(x) & \rho_{S,j}(x) \\ \rho_{S,j}(x) & \omega_{S,j}(x) \end{pmatrix}$$

and $\nu_S(x)$ of the dual. Thus the variable $\nu_{S \cup \{j\}}(x)$ never appears in the dual and we must force it to be 0 in (3.4e), and we do that by enforcing the condition (3.4f). Hence, the dual of (3.3) is

$$\text{maximize} \quad - \sum_{x \in \mathcal{D}_1} \mu(x) + \sum_{x \in \mathcal{D}_1} \nu_\emptyset(x) \quad (3.4a)$$

$$\text{subject to} \quad \sum_{y \in \mathcal{D}_0} \kappa_y \leq 1; \quad (3.4b)$$

$$\phi_{S,j}(x) - \mu(x) = 0 \quad \text{for all } (S, j) \in \mathcal{E} \text{ and } x \in \mathcal{D}_1; \quad (3.4c)$$

$$\sum_{\substack{x \in \mathcal{D}_1 \\ x_S = a}} \omega_{S,j}(x) - \sum_{\substack{y \in \mathcal{D}_1 \\ y_S = a}} \kappa_y \leq 0 \quad \text{for all } (S, j) \in \mathcal{E} \text{ and } a \in \Sigma^S; \quad (3.4d)$$

$$2\rho_{S,j}(x) + \nu_{S \cup \{j\}}(x) - \nu_S(x) = 0 \quad \text{for all } (S, j) \in \mathcal{E} \text{ and } x \in \mathcal{D}_1; \quad (3.4e)$$

$$\nu_S(x) = 0 \quad \text{whenever } x_S \text{ is a 1-certificate}; \quad (3.4f)$$

$$\begin{pmatrix} \phi_{S,j}(x) & \rho_{S,j}(x) \\ \rho_{S,j}(x) & \omega_{S,j}(x) \end{pmatrix} \succeq 0 \quad \text{for all } (S, j) \in \mathcal{E} \text{ and } x \in \mathcal{D}_1; \quad (3.4g)$$

$$\mu(x) \geq 0, \nu_S(x) \in \mathbb{R}, \kappa_y \geq 0 \quad \text{for all } S \subseteq [n], x \in \mathcal{D}_1, \text{ and } y \in \mathcal{D}_0. \quad (3.4h)$$

Here, in (3.4c), (3.4d), (3.4e) we have already expanded the inner product of $\begin{pmatrix} \phi_{S,j}(x) & \rho_{S,j}(x) \\ \rho_{S,j}(x) & \omega_{S,j}(x) \end{pmatrix}$

and, respectively, $\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$, $\begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$, $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$.

To construct a strictly feasible solution of (3.4), first choose any $\kappa_y > 0$ strictly satisfying (3.4b) and any $\mu(x) > 0$. Also choose $\rho_{S,j} = 0$ and $\nu_S(x) = 0$. And, finally, choose any $\omega_{S,j}(x) > 0$ strictly satisfying (3.4d) and $\phi_{S,j}(x) = \mu(x)$. Hence, we have the strong duality and both (3.3) and (3.4) attain the same optimal value.

Notice that (3.4c) requires us to have $\phi_{S,j}(x) = \mu(x) \geq 0$, therefore the conditions (3.4c), (3.4e), and (3.4g) can be replaced by a single condition

$$\mu(x)\omega_{S,j}(x) \geq (\nu_S(x) - \nu_{S \cup \{j\}}(x))^2/4 \quad \text{for all } (S, j) \in \mathcal{E} \text{ and } x \in \mathcal{D}_1,$$

eliminating the variables $\phi_{S,j}(x)$ and $\rho_{S,j}(x)$ in the process. Now, let us substitute

$$\alpha_S(x) := \nu_S(x) / (2\sqrt{\mu(x)}) \quad (3.5)$$

to obtain the following problem equivalent to (3.4):

$$\text{maximize} \quad \sum_{x \in \mathcal{D}_1} (2\sqrt{\mu(x)}\alpha_\emptyset(x) - \mu(x)) \quad (3.6a)$$

$$\text{subject to} \quad \sum_{y \in \mathcal{D}_1} \kappa_y \leq 1; \quad (3.6b)$$

$$\sum_{\substack{x \in \mathcal{D}_1 \\ x_S = a}} \omega_{S,j}(x) \leq \sum_{\substack{y \in \mathcal{D}_1 \\ y_S = a}} \kappa_y \quad \text{for all } (S, j) \in \mathcal{E} \text{ and } a \in \Sigma^S; \quad (3.6c)$$

$$\alpha_S(x) = 0 \quad \text{whenever } x_S \text{ is a 1-certificate}; \quad (3.6d)$$

$$\omega_{S,j}(x) \geq (\alpha_S(x) - \alpha_{S \cup \{j\}}(x))^2 \quad \text{for all } (S, j) \in \mathcal{E} \text{ and } x \in \mathcal{D}_1; \quad (3.6e)$$

$$\mu(x) \geq 0, \alpha_S(x) \in \mathbb{R}, \kappa_y \geq 0 \quad \text{for all } S \subseteq [n], x \in \mathcal{D}_1, \text{ and } y \in \mathcal{D}_0. \quad (3.6f)$$

One can verify that we do not need to be concerned about $\mu(x)$ being 0 in (3.5). In order to maximize (3.6a), we have to choose $\mu(x) = \alpha_\emptyset(x)^2$. And there is also no loss in choosing $\omega_{S,j}(x)$ as small as possible, namely, so that there is equality in (3.6e).

Recall that we obtained (3.3) from (3.1) by squaring the objective value. So, by doing the reverse, that is, taking the square root of (3.6a), we get that the dual of the adaptive learning graph complexity is

$$\text{maximize} \quad \sqrt{\sum_{x \in \mathcal{D}_1} \alpha_\emptyset(x)^2} \quad (3.7a)$$

$$\text{subject to} \quad \sum_{y \in \mathcal{D}_1} \kappa_y \leq 1; \quad (3.7b)$$

$$\sum_{\substack{x \in \mathcal{D}_1 \\ x_S = a}} (\alpha_S(x) - \alpha_{S \cup \{j\}}(x))^2 \leq \sum_{\substack{y \in \mathcal{D}_1 \\ y_S = a}} \kappa_y \quad \text{for all } (S, j) \in \mathcal{E}, \text{ and } a \in \Sigma^S; \quad (3.7c)$$

$$\alpha_S(x) = 0 \quad \text{whenever } x_S \text{ is 1-certificate}; \quad (3.7d)$$

$$\alpha_S(x) \in \mathbb{R}, \kappa_y \geq 0 \quad \text{for all } S \subseteq [n], x \in \mathcal{D}_1, \text{ and } y \in \mathcal{D}_0. \quad (3.7e)$$

3.1.2 SDPs for learning graph complexity of certificate structures

Again, consider a decision problem $\mathcal{P}: \Sigma^n \rightarrow \{0, 1\}$ with a domain $\mathcal{D} = \mathcal{D}_0 \sqcup \mathcal{D}_1$. Just like the adaptive learning graph complexity, the non-adaptive learning graph complexity of \mathcal{P} is given by (3.1), except now we require that $w_{S,j}(a)$ is independent from a ; we simply denote it by $w_{S,j}$. Hence, the 0-complexity $LG_0(\mathcal{G}, y)$ becomes the same for all $y \in \mathcal{D}_0$, and we can assume that the variable K in (3.1) equals it.

Suppose we have fixed all the weights $w_{S,j}$, and recall from Section 2.1.2 the minimal certificate structure $\mathcal{C}_{\mathcal{P}}$ of \mathcal{P} . If $x \in \mathcal{D}_1$ is such that $M_x \in \mathcal{C}_{\mathcal{P}}$, then all sinks of the flow $p(x)$ must be in

M_x . If $x \in \mathcal{D}_1$ is such that $M_x \notin \mathcal{C}_{\mathcal{P}}$, then there is $x' \in \mathcal{D}_1$ with $M_{x'} \subseteq M_x$, and we can choose $p(x) = p(x')$, yielding $LG_1(\mathcal{G}, x) = LG_1(\mathcal{G}, x')$. Hence, to compute the non-adaptive learning graph complexity of \mathcal{P} , it suffices to consider flows ending in each certificate placement of the minimal certificate structure. Because of this fact, it is useful to define the (non-adaptive) learning graph complexity of a certificate structure.

Definition 3.1. The learning graph complexity of a certificate structure \mathcal{C} on n variables is equal to the optimal value of the following semidefinite program:

$$\text{minimize } \sqrt{\sum_{(S,j) \in \mathcal{E}} w_{S,j}} \quad (3.8a)$$

$$\text{subject to } \sum_{(S,j) \in \mathcal{E}} \frac{p_{S,j}(M)^2}{w_{S,j}} \leq 1 \quad \text{for all } M \in \mathcal{C}; \quad (3.8b)$$

$$\sum_{j \in S} p_{S \setminus \{j\}, j}(M) = \sum_{j \notin S} p_{S,j}(M) \quad \text{for all } M \in \mathcal{C} \text{ and } S \in 2^{[n]} \setminus (M \cup \{\emptyset\}); \quad (3.8c)$$

$$\sum_{j \in [n]} p_{\emptyset, j}(M) = 1 \quad \text{for all } M \in \mathcal{C}; \quad (3.8d)$$

$$p_{S,j}(M) \in \mathbb{R}, w_{S,j} \geq 0 \quad \text{for all } (S,j) \in \mathcal{E} \text{ and } M \in \mathcal{C}, \quad (3.8e)$$

where $0/0$ in (3.8b) is defined to be 0.

Claim 3.2. *Suppose \mathcal{C} is a certificate structure of a decision problem \mathcal{P} . The non-adaptive learning graph complexity of \mathcal{P} is at most the learning graph complexity of \mathcal{C} , with equality achieved if $\mathcal{C} = \mathcal{C}_{\mathcal{P}}$.*

Proof. Suppose we are given a learning graph for \mathcal{C} (i.e., a feasible solution of (3.8)). We construct a non-adaptive learning graph for \mathcal{P} by using exactly the same weights and, for $x \in \mathcal{D}_1$, we choose $p(x) = p(M')$, where M' is a certificate placement satisfying $M' \subseteq M_x$ (such M' exists by the definition of certificate structures). Hence, the complexity of the constructed learning graph for \mathcal{P} is at most that of the original learning graph for \mathcal{C} . And we have already considered the $\mathcal{C} = \mathcal{C}_{\mathcal{P}}$ case above Definition 3.1. \square

This claim shows that, for example, THRESHOLD- k , k -DISTINCTNESS, and k -SUM have the same non-adaptive learning graph complexity, and so do TRIANGLE and TRIANGLE-SUM. Analogously to Theorem 2.19, we have

Theorem 3.3 ([Bel12d], [BL11]). *The quantum query complexity of a certificate structure is at most a constant times its learning graph complexity.*

In Section 4.3, using the dual of (3.8), we prove the reverse statement for all certificate structures. We can obtain the dual the same way as we did for adaptive learning graphs in Section 3.1.1, and we omit these derivations. (In fact, now the process of obtaining the dual is simpler, as the primal has less variables.) The dual of (3.8) is

$$\text{maximize } \sqrt{\sum_{M \in \mathcal{C}} \alpha_\emptyset(M)^2} \quad (3.9a)$$

$$\text{subject to } \sum_{M \in \mathcal{C}} (\alpha_S(M) - \alpha_{S \cup \{j\}}(M))^2 \leq 1 \quad \text{for all } (S, j) \in \mathcal{E}; \quad (3.9b)$$

$$\alpha_S(M) = 0 \quad \text{whenever } S \in M; \quad (3.9c)$$

$$\alpha_S(M) \in \mathbb{R} \quad \text{for all } S \subseteq [n] \text{ and } M \in \mathcal{C}. \quad (3.9d)$$

Again, one can see that (3.8) and (3.9) are strongly dual. We call a feasible solution of (3.9) a *dual learning graph of a certificate structure* \mathcal{C} .

3.2 Learning graph complexity of certificate structures

In this section, we construct dual learning graphs for certificate structures considered in Section 2.1.2. Let \mathcal{T} be the target objective value of (3.9), namely, the lower bound we want to prove. We construct $\alpha_S(M)$ for all certificate structures \mathcal{C} , implicitly or explicitly, in the form

$$\alpha_S(M) := \begin{cases} \max\{\mathcal{T} - |S| - \sum_{i=1}^m g_i(S, M), 0\} / \sqrt{|\mathcal{C}|}, & \text{if } S \notin M, \\ 0, & \text{otherwise,} \end{cases} \quad (3.10a)$$

$$(3.10b)$$

where $g_i(S, M)$ is a function satisfying $g_i(S, M) \geq 0$ and $g_i(\emptyset, M) = 0$. The objective value (3.9a) is therefore indeed $\sqrt{\sum_{M \in \mathcal{C}} \mathcal{T}^2 / |\mathcal{C}|} = \mathcal{T}$. In Section 3.3, when constructing lower bounds on the adaptive learning graph complexity, we use a form very similar to (3.10) (in particular, see (3.20) and (3.26)).

In practice, we care only about asymptotic behaviour of the learning graph complexity, and, instead of (3.9b), we use

$$\sum_{M \in \mathcal{C}} (\alpha_S(M) - \alpha_{S \cup \{j\}}(M))^2 = O(1) \quad \text{for all } (S, j) \in \mathcal{E}. \quad (3.11)$$

From (3.10), by the Cauchy–Schwarz inequality, we get

$$\begin{aligned} (\alpha_S(M) - \alpha_{S \cup \{j\}}(M))^2 &\leq \frac{m+2}{|\mathcal{C}|} \left(1 + \sum_{i=1}^m (g_i(S \cup \{j\}, M) - g_i(S, M))^2 \right. \\ &\quad \left. + \begin{cases} \mathcal{T}^2, & \text{if } S \notin M \text{ and } S \cup \{j\} \in M \\ 0, & \text{otherwise} \end{cases} \right), \end{aligned} \quad (3.12)$$

where 1 comes from $(|S \cup \{j\}| - |S|)^2$. Thus, if $m = O(1)$, we can add the term $|S|$ in (3.10a) without loss of generality as $\sqrt{\sum_{M \in \mathcal{C}} 1/|\mathcal{C}|} = 1$. If $m = \omega(1)$, as in Section 3.2.2 below, one requires to scale down all $\alpha_S(M)$ of (3.10) by a factor of \sqrt{m} , thus scaling down the objective value to \mathcal{T}/\sqrt{m} .

The term $|S|$ in (3.10a) ensures that we never have to consider the condition (3.11) for S of size larger than \mathcal{T} . The purpose of this term together with the $g_i(S, M)$ terms is to ensure that not many $M \in \mathcal{C}$ experience the “jump” from the first case (3.10a) for $\alpha_S(M) > 0$ to the second case (3.10b) for $\alpha_{S \cup \{j\}}(M)$. We want to limit the number of such jumps as each jump may contribute to the left hand side of (3.11) as much as \mathcal{T}^2 .

3.2.1 Lower bounds for the k -subset and hidden shift certificate structures

Using the dual learning graph (3.9), let us construct lower bounds on the learning graph complexity of the k -subset and hidden shift certificate structures. For both certificate structures, we construct $\alpha_S(M)$ as in (3.10) without any $g_i(S, M)$ terms (i.e., $m = 0$). As discussed above, this means that the objective value (3.9a) is \mathcal{T} , and all we have to do is to show that (3.11) holds.

Proposition 3.4. *Given a constant k , the learning graph complexity of the k -subset certificate structure is $\Omega(n^{k/(k+1)})$.*

Proof. Let \mathcal{C} be the k -subset certificate structure. Note that $|\mathcal{C}| = \binom{n}{k}$, and let $\mathcal{T} := n^{k/(k+1)}$. Take any $(S, j) \in \mathcal{E}$. If $|S| \geq \mathcal{T} = n^{k/(k+1)}$, then $\alpha_S(M) = \alpha_{S \cup \{j\}}(M) = 0$, and we are done. Thus, we further assume $|S| < n^{k/(k+1)}$. There are at most $\binom{|S|}{k-1} \leq n^{k(k-1)/(k+1)}$ choices of $M \in \mathcal{C}$ such that $S \notin M$ and $S \cup \{j\} \in M$. For each of them, the value of $\alpha_S(M)$ changes by at most $\binom{n}{k}^{-1/2} n^{k/(k+1)}$. Thus, the sum of (3.12) over all $M \in \mathcal{C}$ is

$$\sum_{M \in \mathcal{C}} (\alpha_S(M) - \alpha_{S \cup \{j\}}(M))^2 \leq 2(n^{k(k-1)/(k+1)} \mathcal{T}^2 + |\mathcal{C}| \cdot 1) / |\mathcal{C}| = O(1).$$

□

Proposition 3.5. *The learning graph complexity of the hidden shift (and, hence, the set equality and the collision) certificate structure is $\Omega(n^{1/3})$.*

Proof. Let \mathcal{C} be the hidden shift certificate structure. Note that $|\mathcal{C}| = n$, and let $\mathcal{T} := n^{1/3}$. Take any $(S, j) \in \mathcal{E}$, and again we may assume that $|S| < \mathcal{T} = n^{1/3}$. There are at most $|S| = n^{1/3}$ choices of $M \in \mathcal{C}$ such that $S \notin M$ and $S \cup \{j\} \in M$. Thus,

$$\sum_{M \in \mathcal{C}} (\alpha_S(M) - \alpha_{S \cup \{j\}}(M))^2 \leq 2(n^{1/3} \mathcal{T}^2 + |\mathcal{C}| \cdot 1) / |\mathcal{C}| = O(1).$$

For the set equality and the collision certificate structures, just assign $\alpha_S(M) = 0$ for all M that are not in the hidden shift certificate structure (this is a case of Remark 2.16). \square

The results of Propositions 3.4 and 3.5 are tight. Belovs, Lee, and Zhu show an $O(n^{k/(k+1)})$ upper bound for the k -subset certificate structure [BL11, Zhu12], and an $O(n^{1/3})$ upper bound for the collision (and, hence, the set equality and the hidden shift) certificate structure can be derived by similar methods (we omit here the construction of the corresponding learning graph).

As illustrated by the proofs of the two propositions above, in general, one can choose \mathcal{T} such that $|S| < \mathcal{T}$ ensures that the number of $M \in \mathcal{C}$ such that $S \notin M$ and $S \cup \{j\} \in M$ is at most $|\mathcal{C}|\mathcal{T}^2$. For the triangle certificate structure, this allows us to choose the value of \mathcal{T} no higher than n (recall that the number of input variables for this structure is $\binom{n}{2}$).

3.2.2 Lower bound for the triangle certificate structure

In this section we prove an almost tight lower bound on the learning graph complexity of the triangle certificate structure. The best known upper bound $O(n^{9/7})$ was given by Lee, Magniez, and Santha [LMS13], improving upon an $O(n^{35/27})$ upper bound by Belovs [Bel12d].

Theorem 3.6. *The learning graph complexity of the triangle certificate structure (and, thus, the non-adaptive learning graph complexity of TRIANGLE) is $\Omega(n^{9/7}/\sqrt{\log n})$.*

The proof of this lower bound is rather bulky. It resulted from a weaker, yet non-trivial $\Omega(n^{5/4})$ lower bound. The $\Omega(n^{5/4})$ lower bound can be found in Ref. [Bel13], and it too is constructed in the form (3.10).

Proof of Theorem 3.6. Let $E = \{uu' \mid 1 \leq u < u' \leq n\}$ be the set of input variables (potential edges of the graph); this is not to be confused with \mathcal{E} , the set of arcs of the learning graph. Let \mathcal{C} be the triangle certificate structure. For each $M \in \mathcal{C}$, fix three vertices $a = a(M), b = b(M), c = c(M)$ forming the triangle: $S \in M$ if and only if $ab, ac, bc \in S$. (Technically a, b , and c are functions, but, whenever M is clear from the context, we can think of them as vertices.) Let $L := \{a, b, c\}$, where we think of $v \in L$ as an indicator which function (a, b , or c) to consider (that is, the set L is independent from M , while $\{a(M), b(M), c(M)\}$ is not). All definitions and arguments made for $v = a$ translate to $v = b$ and $v = c$ by symmetry.

We construct the dual learning graph (3.9) in the form

$$\alpha_S(M) := \begin{cases} \max \{n^{-3/14} - \sum_{i=0}^m \sum_{v \in L} g_{i,v}(S, M), 0\}, & S \notin M, \\ 0, & \text{otherwise,} \end{cases}$$

where $g_{i,v}(S, M)$ is a function satisfying $0 \leq g_{i,v}(S, M) \leq n^{-3/14}$ and $g_{i,v}(\emptyset, M) = 0$. Hence the objective value (3.9a) is $\sqrt{\binom{n}{3}} n^{-3/14} = \Omega(n^{9/7})$. The hard part will be to show that (3.9b) holds up to logarithmic factors.

We define

$$g_{0,v}(S, M) := \min\{n^{-3/2}|S|, n^{-3/14}\}$$

(corresponding to the $|S|$ term in (3.10)). Hence, $\alpha_S(M) = 0$ if $|S| \geq n^{9/7}$, and from now on we assume $|S| \leq n^{9/7}$. We will define $g_{i,v}(S, M)$ for $i \in [1..m]$ later.

For $S \subset E$ and $j \in E \setminus S$, let $F(S, j) \subset \mathcal{C}$ denote the set of $M \in \mathcal{C}$ such that $S \notin M$, but $S \cup \{j\} \in M$. We decompose

$$F(S, j) = \bigsqcup_{i=1}^m \bigsqcup_{v \in L} F_{i,v}(S, j)$$

as follows. Let $\deg_S a = \deg_S a$ be the degree of a vertex a in the graph with edge set S . A certificate placement $M \in F(S, j)$ belongs to $F_{1,a}(S, j)$ if $j = bc$ and $\deg a \leq n^{3/7}$ and, for $i \geq 2$, to $F_{i,a}(S, j)$ if $j = bc$ and $2^{i-2}n^{3/7} < \deg a \leq 2^{i-1}n^{3/7}$. Hence, $m \approx (4/7) \log_2 n$.

For all $i \in [1..m]$, we will define $g_{i,v}(S, M)$ so that, for all $v \in L$, $S \subset E$ of size at most $n^{9/7}$, and $j \in E \setminus S$:

$$\sum_{M \in \mathcal{C} \setminus F(S, j)} (g_{i,v}(S, M) - g_{i,v}(S \cup \{j\}, M))^2 = O(1) \quad (3.13)$$

and

$$\sum_{M \in F_{i,v}(S, j)} (n^{-3/14} - g_{i,v}(S, M))^2 = O(1). \quad (3.14)$$

Note that (3.13) also holds for $i = 0$. Even more, we will show that the set $I := I(S, j)$ of $i \in [0..m]$ such that (3.13) is non-zero has size $O(1)$. Thus, for the left hand side of (3.9b), we will have

$$\begin{aligned} \sum_{M \in \mathcal{C}} (\alpha_S(M) - \alpha_{S \cup \{j\}}(M))^2 &= \sum_{M \in \mathcal{C} \setminus F(S, j)} \left(\sum_{i \in I} \sum_{v \in L} (g_{i,v}(S, M) - g_{i,v}(S \cup \{j\}, M)) \right)^2 \\ &\quad + \sum_{i=1}^m \sum_{v \in L} \sum_{M \in F_{i,v}(S, j)} (\alpha_S(M))^2 \\ &\leq |3I| \sum_{i \in I} \sum_{v \in L} \sum_{M \in \mathcal{C} \setminus F(S, j)} (g_{i,v}(S, M) - g_{i,v}(S \cup \{j\}, M))^2 \end{aligned} \quad (3.15)$$

$$+ \sum_{i=1}^m \sum_{v \in L} \sum_{M \in F_{i,v}(S, j)} (n^{-3/14} - g_{i,v}(S, M))^2, \quad (3.16)$$

which comes from the Cauchy–Schwarz inequality and the fact that

$$|\alpha_S(M)| \leq |n^{-3/14} - g_{i,v}(S, M)|.$$

Due to (3.13) and (3.14), the sum (3.15) is $O(1)$ and the sum (3.16) is $O(\log n)$. By scaling all $\alpha_S(M)$ down by a factor of $O(\sqrt{\log n})$, we obtain a feasible solution to (3.9) with the objective value $\Omega(n^{9/7}/\sqrt{\log n})$.

It remains to construct the functions $g_{i,v}(S, M)$ for $i \in [1..m]$ that satisfy (3.13) and (3.14). In the following, let $\mu(x)$ be the median of 0, x , and 1, i.e., $\mu(x) = \max\{0, \min\{x, 1\}\}$.

Case $i = 1$. Let us define

$$g_{1,a}(S, M) = \begin{cases} n^{-3/14} \mu(2 - n^{-3/7} \deg a), & ab, ac \in S, \\ 0, & \text{otherwise.} \end{cases} \quad (3.17)$$

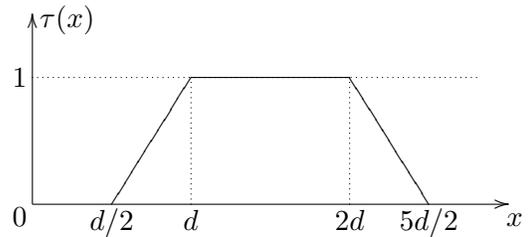
Clearly, $0 \leq g_{1,a}(S, M) \leq n^{-3/14}$ and $g_{1,a}(\emptyset, M) = 0$. We distinguish two cases how $g_{1,a}(S, M)$ may be influenced by ‘‘loading’’ j when $S \cup \{j\} \notin M$. We show that the total contribution to (3.13) is $O(1)$.

- It may happen if $|\{ab, ac\} \cap S| = 1$ and $j \in \{ab, ac\}$, i.e., the transition from the second case of (3.17) to the first one happens. Moreover, $g_{1,a}(S, M)$ changes (i.e., $g_{1,a}(S \cup \{j\}, M) \neq g_{1,a}(S, M)$) only if $\deg a \leq 2n^{3/7}$. Then j identifies two vertices of the triangle, and the third one is among the neighbours of an end-point of j having degree at most $2n^{3/7}$. Thus, the total number of M satisfying this scenario is at most $4n^{3/7}$. The contribution to (3.13) is at most $O(n^{3/7})(n^{-3/14})^2 = O(1)$.
- Another possibility is that $ab, ac \in S$ and $\deg a$ changes. In this case, a is determined as an end-point of j , and b and c are among its at most $2n^{3/7}$ neighbours. The number of M influenced is $O(n^{6/7})$, and the contribution is $O(n^{6/7})(n^{-9/14})^2 = o(1)$, where $n^{-9/14}$ is the coefficient of $\deg a$ in (3.17).

Finally, we have to show that (3.14) holds. If $M \in F_{1,a}$, then $ab, ac \in S$ and $\deg a \leq n^{3/7}$. However, (3.17) implies that, in this case, the left hand side of (3.14) is 0.

Case $i \geq 2$. For $d := 2^{i-2}n^{3/7} \geq n^{3/7}$, define a piece-wise linear function τ as follows

$$\tau(x) = \begin{cases} 0, & x < d/2; \\ (2x - d)/d, & d/2 \leq x < d; \\ 1, & d \leq x < 2d; \\ (5d - 2x)/d, & 2d \leq x \leq 5d/2; \\ 0, & x \geq 5d/2. \end{cases}$$



It can be interpreted as a continuous version of the indicator function that a vertex has a right degree (i.e., a degree between d and $2d$). Define

$$\nu(S, M) = \nu_{i,a}(S, M) := \sum_{u \in N(b) \cap N(c)} \tau(\deg u),$$

where the sum is over the common neighbours of b and c . Let

$$g_{i,a}(S, M) := n^{-3/14} \mu \left(\min \left\{ \frac{2 \deg a}{d}, \frac{\nu(S, M)}{n^{3/7}} \right\} - 1 \right). \quad (3.18)$$

Let us consider how $g_{i,a}(S, M)$ may change by loading j when $S \cup \{j\} \notin M$ and how this contributes to (3.13). Now there are three cases how $g_{i,a}(S, M)$ may be influenced. We again show that the total contribution to (3.13) is $O(1)$.

- It may happen that j is incident to a common neighbour of b and c , and thus $\nu(S, M)$ may change. This means b and c are among the neighbours of an end-point of j of degree at most $5d/2$. Hence, this affects $O(nd^2)$ different M . The contribution is $O(nd^2)(n^{-9/14}/d)^2 = o(1)$, where $n^{-9/14}$ is the coefficient of $\nu(S, M)$ in (3.18) and $\nu(S, M)$ changes by at most $1/d$.
- The set $N(b) \cap N(c)$ may increase. This causes a change in $g_{i,a}(S, M)$ only under the following circumstances. The new edge j is incident to b or c . The second vertex in $\{b, c\}$ is among $\Theta(d)$ neighbours of the second end-point of j . Finally, $\deg a \geq d/2$, that together with $|S| \leq n^{9/7}$ implies that there are $O(n^{9/7}/d)$ choices for a . Altogether, the number of M affected by this is $O(n^{9/7})$, and the change in $g_{i,a}(S, M)$ does not exceed $n^{-9/14}$. The contribution is $O(1)$.
- The degree of a may change. Let us calculate the number P of possible pairs b and c affected by this. There is a change in $g_{i,a}(S, M)$ only if b and c are connected to at least $n^{3/7}$ vertices of degrees between $d/2$ and $5d/2$. Denote the set of these vertices by A . Since $|S| \leq n^{9/7}$, we have $|A| = O(n^{9/7}/d)$.

Let us calculate the number of paths of length 2 in S having the middle vertex in A . On one hand, this number is at least $Pn^{3/7}$, as each pair b and c has at least $n^{3/7}$ common neighbours. On the other hand, it is at most $O(d^2|A|) = O(dn^{9/7})$. Thus, $P = O(dn^{6/7})$. Since a is determined as an end-point of j , the contribution is $O(dn^{6/7})(n^{-3/14}/d)^2 = O(1)$, as $d \geq n^{3/7}$.

Finally, j may be the last edge of the triangle: $S \cup \{j\} \in M$. If $M \in F_{i,a}(S, j)$, then $\deg a > d$, implying that, in this case, $g_{i,a}(S, M) = n^{-3/14} \mu(\nu(S, M)n^{-3/7} - 1)$. Hence, either $n^{-3/14} - g_{i,a}(S, M) = 0$, or $\nu(S, M) \leq 2n^{3/7}$, in which case, there are $O(n^{3/7})$ choices of a satisfying the condition. Hence, the left hand side of (3.14) is $O(n^{3/7})(n^{-3/14})^2 = O(1)$.

Regarding the set $I = I(S, j)$, note that, if $g_{i,a}(S, M) - g_{i,a}(S \cup \{j\}, M) \neq 0$, then either $i = 0$, $i = 1$, or the degree of one of the end-points of j must be between approximately $d/2$ and $5d/2$, where $d = 2^{i-2}n^{3/7}$. Hence, the set I has constant size, as required. \square

3.3 Lower bounds on adaptive learning graph complexity

3.3.1 Adaptive learning graph complexity of the AND function

Proposition 3.7. *The adaptive learning graph complexity of AND is $\Omega(n)$.*

Proof. Let us construct a feasible solution of (3.7), whose objective value gives a lower bound on the adaptive learning graph complexity of AND. The AND function has only one positive input, the all-ones string 1^n , so $\mathcal{D}_1 = \{1^n\}$. Let $Y \subset \mathcal{D}_0$ be the set of negative inputs that contain exactly one zero and $n - 1$ ones. We set $\kappa_y = 1/|Y| = 1/n$ for all $y \in Y$ and $\kappa_y = 0$ for all $y \in \mathcal{D}_0 \setminus Y$ (here we essentially exploit Remark 2.16), thereby saturating the condition (3.7b).

It remains to set $\alpha_S(1^n)$ for all S . Note that, if $a \in \{0, 1\}^S$ contains at least one zero, then the left hand side of (3.7c) equals 0, and the condition (3.7c) is satisfied. On the other hand, if $a = 1^{|S|}$, the condition (3.7c) becomes

$$|\alpha_S(1^n) - \alpha_{S \cup \{j\}}(1^n)| \leq \sqrt{\frac{n - |S|}{n}} \quad \text{for all } (S, j) \in \mathcal{E}, \quad (3.19)$$

because there are exactly $n - |S|$ inputs $y \in Y$ such that $y_S = 1^{|S|}$. We set

$$\alpha_S(1^n) := \begin{cases} 3n/8 - |S|/2, & \text{if } |S| < 3n/4; \\ 0, & \text{otherwise.} \end{cases} \quad (3.20)$$

This clearly satisfies (3.19), as the right hand side of (3.19) is at least $1/2$ whenever $|S| < 3n/4$. The objective value (3.7a) is $3n/8$, which concludes the proof. \square

3.3.2 Adaptive learning graph complexity of k -Distinctness

Recall the k -DISTINCTNESS problem from Section 2.1.1. Let us construct a feasible solution to (3.7), yielding the following lower bound on the adaptive learning graph complexity of k -DISTINCTNESS.

Theorem 3.8. *Given a constant k , the adaptive learning graph complexity of the k -DISTINCTNESS problem is $\Omega(n^{1-2^{k-2}/(2^k-1)})$.*

Notice that this lower bound matches the best known upper bound on the quantum query complexity of k -DISTINCTNESS by Belovs [Bel12c]. This upper bound, however, uses span programs that are more general than adaptive learning graphs. An adaptive learning graph of complexity $O(n^{1-2^{k-2}/(2^k-1)})$ was given by Lee and Belovs assuming that, for all $i < k$, one knows approximately how many characters in Σ appear in an input x exactly i times [BL11]. For the inputs that we consider in the following proof, we have such knowledge, therefore $\Omega(n^{1-2^{k-2}/(2^k-1)})$ is the best possible lower bound one can obtain in this setting.

Proof of Theorem 3.8. Note that both sides of (3.7c) are zero whenever a is a 1-certificate, so we are concerned only with the case when it is not. Since k is constant, for notational convenience, we assume that the length of the input is $(k-1)n$. Without loss of generality, we also assume that the size of the input alphabet Σ is n ; the case $|\Sigma| < n$ is trivial and a lower bound for $|\Sigma| = n$ is also a lower bound whenever $|\Sigma| > n$.

Let $Y := \mathcal{D}_0$ be the set of all negative inputs. Due to our choice of the alphabet size, all negative inputs are the same up to an index permutation—every negative input y contains each character exactly $k-1$ times (we also say: each character has multiplicity $k-1$ in y). Due to this symmetry, we choose $\kappa_y = 1/|Y|$ for all $y \in Y$ to saturate the condition (3.7b).

Let $X \subset \mathcal{D}_1$ be the set of positive inputs that contain one character $2(k-1)$ times, $n-2$ characters $k-1$ times, and does not contain one character at all. We choose $\alpha_S(x) = 0$ for all $x \in \mathcal{D}_1 \setminus X$ and all S . (Essentially, we exploit Remark 2.16 here.)

Cardinalities of various sets of inputs. It is rather simple to compute cardinalities of X and Y , for example,

$$|Y| = \frac{((k-1)n)!}{((k-1)!)^n n!},$$

but, in fact, we will only need to know the ratio of these two cardinalities. Therefore, let us consider the following procedure that from one negative input constructs $n(n-1)$ positive inputs.

Procedure 3.9. Given a negative input $y \in Y$, first choose a character $v \in \Sigma$, then choose a character $v' \neq v$, and then substitute each occurrence of v' in y with v . Let P be the function that takes y as an input and outputs the set of $n(n-1)$ positive $x \in X$ constructed this way.

One can see that, for every positive input $x \in X$, there are exactly $\binom{2(k-1)}{k-1}$ negative inputs $y \in Y$ such that $x \in P(y)$. Hence,

$$\frac{|X|}{|Y|} = \binom{2(k-1)}{k-1}^{-1} n(n-1) = \Theta(n^2). \quad (3.21)$$

Fix arbitrary $(S, j) \in \mathcal{E}$ (i.e., $S \subset [(k-1)n]$ and $j \in [(k-1)n] \setminus S$) and $a \in \Sigma^S$. For $i \geq 0$, let $\ell_i(a)$ be the number of characters that has multiplicity i in a . Note that

$$\ell_0(a) = n - \sum_{i \geq 1} \ell_i(a) \quad (3.22)$$

is the number of characters that do not appear in a . Let us assume that $|S| = o(n)$ and that each character appears in a at most $k-1$ times (as in Section 3.2 for non-adaptive learning graphs, we later define $\alpha_S(x)$ to be 0 whenever S is large).

Let $Y(a) \subseteq Y$ be the set of negative inputs compatible with a and, for $i \in [0..k-2]$, let $Y_i(a, j)$ be its subset consisting of $y \in Y(a)$ such that the character y_j has multiplicity i in a .

Claim 3.10. *For $i \in [0..k-2]$, we have*

$$\frac{|Y_i(a, j)|}{|Y(a)|} = \frac{(k-1-i) \cdot \ell_i(a)}{(k-1)n - |S|} = \Theta\left(\frac{\ell_i(a)}{n}\right). \quad (3.23)$$

Proof. It is clear that the cardinality of $Y_i(a, j)$ is independent from which $j \in [(k-1)n] \setminus S$ one considers. Hence, if we choose any y consistent with a and uniformly at random choose $j \in [(k-1)n] \setminus S$, then the ratio (3.23) is the probability that y_j has multiplicity i in a . \square

Analogously, for $i \in [0..k-1]$, let $X_i(a, j)$ be the set of positive inputs $x \in X$ such that x is compatible with a and the character x_j has multiplicity i in a (i.e., $x_S = a$ and $\ell_{i+1}(x_{S \cup \{j\}}) = \ell_{i+1}(x_S) + 1$). The case $i = k-1$ is special, as it implies that $x_{S \cup \{j\}}$ is a 1-certificate. For all other i , we have the following.

Claim 3.11. *For $i \in [0..k-2]$, we have*

$$\frac{|X_i(a, j)|}{|Y_i(a, j)|} \leq n(n-1).$$

Proof. For every $x \in X_i(a, j)$, there exists $y \in Y_i(a, j)$ such that $x \in P(y)$. Hence,

$$X_i(a, j) \subseteq \bigcup_{y \in Y_i(a, j)} P(y),$$

and $|P(y)| = n(n-1)$ completes the proof. \square

Hence, Claims 3.10 and 3.11 together imply that

$$\frac{|X_i(a, j)|}{|Y(a)|} = O(n\ell_i(a)). \quad (3.24)$$

We are left to consider the special case $i = k-1$, and we have

Claim 3.12.

$$\frac{|X_{k-1}(a, j)|}{|Y(a)|} \leq \ell_{k-1}(a). \quad (3.25)$$

Proof. The character that has multiplicity $2(k-1)$ in $x \in X_{k-1}(a, j)$ has to be one of $\ell_{k-1}(a)$ characters that has multiplicity $k-1$ in a . Fix one such character v . Let $X_{k-1}^{(v)}(a, j)$ be the set of positive inputs $x \in X_{k-1}(a, j)$ that contain v with multiplicity $2(k-1)$. Consider the following two-step procedure:

1. choose $x \in X_{k-1}^{(v)}(a, j)$ and let v' be the character not present in x ;
2. obtain $y(x) \in Y(a)$ by substituting all $k-1$ occurrences of v outside S by v' .

Note that $y(x)_j = v'$ because $x_j = v$, and one can see that $y(x) \neq y(x')$ for all $x, x' \in X_{k-1}^{(v)}(a, j)$ such that $x \neq x'$. Hence, $|X_{k-1}^{(v)}(a, j)| \leq |Y(a)|$, and there are $\ell_{k-1}(a)$ possible choices for v . \square

Choosing values of α 's. Due to symmetry, we choose $\alpha_\emptyset(x)$ to be the same for all $x \in X$. If \mathcal{T} is the target objective value of (3.7), from (3.7a) we get $\alpha_\emptyset(x) = \mathcal{T}/\sqrt{|X|}$. We construct $\alpha_S(x)$ in the form

$$\alpha_S(x) = \begin{cases} 0, & \text{if } x_S \text{ is a 1-certificate;} \\ \max\{\mathcal{T} - \sum_{i=1}^{k-1} \gamma_i \ell_i(x_S), 0\} / \sqrt{|X|}, & \text{otherwise.} \end{cases} \quad (3.26)$$

When we optimize the coefficients \mathcal{T} and γ_i later, we choose

$$1 = \gamma_1 \ll \gamma_2 \ll \dots \ll \gamma_{k-1} \ll \mathcal{T} \ll n,$$

where $\gamma \ll \gamma'$ stands for $\gamma = o(\gamma')$. Hence, $\alpha_S(x) = 0$ whenever $|S| \geq \mathcal{T}$.

Let us now consider the condition (3.7c). Again, fix arbitrary $S \subset [(k-1)n]$ of size $|S| = o(n)$, $j \in [(k-1)n] \setminus S$, and $a \in \Sigma^S$ that is not a 1-certificate. The right hand side of (3.7c) is equal to $|Y(a)|/|Y|$, so let us now consider the left hand side.

First note that we always have $\alpha_{S \cup \{j\}}(x) \leq \alpha_S(x)$ because $\gamma_{i+1} > \gamma_i$ for all i . Indeed, there is a unique i such that $\ell_i(x_{S \cup \{j\}}) = \ell_i(x_S) - 1$, $\ell_{i+1}(x_{S \cup \{j\}}) = \ell_{i+1}(x_S) + 1$, and $\ell_{i'}(x_{S \cup \{j\}}) = \ell_{i'}(x_S)$ for all $i' \notin \{i, i+1\}$. This also implies that, if $\alpha_S(x) = 0$, then $\alpha_{S \cup \{j\}}(x) = 0$ too. So let us assume that $\alpha_S(x) > 0$ and also $x_S = a$. This enforces

$$\ell_i(a)\gamma_i < \mathcal{T} \quad \forall i \in \{1, \dots, k-1\}, \quad \text{and} \quad \ell_0(a) = n - o(n) \quad (3.27)$$

is due to (3.22) and $|S| = o(n)$.

For $i \in [0..k-1]$, there are exactly $|X_i(a, j)|$ positive inputs consistent with a such that $\ell_{i+1}(x_{S \cup \{j\}}) = \ell_{i+1}(x_S) + 1$. Hence, from the definition (3.26) of $\alpha_S(x)$ we get that the left hand side of (3.7c) is at most

$$\begin{aligned} & \sum_{i=0}^{k-2} |X_i(a, j)| \frac{(\gamma_{i+1} - \gamma_i)^2}{|X|} + |X_{k-1}(a, j)| \frac{(\mathcal{T} - \gamma_{k-1})^2}{|X|} \\ &= O\left(\sum_{i=0}^{k-2} n \ell_i(a) |Y(a)| \frac{\gamma_{i+1}^2}{n^2 |Y|} + \ell_{k-1}(a) |Y(a)| \frac{\mathcal{T}^2}{n^2 |Y|}\right) \\ &= \frac{|Y(a)|}{|Y|} \cdot O\left(\max\left\{\gamma_1^2, \frac{\mathcal{T}\gamma_2^2}{\gamma_1 n}, \frac{\mathcal{T}\gamma_3^2}{\gamma_2 n}, \dots, \frac{\mathcal{T}\gamma_{k-1}^2}{\gamma_{k-2} n}, \frac{\mathcal{T}^3}{\gamma_{k-1} n^2}\right\}\right), \end{aligned}$$

where $\gamma_0 = 0$, the first Big-O relation comes from (3.21), (3.24), and (3.25), and the second from (3.27) and the fact that k is constant. In order for (3.7c) to hold, we need to have

$$\max\left\{\gamma_1^2, \frac{\mathcal{T}\gamma_2^2}{\gamma_1 n}, \frac{\mathcal{T}\gamma_3^2}{\gamma_2 n}, \dots, \frac{\mathcal{T}\gamma_{k-1}^2}{\gamma_{k-2} n}, \frac{\mathcal{T}^3}{\gamma_{k-1} n^2}\right\} = O(1),$$

which is saturated by choosing

$$\mathcal{T} = n^{1-2^{k-2}/(2^k-1)} \quad \text{and} \quad \gamma_i = n^{(2^{k-2}-2^{k-1-i})/(2^k-1)} \quad \forall i \in \{1, \dots, k-1\}.$$

This concludes the proof. □

Chapter 4

Adversary bounds using matrix embedding

In Section 2.3.2 we described multiple tools that simplify construction of adversary bounds, i.e., adversary matrices Γ in Theorem 2.15. To various degrees, we will use all of those tools in this chapter. Most notably, rows and columns of adversary matrices Γ will correspond to positive inputs and negative inputs, respectively, and we will embed Γ in larger matrices $\tilde{\Gamma}$.

Suppose we are given a decision problem $\mathcal{P} : \Sigma^n \rightarrow \{0, 1\}$. In this chapter, let q denote the size of the input alphabet. Here we also assume that $\Sigma := [q]$ (except in Section 4.3.5, where we will have to consider a more elaborate input alphabet). For all decision problems considered in this chapter, we have that, if q is large enough, an input string chosen from $[q]^n$ uniformly at random is a negative input of \mathcal{P} with constant probability. All matrices $\tilde{\Gamma}$ considered here will have columns labeled by all input strings in $[q]^n$.

To choose a random input, one can independently and uniformly at random choose each of its symbols. This independence makes the construction of adversary bounds much easier. For example, the adversary bound for ELEMENT DISTINCTNESS with large range (Section 4.1) is much simpler than the adversary bound for ELEMENT DISTINCTNESS with minimal range (Chapter 5). In the latter case, the probability of a random input being negative is the minuscule $n!/n^n$.

Let $\mathcal{H} := \mathbb{C}^{[q]}$ be the space corresponding to all symbols of the input alphabet. Recall that $\{\mathbf{j} : j \in [q]\}$ is the *standard* basis of \mathcal{H} , and, for a vector $v \in \mathcal{H}$ given in the standard basis, $v_j = v[\mathbf{j}] = \mathbf{j}^* v$ denotes its j -th entry. An *e-basis* of \mathcal{H} is an orthonormal basis e_0, e_1, \dots, e_{q-1} satisfying $e_0[\mathbf{j}] = 1/\sqrt{q}$ for all $j \in [q]$. The precise choice of the remaining basis elements is irrelevant (except in Section 4.3.5).

Let $\mathcal{H} = \mathcal{H}_0 \oplus \mathcal{H}_1$, where

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

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, $\Pi_0 = e_0 e_0^* = \mathbb{J}_q/q$ and $\Pi_1 = \mathbb{I}_q - \mathbb{J}_q/q$ (recall that \mathbb{J}_q is the all-ones matrix in the standard basis).

The standard basis vectors of the space $\mathcal{H}^{\otimes n}$ correspond to possible input strings $x \in [q]^n$, and the i -th multiplier in $\mathcal{H}^{\otimes n}$ corresponds to the i -th variable x_i . Suppose \mathcal{H} corresponds to the i -th variable. The difference matrix Δ_i in the standard basis of \mathcal{H} is thus $\mathbb{J}_q - \mathbb{I}_q$, and

$$\Delta_i \circ \Pi_0 = \Pi_0 - \mathbb{I}_q/q \quad \text{and} \quad \Delta_i \circ \Pi_1 = -\Pi_0 + \mathbb{I}_q/q. \quad (4.1)$$

As described in Section 2.3.2, we can approximate this Δ_i -action as

$$\Delta_i \diamond \Pi_0 := \Pi_0 \quad \text{and} \quad \Delta_i \diamond \Pi_1 := -\Pi_0, \quad (4.2)$$

because $\Delta_i \circ \mathbb{I}_{\mathcal{H}} = 0$. We choose to use these approximations because they have better “orthogonality properties” than (4.1). All approximations used in this chapter are essentially based on (4.2) or the *trivial approximation* $\Delta_i \diamond A := A$, where rows and columns of A correspond to input strings.

Similarly as for \mathcal{H} , an e -basis of $\mathcal{H}^{\otimes n}$ consists of n -fold tensor products of the vectors in $\{e_i\}$. In Sections 4.1 and 4.2 for the ELEMENT DISTINCTNESS, COLLISION, and SET EQUALITY problems, the 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. We use the decomposition $\mathcal{H}^{\otimes n} = \bigoplus_{k=0}^n \mathcal{H}_k^{(n)}$, where

$$\mathcal{H}_k^{(n)} := \bigoplus_{c \in \{0,1\}^n, |c|=k} \mathcal{H}_{c_1} \otimes \dots \otimes \mathcal{H}_{c_n} \quad (4.3)$$

is the space spanned by all the e -basis elements of weight k . (For $m \neq n$, we define the subspace $\mathcal{H}_k^{(m)} \subset \mathcal{H}^{\otimes m}$ the same way.) Note that, for $k \in [0..n]$, $\Pi_k^{(n)}$ is a projector on an eigenspace of the Hamming scheme (see Section 1.5.1, in particular, (1.16)). Let us also define $\Pi_{-1}^{(n)} := 0$ to avoid exception handling.

In Section 4.3 for the CERTIFICATE-SUM and ORTHOGONAL ARRAY problems, we use the decomposition $\mathcal{H}^{\otimes n} = \bigoplus_{S \subseteq [n]} \mathcal{H}_S$, where, given that $s_j = 1$ if $j \in S$ and $s_j = 0$ otherwise, we define $\mathcal{H}_S := \bigotimes_{j \in [n]} \mathcal{H}_{s_j}$. Therefore,

$$\Pi_S = \bigotimes_{j \in [n]} \Pi_{s_j}. \quad (4.4)$$

4.1 Adversary bound for Element Distinctness

Let us start by constructing the adversary bound for the ELEMENT DISTINCTNESS problem when the size of the input alphabet q is at least $\Omega(n^2)$. The adversary matrix that we construct here is

almost exactly the same as the one originally given by Belovs [Bel12b] (Remark 4.1 will explain the difference). Nevertheless, we choose to present it here because it provides some intuition behind adversary bounds for COLLISION and SET EQUALITY (Section 4.2), CERTIFICATE-SUM and ORTHOGONAL ARRAY (Section 4.3), and ELEMENT DISTINCTNESS with small range (Chapter 5).

4.1.1 Construction of the adversary matrix

For every positive input x of ELEMENT DISTINCTNESS, there exists a pair of indices $i \neq j$ such that $x[[i]] = x[[j]]$. For the sake of constructing the adversary matrix, we represent each such pair as follows. Let

$$\mu := ((\mu_{1,1}, \mu_{1,2}), \mu_2, \mu_3, \dots, \mu_{n-1})$$

be a tuple such that

$$\{\mu_{1,1}, \mu_{1,2}, \mu_2, \mu_3, \dots, \mu_{n-1}\} = [n],$$

$\mu_{1,1} < \mu_{1,2}$, and $\mu_i < \mu_{i+1}$ for all $i \in [2..n-2]$. We call μ a *pair*, having in mind $\widehat{\mu} := \{\mu_{1,1}, \mu_{1,2}\}$. The map $\mu \mapsto \widehat{\mu}$ is a bijection between the set of all such pairs μ and the set of size-two subsets of $[n]$. Due to this bijection and by abuse of notation, let N denote both of these sets.

As described in Section 2.3.2, we initially embed the adversary matrix Γ into a larger $|N|q^{n-1} \times q^n$ matrix $\tilde{\Gamma}$. Columns of $\tilde{\Gamma}$ are labeled by all possible inputs in $[q]^n$. The rows of $\tilde{\Gamma}$ are split into blocks corresponding to the pairs in N . Inside a block corresponding to a pair μ , the rows correspond to all possible inputs $x \in [q]^n$ such that $x[[\mu_{1,1}]] = x[[\mu_{1,2}]]$. We label rows by specifying both the input and the block, i.e., like (x, μ) .

Let $\mathcal{N} := \mathbb{C}^N$. Then, $\tilde{\Gamma}$ can be considered as a linear map from $\mathcal{H}^{\otimes n}$ to $\mathcal{N} \otimes \mathcal{H}^{\otimes(n-1)}$ if we identify a standard basis element $(\boldsymbol{\mu}, \boldsymbol{z}) \in \mathcal{N} \otimes \mathcal{H}^{\otimes(n-1)}$ with the row in the μ -block of $\tilde{\Gamma}$ corresponding to positive input x such that $x[[\mu_{1,1}]] = x[[\mu_{1,2}]] = z_1$ and $x[[\mu_i]] = z_i$ for $i \in [2..n-1]$.

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

$$\tilde{\Gamma} := \sum_k \alpha_k W_k, \tag{4.5}$$

where, for each k , W_k is a linear map from $\mathcal{H}_k^{(n)}$ to $\mathcal{N} \otimes \mathcal{H}_k^{(n-1)}$ and we optimize the coefficients α_k later. Recall that the matrix $\tilde{\Gamma}$ can be decomposed into blocks corresponding to different pairs $\mu \in N$. We first define one block of the matrix. Let

$$\Psi_0 := \Pi_0 \otimes e_0^* = e_0^* \otimes \Pi_0 \quad \text{and} \quad \Psi_1 := e_0^* \otimes \Pi_1 + \Pi_1 \otimes e_0^* \tag{4.6}$$

be two linear maps from $\mathcal{H} \otimes \mathcal{H}$ to \mathcal{H} . In the standard basis, we think of the rows of Ψ_0 and Ψ_1 to correspond to (a, a) , where $a \in [q]$, and columns to $(a, b) \in [q]^2$. For example, for $q = 3$, the

two matrices constituting Ψ_1 are

$$e_0^* \otimes \Pi_1 = \begin{matrix} & \begin{matrix} (1,1) & (1,2) & (1,3) & (2,1) & (2,2) & (2,3) & (3,1) & (3,2) & (3,3) \end{matrix} \\ \begin{matrix} (1,1) \\ (2,2) \\ (3,3) \end{matrix} & \begin{pmatrix} 2 & -1 & -1 & 2 & -1 & -1 & 2 & -1 & -1 \\ -1 & 2 & -1 & -1 & 2 & -1 & -1 & 2 & -1 \\ -1 & -1 & 2 & -1 & -1 & 2 & -1 & -1 & 2 \end{pmatrix} \end{matrix} / 3^{3/2}, \quad (4.7)$$

$$\Pi_1 \otimes e_0^* = \begin{matrix} & \begin{matrix} (1,1) & (1,2) & (1,3) & (2,1) & (2,2) & (2,3) & (3,1) & (3,2) & (3,3) \end{matrix} \\ \begin{matrix} (1,1) \\ (2,2) \\ (3,3) \end{matrix} & \begin{pmatrix} 2 & 2 & 2 & -1 & -1 & -1 & -1 & -1 & -1 \\ -1 & -1 & -1 & 2 & 2 & 2 & -1 & -1 & -1 \\ -1 & -1 & -1 & -1 & -1 & -1 & 2 & 2 & 2 \end{pmatrix} \end{matrix} / 3^{3/2}. \quad (4.8)$$

For every $k \in [n]$ and every $\mu \in N$, define the linear map $W_{\mu,k} \in \mathbb{L}(\mathcal{H}^{\otimes n}, \mathcal{H}^{\otimes(n-1)})$ by

$$W_{\mu,k} := \sum_{c \in \{0,1\}^{n-1}, |c|=k} \Psi_{c_1} \otimes \Pi_{c_2} \dots \otimes \Pi_{c_{n-1}} = \Psi_0 \otimes \Pi_k^{(n-2)} + \Psi_1 \otimes \Pi_{k-1}^{(n-2)}, \quad (4.9)$$

where Ψ_{c_1} maps the $\mu_{1,1}$ -th and the $\mu_{1,2}$ -th multiplier in $\mathcal{H}^{\otimes n}$ to the 1-st multiplier in $\mathcal{H}^{\otimes(n-1)}$ and, for $i \in [2..n-1]$, Π_{c_i} maps the μ_i -th multiplier in $\mathcal{H}^{\otimes n}$ to the i -th multiplier in $\mathcal{H}^{\otimes(n-1)}$.

To gain intuition on the matrix $W_{\mu,k}$, it is best to consider $\mu = \{\{1,2\}, 3, 4, \dots, n\}$. Let the rows and columns of $W_{\mu,k}$ be labelled according to the lexicographical order: columns by all strings $y \in [q]^n$ and rows by all strings $x \in [q]^n$ such that $x[[1]] = x[[2]]$. Then, for this μ , the tensor products in (4.9) can be viewed as Kronecker products. The matrices $W_{\mu,k}$ for other μ can be obtained by symmetry.

The matrix $W_{\mu,k}$ is also closely related to the Hamming scheme. In particular, the matrices $\Pi_k^{(n-2)}$ and $\Pi_{k-1}^{(n-2)}$ in (4.9) project on eigenspaces of the Hamming scheme (see (1.16)). And Ψ_0 and Ψ_1 are obtained from $\sqrt{q}\Pi_0^{(2)}$ and $\sqrt{q}\Pi_1^{(2)}$, respectively, by removing all the rows that correspond to pairs $(a, a') \in [q]^2$ such that $a \neq a'$.

We define W_k by specifying each of its blocks:

$$\text{the block of the matrix } W_k \text{ corresponding to } \mu \in N \text{ is defined by } \frac{1}{\sqrt{|N|}} W_{\mu,k} \quad (4.10)$$

(see Figure 4.1). Let us split $W_k = W_{\mathcal{A},k} + W_{\mathcal{B},k}$, where, as in (4.10), we define $W_{\mathcal{A},k}$ and $W_{\mathcal{B},k}$ via blocks

$$W_{\mathcal{A},\mu,k} := \Psi_0 \otimes \Pi_k^{(n-2)} \quad \text{and} \quad W_{\mathcal{B},\mu,k} := \Psi_1 \otimes \Pi_{k-1}^{(n-2)},$$

respectively. In accordance with (4.5), this also splits $\tilde{\Gamma} = \tilde{\Gamma}_{\mathcal{A}} + \tilde{\Gamma}_{\mathcal{B}}$.

Remark 4.1. The matrix $\tilde{\Gamma}$ considered here differs only slightly from the one used by Belovs in [Bel12b]. Here, when constructing it as a linear combination of $W_{\mathcal{A},k}$ and $W_{\mathcal{B},k}$ (see (4.5)), we choose the same coefficient for $W_{\mathcal{A},k}$ and $W_{\mathcal{B},k}$, while Belovs chooses the same coefficient for $W_{\mathcal{A},k}$ and $W_{\mathcal{B},k+1}$. (Notice that $W_{\mathcal{A},\mu,k} + W_{\mathcal{B},\mu,k+1} = (\Psi_0 + \Psi_1) \otimes \Pi_k^{(n-2)}$.)

$$\begin{array}{c}
\boxed{W_k} \\
:= \\
\left. \begin{array}{c}
\overbrace{\left[\begin{array}{c}
\frac{1}{\sqrt{|N|}} W_{\{1,2\},k} \\
\frac{1}{\sqrt{|N|}} W_{\{1,3\},k} \\
\vdots \\
\frac{1}{\sqrt{|N|}} W_{\{n-1,n\},k}
\end{array} \right]}^{q^n} \\
\left. \right\}^{q^{n-1}}
\end{array}
\right.
\end{array}$$

Figure 4.1: The decomposition of W_k into blocks. Here, by abuse of notation, we write $W_{\widehat{\mu},k}$ instead of $W_{\mu,k}$.

Suppose we fix i and we want to study the action of Δ_i on $\tilde{\Gamma}$ and W_k . Let $\tilde{\Gamma}'$ and $\tilde{\Gamma}''$ be the part of $\tilde{\Gamma}$ corresponding to all μ such that $i \in \widehat{\mu}$ and all μ such that $i \notin \widehat{\mu}$, respectively. We use analogous single and double prime notation for $\tilde{\Gamma}_{\mathcal{A}}$, $\tilde{\Gamma}_{\mathcal{B}}$, W_k , $W_{\mathcal{A},k}$, and $W_{\mathcal{B},k}$. We exploit the fact that $\|\Delta_i \circ \tilde{\Gamma}\| = O(1)$ if and only if both $\|\Delta_i \circ \tilde{\Gamma}'\| = O(1)$ and $\|\Delta_i \circ \tilde{\Gamma}''\| = O(1)$. Also, $\|\Delta_i \circ \tilde{\Gamma}''\| = O(1)$ if both $\|\Delta_i \circ \tilde{\Gamma}''_{\mathcal{A}}\| = O(1)$ and $\|\Delta_i \circ \tilde{\Gamma}''_{\mathcal{B}}\| = O(1)$. (See Figure 4.2 for the case when $i = 1$.)

$$\begin{array}{c}
\boxed{\Delta_1 \circ \tilde{\Gamma}} \\
= \\
\left. \begin{array}{c}
\left[\begin{array}{c}
\Delta_1 \circ \tilde{\Gamma}' \\
\hline
\Delta_1 \circ \tilde{\Gamma}''_{\mathcal{A}} \\
+ \\
\Delta_1 \circ \tilde{\Gamma}''_{\mathcal{B}}
\end{array} \right] \\
\left. \right\} \begin{array}{l}
\widehat{\mu} = \{1, 2\}, \{1, 3\}, \dots, \{1, n\} \\
\widehat{\mu} = \{2, 3\}, \{2, 4\}, \dots, \{n-1, n\}
\end{array}
\end{array}
\right.
\end{array}$$

Figure 4.2: The decomposition of $\Delta_1 \circ \tilde{\Gamma}$ for the sake of estimating its norm. The top part of the matrix on the right consists of $n - 1$ blocks, the bottom part to $\binom{n-1}{2}$ blocks.

Due to symmetry, $\|\Delta_i \circ \tilde{\Gamma}'\| = \|\Delta_j \circ \tilde{\Gamma}'\|$ and $\|\Delta_i \circ \tilde{\Gamma}''\| = \|\Delta_j \circ \tilde{\Gamma}''\|$ for all i, j . For notational convenience, let us consider $\|\Delta_1 \circ \tilde{\Gamma}'\|$ and $\|\Delta_n \circ \tilde{\Gamma}''\|$.

4.1.2 Bounding $\|\Delta_1 \circ \tilde{\Gamma}'\|$

For μ such that $1 \in \widehat{\mu}$ (that is, $\mu_{1,1} = 1$), let $W_{\mu,k} = X'_{\mu,k} + Y'_{\mu,k} + Z'_{\mu,k}$, where

$$X'_{\mu,k} := \Psi_0 \otimes \Pi_k^{(n-2)}, \quad Y'_{\mu,k} := (e_0^* \otimes \Pi_1) \otimes \Pi_{k-1}^{(n-2)}, \quad Z'_{\mu,k} := (\Pi_1 \otimes e_0^*) \otimes \Pi_{k-1}^{(n-2)}.$$

We define $X'_k, Y'_k,$ and Z'_k similarly to (4.10), but consisting of $n-1$ blocks like $\tilde{\Gamma}'$. Due to (4.2), we can choose

$$X'_k \xrightarrow{\Delta_1} X'_k, \quad Y'_k \xrightarrow{\Delta_1} Y'_k, \quad Z'_k \xrightarrow{\Delta_1} -X'_{k-1}.$$

(Example (4.8) helps to illustrate that $\Delta_1 \circ (\Pi_1 \otimes e_0^*) = -\Delta_1 \circ \Psi_0$, justifying the last approximation.) By linearity, we therefore have

$$\Delta_1 \circ \tilde{\Gamma}' = \sum_k \alpha_k (X'_k + Y'_k - X'_{k-1}) = \sum_k (\alpha_{k-1} - \alpha_k) X'_{k-1} + \sum_k \alpha_k Y'_k.$$

$(Y'_k)^* Y'_k$ equals $1/|N|$ times the sum of $(Y'_{\mu,k})^* Y'_{\mu,k}$ over all $n-1$ pairs μ satisfying $1 \in \widehat{\mu}$, and it has to be proportional to $\Pi_0 \otimes \Pi_k^{(n-1)}$. Since

$$\text{Tr}(\Pi_0 \otimes \Pi_k^{(n-1)}) = \binom{n-1}{k} (q-1)^k \quad \text{and} \quad \text{Tr}((Y'_{\mu,k})^* Y'_{\mu,k}) = \binom{n-2}{k-1} (q-1)^k,$$

we have

$$(Y'_k)^* Y'_k = \frac{n-1}{|N|} \frac{\binom{n-2}{k-1}}{\binom{n-1}{k}} \Pi_0 \otimes \Pi_k^{(n-1)} = \frac{k}{|N|} \Pi_0 \otimes \Pi_k^{(n-1)},$$

implying $\|Y'_k\| = \Theta(\sqrt{k}/n)$. The same way we show $\|X'_k\| = \Theta(1/\sqrt{n})$.

Since $X'_k \perp Y'_{k'}$ for all k, k' and $X'_k \perp X'_{k'}$ and $Y'_k \perp Y'_{k'}$ whenever $k \neq k'$, the requirement $\|\Delta_1 \circ \tilde{\Gamma}'\| = O(1)$ imposes two conditions on the coefficients α_k :

$$|\alpha_{k-1} - \alpha_k| = O(\sqrt{n}) \quad \text{and} \quad |\alpha_k| = O(n/\sqrt{k}). \quad (4.11)$$

4.1.3 Bounding $\|\Delta_n \circ \tilde{\Gamma}''\|$.

For a label $c \in \{\mathcal{A}, \mathcal{B}\}$ and for μ such that $n \notin \widehat{\mu}$ (that is, $\mu_{n-1} = n$), let

$$W_{c,\mu,k} = X''_{c,\mu,k} + Z''_{c,\mu,k},$$

where

$$\begin{aligned} X''_{\mathcal{A},\mu,k} &:= \Psi_0 \otimes \Pi_k^{(n-3)} \otimes \Pi_0, & Z''_{\mathcal{A},\mu,k} &:= \Psi_0 \otimes \Pi_{k-1}^{(n-3)} \otimes \Pi_1, \\ X''_{\mathcal{B},\mu,k} &:= \Psi_1 \otimes \Pi_{k-1}^{(n-3)} \otimes \Pi_0, & Z''_{\mathcal{B},\mu,k} &:= \Psi_1 \otimes \Pi_{k-2}^{(n-3)} \otimes \Pi_1. \end{aligned}$$

Similarly to (4.10), we define X''_k and Z''_k , each consisting of $\binom{n-1}{2}$ blocks. We can choose

$$X''_{c,k} \xrightarrow{\Delta_n} X''_{c,k} \quad \text{and} \quad Z''_{c,k} \xrightarrow{\Delta_n} -X''_{c,k-1}$$

because of (4.2), therefore we have

$$\Delta_n \diamond \tilde{\Gamma}''_c = \sum_k \alpha_k (X''_{c,k} - X''_{c,k-1}) = \sum_k (\alpha_{k-1} - \alpha_k) X''_{c,k-1}.$$

Note that $X''_{c,k} \perp X''_{c,k'}$ whenever $k \neq k'$ and $X''_{\mathcal{A},k} \perp X''_{\mathcal{B},k'}$ for all k, k' . Since

$$\|X''_{\mathcal{A},\mu,k}\| = \|\Psi_0\| = 1 \quad \text{and} \quad \|X''_{\mathcal{B},\mu,k}\| = \|\Psi_1\| = \sqrt{2}$$

and since in total we have $\binom{n-1}{2} < |N|$ pairs μ to consider, we have $\|X''_{c,k}\| < \sqrt{2}$. The requirement $\|\Delta_n \diamond \tilde{\Gamma}''_c\| = O(1)$ thus imposes the condition

$$|\alpha_{k-1} - \alpha_k| = O(1), \tag{4.12}$$

which is stricter than the left condition in (4.11).

4.1.4 Removal of illegal columns

We obtain Γ by removing all the illegal columns of $\tilde{\Gamma}$. Note that none of the rows are illegal, though multiple rows may correspond to the same positive input, which is fine according to Section 2.3.2. Recall from the same section that $\|\Delta_i \circ \Gamma\| \leq \|\Delta_i \circ \tilde{\Gamma}\|$. Now we need to show that $\|\Gamma\|$ is not much smaller than $\|\tilde{\Gamma}\|$, in particular, not much smaller than α_0 . (The automorphism principle essentially implies that, when constructing the adversary matrix Γ , we can aim for α_0 to yield its principal singular value.)

We have $\Gamma = \sum_k \alpha_k \check{W}_k$, where \check{W}_k is obtained from W_k by removing all the illegal columns. Let us assume that $q = \Omega(n^2)$, so that a constant ratio of columns of $\tilde{\Gamma}$ remains in Γ (this is due to the birthday problem, see [KL07], for example). In particular, since W_0 is the matrix of all entries equal and $\|W_0\| = 1$, we have $\|\check{W}_0\| = \Omega(1)$ and its principal left-singular vector is the all-ones vector $\vec{1}_{|N|q^{n-1}}$. As each column of \check{W}_k is also a column of W_k , we have $\vec{1}_{|N|q^{n-1}}^* \check{W}_k = 0$ whenever $k \neq 0$. Hence, $\|\Gamma\| = \Omega(\alpha_0)$.

We maximize α_0 subject to the conditions (4.11) and (4.12). The right condition in (4.11) imposes $\alpha_k = O(n^{2/3})$ for all $k \geq n^{2/3}$, and, in turn, (4.12) imposes this bound on α_k for all k . Thus, up to a constant scalar, the optimum is obtained by choosing $\alpha_k := \max\{n^{2/3} - k, 0\}$.

Theorem 4.2. *Let W_k be defined via (4.9) and (4.10), let*

$$\tilde{\Gamma} := \sum_{k=0}^{n^{2/3}} (n^{2/3} - k) W_k,$$

and let Γ be obtained from $\tilde{\Gamma}$ by removing all the illegal columns. Given that $q = \Omega(n^2)$, Γ is an adversary matrix for ELEMENT DISTINCTNESS giving an $\Omega(n^{2/3})$ lower bound on the quantum query complexity of the problem.

The certificate complexity barrier (see Section 2.3.4) implies that the positive-weights adversary method cannot yield a lower bound for ELEMENT DISTINCTNESS better than $\Omega(n^{1/2})$. Hence, the adversary matrix Γ given in Theorem 4.2 contains negative weights. Nevertheless, we do not explicitly calculate the weights in Γ , therefore we do not know which of them are negative and which positive.

Aside from removing the illegal columns of $\tilde{\Gamma}$, in his construction, Belovs also removed all rows corresponding to positive inputs x that have more than one collision (i.e., x having multiple pairs $i \neq j$ such that $x[[i]] = x[[j]]$). This way, none of the positive inputs correspond to multiple rows and many positive inputs have no corresponding rows at all, which is fine according to Remark 2.16. In fact, when constructing the adversary bound for ELEMENT DISTINCTNESS with minimal input alphabet in Chapter 5, we will also consider only positive inputs having a unique collision.

4.2 Adversary lower bounds for the Collision and Set Equality problems

In this section, we construct optimal adversary bounds for the COLLISION and SET EQUALITY problems. The proofs for the both problems are almost identical, so we present them in parallel.¹

At the beginning, we proceed to construct adversary matrices for COLLISION and SET EQUALITY in a similar manner as for ELEMENT DISTINCTNESS above. This process, however, fails for specific reasons later mentioned in Section 4.3.3. Therefore, to achieve the desired lower bound, we then modify the adversary matrices using the representation theory of the symmetric group.

4.2.1 Preliminaries

In the following, we 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.

¹An adversary matrix for SET EQUALITY is also, of course, an adversary matrix for COLLISION yielding the same lower bound. However, COLLISION has more symmetry than SET EQUALITY, and we will construct an adversary matrix that respects this symmetry.

Recall that for the COLLISION and SET EQUALITY problems, n denotes half of the input length. Positive inputs of both problems naturally give rise to corresponding matchings. A *matching* μ on $[2n]$ is a decomposition

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

of the set $[2n]$ into n pairwise disjoint pairs of elements. For concreteness, we will usually assume that $\mu_{i,1} < \mu_{i,2}$ for all $i \in [n]$, and $\mu_{1,1} < \mu_{2,1} < \cdots < \mu_{n,1}$. In particular, $\mu_{1,1} = 1$. Clearly, this assumption is without loss of generality. Let N_{CP} denote the set of all matchings on $[2n]$, and let N_{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 .

Our aim is to construct adversary matrices Γ_{CP} and Γ_{SE} for the COLLISION and SET EQUALITY problems, respectively. As for ELEMENT DISTINCTNESS above, we embed the adversary matrix Γ_{Q} into a larger $|N_{\text{Q}}|q^n \times q^{2n}$ matrix $\tilde{\Gamma}_{\text{Q}}$. Columns of $\tilde{\Gamma}_{\text{Q}}$ are labeled by all possible inputs in $[q]^{2n}$. The rows of $\tilde{\Gamma}_{\text{Q}}$ are split into blocks corresponding to the matchings in N_{Q} . Inside a block corresponding to μ , the rows correspond to all possible inputs $x \in [q]^{2n}$ such that $x[\mu_{i,1}] = x[\mu_{i,2}]$ for all i . We label rows by specifying both the input and the block, i.e., like (x, μ) .

Note that now $\tilde{\Gamma}_{\text{Q}}$ contains both illegal rows and illegal columns. A column is illegal if its label $y \in [q]^{2n}$ contains two equal elements. A row labeled by (x, μ) is illegal if $x[\mu_{i,1}] = x[\mu_{j,1}]$ for some $i \neq j$.

Let $\mathcal{N}_{\text{Q}} := \mathbb{C}^{N_{\text{Q}}}$. Then, $\tilde{\Gamma}_{\text{Q}}$ can be considered as a linear map from $\mathcal{H}^{\otimes 2n}$ to $\mathcal{N}_{\text{Q}} \otimes \mathcal{H}^{\otimes n}$ if we identify a basis standard basis element $(\boldsymbol{\mu}, \boldsymbol{z}) \in \mathcal{N}_{\text{Q}} \otimes \mathcal{H}^{\otimes n}$ with the row label (x, μ) of $\tilde{\Gamma}_{\text{Q}}$ where the positive input x is defined by $x[\mu_{i,a}] = z_i$.

Symmetry. Recall that \mathbb{S}_L denotes the symmetric group of a finite set L , and, for $m \in \mathbb{N}$, \mathbb{S}_m denotes the isomorphism class of all symmetric groups \mathbb{S}_L with $|L| = m$. The group $\mathbb{S}_{\text{CP}} := \mathbb{S}_{[2n]}$ and its subgroup $\mathbb{S}_{\text{SE}} := \mathbb{S}_{[1..n]} \times \mathbb{S}_{[n+1..2n]}$ are automorphisms of COLLISION and SET EQUALITY, respectively. Hence, the automorphism principle describes that we may construct Γ_{Q} that are invariant under the permutations of these groups. We extend this symmetry to $\tilde{\Gamma}_{\text{Q}}$ by requiring that, for each $\pi \in \mathbb{S}_{\text{Q}}$, and labels (x, μ) and y , we have

$$\tilde{\Gamma}_{\text{Q}}[(x, \mu), y] = \tilde{\Gamma}_{\text{Q}}[(\pi x, \pi \mu), \pi y], \quad (4.13)$$

where

$$(\pi x)_i = x_{\pi^{-1}(i)} \quad \text{and} \quad \pi \mu = \{\{\pi(\mu_{1,1}), \pi(\mu_{1,2})\}, \dots, \{\pi(\mu_{n,1}), \pi(\mu_{n,2})\}\}. \quad (4.14)$$

Let $V_{\pi} \in \text{U}(\mathcal{H}^{\otimes 2n})$ and $V'_{\text{Q},\pi} \in \text{U}(\mathcal{N}_{\text{Q}} \otimes \mathcal{H}^{\otimes n})$ and be the permutation representations corresponding to the action of \mathbb{S}_{Q} on the column and row labels of $\tilde{\Gamma}_{\text{Q}}$, respectively, defined according to (4.14). Then (4.13) is equivalent to

$$V'_{\text{Q},\pi} \tilde{\Gamma}_{\text{Q}} V_{\pi}^{-1} = \tilde{\Gamma}_{\text{Q}} \quad (4.15)$$

for all $\pi \in \mathbb{S}_Q$.

Because of this symmetry, we may use the representation theory in the construction of $\tilde{\Gamma}_Q$. The COLLISION and SET EQUALITY problems also have symmetries associated with the permutations in $\mathbb{S}_{[q]}$, i.e., the permutations of the symbols of the alphabet. We make use of this symmetry only in Appendices A.2 and A.3, where we prove one of the lemmas used below.

Intended form of the adversary matrix. Similarly to ELEMENT DISTINCTNESS, the adversary matrix $\tilde{\Gamma}_Q$ is constructed as a linear combination

$$\tilde{\Gamma}_Q := \sum_k \alpha_k \bar{W}_{Q,k}, \quad (4.16)$$

where, for each k , $\bar{W}_{Q,k}$ is a linear map from $\mathcal{H}_k^{(2n)}$ to $\mathcal{N}_Q \otimes \mathcal{H}_k^{(n)}$. The coefficients α_k are given by $\alpha_k := \max\{n^{1/3} - k, 0\}$. We again assume that $\bar{W}_{Q,k}$ is fixed under the action of \mathbb{S}_Q (in the sense of (4.13) and (4.15)).

In Section 4.2.2, we define matrices $W_{Q,k}$ very similarly as we defined W_k in (4.5) for ELEMENT DISTINCTNESS. We show that the construction of $\tilde{\Gamma}_Q$ that uses $\bar{W}_{Q,k} := W_{Q,k}$ in (4.16) unfortunately fails to yield non-trivial lower bounds. Fortunately, it is possible to modify the $\bar{W}_{Q,k}$ matrices so that (4.16) gives an optimal adversary matrix. We describe this in Section 4.2.3. Finally, in Section 4.2.4, we show that the removal of illegal rows and columns from $\tilde{\Gamma}_Q$ in order to obtain a valid adversary matrix Γ_Q does not change the value of the adversary bound by more than a constant factor.

4.2.2 Simple yet unsuccessful construction

In this section, we define matrices $W_{Q,k}$ that may seem as the most natural (see Section 4.3.3) choice for the decomposition (4.16). As we show below, they do not work well enough, yet we will use them in Section 4.2.3 to construct matrices $\bar{W}_{Q,k}$ that do work.

Recall that the matrix $\tilde{\Gamma}_Q$ can be decomposed into blocks corresponding to different matchings $\mu \in N_Q$. We first define one block of the matrix. Recall the maps $\Psi_0, \Psi_1 \in \mathcal{L}(\mathcal{H}^{\otimes 2}, \mathcal{H})$ from (4.6). For every $k \in [n]$ and every $\mu \in N_Q$, define the linear map $W_k^\mu \in \mathcal{L}(\mathcal{H}^{\otimes 2n}, \mathcal{H}^{\otimes n})$ by

$$W_k^\mu := \sum_{c \in \{0,1\}^n, |c|=k} \Psi_{c_1} \otimes \dots \otimes \Psi_{c_n}, \quad (4.17)$$

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 $\mathcal{H}^{\otimes n}$. The block of the matrix $W_{Q,k}$ corresponding to $\mu \in N_Q$ is defined by $\frac{1}{\sqrt{|N_Q|}} W_k^\mu$. (Figure 4.1 illustrates an analogous block structure of W_k for ELEMENT DISTINCTNESS.)

Suppose we use $\bar{W}_{Q,k} := W_{Q,k}$ in (4.16). One can see that $W_{Q,k}$ satisfy the symmetry (4.13). 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 define the following decomposition:

$$W_{Q,k} = X_{Q,k} + Y_{Q,k} + Z_{Q,k}, \quad (4.18)$$

where $X_{Q,k}$, $Y_{Q,k}$, and $Z_{Q,k}$ are defined similarly to $W_{Q,k}$ via

$$\begin{aligned} X_k^\mu &= \Psi_0 \otimes \sum_{c \in \{0,1\}^{n-1}, |c|=k} \Psi_{c_2} \otimes \dots \otimes \Psi_{c_n}, \\ Y_k^\mu &= (e_0^* \otimes \Pi_1) \otimes \sum_{c \in \{0,1\}^{n-1}, |c|=k-1} \Psi_{c_2} \otimes \dots \otimes \Psi_{c_n}, \\ Z_k^\mu &= (\Pi_1 \otimes e_0^*) \otimes \sum_{c \in \{0,1\}^{n-1}, |c|=k-1} \Psi_{c_2} \otimes \dots \otimes \Psi_{c_n}. \end{aligned} \quad (4.19)$$

Recall that we always have $\mu(1,1) = 1$. Again, one can see that $X_{Q,k}$, $Y_{Q,k}$, and $Z_{Q,k}$ are symmetric under the action of \mathbb{S}'_Q (in the sense of (4.13) and (4.15)), where $\mathbb{S}'_{\text{CP}} := \mathbb{S}_{[2..2n]}$ and $\mathbb{S}'_{\text{SE}} := \mathbb{S}_{[2..n]} \times \mathbb{S}_{[n+1..2n]}$.

As for ELEMENT DISTINCTNESS above, we can choose $\Delta_1 \diamond X_{Q,k} := X_{Q,k}$, $\Delta_1 \diamond Y_{Q,k} := Y_{Q,k}$, and $\Delta_1 \diamond Z_{Q,k} := -X_{Q,k-1}$. Hence, by linearity,

$$\Delta_1 \diamond \tilde{\Gamma}_Q = \sum_k (\alpha_{k-1} - \alpha_k) X_{Q,k-1} + \sum_k \alpha_k Y_{Q,k}, \quad (4.20)$$

if $\tilde{\Gamma}_Q$ is defined by (4.16) with $\bar{W}_{Q,k} = W_{Q,k}$. However, it is not hard to show that

$$\|W_{Q,k}\| = \Theta(2^{k/2}), \quad \|X_{Q,k}\| = \Theta(2^{k/2}), \quad \text{and} \quad \|Y_{Q,k}\| = \Theta(2^{k/2} \sqrt{k/n}) \quad (4.21)$$

if $k = o(\sqrt{n})$ (see [BR13b]). Thus, this construction only gives a trivial lower bound. It is also possible to show that this problem cannot be fixed by a mere modification of the coefficients α_k .

This failure can be explained by the results in Section 4.3 below: this construction only uses that the non-adaptive learning graph complexity of the COLLISION problem is $\Omega(n^{1/3})$. On the other hand, as we saw in Proposition 3.5, the non-adaptive learning graph complexity of the HIDDEN SHIFT problem is also $\Omega(n^{1/3})$. Thus, if this construction worked for COLLISION, we most likely would also be able to prove an $\Omega(n^{1/3})$ lower bound for HIDDEN SHIFT, that is in contradiction with the fact that the query complexity of this problem is logarithmic [EHK04].

4.2.3 Successful construction

Our aim is to get rid of the $2^{k/2}$ factor in (4.21) while preserving an analogue of (4.20). Recall that the operator $W_{Q,k}$ is symmetric with respect to \mathbb{S}_Q , hence, Schur's lemma (Lemma 1.2) implies that $W_{Q,k}$ can be subdivided into parts corresponding to different irreps of \mathbb{S}_Q . We define the

operator $\bar{W}_{Q,k}$ by taking the part of $W_{Q,k}$ corresponding to Young diagrams with many boxes below the first row.

From (4.3) and (4.17), one can see that $W_k^\mu \in \mathcal{L}(\mathcal{H}_k^{(2n)}, \mathcal{H}_k^{(n)})$, so we have $W_{Q,k}\Pi_k^{(2n)} = W_{Q,k}$. The space $\mathcal{H}_k^{(2n)}$ is stable under all permutations in \mathbb{S}_Q , therefore it can be decomposed into irreps of \mathbb{S}_Q . In Appendix A.2 (see Lemma A.2, in particular) we show that $\mathcal{H}_k^{(2n)}$ contains irreps of \mathbb{S}_{CP} whose corresponding Young diagram has from $2n$ to $2n - k$ boxes in the first row. Similarly, $\mathcal{H}_k^{(2n)}$ contains irreps of \mathbb{S}_{SE} such that the sum of the number of boxes in the first rows of the two Young diagrams defining the irrep is between $2n$ and $2n - k$.

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 $\mathcal{S}^{(m-k,\lambda)}$, where $\lambda \vdash k$. We will also use the subspace $\bar{\bar{\mathcal{H}}}_k^{(m)}$ of $\mathcal{H}_k^{(m)}$ that is spanned by the irreps having exactly $k - 1$ boxes below the first row.

We restrict the operator $W_{Q,k}$ by

$$\bar{W}_{Q,k} := W_{Q,k}\bar{\Pi}_{Q,k}, \quad (4.22)$$

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

$$\bar{\mathcal{H}}_{\text{CP},k} := \bar{\mathcal{H}}_k^{(2n)} \quad \text{or} \quad \bar{\mathcal{H}}_{\text{SE},k} := \sum_{\ell=0}^k \bar{\mathcal{H}}_{k-\ell}^{(n)} \otimes \bar{\mathcal{H}}_\ell^{(n)}. \quad (4.23)$$

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. Note that all entries of $\bar{\Pi}_{Q,k}$ are real in the standard basis (see Claim 2.18). (Essentially, $W_{\text{CP},k}$ corresponds to the columns $0, 1, 2, \dots, k$ in Figure 2.8, and $\bar{W}_{\text{CP},k}$ is its restriction to the column k . Similarly for $W_{\text{SE},k}$ and $\bar{W}_{\text{SE},k}$).

In order to define the action of Δ_1 , we need the following decomposition result. Its proof is rather technical, it uses multiple concepts of the representation theory of the symmetric and the unitary group from Section 1.4, and it is given in Appendix A.2.

Lemma 4.3. *If $k = o(m)$, then*

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

where $\|\Phi_k^{(m)}\| = O(1/\sqrt{m})$ and the support of $\Phi_k^{(m)}$ is contained in $\bar{\bar{\mathcal{H}}}_k^{(m)}$.

With $\Phi_k^{(m)}$ as in Lemma 4.3, let us denote

$$\begin{aligned} \bar{\Pi}'_{\text{CP},k} &:= \bar{\Pi}_k^{(2n-1)}, & \Phi_{\text{CP},k} &:= \Phi_k^{(2n)}, & \bar{\bar{\Pi}}_{\text{CP},k} &:= \bar{\bar{\Pi}}_k^{(2n)}, \\ \bar{\Pi}'_{\text{SE},k} &:= \sum_{\ell=0}^k (\bar{\Pi}_{k-\ell}^{(n-1)} \otimes \bar{\Pi}_\ell^{(n)}), & \Phi_{\text{SE},k} &:= \sum_{\ell=0}^{k-1} (\Phi_{k-\ell}^{(n)} \otimes \bar{\Pi}_\ell^{(n)}), & \bar{\bar{\Pi}}_{\text{SE},k} &:= \sum_{\ell=0}^{k-1} (\bar{\bar{\Pi}}_{k-\ell}^{(n)} \otimes \bar{\Pi}_\ell^{(n)}). \end{aligned} \quad (4.24)$$

Note that $\bar{\Pi}_{Q,k}$ and $\Phi_{Q,k}$ act on $\mathcal{H}^{\otimes 2n}$ while $\bar{\Pi}'_{Q,k}$ acts on $\mathcal{H}^{\otimes (2n-1)}$. Also, $\Phi_{Q,k} = \bar{\Pi}_{Q,k}\bar{\Phi}_{Q,k}$. From Lemma 4.3, we have

$$\bar{\Pi}_{Q,k} = \Pi_0 \otimes \bar{\Pi}'_{Q,k} + \Pi_1 \otimes \bar{\Pi}'_{Q,k-1} + \Phi_{Q,k}. \quad (4.25)$$

With $X_{Q,k}$, $Y_{Q,k}$ and $Z_{Q,k}$ as in Section 4.2.2, 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 (4.18), (4.19) and (4.25), 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

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

The validity of the last action follows from $\Pi_1 \xrightarrow{\Delta_1} -\Pi_0$. Thus, for $\tilde{\Gamma}_Q$ as defined in (4.16), 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 (4.20). The main difference between this construction and the one in Section 4.2.2 is given by the following result.

Lemma 4.4. *In the above notations, we have:*

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

Note the difference with (4.21). We prove Lemma 4.4 in Appendix A.1.

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 (4.16). Indeed, for different k , all the operators $\bar{X}_{Q,k}$ are 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 $\alpha_k = \max\{n^{1/3} - k, 0\}$ satisfy these conditions, giving us

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

4.2.4 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 $\tilde{\Gamma}_Q$. We already have $\|\Delta_i \circ \Gamma\| \leq \|\Delta_i \circ \tilde{\Gamma}\|$, and it is left to show that $\|\Gamma_Q\|$ is not much smaller than $\|\tilde{\Gamma}_Q\|$, in particular, not much smaller than α_0 .

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. Let us assume that $q = \Omega(n^2)$, so that a constant ratio of rows and columns of $\tilde{\Gamma}_Q$ remains in Γ_Q . 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 $\vec{1}_{q!/(q-2n)!}$. All that is left to show is that $\check{W}_{Q,k} \vec{1}_{q!/(q-2n)!} = 0$ whenever $k \neq 0$.

Let us split $\mathcal{H}^{\otimes 2n} = \mathcal{H}_{\text{legal}} \oplus \mathcal{H}_{\text{illegal}}$, where $\mathcal{H}_{\text{legal}}$ and $\mathcal{H}_{\text{illegal}}$ are the spaces spanned by standard basis vectors corresponding to legal and illegal negative inputs, respectively. Let us further decompose

$$\mathcal{H}_{\text{legal}} = \bigoplus_{L \subset [q], |L|=2n} \mathcal{H}_{\text{legal},L},$$

where $\mathcal{H}_{\text{legal},L}$ is the space spanned by standard basis vectors corresponding to negative inputs $x \in [q]^{2n}$ such that entries of x form the set L of size $2n$. It is easy to see that, for all L , $\mathcal{H}_{\text{legal},L}$ is invariant under the action of $\mathbb{S}_{[2n]}$. Since $\bar{\mathcal{H}}_{\text{SE},k} \subseteq \bar{\mathcal{H}}_{\text{CP},k} = \bar{\mathcal{H}}_k^{(2n)}$, it suffices to show $\bar{\Pi}_k^{(2n)} \Pi_{\text{legal},L} \vec{1}_{q^{2n}} = 0$.

Note that the action of $\mathbb{S}_{[2n]}$ on $\mathcal{H}_{\text{legal},L}$ is isomorphic to the regular representation of $\mathbb{S}_{[2n]}$, and we can decompose

$$\mathcal{H}_{\text{legal},L} = \bigoplus_{\sigma \vdash 2n} \mathcal{H}_{\text{legal},L,\sigma},$$

where $\mathcal{H}_{\text{legal},L,\sigma}$ is the subspace corresponding to $\dim \sigma$ copies of the irrep \mathcal{S}^σ in the regular representation. Since we consider the regular representation, we have $\Pi_{\text{legal},L,\sigma} \vec{1}_{q^{2n}} = 0$ whenever $\sigma \neq (2n)$. So it is enough to consider $\Pi_{\text{legal},L,(2n)}$, but from the definition of $\bar{\Pi}_k^{(2n)}$ we have $\bar{\Pi}_k^{(2n)} \Pi_{\text{legal},L,(2n)} = 0$ whenever $k \neq 0$. Hence $\check{W}_{Q,k} \vec{1}_{q!/(q-2n)!} = 0$ whenever $k \neq 0$, and $\|\Gamma_Q\| = \Omega(\alpha_0)$. This gives us the desired lower bound.

Theorem 4.5. *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 (4.22), and let Γ_Q be obtained from $\tilde{\Gamma}_Q$ by removing all the illegal rows and columns. Given that $q = \Omega(n^2)$, 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.

4.3 Adversary bounds for certificate structures

In this section, we consider certificate structures, and we prove that

Theorem 4.6. *For every certificate structure, its quantum query and learning graph complexities differ by at most a constant multiplicative factor.*

Namely, for every certificate structure \mathcal{C} , we construct a Boolean-valued function \mathcal{P} such that \mathcal{P} has \mathcal{C} as a certificate structure and the quantum query complexity of \mathcal{P} is no less than a constant times the learning graph complexity of \mathcal{C} . The opposite direction follows from Claim 3.2 and Theorem 3.3.

Although Theorem 4.6 is a very general result, it is unsatisfactory in the sense that the function \mathcal{P} having the required quantum query complexity is rather artificial, and the size of the alphabet is enormous. However, for boundedly generated certificate structures (see Definition 2.9), it is possible to construct a relatively natural problem with a modestly-sized alphabet having the required quantum query complexity.

In order to define the function with the desired complexity, we first have to introduce the following special case of a well-studied combinatorial object.

Definition 4.7 (Orthogonal Array). Assume T is a subset of $[q]^k$. We say that T is an *orthogonal array* over alphabet $[q]$ if, for every index $i \in [k]$ and for every sequence $z_1, \dots, z_{i-1}, z_{i+1}, \dots, z_k$ of elements in $[q]$, there exist exactly $|T|/q^{k-1}$ choices of $z_i \in [q]$ such that $(z_1, \dots, z_k) \in T$. We call $|T|$ the *size* of the array, and k its *length*.

Compared to a standard definition of orthogonal arrays (cf. [HSS99]), we always require that the so-called *strength* of the array equals $k - 1$.² Recall from Definition 2.9 that a certificate structure \mathcal{C} is boundedly generated if each $M \in \mathcal{C}$ consists exactly of all supersets of some subset $A_M \subset [n]$ of size $O(1)$.

Theorem 4.8. *Assume a certificate structure \mathcal{C} is boundedly generated, and let A_M be like in Definition 2.9. Assume the alphabet is $[q]$ for some $q \geq 2|\mathcal{C}|$, and each A_M is equipped with an orthogonal array T_M over alphabet $[q]$ of length $|A_M|$ and size $q^{|A_M|-1}$. Consider a function $\mathcal{P}: [q]^n \rightarrow \{0, 1\}$ defined by $\mathcal{P}(x) = 1$ iff there exists $M \in \mathcal{C}$ such that $x_{A_M} \in T_M$. Then, the quantum query complexity of f is at least a constant times the learning graph complexity of \mathcal{C} .*

Note that the \mathcal{C} -SUM problem (see Section 2.1.2) is a special case of such a function. So, if $q \geq 2|\mathcal{C}|$, Theorem 4.8 implies that the quantum query complexity of \mathcal{C} -SUM is at least a constant

²For certain type of problems, Špalek constructs adversary lower bounds based on orthogonal arrays with arbitrary strengths [Špa13].

times the learning graph complexity of \mathcal{C} . In particular, Theorem 4.8 together with Propositions 3.4 and 3.5 give $\Omega(n^{k/(k+1)})$ and $\Omega(n^{9/7}/\sqrt{\log n})$ lower bounds on the quantum query complexity of the k -SUM and TRANGLE-SUM problems, respectively, assuming that the alphabet size is sufficiently large.

Theorem 4.8 is a generalization of the lower bound for the k -SUM problem from [BŠ13], and provides additional intuition on the construction, by linking it to learning graphs. Much of the discussion in [BŠ13] applies here as well.

4.3.1 Outline of the lower bound

Let us now outline how Theorems 4.6 and 4.8 are proven. Both theorems are strongly connected: In the second one we prove a stronger statement from stronger premisses. As a consequence, the proofs also have many common elements.

Recall the dual formulation (3.9) of the learning graph complexity of a certificate structure. Given a certificate structure \mathcal{C} , let $\alpha_S(M)$ satisfy (3.9), and be such that (3.9a) equals the learning graph complexity of \mathcal{C} . We define an explicit function $\mathcal{P}: \mathcal{D} \rightarrow \{0, 1\}$ with $\mathcal{D} \subseteq [q]^n$ having \mathcal{C} as a certificate structure and having the objective value (3.9a) of program (3.9) as a lower bound on its quantum query complexity. We prove the latter using the adversary bound.

Function. Let M be a certificate placement in the certificate structure \mathcal{C} . Let $A_M^{(1)}, \dots, A_M^{(\ell(M))}$ be all the inclusion-wise minimal elements of M . (In a boundedly generated certificate structure, M has only one inclusion-wise minimal element A_M .) For each $A_M^{(i)}$, we choose an orthogonal array $T_M^{(i)}$ of length $|A_M^{(i)}|$ over the alphabet $[q]$, and define

$$X_M := \left\{ x \in [q]^n : x_{A_M^{(i)}} \in T_M^{(i)} \text{ for all } i \in [\ell(M)] \right\}. \quad (4.26)$$

The orthogonal arrays are chosen so that X_M is non-empty and satisfies the following *orthogonality property*:

$$\forall S \in 2^{[n]} \setminus M \quad \forall z \in [q]^S : \left| \{x \in X_M : x_S = z\} \right| = |X_M|/q^{|S|}. \quad (4.27)$$

For boundedly generated certificate structures, this property is satisfied automatically.

The set of positive inputs is defined by $\mathcal{D}_1 := \bigcup_{M \in \mathcal{C}} X_M$. The set of negative inputs $Y = \mathcal{D}_0$ is defined by

$$Y := \left\{ x \in [q]^n : x_{A_M^{(i)}} \notin T_M^{(i)} \text{ for all } M \in \mathcal{C} \text{ and } i \in [\ell(M)] \right\}. \quad (4.28)$$

It is easy to see that \mathcal{P} has \mathcal{C} as a certificate structure. We call a function \mathcal{P} defined this way an ORTHOGONAL ARRAY problem. The parameters will be chosen so that $|\mathcal{D}_0| = \Omega(q^n)$. One can see that, if \mathcal{C} is boundedly generated, the function \mathcal{P} is total. (Note: unlike in Sections 4.1 and 4.2, in the current section we use X and Y for denoting sets of inputs.)

Significant matrices. Like for ELEMENT DISTINCTNESS, COLLISION, and SET EQUALITY in Sections 4.1 and 4.2, we initially embed the adversary matrix Γ into a larger matrix $\tilde{\Gamma}$. Columns of $\tilde{\Gamma}$ are labeled by all input strings $y \in [q]^n$, and columns labeled by $y \notin Y$ are illegal. We construct $\tilde{\Gamma}$ so that all its rows correspond to positive inputs, and are thus legal. Γ is obtained from $\tilde{\Gamma}$ by removing the illegal columns. In Sections 4.1 and 4.2, the matrix $\tilde{\Gamma}$ was divided into blocks corresponding to $\mu \in N$. Similarly, here $\tilde{\Gamma}$ is divided into blocks \tilde{G}_M corresponding to $M \in \mathcal{C}$, and the rows of the block \tilde{G}_M are labeled by (x, M) such that $x \in X_M$.

Unlike in Sections 4.1 and 4.2, however, here we obtain $\tilde{\Gamma}$ from an even larger matrix $\hat{\Gamma}$, whose blocks \hat{G}_M are $[q]^n \times [q]^n$ -matrices. Assuming $\mathcal{C} = \{M_1, \dots, M_k\}$,

$$\hat{\Gamma} = \begin{pmatrix} \hat{G}_{M_1} \\ \hat{G}_{M_2} \\ \vdots \\ \hat{G}_{M_k} \end{pmatrix}. \quad (4.29)$$

The block \tilde{G}_M of $\tilde{\Gamma}$ is obtained from the block \hat{G}_M of $\hat{\Gamma}$ by both scaling it up $\sqrt{q^n/|X_M|}$ times and removing all rows corresponding to $x \notin X_M$. Hence, Γ consists of blocks G_M , like in (4.29), where

$$G_M = \sqrt{q^n/|X_M|} \hat{G}_M[[X_M, Y]]$$

(here, the latter notation stands for the submatrix formed by the specified rows and columns).

We construct $\hat{\Gamma}$ so that $\|\hat{\Gamma}\|$ is at least the objective value (3.9a) and, for each $j \in [n]$, there exists $\hat{\Gamma}^\Delta$ such that $\hat{\Gamma} \xrightarrow{\Delta_j} \hat{\Gamma}^\Delta$ and $\|\hat{\Gamma}^\Delta\| \leq 1$. The matrix $\hat{\Gamma}^\Delta$ has a decomposition into blocks \hat{G}_M^Δ similar to (4.29). The matrices Γ^Δ and $\tilde{\Gamma}^\Delta$ are obtained from $\hat{\Gamma}^\Delta$ the same way as Γ and $\tilde{\Gamma}$ from $\hat{\Gamma}$. It is clear that $\hat{\Gamma} \xrightarrow{\Delta_j} \hat{\Gamma}^\Delta$ implies $\Gamma \xrightarrow{\Delta_j} \Gamma^\Delta$.

Let us define $X := \{(x, M) \in [q]^n \times \mathcal{C} : x \in X_M\}$. So Γ is an $X \times Y$ matrix satisfying

$$\Gamma[(x, M), y] = \sqrt{\frac{q^n}{|X_M|}} \hat{\Gamma}[(x, M), y].$$

We show that $\|\Gamma\|$ is not much smaller than $\|\hat{\Gamma}\|$, and we also show that the norm of Γ^Δ is small by showing that $\|\tilde{\Gamma}^\Delta\| = O(\|\hat{\Gamma}^\Delta\|)$. We denote the blocks of $\tilde{\Gamma}^\Delta$ by \tilde{G}_M^Δ , that is,

$$\tilde{G}_M^\Delta = \sqrt{\frac{q^n}{|X_M|}} \hat{G}_M^\Delta[[X_M, [q]^n]]. \quad (4.30)$$

4.3.2 Common parts of the proofs

Let e_0, \dots, e_{q-1} be an e -basis of $\mathcal{H} = \mathbb{C}^{[q]}$, and recall that $\Pi_0 = e_0 e_0^* = \mathbb{J}_q/q$ and $\Pi_1 = \mathbb{I}_q - \mathbb{J}_q/q$. Also, recall from (4.4) that, if $S \subseteq [n]$ and (s_j) is the corresponding characteristic vector, $\Pi_S =$

$\otimes_{j \in [n]} \Pi_{s_j}$. Note that

$$\Pi_S \Pi_{S'} = 0 \quad \text{whenever } S \neq S'. \quad (4.31)$$

We define the matrices \widehat{G}_M from (4.29) by

$$\widehat{G}_M = \sum_{S \subseteq [n]} \alpha_S(M) \Pi_S, \quad (4.32)$$

where $\alpha_S(M)$ give an optimal solution to (3.9).

Lemma 4.9. *If $\widehat{\Gamma}$ and Γ are defined as in Section 4.3.1, all X_M satisfy the orthogonality property (4.27), and $|Y| = \Omega(q^n)$, then*

$$\|\Gamma\| = \Omega\left(\sqrt{\sum_{M \in \mathcal{C}} \alpha_\emptyset(M)^2}\right). \quad (4.33)$$

Proof. Recall that $G_M = \sqrt{q^n/|X_M|} \widehat{G}_M \llbracket X_M, Y \rrbracket$, hence, from (4.32), we get that

$$G_M = \sqrt{\frac{q^n}{|X_M|}} \alpha_\emptyset(M) \Pi_\emptyset^{\otimes n} \llbracket X_M, Y \rrbracket + \sqrt{\frac{q^n}{|X_M|}} \sum_{S \neq \emptyset} \alpha_S(M) \Pi_S \llbracket X_M, Y \rrbracket.$$

Let us calculate the sum $s(G_M) := \vec{1}_{|X_M|}^* G_M \vec{1}_{|Y|}$ of the entries of G_M . In the first term, each entry of $\Pi_\emptyset^{\otimes n}$ equals q^{-n} . There are $|X_M|$ rows and $|Y|$ columns in the matrix, hence, the sum of the entries of the first term is $\sqrt{|X_M|/q^n} |Y| \alpha_\emptyset(M)$.

In the second term, $s(\alpha_S(M) \Pi_S \llbracket X_M, Y \rrbracket) = 0$ for all $S \neq \emptyset$. Indeed, if $S \in M$, then $\alpha_S(M) = 0$ by (3.9c). Otherwise,

$$\begin{aligned} s(\Pi_S \llbracket X_M, Y \rrbracket) &= \sum_{y \in Y} \sum_{x \in X_M} \Pi_S \llbracket x, y \rrbracket = q^{|S|-n} \sum_{y \in Y} \sum_{x \in X_M} \Pi_1^{\otimes |S|} \llbracket x_S, y_S \rrbracket \\ &= \frac{|X_M|}{q^n} \sum_{y \in Y} \sum_{z \in [q]^S} \Pi_1^{\otimes |S|} \llbracket z, y_S \rrbracket = 0. \end{aligned}$$

(For the third equality, the orthogonality condition (4.27) is used. For the last one, we use that the sum of the entries of every column of $\Pi_1^{\otimes k}$ is zero if $k > 0$.) Summing up,

$$s(G_M) = \sqrt{\frac{|X_M|}{q^n}} |Y| \alpha_\emptyset(M).$$

We are now ready to estimate $\|\Gamma\|$. Define two unit vectors $u \in \mathbb{R}^X$ and $v \in \mathbb{R}^Y$ by

$$u \llbracket (x, M) \rrbracket := \frac{\alpha_\emptyset(M)}{\sqrt{|X_M| \sum_{M \in \mathcal{C}} \alpha_\emptyset(M)^2}} \quad \text{and} \quad v \llbracket y \rrbracket := \frac{1}{\sqrt{|Y|}}$$

for all $(x, M) \in X$ and $y \in Y$. Then,

$$\|\Gamma\| \geq u^* \Gamma v = \frac{\sum_{M \in \mathcal{C}} \alpha_\emptyset(M) s(G_M)}{\sqrt{|X_M| |Y| \sum_{M \in \mathcal{C}} \alpha_\emptyset(M)^2}} = \sqrt{\frac{|Y|}{q^n} \sum_{M \in \mathcal{C}} \alpha_\emptyset(M)^2} = \Omega\left(\sqrt{\sum_{M \in \mathcal{C}} \alpha_\emptyset(M)^2}\right).$$

□

Let us define the transformation $\widehat{\Gamma} \xrightarrow{\Delta_j} \widehat{\Gamma}^\Delta$ and state some of the properties of $\widehat{\Gamma}^\Delta$ that will be used in the subsequent sections. By using the approximation (4.2), we choose $\Pi_S \xrightarrow{\Delta_j} \Pi_S$ if $j \notin S$ and $\Pi_S \xrightarrow{\Delta_j} -\Pi_{S \setminus \{j\}}$ if $j \in S$. We extend this approximation to $\widehat{G}_M \xrightarrow{\Delta_j} \widehat{G}_M^\Delta$ by linearity: from (4.32), we see that

$$\widehat{G}_M^\Delta = \sum_{S \subseteq [n]} \beta_S(M) \Pi_S, \quad (4.34)$$

where $\beta_S(M) = \alpha_S(M) - \alpha_{S \cup \{j\}}(M)$. In particular, $\beta_S(M) = 0$ if $j \in S$ or $S \in M$. The matrix $\widehat{\Gamma}^\Delta$ is of the form (4.29), but with each \widehat{G}_M replaced by \widehat{G}_M^Δ . Thus,

$$(\widehat{\Gamma}^\Delta)^* \widehat{\Gamma}^\Delta = \sum_{M \in \mathcal{C}} (\widehat{G}_M^\Delta)^* \widehat{G}_M^\Delta = \sum_{S \in 2^{[n]}} \left(\sum_{M \in \mathcal{C}} \beta_S(M)^2 \right) \Pi_S. \quad (4.35)$$

In particular, we obtain from (3.9b) that $\|\widehat{\Gamma}^\Delta\| \leq 1$.

4.3.3 Comparison to adversary constructions of Sections 4.1 and 4.2

Let us see how the adversary matrices defined here compare to ones defined earlier in Sections 4.1 and 4.2. First, suppose \mathcal{C} is the 2-subset certificate structure. Each $M \in \mathcal{C}$ is specified by $A_M^{(1)} = \{j, j'\}$, and let $T_M^{(1)}$ be given by $x[[j]] = x[[j']]$. This certificate structure together with these orthogonal arrays correspond to the ELEMENT DISTINCTNESS problem. Just like in the proof of Proposition 3.4, let

$$\alpha_S(M) := \binom{n}{2}^{-1/2} \max\{n^{2/3} - |S|, 0\}$$

if $\{j, j'\} \not\subseteq S$, and $\alpha_S(M) := 0$ otherwise. Then, the adversary matrix Γ that we obtain according to definitions in this section is exactly the same as Γ given in Section 4.1.

Now suppose \mathcal{C} is the collision, the set equality, or the hidden shift certificate structure (see Definition 2.6). Generalizing the proof of Proposition 3.5, we define

$$\alpha_S(M) := \max\{n^{1/3} - |S|, 0\} / \sqrt{|\mathcal{C}|}$$

if $S \notin M$, and $\alpha_S(M) := 0$ if $S \in M$. One can see that this choice of $\alpha_S(M)$ reproves Proposition 3.5: the learning graph complexity of \mathcal{C} is $\Omega(n^{1/3})$. Each $M \in \mathcal{C}$ corresponds to some matching μ on $[2n]$, and the inclusion-wise minimal elements of M are $A_M^{(1)} = \{\mu_{1,1}, \mu_{1,2}\}$, $A_M^{(2)} = \{\mu_{2,1}, \mu_{2,2}\}$, \dots , $A_M^{(n)} = \{\mu_{n,1}, \mu_{n,2}\}$. For each $A_M^{(i)}$, let the orthogonal array $T_M^{(i)}$ be given by $x[\mu_{i,1}] = x[\mu_{i,2}]$. For this collection of orthogonal arrays, the orthogonality property (4.27) holds. For the collision and set equality certificate structures, the matrices $\tilde{\Gamma}$ that we obtain according to definitions of this section are exactly the same as, respectively, the matrices $\tilde{\Gamma}_{\text{CP}}$ and $\tilde{\Gamma}_{\text{SE}}$ of Section 4.2.2. (That is, the construction of $\tilde{\Gamma}_{\text{Q}}$ where one uses $\tilde{W}_{\text{Q},k} := W_{\text{Q},k}$ in (4.16).) The failure of these matrices to yield non-trivial lower bounds for COLLISION and SET EQUALITY suggests that a straightforward generalization of Theorem 4.8 for all certificate structures is not true.

4.3.4 Boundedly generated certificate structures

In this section, we finish the proof of Theorem 4.8. In the settings of the theorem, the orthogonal arrays $T_M^{(i)}$ in (4.26) are already specified. Since each $M \in \mathcal{C}$ has only one inclusion-wise minimal element A_M , we drop all upper indices (i) in this section.

From the statement of the theorem, we have $|X_M| = q^{n-1}$, and, in particular, they are non-empty. Also, X_M satisfies the orthogonality property (4.27), and, by (4.28), we have

$$|Y| = \left| [q]^n \setminus \bigcup_{M \in \mathcal{C}} X_M \right| \geq q^n - \sum_{M \in \mathcal{C}} |X_M| = q^n - |\mathcal{C}|q^{n-1} \geq \frac{q^n}{2}. \quad (4.36)$$

Thus, the conditions of Lemma 4.9 are satisfied, and (4.33) holds.

As Γ^Δ is a submatrix of $\tilde{\Gamma}^\Delta$, it suffices to estimate $\|\tilde{\Gamma}^\Delta\|$. Let $k := \max_{M \in \mathcal{C}} |A_M|$. By the definition of boundedly generated certificate structures (Definition 2.9), $k = O(1)$.

Fix some order of elements in each $A_M = \{a_{M,1}, \dots, a_{M,|A_M|}\}$, and let $L_{M,i}$, where $M \in \mathcal{C}$ and $i \in [k]$, be subsets of $2^{[n]}$ satisfying the following properties:

- for each M , the set $2^{[n]} \setminus M$ is the disjoint union $L_{M,1} \sqcup \dots \sqcup L_{M,k}$;
- for each M and each $i \leq |A_M|$, all elements of $L_{M,i}$ omit $a_{M,i}$;
- for each M and each i such that $|A_M| < i \leq k$, the set $L_{M,i}$ is empty.

Recall that, if $S \subseteq [n]$ and (s_j) is the corresponding characteristic vector, $\Pi_S = \bigotimes_{j \in [n]} \Pi_{s_j}$. The main idea behind defining $L_{M,i}$'s is as follows.

Claim 4.10. *If $S, S' \in L_{M,i}$, then*

$$(\Pi_S \llbracket X_M, [q]^n \rrbracket)^* (\Pi_{S'} \llbracket X_M, [q]^n \rrbracket) = \begin{cases} \Pi_S/q, & S = S'; \\ 0, & \text{otherwise.} \end{cases}$$

Proof. If we remove the $a_{M,i}$ -th entry in all inputs in X_M , we obtain $[q]^{n-1}$ by the definition of an orthogonal array. All elements of $L_{M,i}$ omit $a_{M,i}$, hence, Π_S has Π_0 in the $a_{M,i}$ -th position for all $S \in L_{M,i}$. Thus, the $a_{M,i}$ -th entries of x and y have no impact on the value of $\Pi_S \llbracket x, y \rrbracket$.

Consider a column of Π_S ; it is of the form $\psi \otimes e_0$, where e_0 is on the $a_{M,i}$ -th element of $[q]^n$ and q^{n-1} -dimensional vector ψ is on the others. For every $z \in [q]^{[n] \setminus \{a_{M,i}\}}$, there is a unique $x \in [q]^n$ such that $x_{[n] \setminus \{a_{M,i}\}} = z$ and $x \in X_M$. Therefore, $(\psi \otimes e_0) \llbracket X_M \rrbracket = \psi / \sqrt{q}$. Let (s_j) be the characteristic vector of S . Then,

$$\Pi_S \llbracket X_M, [q]^n \rrbracket = \left(\bigotimes_{j \in [n] \setminus \{a_{M,i}\}} \Pi_{s_j} \right) \otimes \frac{e_0^*}{\sqrt{q}}.$$

Similarly for S' , and the claim follows from (4.31). \square

For each M , decompose \widehat{G}_M^Δ from (4.34) into $\sum_{i \in [k]} \widehat{G}_{M,i}^\Delta$, where

$$\widehat{G}_{M,i}^\Delta := \sum_{S \in L_{M,i}} \beta_S(M) \Pi_S.$$

Define similarly to Section 4.3.1,

$$\widetilde{G}_{M,i}^\Delta := \sqrt{\frac{q^n}{|X_M|}} \widehat{G}_{M,i}^\Delta \llbracket X_M, [q]^n \rrbracket = \sqrt{q} \sum_{S \in L_{M,i}} \beta_S(M) \Pi_S \llbracket X_M, [q]^n \rrbracket,$$

and let $\widetilde{\Gamma}_i^\Delta$ be the matrix consisting of $\widetilde{G}_{M,i}^\Delta$, for all $M \in \mathcal{C}$, stacked one on another like in (4.29). Then, $\widetilde{\Gamma}^\Delta = \sum_{i \in [k]} \widetilde{\Gamma}_i^\Delta$. We have

$$(\widetilde{\Gamma}_i^\Delta)^* \widetilde{\Gamma}_i^\Delta = \sum_{M \in \mathcal{C}} (\widetilde{G}_{M,i}^\Delta)^* \widetilde{G}_{M,i}^\Delta = \sum_{M \in \mathcal{C}} \sum_{S \in L_{M,i}} \beta_S(M)^2 \Pi_S,$$

by Claim 4.10. Similarly to (4.35), we get $\|\widetilde{\Gamma}_i^\Delta\| \leq 1$. By the triangle inequality, $\|\widetilde{\Gamma}^\Delta\| \leq k$, hence, $\|\Gamma^\Delta\| \leq k = O(1)$. Combining this with (4.33), and using the adversary bound (Theorem 2.15), we obtain the necessary lower bound. This finishes the proof of Theorem 4.8.

4.3.5 General certificate structures

In this section, we finish the proof of Theorem 4.6. There are two main reasons why it is likely not possible to prove a general result like Theorem 4.8 for arbitrary certificate structures.

First of all, the proof in Section 4.3.4 cannot be applied here, because k in the decomposition of \widehat{G}_M^Δ into $\sum_{i \in [k]} \widehat{G}_{M,i}^\Delta$ would not be bounded by a constant. This too indicates why the adversary constructions of Section 4.2.2 for the COLLISION and SET EQUALITY problems fail.

Next, the orthogonality property (4.27) is not satisfied automatically for general certificate structures. For instance, assume $A_M^{(1)} = \{1, 2\}$, $A_M^{(2)} = \{2, 3\}$, and the orthogonal arrays are given by the conditions $x_1 = x_2$ and $x_2 = x_3$, respectively. Then, for any input x satisfying both conditions, we have $x_1 = x_3$, and the orthogonality condition fails for $S = \{1, 3\}$.

The problem in the last example is that the orthogonal arrays are not independent because $A_M^{(1)}$ and $A_M^{(2)}$ intersect. We cannot avoid that $A_M^{(i)}$ s intersect, but we still can have $T_M^{(i)}$ s independent by defining them on independent parts of the input alphabet.

Fourier basis. At the beginning of this chapter, we defined an e -basis as an arbitrary orthonormal basis satisfying the requirement that e_0 has all its entries equal to $1/\sqrt{q}$. In this section, however, we consider a concrete choice for e_i . Its construction is based on the Fourier basis.

Let p be a positive integer, and \mathbb{Z}_p be the cyclic group of order p , formed by the integers modulo p . Consider the complex vector space $\mathbb{C}^{\mathbb{Z}_p}$. The vectors $(\chi_a)_{a \in \mathbb{Z}_p}$, defined by $\chi_a[b] := e^{2\pi i ab/p} / \sqrt{p}$, form its orthonormal basis, called the *Fourier basis* of $\mathbb{C}^{\mathbb{Z}_p}$. Note that the value of $\chi_a[b]$ is well-defined because $e^{2\pi i} = 1$.

If $\mathfrak{U} \subseteq \mathbb{Z}_p$, then the *Fourier bias* [TV06] of \mathfrak{U} is defined by

$$\|\mathfrak{U}\|_{\mathfrak{u}} := \frac{1}{p} \left| \max_{a \in \mathbb{Z}_p \setminus \{0\}} \sum_{u \in \mathfrak{U}} e^{2\pi i au/p} \right|. \quad (4.37)$$

It is a real number between 0 and $|\mathfrak{U}|/p$. We need the following result stating the existence of sets with small Fourier bias and arbitrary density.

Theorem 4.11. *For any real $0 < \delta < 1$, it is possible to construct $\mathfrak{U} \subseteq \mathbb{Z}_p$ such that $|\mathfrak{U}| \approx \delta p$ (e.g., $|\mathfrak{U}| = \lceil \delta p \rceil$), $\|\mathfrak{U}\|_{\mathfrak{u}} = O(\text{polylog}(p)/\sqrt{p})$, and p is arbitrary large. In particular, $\|\mathfrak{U}\|_{\mathfrak{u}} = o(1)$.*

For instance, one may prove that a random subset satisfies these properties with high probability [TV06, Lemma 4.16]. There also exist explicit constructions [Gil10].

Input alphabet and orthogonal arrays. Let $\ell := \max_{M \in \mathcal{C}} \ell(M)$, where $\ell(M)$ is defined in Section 4.3.1 as the number of inclusion-wise minimal elements of M . We define the input alphabet as $\Sigma := \mathbb{Z}_p^\ell$ for some p to be defined later (note: the size of the alphabet is $q = p^\ell$). Hence, every input string $x \in \Sigma^n$ can be expressed in the form

$$x = \begin{pmatrix} x_1^{(1)} & \cdots & x_n^{(1)} \\ \vdots & \ddots & \vdots \\ x_1^{(\ell)} & \cdots & x_n^{(\ell)} \end{pmatrix}, \quad (4.38)$$

where $x_j = (x_j^{(1)}, \dots, x_j^{(\ell)}) \in \Sigma$ is j -th entry of x and we call $x^{(i)} := (x_1^{(i)}, \dots, x_n^{(i)}) \in \mathbb{Z}_p^n$ the i -th component of x .

Let $Q_M^{(i)}$ be an orthogonal array of length $|A_M^{(i)}|$ over the alphabet \mathbb{Z}_p . We will specify a concrete choice in a moment. From $Q_M^{(i)}$, we define $T_M^{(i)}$ in (4.26) by requiring that the i -th component of $z \in \Sigma^{A_M^{(i)}}$ satisfy $Q_M^{(i)}$ (other components can be arbitrary). The sets X_M are defined as in (4.26). We additionally define

$$X_M^{(i)} := \{x^{(i)} \in \mathbb{Z}_p^n : x_{A_M^{(i)}}^{(i)} \in Q_M^{(i)}\},$$

for $i \leq \ell(M)$, and $X_M^{(i)} = \mathbb{Z}_p^n$ otherwise. Note that $X_M = \prod_{i=1}^{\ell} X_M^{(i)}$ in the sense that, for each sequence $x^{(i)} \in X_M^{(i)}$ with $i = 1, \dots, \ell$, there is a corresponding element $x \in X_M$ with $x_j = (x_j^{(1)}, \dots, x_j^{(\ell)})$.

Now we make our choice for $Q_M^{(i)}$. Let $\delta := 1/(2\ell|\mathcal{C}|)$ and let $\mathfrak{U} \subseteq \mathbb{Z}_p$ be a set with small Fourier bias and size $|\mathfrak{U}| \approx \delta p$, which exists due to Theorem 4.11. We define $Q_M^{(i)}$ as consisting of all $x \in \mathbb{Z}_p^{A_M^{(i)}}$ such that the sum of the elements of x belongs to \mathfrak{U} . With this definition,

$$|X_M^{(i)}| = \delta p^n. \quad (4.39)$$

Hence, there are exactly δq^n elements $x \in \Sigma^n$ such that $x_{A_M^{(i)}}^{(i)} \in T_M^{(i)}$. Since $\delta = 1/(2\ell|\mathcal{C}|)$, a calculation similar to (4.36) shows that $|Y| \geq q^n/2$. Also, by considering each $i \in [\ell]$ independently, it is easy to see that all X_M satisfy the orthogonality condition (4.27). Thus, Lemma 4.9 applies, and (4.33) holds.

Now it remains to estimate $\|\Gamma^\Delta\|$, and it is done by considering the matrix $\tilde{\Gamma}^\Delta$ as described in Section 4.3.1, and performed once in Section 4.3.4. If $\hat{\Gamma}^\Delta = 0$, then also $\Gamma^\Delta = 0$, and we are done. Thus, we further assume $\hat{\Gamma}^\Delta \neq 0$. Recall that $(\chi_a)_{a \in \mathbb{Z}_p}$ denotes the Fourier basis of $\mathbb{C}^{\mathbb{Z}_p}$. The e -basis that we consider is defined as the Fourier basis of \mathbb{C}^Σ . It consists of the elements of the form $e_a = \bigotimes_{i=1}^{\ell} \chi_{a^{(i)}}$ where $a = (a^{(i)}) \in \Sigma$. Note that e_0 has the required value, where 0 is interpreted as the identity element 0^ℓ of the additive group Σ .

Given $v = (v_j^{(i)}) \in \Sigma^n$, we define $v_j \in \Sigma$ and $v^{(i)} \in \mathbb{Z}_p^n$ as in (4.38). Let $e_v := \bigotimes_{j=1}^n e_{v_j}$, which form an e -basis of $\mathcal{H} = \mathbb{C}^{\Sigma^n}$. Components of these vectors, for $w = (w_j) \in \mathbb{Z}_p^n$, are defined as $\chi_w := \bigotimes_{j=1}^n \chi_{w_j}$.

Fix an arbitrary $M \in \mathcal{C}$. Let $\widehat{B}_M = (\widehat{G}_M^\Delta)^* \widehat{G}_M^\Delta$ and $\widetilde{B}_M = (\widetilde{G}_M^\Delta)^* \widetilde{G}_M^\Delta$. We aim to show that

$$\|\widehat{B}_M - \widetilde{B}_M\| \rightarrow 0 \quad \text{as } p \rightarrow \infty, \quad (4.40)$$

because this implies

$$\|(\widehat{\Gamma}^\Delta)^* \widehat{\Gamma}^\Delta - (\widetilde{\Gamma}^\Delta)^* \widetilde{\Gamma}^\Delta\| = \left\| \sum_{M \in \mathcal{C}} (\widehat{B}_M - \widetilde{B}_M) \right\| \leq \sum_{M \in \mathcal{C}} \|\widehat{B}_M - \widetilde{B}_M\| \rightarrow 0$$

as $p \rightarrow \infty$. As $\|\widehat{\Gamma}^\Delta\| > 0$, this implies that $\|\Gamma^\Delta\| \leq 2\|\widehat{\Gamma}^\Delta\|$ for p large enough, and together with (4.33) and the adversary bound (Theorem 2.15), this implies Theorem 4.6.

Comparison of \widehat{B}_M and \widetilde{B}_M . From (4.34), we conclude that the eigenbasis of \widehat{B}_M consists of the vectors e_v , with $v \in \Sigma^n$, defined above. Hence \widehat{B}_M is diagonal in the e -basis. We prove (4.40) by showing that, in the e -basis,

- \widehat{B}_M and \widetilde{B}_M have the same diagonal entries,
- \widetilde{B}_M is block diagonal with each block having size independent from p ,
- off-diagonal entries of \widetilde{B}_M goes to 0 as $p \rightarrow \infty$.

In order to understand \widetilde{B}_M better, we have to understand how $e_v[X_M]$ behave. We have

$$(e_v[X_M])^* (e_{v'}[X_M]) = \prod_{i=1}^{\ell} (\chi_{v^{(i)}}[X_M^{(i)}])^* (\chi_{v'^{(i)}}[X_M^{(i)}]). \quad (4.41)$$

Hence, it suffices to understand the behaviour of $\chi_w[X_M^{(i)}]$. For $w \in \mathbb{Z}_p^n$, $A \subseteq [n]$, and $c \in \mathbb{Z}_p$, we write $w + cA$ for the sequence $w' \in \mathbb{Z}_p^n$ defined by

$$w'_j := \begin{cases} w_j + c, & j \in A; \\ w_j, & \text{otherwise.} \end{cases}$$

In this case, we say that w and w' are obtained from each other by a *shift on A* .

Claim 4.12. *Assume that w and w' are elements of \mathbb{Z}_p^n , and let $\xi := (\chi_w[X_M^{(i)}])^* (\chi_{w'}[X_M^{(i)}])$. If $w = w'$, then $\xi = \delta$. If $w \neq w'$, but w can be obtained from w' by a shift on $A_M^{(i)}$, then $|\xi| \leq \|\mathfrak{U}\|_{\mathfrak{u}}$. Finally, if w cannot be obtained from w' by a shift on $A_M^{(i)}$, then $\xi = 0$.*

Proof. Arbitrarily enumerate the elements of $\mathfrak{U} = \{u_1, \dots, u_m\}$ where $m = \delta p$. Denote, for the sake of brevity, $A = A_M^{(i)}$. Consider the decomposition $X_M^{(i)} = \bigsqcup_{k=1}^m X_k$, where

$$X_k := \left\{ w \in \mathbb{Z}_p^n : \sum_{j \in A} w_j = u_k \right\}.$$

Fix an arbitrary element $a \in A$ and denote $\bar{w} := w - w_a A$ and $\bar{w}' := w' - w'_a A$. In both of them, $\bar{w}_a = \bar{w}'_a = 0$, and by an argument similar to Claim 4.10, we get that

$$(\chi_{\bar{w}} \llbracket X_k \rrbracket)^* (\chi_{\bar{w}'} \llbracket X_k \rrbracket) = \begin{cases} 1/p, & \bar{w} = \bar{w}'; \\ 0, & \text{otherwise.} \end{cases} \quad (4.42)$$

If $x^{(i)} \in X_k$, then

$$\begin{aligned} \chi_w \llbracket x^{(i)} \rrbracket &= \prod_{j=1}^n \chi_{w_j} \llbracket x_j^{(i)} \rrbracket = \frac{1}{\sqrt{p^n}} \exp \left[\frac{2\pi i}{p} \sum_{j=1}^n w_j x_j^{(i)} \right] \\ &= \frac{1}{\sqrt{p^n}} \exp \left[\frac{2\pi i}{p} \left(\sum_{j=1}^n \bar{w}_j x_j^{(i)} + w_a \sum_{j \in A} x_j^{(i)} \right) \right] = \exp \left(\frac{2\pi i}{p} w_a u_k \right) \chi_{\bar{w}} \llbracket x^{(i)} \rrbracket. \end{aligned}$$

Hence,

$$(\chi_w \llbracket X_M^{(i)} \rrbracket)^* (\chi_{w'} \llbracket X_M^{(i)} \rrbracket) = \sum_{k=1}^m (\chi_w \llbracket X_k \rrbracket)^* (\chi_{w'} \llbracket X_k \rrbracket) = \sum_{k=1}^m e^{2\pi i (w'_a - w_a) u_k / p} (\chi_{\bar{w}} \llbracket X_k \rrbracket)^* (\chi_{\bar{w}'} \llbracket X_k \rrbracket). \quad (4.43)$$

If w' cannot be obtained from w by a shift on A , then $\bar{w} \neq \bar{w}'$ and (4.43) equals zero by (4.42). If $w = w'$, then (4.43) equals $m/p = \delta$. Finally, if w' can be obtained from w by a shift on A but $w \neq w'$, then $\bar{w} = \bar{w}'$ and $w_a \neq w'_a$. By (4.42) and (4.37), we get that (4.43) does not exceed $\|\mathfrak{U}\|_{\mathfrak{u}}$ in absolute value. \square

Let $v \in \Sigma^n$, and $S := \{j \in [n] : v_j \neq 0\}$. Let $v' \in \Sigma^n$, and define S' similarly. We have

$$\begin{aligned} e_v^* \tilde{B}_M e_{v'} &= \frac{q^n \beta_S(M) \beta_{S'}(M)}{|X_M|} (e_v \llbracket X_M \rrbracket)^* (e_{v'} \llbracket X_M \rrbracket) \\ &= \frac{\beta_S(M) \beta_{S'}(M)}{\delta^\ell} \prod_{i=1}^{\ell} (\chi_{v^{(i)}} \llbracket X_M^{(i)} \rrbracket)^* (\chi_{v'^{(i)}} \llbracket X_M^{(i)} \rrbracket), \end{aligned} \quad (4.44)$$

where the first equality is by (4.30) and (4.34) and the second by (4.39) and (4.41). By this and Claim 4.12, we have that

$$e_v^* \tilde{B}_M e_{v'} = \beta_S(M)^2 = e_v^* \widehat{B}_M e_{v'}. \quad (4.45)$$

Call v and v' *equivalent*, if $\beta_S(M)$ and $\beta_{S'}(M)$ are both non-zero and, for each $i \in [\ell]$, $v^{(i)}$ can be obtained from $v'^{(i)}$ by a shift on $A_M^{(i)}$. By (4.44) and Claim 4.12, we have that $e_v^* \tilde{B}_M e_{v'}$ is non-zero only if v and v' are equivalent.

Note that, if $v_j^{(i)} \neq 0$ for all $j \in A_M^{(i)}$, then S' is such that $A_M^{(i)} \subseteq S'$, therefore $S' \in M$ and $\beta_{S'}(M) = 0$ by (3.9c). For each $i \in [\ell]$, there are at most $|A_M^{(i)}| \leq n$ shifts of $v^{(i)}$ on $A_M^{(i)}$ that have an element with an index in $A_M^{(i)}$ equal to 0. Hence, for each $v \in \Sigma^n$, there are at most n^ℓ elements of Σ^n equivalent to it.

Thus, in the e -basis, the matrix \tilde{B}_M has the required properties. Namely, by (4.45), its diagonal entries equal the diagonal entries of \widehat{B}_M . Next, \tilde{B}_M is block-diagonal with the blocks of size at most n^ℓ . By (4.44) and Claim 4.12, the off-diagonal elements satisfy

$$|e_v^* \tilde{B}_M e_{v'}| \leq \frac{\|\mathfrak{U}\|_{\mathfrak{u}}}{\delta} |\beta_S(M) \beta_{S'}(M)|,$$

because $\|\mathfrak{U}\|_{\mathfrak{u}} \leq \delta$. Since the values of $\beta_S(M)$ do not depend on p , and by Theorem 4.11, the off-diagonal elements of \tilde{B}_M tend to zero as p tends to infinity. Since the sizes of the blocks also do not depend on p , the norm of $\widehat{B}_M - \tilde{B}_M$ also tends to 0, as required in (4.40). This finishes the proof of Theorem 4.6.

Chapter 5

Adversary bound for Element Distinctness with small range

Recall that, given an input string $z \in \Sigma^n$, the ELEMENT DISTINCTNESS problem is to decide whether z contains a collision or not, namely, whether there exist $i, j \in [n]$ such that $i \neq j$ and $z_i = z_j$. In this chapter, we only consider a special case of the problem where we are given a promise that the input contains at most one collision. This promise does not change the complexity of the problem [Amb07].

We construct a tight $\Omega(n^{2/3})$ adversary lower bound for ELEMENT DISTINCTNESS with minimal alphabet such that the problem is still non-trivial. Due to Remark 2.16, a lower bound for a minimal alphabet is also a lower bound for any larger alphabet. We also provide certain “tight” conditions that every optimal adversary matrix for ELEMENT DISTINCTNESS must satisfy,¹ therefore suggesting that every optimal adversary matrix for ELEMENT DISTINCTNESS might have to be, in some sense, close to the adversary matrix that we have constructed.

5.1 Preliminaries

As before, let \mathcal{D}_1 and \mathcal{D}_0 denote the sets of positive and negative inputs, respectively, that is, inputs with a unique collision and inputs without a collision. If $|\Sigma| < n$, then $\mathcal{D}_0 = \emptyset$, and the problem becomes trivial. Therefore we consider the case when $|\Sigma| = n$. We have

$$|\mathcal{D}_1| = \binom{n}{2} \frac{|\Sigma|!}{(|\Sigma| - n + 1)!} = \binom{n}{2} n! \quad \text{and} \quad |\mathcal{D}_0| = \frac{|\Sigma|!}{(|\Sigma| - n)!} = n!.$$

¹Assuming, without loss of generality, that the adversary matrix has the symmetry given by the automorphism principle.

Adversary method. As ELEMENT DISTINCTNESS is a decision problem, we assume that the rows of the adversary matrix Γ are labeled by the positive inputs \mathcal{D}_1 and columns by the negative inputs \mathcal{D}_0 . Recall the difference matrices Δ_i and $\overline{\Delta}_i$ from (2.24). To apply the adversary bound (Theorem 2.15), our goal is to construct Γ such that $\|\Gamma\| = \Omega(n^{2/3})$ and $\|\Delta_i \circ \Gamma\| = O(1)$ for all $i \in [n]$. Recall that $\Delta_i \circ \Gamma = \Gamma - \overline{\Delta}_i \circ \Gamma$.

Symmetries of the adversary matrix. The automorphism principle (see Section 2.3.2) implies that, without loss of generality, we can assume that Γ is fixed under all index and all alphabet permutations. Namely, as defined in Section 2.2.2, index permutations $\pi \in \mathbb{S}_{[n]}$ and alphabet permutations $\tau \in \mathbb{S}_\Sigma$ act on input strings $z \in \Sigma^n$ in the natural way:

$$\begin{aligned} \pi \in \mathbb{S}_{[n]} : z = (z_1, \dots, z_n) &\mapsto z_\pi = (z_{\pi^{-1}(1)}, \dots, z_{\pi^{-1}(n)}), \\ \tau \in \mathbb{S}_\Sigma : z = (z_1, \dots, z_n) &\mapsto z^\tau = (\tau(z_1), \dots, \tau(z_n)). \end{aligned}$$

The actions of π and τ commute: we have $(z_\pi)^\tau = (z^\tau)_\pi$, which we denote by z_π^τ for short. The automorphism principle (see (2.28), in particular) implies that we can assume

$$\Gamma[x, y] = \Gamma[x_\pi^\tau, y_\pi^\tau] \tag{5.1}$$

for all $x \in \mathcal{D}_1$, $y \in \mathcal{D}_0$, $\pi \in \mathbb{S}_{[n]}$, and $\tau \in \mathbb{S}_\Sigma$.

Let us state the same symmetry via representations of $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$. Let $\mathcal{X} := \mathbb{R}^{\mathcal{D}_1}$ and $\mathcal{Y} := \mathbb{R}^{\mathcal{D}_0}$ be the vector spaces corresponding to the positive and the negative inputs, respectively. (We can view Γ as a linear map from \mathcal{Y} to \mathcal{X} .) Let U_π^τ and V_π^τ be the permutation matrices that respectively act on the spaces \mathcal{X} and \mathcal{Y} and that map every $x \in \mathcal{D}_1$ to x_π^τ and every $y \in \mathcal{D}_0$ to y_π^τ . Then (5.1) is equivalent to

$$U_\pi^\tau \Gamma = \Gamma V_\pi^\tau \tag{5.2}$$

for all $\pi \in \mathbb{S}_{[n]}$, and $\tau \in \mathbb{S}_\Sigma$. Both U and V are permutation representations of $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$.

Representation theory of the symmetric group. Let us recall basics of the representation theory of the symmetric group from Section 1.4. We represent each partition of m by an m -box Young diagram, and we use these terms interchangeably. In this chapter, we use ζ , η , and θ to denote Young diagrams having $o(n)$ boxes, λ , μ , and ν to denote Young diagrams having n , $n-1$, and $n-2$ boxes, respectively, and ρ and σ to denote Young diagrams for general statements and other purposes. Also, given a finite set A , recall that \mathcal{S}^ρ denotes the irrep of \mathbb{S}_A corresponding to $\rho \vdash |A|$, and the dimension of this irrep, denoted $\dim \rho$, is given by the hook-length formula (1.8). For a set $\{a, b\}$, let $\mathcal{S}^{\text{id}} := \mathcal{S}^{(2)}$ and $\mathcal{S}^{\text{sgn}} := \mathcal{S}^{(1,1)}$ be, respectively, the trivial and the sign representation of $\mathbb{S}_{\{a,b\}}$.

In this chapter, for a Young diagram ρ , let $\rho(i)$ and $\rho^\top(j)$ denote the number of boxes in the i -th row and j -th column of ρ , respectively. Recall the shorthand $(m, \rho) := (m, \rho(1), \rho(2), \dots, \rho(r))$, where $m \geq \rho(1)$.

Given $\ell \in \{0, 1, 2, 3\}$, a set $A = [n]$ or $A = \Sigma$, its subset $A \setminus \{a_1, \dots, a_\ell\}$, and $\rho \vdash n - \ell$, let us write $\rho_{a_1 \dots a_\ell}$ if we want to stress that we think of \mathcal{S}^ρ as an irrep of $\mathbb{S}_{A \setminus \{a_1, \dots, a_\ell\}}$. We omit the subscript if $\ell = 0$ or when $\{a_1, \dots, a_\ell\}$ is clear from the context. To lighten the notations, given $k = o(n)$ and $\eta \vdash k$, let $\bar{\eta}_{a_1 \dots a_\ell} = (n - \ell - k, \eta)_{a_1 \dots a_\ell} \vdash n - \ell$; here we omit the subscript if and only if $\ell = 0$.

Also recall that $\sigma \subset \rho$ and $\sigma \subset\subset \rho$ denotes that a Young diagram σ is obtained from ρ by removing exactly one box and exactly two boxes, respectively. And, given $\sigma \subset\subset \rho$, we write $\sigma \subset\subset_r \rho$ or $\sigma \subset\subset_c \rho$ if the two boxes removed from ρ to obtain σ are, respectively, in different rows or different columns. And $\sigma \subset\subset_{rc} \rho$ is a shorthand for $(\sigma \subset\subset_r \rho) \& (\sigma \subset\subset_c \rho)$. Given $\sigma \subset\subset_{rc} \rho$, let $d_{\rho, \sigma} \geq 2$ denote the distance between the two boxes that we remove from ρ to obtain σ .

Transporters. We will construct the adversary matrix Γ using transporters between isomorphic irreps. In this chapter we only consider real vector spaces, which we can do because of Claim 2.18, therefore each transporter is unique up to a global phase ± 1 . We always choose the global phases so that they respect composition and inversion, as described in Section 1.3.4.

Structure of the chapter. In Section 5.2 we show that the adversary matrix Γ can be expressed as a linear combination of specific matrices. In this section we also present Claim 5.2, which states what conditions every optimal adversary matrix for ELEMENT DISTINCTNESS must satisfy; we prove this claim in the Appendix B. In Section 5.3 we show how to specify the adversary matrix Γ via its submatrix $\Gamma_{1,2}$, which will make the analysis of the adversary matrix simpler. In Section 5.4 we present tools for estimating the norm $\|\Delta_i \circ \Gamma\|$. Finally, in Section 5.5 we use the conditions given by Claim 5.2 to construct an adversary matrix for ELEMENT DISTINCTNESS with the alphabet size n , and we show that this matrix indeed yields the desired $\Omega(n^{2/3})$ lower bound.

5.2 Building blocks of Γ

5.2.1 Decomposition of U and V into irreps

Without loss of generality, we assume that the adversary matrix Γ satisfy the symmetry (5.2) given by the automorphism principle. Both U and V are representations of $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$ and, in order to use Schur's lemma (Lemma 1.2), we want to see what irreps of $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$ occur in both U and V . It is also convenient to consider U and V as representations of just $\mathbb{S}_{[n]}$ or just \mathbb{S}_Σ .

Claim 5.1. V decomposes into irreps of $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$ as $V \cong \bigoplus_{\lambda \vdash n} \mathcal{S}^\lambda \times \mathcal{S}^\lambda$.

Proof. As a representation of $\mathbb{S}_{[n]}$ and \mathbb{S}_Σ , respectively, V is isomorphic to the regular representation of $\mathbb{S}_{[n]}$ and \mathbb{S}_Σ (see Section 1.3.6). For every $y \in \mathcal{D}_0$ and every $\pi \in \mathbb{S}_{[n]}$, there is a unique $\tau \in \mathbb{S}_\Sigma$ such that $y_\pi = y^\tau$, and π and τ belong to isomorphic conjugacy classes. Thus, Theorem 1.9 implies that, for every $\lambda \vdash n$, the \mathcal{S}^λ -isotypical subspace of \mathcal{Y} is the same for both $\mathbb{S}_{[n]}$ and \mathbb{S}_Σ . Since V is isomorphic to the regular representation, the dimension of this subspace is $(\dim \lambda)^2$, which is exactly the dimension of the irrep $\mathcal{S}^\lambda \times \mathcal{S}^\lambda$ of $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$. \square

Now let us address U , which acts on the space \mathcal{X} corresponding to the positive inputs $x \in \mathcal{D}_1$. Let us decompose \mathcal{D}_1 as a disjoint union of $\binom{n}{2}$ sets $\mathcal{D}_{i,j}$, where $\{i, j\} \subset [n]$ and $\mathcal{D}_{i,j}$ is the set of all $x \in \mathcal{D}_1$ such that $x_i = x_j$ (note: we used an analogous decomposition in Section 4.1 for the ELEMENT DISTINCTNESS problem with large range). Let us further decompose $\mathcal{D}_{i,j}$ as a disjoint union of $\binom{n}{2}$ sets $\mathcal{D}_{i,j}^{s,t}$, where $\{s, t\} \subset \Sigma$ and $\mathcal{D}_{i,j}^{s,t}$ is the set of all $x \in \mathcal{D}_{i,j}$ that does not contain s and contains t twice or vice versa. Let $\mathcal{X}_{i,j}$ and $\mathcal{X}_{i,j}^{s,t}$ be the subspaces of \mathcal{X} that correspond to the sets $\mathcal{D}_{i,j}$ and $\mathcal{D}_{i,j}^{s,t}$, respectively. The space $\mathcal{X}_{i,j}^{s,t}$ is stable under the action of

$$\mathbb{S}_{i,j}^{s,t} := (\mathbb{S}_{\{i,j\}} \times \mathbb{S}_{[n] \setminus \{i,j\}}) \times (\mathbb{S}_{\{s,t\}} \times \mathbb{S}_{\Sigma \setminus \{s,t\}}),$$

namely, $U_\pi^\tau \mathcal{X}_{i,j}^{s,t} = \mathcal{X}_{i,j}^{s,t}$ for all $(\pi, \tau) \in \mathbb{S}_{i,j}^{s,t}$. Therefore, U restricted to the subspace $\mathcal{X}_{i,j}^{s,t}$ is a representation of $\mathbb{S}_{i,j}^{s,t}$, and, similarly to Claim 5.1, it decomposes into irreps as

$$\bigoplus_{\nu \vdash n-2} (\mathcal{S}^{\text{id}} \times \mathcal{S}^\nu) \times ((\mathcal{S}^{\text{id}} \oplus \mathcal{S}^{\text{sgn}}) \times \mathcal{S}^\nu). \quad (5.3)$$

To see how U decomposes into irreps of $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$, we induce the representation (5.3) from $\mathbb{S}_{i,j}^{s,t}$ to $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$.

The Littlewood–Richardson rule (1.11) implies that an irrep of $\mathbb{S}_{[n]} \times \mathbb{S}_\Sigma$ isomorphic to $\mathcal{S}^\lambda \times \mathcal{S}^\lambda$ can occur in U due to one of the following scenarios.

- If $\nu \subset_c \lambda$ and $\nu \not\subset_r \lambda$ (i.e., ν is obtained from λ by removing two boxes in the same row), then $\mathcal{S}^\lambda \times \mathcal{S}^\lambda$ occurs once in the induction of $(\mathcal{S}^{\text{id}} \times \mathcal{S}^\nu) \times (\mathcal{S}^{\text{id}} \times \mathcal{S}^\nu)$. Let $\mathcal{X}_{\text{id},\nu}^\lambda$ denote the subspace of \mathcal{X} corresponding to this instance of $\mathcal{S}^\lambda \times \mathcal{S}^\lambda$.
- If $\nu \subset_{rc} \lambda$, then $\mathcal{S}^\lambda \times \mathcal{S}^\lambda$ occurs once in the induction of $(\mathcal{S}^{\text{id}} \times \mathcal{S}^\nu) \times (\mathcal{S}^{\text{id}} \times \mathcal{S}^\nu)$ and once in the induction of $(\mathcal{S}^{\text{id}} \times \mathcal{S}^\nu) \times (\mathcal{S}^{\text{sgn}} \times \mathcal{S}^\nu)$. Let $\mathcal{X}_{\text{id},\nu}^\lambda$ and $\mathcal{X}_{\text{sgn},\nu}^\lambda$ denote the respective subspaces of \mathcal{X} corresponding to these instances of $\mathcal{S}^\lambda \times \mathcal{S}^\lambda$.

Note: from the definition of the induction in Section 1.3.7, one can see that the subspaces $\mathcal{X}_{\text{id},\nu}^\lambda$ and $\mathcal{X}_{\text{sgn},\nu}^\lambda$ are independent from the choice of $\{i, j\} \subset [n]$ and $\{s, t\} \subset \Sigma$.

5.2.2 Γ as a linear combination of transporters

Let $\Xi_{\text{id},\nu}^\lambda$ and $\Xi_{\text{sgn},\nu}^\lambda$ denote the transporters from the unique instance of $\mathcal{S}^\lambda \times \mathcal{S}^\lambda$ in \mathcal{Y} to the subspaces $\mathcal{X}_{\text{id},\nu}^\lambda$ and $\mathcal{X}_{\text{sgn},\nu}^\lambda$, respectively. We will specify the global phases of these transporters in Section 5.3.3. We consider $\Xi_{\text{id},\nu}^\lambda$ and $\Xi_{\text{sgn},\nu}^\lambda$ as matrices of the same dimensions as Γ , namely, $\binom{n}{2}n! \times n!$. Schur’s lemma implies that, due to (5.2), we can express Γ as a linear combination of these transporters. Namely,

$$\Gamma = \sum_{\lambda \vdash n} \left(\sum_{\nu \subset_c \lambda} \beta_{\text{id},\nu}^\lambda \Xi_{\text{id},\nu}^\lambda + \sum_{\nu \subset_{rc} \lambda} \beta_{\text{sgn},\nu}^\lambda \Xi_{\text{sgn},\nu}^\lambda \right), \quad (5.4)$$

where the coefficients $\beta_{\text{id},\nu}^\lambda$ and $\beta_{\text{sgn},\nu}^\lambda$ are real.

Thus we have reduced the construction of the adversary matrix Γ to choosing the coefficients β of the transporters in (5.4). To illustrate what are the available transporters, let us consider the last four $(n-2)$ -box Young diagrams ν of the lexicographical order— $(n-2)$, $(n-3,1)$, $(n-4,2)$, and $(n-4,1,1)$ —and all λ that are obtained from these ν by adding two boxes in different columns. Table 5.1 shows pairs of λ and ν for which we have both $\Xi_{\text{id},\nu}^\lambda$ and $\Xi_{\text{sgn},\nu}^\lambda$ available for the construction of Γ (double check mark “ $\checkmark\checkmark$ ”) or just $\Xi_{\text{id},\nu}^\lambda$ available (single check mark “ \checkmark ”).

$\lambda \backslash \nu$	$(n-2)$	$(n-3,1)$	$(n-4,2)$	$(n-4,1,1)$
(n)	\checkmark_0			
$(n-1,1)$	$\checkmark\checkmark_1$	\checkmark_0		
$(n-2,2)$	\checkmark_2	$\checkmark\checkmark_1$	\checkmark_0	
$(n-2,1,1)$		$\checkmark\checkmark_1$		\checkmark_0
$(n-3,3)$		\checkmark_2	$\checkmark\checkmark_1$	
$(n-3,2,1)$		$\checkmark\checkmark_2$	$\checkmark\checkmark_1$	$\checkmark\checkmark_1$
$(n-3,1,1,1)$				$\checkmark\checkmark_1$
$(n-4,4)$			\checkmark_2	
$(n-4,3,1)$			$\checkmark\checkmark_2$	\checkmark_2
$(n-4,2,2)$			\checkmark_2	
$(n-4,2,1,1)$				$\checkmark\checkmark_2$

Table 5.1: Available operators for the construction of Γ . We distinguish three cases: both λ and ν are the same below the first row (label “ \checkmark_0 ”), λ has one box more below the first row than ν (label “ $\checkmark\checkmark_1$ ”), λ has two boxes more below the first row than ν (labels “ \checkmark_2 ” and “ $\checkmark\checkmark_2$ ”).

Due to the symmetry, $\|\Delta_i \circ \Gamma\|$ is the same for all $i \in [n]$, so, from now on, let us only consider $\Delta_1 \circ \Gamma$. We want to choose the coefficients β so that $\|\Gamma\| = \Omega(n^{2/3})$ and $\|\Delta_1 \circ \Gamma\| = O(1)$. The automorphism principle also implies that we can assume that the principal left and right

singular vectors of Γ are the all-ones vectors, which correspond to $\Xi_{\text{id},(n-2)}^{(n)}$. We thus choose $\beta_{\text{id},(n-2)}^{(n)} = \Theta(n^{2/3})$.

In order to understand how to choose the coefficients β , later in Appendix B we prove the following claim, which relates all the coefficients of transporters of Table 5.1 and more.

Claim 5.2. *Suppose Γ is given as in (5.4) and $\beta_{\text{id},(n-2)}^{(n)} = n^{2/3}$. Consider $\lambda \vdash n$ that has $O(1)$ boxes below the first row and $\nu \subset_c \lambda$. In order for $\|\Delta_1 \circ \Gamma\| = O(1)$ to hold, we need to have*

1. $\beta_{\text{id},\nu}^\lambda = n^{2/3} + O(1)$ if λ and ν are the same below the first row,
2. $\beta_{\text{id},\nu}^\lambda, \beta_{\text{sgn},\nu}^\lambda = c_\nu^\lambda n^{1/6} + O(1)$ if λ has one box more below the first row than ν , where c_ν^λ is a constant depending only on the part of λ and ν below the first row,²
3. $\beta_{\text{id},\nu}^\lambda, \beta_{\text{sgn},\nu}^\lambda = O(1)$ if λ has two boxes more below the first row than ν .

Note that we always have the freedom of changing (a constant number of) coefficients β up to an additive term of $O(1)$ because of the fact that $\|\Delta_j \circ B\| \leq 2\|B\|$ for all matrices B and the triangle inequality.

5.3 Specification of Γ via $\Gamma_{1,2}$

Due to the symmetry (5.1), it suffices to specify a single row of the adversary matrix Γ in order to specify the whole matrix. For the convenience, let us instead specify Γ via specifying its $n! \times n!$ submatrix $\Gamma_{1,2}$ —for $\{i, j\} \subset [n]$, we define $\Gamma_{i,j}$ to be the submatrix of Γ that corresponds to the rows labeled by $x \in \mathcal{D}_{i,j}$, that is, positive inputs x with $x_i = x_j$. We think of $\Gamma_{i,j}$ both as an $n! \times n!$ square matrix and as a matrix of the same dimensions as Γ that is obtained from Γ by setting to zero all the $\binom{n}{2} - 1$ rows that correspond to $x \notin \mathcal{D}_{i,j}$.

5.3.1 Necessary and sufficient symmetries of $\Gamma_{1,2}$

For all $(\pi, \tau) \in (\mathbb{S}_{\{1,2\}} \times \mathbb{S}_{[3..n]}) \times \mathbb{S}_\Sigma$, we have $U_\pi^\tau \mathcal{X}_{1,2} = \mathcal{X}_{1,2}$ and, therefore, $U_\pi^\tau \Gamma_{1,2} = \Gamma_{1,2} V_\pi^\tau$. This is the necessary and sufficient symmetry that $\Gamma_{1,2}$ must satisfy in order for Γ to be fixed under all index and alphabet permutations. Since $U_{(12)} \Gamma_{1,2} = \Gamma_{1,2}$, where $\pi = (12)$ denotes the transposition of indices 1 and 2, we also have $\Gamma_{1,2} V_{(12)} = \Gamma_{1,2}$. We have

$$\Gamma = \sum_{\{i,j\} \subset [n]} \Gamma_{i,j} = \sum_{\pi \in R} U_\pi \Gamma_{1,2} V_{\pi^{-1}} = \binom{n}{2} \frac{1}{n!} \sum_{\pi \in \mathbb{S}_{[n]}} U_\pi \Gamma_{1,2} V_{\pi^{-1}}, \quad (5.5)$$

²Let $\hat{\lambda}$ and $\hat{\nu}$ be the part of λ and ν below the first row, respectively, and let $h(\cdot)$ be the product of hook lengths as in (1.8). Then $c_\nu^\lambda = \sqrt{h(\hat{\lambda})/h(\hat{\nu})} = \sqrt{n \dim \nu / \dim \lambda} + O(1/n)$.

where $R := \text{Rep}(\mathbb{S}_{[n]} / (\mathbb{S}_{\{1,2\}} \times \mathbb{S}_{[3..n]}))$ is a transversal of the left cosets of $\mathbb{S}_{\{1,2\}} \times \mathbb{S}_{[3..n]}$ in $\mathbb{S}_{[n]}$.

Let f be a bijection between \mathcal{D}_0 and $\mathcal{D}_{1,2}$ defined as

$$f : \mathcal{D}_0 \rightarrow \mathcal{D}_{1,2} : (y_1, y_2, y_3, \dots, y_n) \mapsto (y_1, y_1, y_3, \dots, y_n),$$

and let F be the corresponding permutation matrix mapping \mathcal{Y} to $\mathcal{X}_{1,2}$. Let us order rows and columns of $\Gamma_{1,2}$ so that they correspond to $f(y)$ and y , respectively, where we take $y \in \mathcal{D}_0$ in the same order for both (see Figure 5.1). Hence, F becomes the identity matrix on \mathcal{Y} , and, from this point onward, we essentially think of $\mathcal{X}_{1,2}$ and \mathcal{Y} as the same space. Let us denote this identity matrix simply by \mathbb{I} .

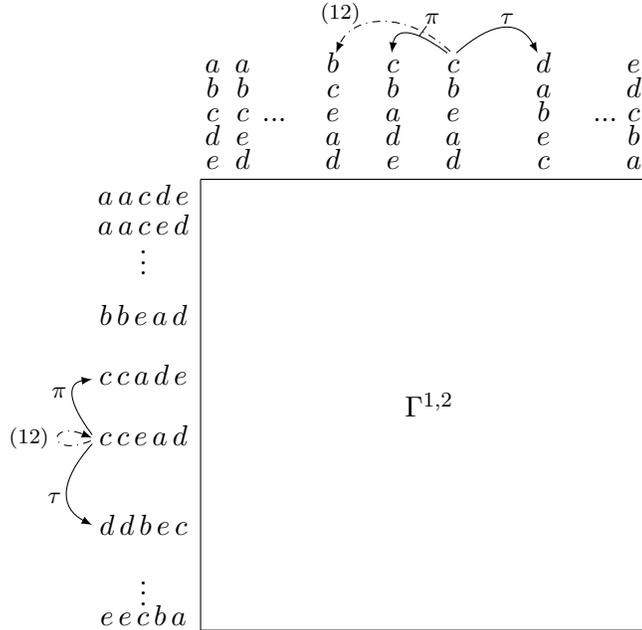


Figure 5.1: Symmetries of $\Gamma_{1,2}$ for $n = 5$ and $\Sigma = \{a, b, c, d, e\}$. With respect to the bijection f , the order of rows and columns matches. The solid arrows show that U^τ and V^τ act symmetrically on $\Gamma_{1,2}$ (here we use $\tau = (aeb)(cd) \in \mathbb{S}_\Sigma$), and so do U_π and V_π for $\pi \in \mathbb{S}_{[3..n]}$ (here we use $\pi = (354)$). However, as shown by the dash-dotted arrows, $U_{(12)}$ acts as the identity on the rows, while $V_{(12)}$ transposes the columns.

For all $(\pi, \tau) \in \mathbb{S}_{[3..n]} \times \mathbb{S}_\Sigma$ we have $f(y_\pi^\tau) = (f(y))_\pi^\tau$ and, thus, $V_\pi^\tau = FV_\pi^\tau = U_\pi^\tau F = U_\pi^\tau$, where we consider the restriction of U_π^τ to $\mathcal{X}_{1,2}$. Note that $U_{(12)} = \mathbb{I}$ on $\mathcal{X}_{1,2}$, while $V_{(12)} \neq \mathbb{I}$. Hence now the two necessary and sufficient symmetries that $\Gamma_{1,2}$ must satisfy are

$$V_\pi^\tau \Gamma_{1,2} = \Gamma_{1,2} V_\pi^\tau \quad \text{for all } (\pi, \tau) \in \mathbb{S}_{[3..n]} \times \mathbb{S}_\Sigma \quad \text{and} \quad \Gamma_{1,2} V_{(12)} = \Gamma_{1,2}. \quad (5.6)$$

Figure 5.1 illustrates these symmetries.

5.3.2 Labeling of projectors and transporters

We use Π , with some sub- and superscripts, to denote operators acting on \mathcal{Y} ; we use subscripts for irreps of index permutations and superscripts for irreps of alphabet permutations. We also think of each such an operator Π to map \mathcal{Y} to $\mathcal{X}_{1,2}$ and vice versa (technically, $F\Pi$ and ΠF^* , respectively).

Let $\Pi_{\text{id}} := (\mathbb{I} + V_{(12)})/2$ and $\Pi_{\text{sgn}} := (\mathbb{I} - V_{(12)})/2$ denote the projectors on the isotypical subspaces of \mathcal{Y} corresponding to irreps \mathcal{S}^{id} and \mathcal{S}^{sgn} of $\mathbb{S}_{\{1,2\}}$, respectively (see Theorem 1.9). Let $\Pi_{\rho_{i_1 \dots i_\ell}}$ and $\Pi^{\sigma_{s_1 \dots s_m}}$ denote the projectors on the isotypical subspaces corresponding to an irrep \mathcal{S}^ρ of $\mathbb{S}_{[n] \setminus \{i_1, \dots, i_\ell\}}$ and an irrep \mathcal{S}^σ of $\mathbb{S}_{\Sigma \setminus \{s_1, \dots, s_m\}}$, respectively. Note that $\Pi_{\rho_{i_1 \dots i_\ell}}$ and $\Pi^{\sigma_{s_1 \dots s_m}}$ commute, and let

$$\Pi_{\rho_{i_1 \dots i_\ell}}^{\sigma_{s_1 \dots s_m}} := \Pi_{\rho_{i_1 \dots i_\ell}} \Pi^{\sigma_{s_1 \dots s_m}} = \Pi^{\sigma_{s_1 \dots s_m}} \Pi_{\rho_{i_1 \dots i_\ell}},$$

which is the projector on the isotypical subspace corresponding to the irrep $\mathcal{S}^\rho \times \mathcal{S}^\sigma$ of

$$\mathbb{S}_{[n] \setminus \{i_1, \dots, i_\ell\}} \times \mathbb{S}_{\Sigma \setminus \{s_1, \dots, s_m\}}$$

(note: this subspace may contain multiple instances of the irrep). In general, when multiple such projectors mutually commute, we denote their product with a single Π whose sub- and superscript is, respectively, a concatenation of the sub- and superscripts of these projectors. For example, $\Pi_{\text{id}, \nu_{12}}^\lambda := \Pi_{\text{id}}^\lambda \Pi_{\nu_{12}}^\lambda \Pi^\lambda$ (note: Π^λ corresponds to an irrep \mathcal{S}^λ of $\mathbb{S}_{\Sigma \setminus \emptyset} = \mathbb{S}_\Sigma$).

Suppose that Π_{sub}^λ and $\Pi_{\text{sub}'}^\lambda$ are two projectors each projecting onto a single instance of an irrep $\mathcal{S}^{\rho_{i_1 \dots i_\ell}} \times \mathcal{S}^\lambda$ of $\mathbb{S}_{[n] \setminus \{i_1, \dots, i_\ell\}} \times \mathbb{S}_\Sigma$, where sub and sub' are subscripts determining these instances. Then let $\Pi_{\text{sub}' \leftarrow \text{sub}}^\lambda$ denote the transporter from the instance corresponding to Π_{sub}^λ to one corresponding to $\Pi_{\text{sub}'}^\lambda$. Let $\Pi_{\text{sub}' \leftrightarrow \text{sub}}^\lambda := \Pi_{\text{sub}' \leftarrow \text{sub}}^\lambda + \Pi_{\text{sub} \leftarrow \text{sub}'}^\lambda$ for short.

5.3.3 Decomposition of $\Gamma_{1,2}$ into projectors and transporters

Due to (5.6), we can express $\Gamma_{1,2}$ as a linear combination of projectors onto irreps and transporters between isomorphic irreps of $\mathbb{S}_{[3..n]} \times \mathbb{S}_\Sigma$. Due to (5.6) we also have $\Gamma_{1,2} \Pi_{\text{id}} = \Gamma_{1,2}$ and $\Gamma_{1,2} \Pi_{\text{sgn}} = 0$. Claim 5.1 states that $\mathbb{I} = \sum_{\lambda \vdash n} \Pi_\lambda^\lambda$, and we have $\Pi_\lambda^\lambda = \sum_{\nu \subset \lambda} \Pi_{\nu_{12}}^\lambda$. If the two boxes removed from λ to obtain ν are in the same row or the same column, then $\Pi_{\nu_{12}}^\lambda$ projects onto the unique instance of the irrep $\mathcal{S}^\nu \times \mathcal{S}^\lambda$ in V , and $\Pi_{\nu_{12}}^\lambda = \Pi_{\text{id}, \nu_{12}}^\lambda$ or $\Pi_{\nu_{12}}^\lambda = \Pi_{\text{sgn}, \nu_{12}}^\lambda$, respectively. On the other hand, if they are in different rows and columns, then $\Pi_{\nu_{12}}^\lambda = \Pi_{\text{id}, \nu_{12}}^\lambda + \Pi_{\text{sgn}, \nu_{12}}^\lambda$, where each $\Pi_{\text{id}, \nu_{12}}^\lambda$ and $\Pi_{\text{sgn}, \nu_{12}}^\lambda$ projects onto an instance of the irrep $\mathcal{S}^\nu \times \mathcal{S}^\lambda$. Hence, similarly to (5.4), we can express $\Gamma_{1,2}$ as a linear combination

$$\Gamma_{1,2} = \sum_{\lambda \vdash n} \left(\sum_{\nu \subset \lambda} \alpha_{\text{id}, \nu}^\lambda \Pi_{\text{id}, \nu_{12}}^\lambda + \sum_{\nu \subset \lambda} \alpha_{\text{sgn}, \nu}^\lambda \Pi_{\text{sgn}, \nu_{12} \leftarrow \text{id}, \nu_{12}}^\lambda \right). \quad (5.7)$$

If $\nu \subset_{rc} \lambda$, then there exist two distinct $\mu, \mu' \vdash n-1$ such that $\nu \subset \mu \subset \lambda$ and $\nu \subset \mu' \subset \lambda$, and let μ appear in the lexicographic order after μ' . Note that $\Pi_{\nu_{12}, \mu_1}^\lambda$ projects onto a single instance of $\mathcal{S}^\nu \times \mathcal{S}^\lambda$. We have

$$\Pi_{\text{sgn}, \nu_{12} \leftarrow \text{id}, \nu_{12}}^\lambda \propto \Pi_{\text{sgn}, \nu_{12}}^\lambda \Pi_{\nu_{12}, \mu_1}^\lambda \Pi_{\text{id}, \nu_{12}}^\lambda,$$

and we specify the global phase of the transporter $\Pi_{\text{sgn}, \nu_{12} \leftarrow \text{id}, \nu_{12}}^\lambda$ by assuming that the coefficient of this proportionality is positive. We present the value of this coefficient in Section 5.4.3.

Let us relate (5.4) and (5.7), the two ways in which we can specify the adversary matrix. One can see that the $2(n-2)! \times n!$ submatrix of $\Xi_{\text{id}, \nu_{12}}^\lambda$ and $\Xi_{\text{sgn}, \nu_{12}}^\lambda$ corresponding to $\mathcal{D}_{1,2}^{s,t}$ is proportional, respectively, to the $2(n-2)! \times n!$ submatrix of $\Pi_{\text{id}, \nu_{12}}^\lambda$ and $\Pi_{\text{sgn}, \nu_{12} \leftarrow \text{id}, \nu_{12}}^\lambda$ corresponding to $\mathcal{D}_{1,2}^{s,t}$. Hence, just like in (5.5), we have

$$\Xi_{\text{id}, \nu}^\lambda = \frac{1}{\gamma_{\text{id}, \nu}^\lambda} \sum_{\pi \in R} U_\pi \Pi_{\text{id}, \nu_{12}}^\lambda V_{\pi^{-1}} \quad \text{and} \quad \Xi_{\text{sgn}, \nu}^\lambda = \frac{1}{\gamma_{\text{sgn}, \nu}^\lambda} \sum_{\pi \in R} U_\pi \Pi_{\text{sgn}, \nu_{12} \leftarrow \text{id}, \nu_{12}}^\lambda V_{\pi^{-1}},$$

and we specify the global phase of the transporters Ξ by assuming that the normalization scalars γ are positive. Note that

$$\begin{aligned} (\gamma_{\text{id}, \nu}^\lambda)^2 \Pi_\lambda^\lambda &= (\gamma_{\text{id}, \nu}^\lambda \Xi_{\text{id}, \nu}^\lambda)^* (\gamma_{\text{id}, \nu}^\lambda \Xi_{\text{id}, \nu}^\lambda) = \left(\sum_{\pi \in R} U_\pi \Pi_{\text{id}, \nu_{12}}^\lambda V_{\pi^{-1}} \right)^* \sum_{\pi \in R} U_\pi \Pi_{\text{id}, \nu_{12}}^\lambda V_{\pi^{-1}} \\ &= \binom{n}{2} \frac{1}{n!} \sum_{\pi \in \mathbb{S}_{[n]}} V_\pi \Pi_{\text{id}, \nu_{12}}^\lambda V_{\pi^{-1}} = \binom{n}{2} \frac{\dim \nu}{\dim \lambda} \Pi_\lambda^\lambda, \end{aligned}$$

where the last equality holds because V_π and Π_λ^λ commute (thus the sum has to be proportional to Π_λ^λ) and $\text{Tr}(\Pi_{\text{id}, \nu_{12}}^\lambda) / \text{Tr}(\Pi_\lambda^\lambda) = \dim \nu / \dim \lambda$. The same way we calculate $\gamma_{\text{sgn}, \nu}^\lambda$, and we have

$$\gamma_{\text{id}, \nu}^\lambda = \frac{\beta_{\text{id}, \nu}^\lambda}{\alpha_{\text{id}, \nu}^\lambda} = \gamma_{\text{sgn}, \nu}^\lambda = \frac{\beta_{\text{sgn}, \nu}^\lambda}{\alpha_{\text{sgn}, \nu}^\lambda} = \sqrt{\binom{n}{2} \frac{\dim \nu}{\dim \lambda}}.$$

5.4 Tools for estimating $\|\Delta_1 \circ \Gamma\|$

5.4.1 Division of $\Delta_1 \circ \Gamma$ into two parts

For all $j \in [2..n]$, $\Delta_1 \circ \Gamma_{1,j}$ is essentially the same as $\Delta_1 \circ \Gamma_{1,2}$. And, for all $\{i, j\} \subset [2..n]$, $\Delta_1 \circ \Gamma_{i,j}$ is essentially the same as $\Delta_1 \circ \Gamma_{2,3}$, which, in turn, is essentially the same as $\Delta_3 \circ \Gamma_{1,2}$. Just like in Section 4.1 for ELEMENT DISTINCTNESS with large range (see Figure 4.2), let us distinguish these two cases by dividing Γ into two parts: let Γ' be the $(n-1)n! \times n!$ submatrix of Γ corresponding to $x \in \mathcal{D}_{1,j}$, where $j \in [2..n]$, and let Γ'' be the $\binom{n-1}{2}n! \times n!$ submatrix of Γ corresponding to $x \in \mathcal{D}_{i,j}$, where $\{i, j\} \in [2..n]$.

Claim 5.3. *We have $\|\Delta_1 \circ \Gamma\| = O(1)$ if and only if both $\|\Delta_1 \circ \Gamma'\| = O(1)$ and $\|\Delta_1 \circ \Gamma''\| = O(1)$.*

Let $R' := \text{Rep}(\mathbb{S}_{[2..n]}/\mathbb{S}_{[3..n]})$ and $R'' := \text{Rep}(\mathbb{S}_{[n]\setminus\{3\}}/(\mathbb{S}_{\{1,2\}} \times \mathbb{S}_{[4..n]}))$ be transversals of the left cosets of $\mathbb{S}_{[3..n]}$ in $\mathbb{S}_{[2..n]}$ and of $\mathbb{S}_{\{1,2\}} \times \mathbb{S}_{[4..n]}$ in $\mathbb{S}_{[n]\setminus\{3\}}$, respectively. Similarly to (5.5), we have

$$\Delta_1 \circ \Gamma' = \sum_{\pi \in R'} U_\pi(\Delta_1 \circ \Gamma_{1,2}) V_{\pi^{-1}} \quad \text{and} \quad \Delta_1 \circ \Gamma'' = U_{(13)} \left(\sum_{\pi \in R''} U_\pi(\Delta_3 \circ \Gamma_{1,2}) V_{\pi^{-1}} \right) V_{(13)}, \quad (5.8)$$

which imply

$$\|\Delta_1 \circ \Gamma'\|^2 = \|(\Delta_1 \circ \Gamma')^*(\Delta_1 \circ \Gamma')\| = \left\| \sum_{\pi \in R'} V_\pi(\Delta_1 \circ \Gamma_{1,2})^*(\Delta_1 \circ \Gamma_{1,2}) V_{\pi^{-1}} \right\|, \quad (5.9)$$

$$\|\Delta_1 \circ \Gamma''\|^2 = \|(\Delta_1 \circ \Gamma'')^*(\Delta_1 \circ \Gamma'')\| = \left\| \sum_{\pi \in R''} V_\pi(\Delta_3 \circ \Gamma_{1,2})^*(\Delta_3 \circ \Gamma_{1,2}) V_{\pi^{-1}} \right\|. \quad (5.10)$$

Therefore, we have to consider $\Delta_1 \circ \Gamma_{1,2}$ and $\Delta_3 \circ \Gamma_{1,2}$.

5.4.2 Commutativity with the action of Δ_i

Instead of Δ_i , let us first consider the action of $\bar{\Delta}_i$. For $i \in [n]$ and $s \in \Sigma$, let $\hat{\Pi}_i^s$ be the projector on all $y \in \mathcal{D}_0$ such that $y_i = s$. Then, due to the particular way we define the bijection f , we have

$$\bar{\Delta}_i \circ \Gamma_{1,2} = \sum_{s \in \Sigma} \hat{\Pi}_i^s \Gamma_{1,2} \hat{\Pi}_i^s \quad \text{whenever } i \neq 2 \quad \text{and} \quad \bar{\Delta}_2 \circ \Gamma_{1,2} = \sum_{s \in \Sigma} \hat{\Pi}_1^s \Gamma_{1,2} \hat{\Pi}_2^s. \quad (5.11)$$

Note that $\hat{\Pi}_i^s$ commutes with every $\Pi_{\rho_{j_1 \dots j_m}}$ whenever $i \in \{j_1, \dots, j_m\}$. Hence, for

$$i \in \{j_1, \dots, j_m\} \setminus \{2\}$$

and every $n! \times n!$ matrix A , we have

$$\Delta_i \circ (\Pi_{\rho_{j_1 \dots j_m}} A) = \Pi_{\rho_{j_1 \dots j_m}} (\Delta_i \circ A) \quad \text{and} \quad \Delta_i \circ (A \Pi_{\rho_{j_1 \dots j_m}}) = (\Delta_i \circ A) \Pi_{\rho_{j_1 \dots j_m}}. \quad (5.12)$$

5.4.3 Relations among irreps of $\mathbb{S}_{[3..n]} \times \mathbb{S}_\Sigma$ within an isotypical subspace

We are interested to see how Δ_1 acts on $\Gamma_{1,2}$, which requires us to consider how it acts on $\Pi_{\text{id}, \nu_{12}}^\lambda$ and $\Pi_{\text{sgn}, \nu_{12} \leftarrow \text{id}, \nu_{12}}^\lambda$. Unfortunately, this action is hard to calculate directly, therefore we express $\Pi_{\text{id}, \nu_{12}}^\lambda$ and $\Pi_{\text{sgn}, \nu_{12} \leftarrow \text{id}, \nu_{12}}^\lambda$ as linear combinations of certain operators on which the action of Δ_1 is easier to calculate.

Consider $\lambda \vdash n$ and $\nu \subset_{rc} \lambda$. The projector $\Pi_{\nu_{12}}^\lambda$ projects onto the isotypical subspace of \mathcal{Y} corresponding to the irrep $\mathcal{S}^\nu \times \mathcal{S}^\lambda$ of $\mathbb{S}_{[3..n]} \times \mathbb{S}_\Sigma$, and this subspace contains two instances of this irrep. There are as many degrees of freedom in splitting this subspace in half so that each half corresponds to a single instance of the irrep as in splitting \mathbb{R}^2 in orthogonal one-dimensional subspaces. We already considered one such split, $\Pi_{\nu_{12}}^\lambda = \Pi_{\text{id}, \nu_{12}}^\lambda + \Pi_{\text{sgn}, \nu_{12}}^\lambda$, and now let us relate it to another.

Let $\mu, \mu' \vdash n-1$ be such that $\nu \subset \mu \subset \lambda$, $\nu \subset \mu' \subset \lambda$, and μ appears after μ' in the lexicographical order. Then $\Pi_{\nu_{12}, \mu_1}^\lambda$ and $\Pi_{\nu_{12}, \mu'_1}^\lambda$ project onto two orthogonal instances of the irrep $\mathcal{S}^\nu \times \mathcal{S}^\lambda$, and $\Pi_{\nu_{12}}^\lambda = \Pi_{\nu_{12}, \mu_1}^\lambda + \Pi_{\nu_{12}, \mu'_1}^\lambda$. Note that $V_{(12)}$ commutes with $\Pi_{\nu_{12}}^\lambda$ and that $\Pi^\lambda = \Pi_\lambda$. The orthogonal form (see Section 1.4.4) of the irrep \mathcal{S}^λ tells us that $V_{(12)}$ restricted to the isotypical subspace corresponding to $\mathcal{S}^\nu \times \mathcal{S}^\lambda$ is

$$V_{(12)} \Big|_{\nu_{12} \times \lambda} = \frac{1}{d_{\lambda, \nu}} \left(\Pi_{\nu_{12}, \mu'_1}^\lambda - \Pi_{\nu_{12}, \mu_1}^\lambda + \sqrt{d_{\lambda, \nu}^2 - 1} \Pi_{\nu_{12}, \mu'_1 \leftrightarrow \nu_{12}, \mu_1}^\lambda \right). \quad (5.13)$$

In effect, (5.13) defines the global phase of the transporters $\Pi_{\nu_{12}, \mu'_1 \leftarrow \nu_{12}, \mu_1}^\lambda$ and $\Pi_{\nu_{12}, \mu'_1 \leftrightarrow \nu_{12}, \mu_1}^\lambda$.

Recall that $\Pi_{\text{id}} = (\mathbb{I} + V_{(12)})/2$, and therefore

$$\Pi_{\text{id}, \nu_{12}}^\lambda = \frac{\Pi_{\nu_{12}}^\lambda + V_{(12)} \Big|_{\nu_{12} \times \lambda}}{2} = \frac{d_{\lambda, \nu} - 1}{2d_{\lambda, \nu}} \Pi_{\nu_{12}, \mu_1}^\lambda + \frac{d_{\lambda, \nu} + 1}{2d_{\lambda, \nu}} \Pi_{\nu_{12}, \mu'_1}^\lambda + \frac{\sqrt{d_{\lambda, \nu}^2 - 1}}{2d_{\lambda, \nu}} \Pi_{\nu_{12}, \mu'_1 \leftrightarrow \nu_{12}, \mu_1}^\lambda \quad (5.14)$$

and

$$\begin{aligned} \Pi_{\text{sgn}, \nu_{12} \leftarrow \text{id}, \nu_{12}}^\lambda &= \frac{2d_{\lambda, \nu}}{\sqrt{d_{\lambda, \nu}^2 - 1}} \Pi_{\text{sgn}, \nu_{12}}^\lambda \Pi_{\nu_{12}, \mu_1}^\lambda \Pi_{\text{id}, \nu_{12}}^\lambda \\ &= \frac{\sqrt{d_{\lambda, \nu}^2 - 1}}{2d_{\lambda, \nu}} \Pi_{\nu_{12}, \mu_1}^\lambda - \frac{\sqrt{d_{\lambda, \nu}^2 - 1}}{2d_{\lambda, \nu}} \Pi_{\nu_{12}, \mu'_1}^\lambda + \frac{d_{\lambda, \nu} + 1}{2d_{\lambda, \nu}} \Pi_{\nu_{12}, \mu_1 \leftarrow \nu_{12}, \mu'_1}^\lambda - \frac{d_{\lambda, \nu} - 1}{2d_{\lambda, \nu}} \Pi_{\nu_{12}, \mu'_1 \leftarrow \nu_{12}, \mu_1}^\lambda. \end{aligned} \quad (5.15)$$

5.4.4 Relations among irreps of $\mathbb{S}_{[4..n]} \times \mathbb{S}_\Sigma$ within an isotypical subspace

We are also interested to see how Δ_3 acts on $\Gamma_{1,2}$, which will require us to consider irreps of $\mathbb{S}_{[4..n]} \times \mathbb{S}_\Sigma$. Let us now consider $k = o(n)$, $\eta \vdash k$, and $\theta \subset \eta$. Recall that, according to our notation, $\bar{\eta} = (n-k, \eta) \vdash n$ and $\bar{\theta}_{123} = (n-k-2, \theta)_{123} \vdash n-3$ is obtained from $\bar{\eta}$ by removing two boxes in the first row and one box below the first row.

V contains three instances of the irrep $\mathcal{S}^{\bar{\theta}_{123}} \times \mathcal{S}^{\bar{\eta}}$ of $\mathbb{S}_{[4..n]} \times \mathbb{S}_\Sigma$: we have

$$\Pi_{\bar{\theta}_{123}}^{\bar{\eta}} = \Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}, (\bar{\eta}_1)}^{\bar{\eta}} + \Pi_{\bar{\theta}_{123}, \bar{\theta}_{12}, \bar{\eta}_1}^{\bar{\eta}} + \Pi_{\bar{\theta}_{123}, (\bar{\theta}_{12}), \bar{\theta}_1}^{\bar{\eta}} = \Pi_{\text{id}, \bar{\theta}_{123}, \bar{\eta}_3}^{\bar{\eta}} + \Pi_{\text{sgn}, \bar{\theta}_{123}, \bar{\eta}_3}^{\bar{\eta}} + \Pi_{(\text{id}), \bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}},$$

where each projector (other than $\Pi_{\bar{\theta}_{123}}^{\bar{\eta}}$) projects on a single instance of the irrep and the subscripts in parenthesis are optional. These two decompositions follow essentially the chain of restrictions $\mathbb{S}_{[n]} \rightarrow \mathbb{S}_{[2..n]} \rightarrow \mathbb{S}_{[3..n]} \rightarrow \mathbb{S}_{[4..n]}$ and $\mathbb{S}_{[n]} \rightarrow \mathbb{S}_{[n] \setminus \{3\}} \rightarrow \mathbb{S}_{\{1,2\}} \times \mathbb{S}_{[4..n]} \rightarrow \mathbb{S}_{[4..n]}$, respectively.

From the orthogonal form of the irrep $\bar{\eta}$, we get that the restriction of $V_{(12)}$ and $V_{(23)}$ to the isotypical subspace corresponding to $\mathcal{S}^{\bar{\theta}_{123}} \times \mathcal{S}^{\bar{\eta}}$ is, respectively,

$$\begin{aligned} V_{(12)} \Big|_{\bar{\theta}_{123} \times \bar{\eta}} &= \Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}}^{\bar{\eta}} + \frac{1}{d_{\bar{\eta}, \bar{\theta}_{12}}} \left(\Pi_{\bar{\theta}_{123}, \bar{\theta}_{12}, \bar{\eta}_1}^{\bar{\eta}} - \Pi_{\bar{\theta}_{123}, \bar{\theta}_1}^{\bar{\eta}} + \sqrt{d_{\bar{\eta}, \bar{\theta}_{12}}^2 - 1} \Pi_{\bar{\theta}_{123}, \bar{\theta}_{12}, \bar{\eta}_1 \leftrightarrow \bar{\theta}_{123}, \bar{\theta}_1}^{\bar{\eta}} \right), \\ V_{(23)} \Big|_{\bar{\theta}_{123} \times \bar{\eta}} &= \frac{1}{d_{\bar{\eta}, \bar{\theta}_{12}} - 1} \left(\Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}}^{\bar{\eta}} - \Pi_{\bar{\theta}_{123}, \bar{\theta}_{12}, \bar{\eta}_1}^{\bar{\eta}} + \sqrt{(d_{\bar{\eta}, \bar{\theta}_{12}} - 1)^2 - 1} \Pi_{\bar{\theta}_{123}, \bar{\eta}_{12} \leftrightarrow \bar{\theta}_{123}, \bar{\theta}_{12}, \bar{\eta}_1}^{\bar{\eta}} \right) + \Pi_{\bar{\theta}_{123}, \bar{\theta}_1}^{\bar{\eta}}, \end{aligned}$$

where the global phases of the transporters in the expression for $V_{(12)} \Big|_{\bar{\theta}_{123} \times \bar{\eta}}$ are consistent with (5.13). Therefore we can calculate the overlap of $\Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}}^{\bar{\eta}}$ and

$$\begin{aligned} \Pi_{\text{id}, \bar{\theta}_{123}, \bar{\eta}_3}^{\bar{\eta}} &= V_{(13)} (\mathbb{I} + V_{(23)}) \Pi_{\bar{\theta}_{123}, \bar{\eta}_1}^{\bar{\eta}} V_{(13)} / 2 \\ &= V_{(23)} V_{(12)} (\mathbb{I} + V_{(23)}) \left(\Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}}^{\bar{\eta}} + \Pi_{\bar{\theta}_{123}, \bar{\theta}_{12}, \bar{\eta}_1}^{\bar{\eta}} \right) V_{(12)} V_{(23)} / 2 \end{aligned}$$

to be

$$\frac{\text{Tr} \left(\Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}}^{\bar{\eta}} \Pi_{\text{id}, \bar{\theta}_{123}, \bar{\eta}_3}^{\bar{\eta}} \right)}{\dim \bar{\theta}_{123} \dim \bar{\eta}} = \frac{2}{d_{\bar{\eta}, \bar{\theta}_{12}} (d_{\bar{\eta}, \bar{\theta}_{12}} - 1)}. \quad (5.16)$$

Since $\Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}}^{\bar{\eta}} = \Pi_{\text{id}} \Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}}^{\bar{\eta}}$, we have

$$\begin{aligned} \Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}}^{\bar{\eta}} &= \Pi_{\bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}} + \frac{2}{d_{\bar{\eta}, \bar{\theta}_{12}}^2 - d_{\bar{\eta}, \bar{\theta}_{12}}} \left(\Pi_{\text{id}, \bar{\theta}_{123}, \bar{\eta}_3}^{\bar{\eta}} - \Pi_{\bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}} \right) \\ &\quad + \frac{\sqrt{2(d_{\bar{\eta}, \bar{\theta}_{12}}^2 - d_{\bar{\eta}, \bar{\theta}_{12}} - 2)}}{d_{\bar{\eta}, \bar{\theta}_{12}}^2 - d_{\bar{\eta}, \bar{\theta}_{12}}} \Pi_{\bar{\theta}_{123}, \bar{\theta}_3 \leftrightarrow \text{id}, \bar{\theta}_{123}, \bar{\eta}_3}^{\bar{\eta}}. \end{aligned} \quad (5.17)$$

5.4.5 Summing the permutations of $(\Delta_1 \circ \Gamma_{1,2})^* (\Delta_1 \circ \Gamma_{1,2})$

We will express $(\Delta_1 \circ \Gamma_{1,2})^* (\Delta_1 \circ \Gamma_{1,2})$ as a linear combination of projectors $\Pi_{\nu_{12}, \mu_1}^\lambda$ and transporters $\Pi_{\nu_{12}, \mu'_1 \leftarrow \nu_{12}, \mu_1}^\lambda$, where $\lambda \vdash n$, $\nu \subset_c \lambda$, and $\mu, \mu' \vdash n-1$ are such that $\nu \subset \mu \subset \lambda$ and $\nu \subset \mu' \subset \lambda$ (we consider transporters only if $\nu \subset_{rc} \lambda$, and thus $\mu \neq \mu'$). In order to calculate $\|\Delta_1 \circ \Gamma'\|$ via

(5.9), we use

$$\begin{aligned} \frac{1}{n-1} \sum_{\pi \in R'} V_\pi \Pi_{\nu_{12}, \mu_1}^\lambda V_{\pi^{-1}} &= \frac{1}{(n-1)!} \sum_{\pi \in \mathbb{S}_{[2..n]}} V_\pi \Pi_{\nu_{12}, \mu_1}^\lambda V_{\pi^{-1}} = \frac{\text{Tr}(V_\pi \Pi_{\nu_{12}, \mu_1}^\lambda V_{\pi^{-1}})}{\text{Tr}(\Pi_{\mu_1}^\lambda)} \Pi_{\mu_1}^\lambda = \frac{\dim \nu}{\dim \mu} \Pi_{\mu_1}^\lambda, \\ \frac{1}{n-1} \sum_{\pi \in R'} V_\pi \Pi_{\nu_{12}, \mu_1' \leftarrow \nu_{12}, \mu_1}^\lambda V_{\pi^{-1}} &= \frac{1}{(n-1)!} \sum_{\pi \in \mathbb{S}_{[2..n]}} V_\pi \Pi_{\nu_{12}, \mu_1' \leftarrow \nu_{12}, \mu_1}^\lambda V_{\pi^{-1}} = 0. \end{aligned} \tag{5.18}$$

The equalities in (5.18) hold because, first of all, $\Pi_{\nu_{12}, \mu_1}^\lambda$ and $\Pi_{\nu_{12}, \mu_1' \leftarrow \nu_{12}, \mu_1}^\lambda$ are fixed under $\mathbb{S}_{[3..n]} \times \mathbb{S}_\Sigma$. Second, V as a representation of $\mathbb{S}_{[2..n]} \times \mathbb{S}_\Sigma$ is multiplicity-free, and thus every operator on \mathcal{Y} that is fixed under $\mathbb{S}_{[2..n]} \times \mathbb{S}_\Sigma$ can be expressed as a linear combination of projectors $\Pi_{\mu_1}^{\lambda'}$, where $\lambda' \vdash n$ and $\mu'' \subset \lambda'$. And third, for $\pi \in \mathbb{S}_{[2..n]}$, V_π commutes with both $\Pi_{\mu_1}^\lambda$ and $\Pi_{\mu_1}^{\lambda'}$.

5.5 Construction of the optimal adversary matrix

In Section 5.3.3 we showed that

$$\beta_{\text{id}, \nu}^\lambda / \alpha_{\text{id}, \nu}^\lambda = \beta_{\text{sgn}, \nu}^\lambda / \alpha_{\text{sgn}, \nu}^\lambda = \sqrt{\binom{n}{2} \frac{\dim \nu}{\dim \lambda}}.$$

We calculate $\dim \nu$ and $\dim \lambda$ using the hook-length formula (1.8), and recall that, given a fixed $\zeta \vdash k$, $\dim \bar{\zeta}$ can be expressed as a polynomial in n of degree k and having the leading coefficient $1/h(\zeta)$ (see Table 1.1 for examples). Therefore we get that Claim 5.2 is equivalent to the following claim, which we prove in Appendix B.

Claim 5.4. *Suppose $\Gamma_{1,2}$ is given as in (5.7), $\alpha_{\text{id}, (n-2)}^{(n)} = n^{-1/3}$, and Γ is obtained from $\Gamma_{1,2}$ via (5.5). Consider $\lambda \vdash n$ that has $O(1)$ boxes below the first row and $\nu \subset_c \lambda$. In order for $\|\Delta_1 \circ \Gamma\| = O(1)$ to hold, we need to have*

1. $\alpha_{\text{id}, \nu}^\lambda = n^{-1/3} + O(1/n)$ if λ and ν are the same below the first row,
2. $\alpha_{\text{id}, \nu}^\lambda, \alpha_{\text{sgn}, \nu}^\lambda = n^{-1/3} + O(1/\sqrt{n})$ if λ has one box more below the first row than ν ,
3. $\alpha_{\text{id}, \nu}^\lambda, \alpha_{\text{sgn}, \nu}^\lambda = O(1)$ if λ has two boxes more below the first row than ν .

(Note that $\alpha_{\text{id}, (n-2)}^{(n)} = n^{-1/3}$ implies $\|\Gamma\| \geq \beta_{\text{id}, (n-2)}^{(n)} = \Theta(n^{2/3})$.)

Consider $k = o(n)$ and $\eta \vdash k$. Claims 5.2 and 5.4 hint that for the optimal adversary matrix we could choose coefficients $\alpha_{\text{id}, \bar{\eta}_{12}}^{\bar{\eta}} \approx \alpha_{\text{id}, \bar{\eta}_{12}}^{\bar{\zeta}} \approx \alpha_{\text{sgn}, \bar{\eta}_{12}}^{\bar{\zeta}}$ whenever $\zeta \supset \eta$ and $\alpha_{\text{id}, \bar{\eta}_{12}}^{\bar{\zeta}} = \alpha_{\text{sgn}, \bar{\eta}_{12}}^{\bar{\zeta}} = 0$ whenever $\zeta \supset \eta$. Let us do that. For $\zeta \supset \eta$, note that $\bar{\eta}_{12} \subset \bar{\eta}_1 \subset \bar{\zeta}$, $\bar{\eta}_{12} \subset \bar{\zeta}_1 \subset \bar{\zeta}$, and $\bar{\eta}_1$ appears after $\bar{\zeta}_1$ in the lexicographic order, and also note that $d_{\bar{\zeta}, \bar{\eta}_{12}} \geq n - 2k - 1$ (equality is achieved by $\eta = (k)$ and $\zeta = (k + 1)$). Therefore, according to (5.14) and (5.15), we have

$$\begin{aligned} \Pi_{\text{id}, \bar{\eta}_{12}}^{\bar{\eta}} + \sum_{\zeta \supset \eta} (\Pi_{\text{id}, \bar{\eta}_{12}}^{\bar{\zeta}} + \Pi_{\text{sgn}, \bar{\eta}_{12} \leftarrow \text{id}, \bar{\eta}_{12}}^{\bar{\zeta}}) &= \Pi_{\bar{\eta}_{12}}^{\bar{\eta}} + \sum_{\zeta \supset \eta} (\Pi_{\bar{\eta}_{12}, \bar{\eta}_1}^{\bar{\zeta}} + \Pi_{\bar{\eta}_{12}, \bar{\eta}_1 \leftarrow \bar{\eta}_{12}, \bar{\zeta}_1}^{\bar{\zeta}}) + O(1/n) \\ &= \Pi_{\bar{\eta}_{12}}^{\bar{\eta}} + \sum_{\zeta \supset \eta} 2\Pi_{\bar{\eta}_{12}, \bar{\eta}_1}^{\bar{\zeta}} \Pi_{\text{id}} + O(1/n) = 2\Pi_{\bar{\eta}_{12}, \bar{\eta}_1} \Pi_{\text{id}} - \Pi_{\bar{\eta}_{12}}^{\bar{\eta}} + O(1/n), \end{aligned}$$

where the last equality is due to $\Pi_{\bar{\eta}_{12}}^{\bar{\eta}} = \Pi_{\bar{\eta}_{12}, \bar{\eta}_1}^{\bar{\eta}} = \Pi_{\text{id}, \bar{\eta}_{12}}^{\bar{\eta}}$ and $\mathcal{S}^{\bar{\eta}_1} \uparrow \mathbb{S}_{[n]} \cong \mathcal{S}^{\bar{\eta}} \oplus \bigoplus_{\zeta \supset \eta} \mathcal{S}^{\bar{\zeta}}$, that is, the branching rule (1.9). Thus we choose to construct $\Gamma_{1,2}$ as a linear combination of matrices

$$2\Pi_{\bar{\eta}_{12}, \bar{\eta}_1} \Pi_{\text{id}} - \Pi_{\bar{\eta}_{12}}^{\bar{\eta}} = \Pi_{\bar{\eta}_{12}}^{\bar{\eta}} + \sum_{\zeta \supset \eta} \left(\frac{d_{\bar{\zeta}, \bar{\eta}_{12}} - 1}{d_{\bar{\zeta}, \bar{\eta}_{12}}} \Pi_{\text{id}, \bar{\eta}_{12}}^{\bar{\zeta}} + \frac{\sqrt{d_{\bar{\zeta}, \bar{\eta}_{12}}^2 - 1}}{d_{\bar{\zeta}, \bar{\eta}_{12}}} \Pi_{\text{sgn}, \bar{\eta}_{12} \leftarrow \text{id}, \bar{\eta}_{12}}^{\bar{\zeta}} \right).$$

(At first glance, it may seem that the matrix on the left hand side does not “treat” indices 1 and 2 equally, but that is an illusion due to the way we define the bijection f .)

Theorem 5.5. *Let Γ be constructed via (5.5) from*

$$\Gamma_{1,2} := \sum_{k=0}^{n^{2/3}} \frac{n^{2/3} - k}{n} \sum_{\eta \vdash k} (2\Pi_{\bar{\eta}_{12}, \bar{\eta}_1} \Pi_{\text{id}} - \Pi_{\bar{\eta}_{12}}^{\bar{\eta}}).$$

Then $\|\Gamma\| = \Omega(n^{2/3})$ and $\|\Delta_1 \circ \Gamma\| = O(1)$, and therefore Γ is, up to constant factors, an optimal adversary matrix for ELEMENT DISTINCTNESS.

For $\Gamma_{1,2}$ of Theorem 5.5 expressed in the form (5.7), we have $\alpha_{\text{id}, (n-2)}^{(n)} = n^{-1/3}$, and therefore $\|\Gamma\| = \Omega(n^{2/3})$. In the remainder of this section, let us prove $\|\Delta_1 \circ \Gamma'\| = O(1)$ and $\|\Delta_1 \circ \Gamma''\| = O(1)$, which is sufficient due to Claim 5.3.

5.5.1 Approximate action of Δ_i

The precise calculation of $\Delta_1 \circ \Gamma$ is tedious; we consider it in Appendix B. For that reason, as described in Section 2.3.2, we can use any valid approximation $\Delta_1 \diamond \Gamma$ instead. It suffices to show that $\|\Delta_1 \diamond \Gamma'\| = O(1)$ and $\|\Delta_1 \diamond \Gamma''\| = O(1)$ for any valid $\Delta_1 \diamond \Gamma'$ and $\Delta_1 \diamond \Gamma''$. That is, it suffices to show that we can change entries of Γ' and Γ'' corresponding to (x, y) with $x_1 = y_1$ in a way that the spectral norms of the resulting matrices are constantly bounded.

We will express $\Gamma_{1,2}$ as a linear combination of certain $n! \times n!$ matrices and, for every such matrix A , we will choose $\Delta_i \diamond A = A$, except for the following three, for which we calculate the action of Δ_1 or Δ_3 precisely. We have

$$\Delta_1 \circ \Pi_{\text{id}} = V_{(12)}/2, \quad \Delta_3 \circ \Pi_{\bar{\theta}_{123}, \bar{\theta}_3} = 0, \quad \text{and} \quad \Delta_3 \circ \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}} = 0$$

due to $\Delta_1 \circ \mathbb{I} = \Delta_3 \circ \mathbb{I} = 0$ and the commutativity relation (5.12).

Due to (5.12), we also have $\Delta_3 \circ (A\Pi_{\text{id}}) = (\Delta_3 \circ A)\Pi_{\text{id}}$ for every $n! \times n!$ matrix A . One can see that, given any choice of $\Delta_3 \diamond A$, we can choose $\Delta_3 \circ (A\Pi_{\text{id}}) = (\Delta_3 \diamond A)\Pi_{\text{id}}$.

5.5.2 Bounding $\|\Delta_1 \circ \Gamma'\|$

For $k \leq n^{2/3}$ and $\eta \vdash k$, define $n! \times n!$ matrices $(\Gamma_\eta)_{1,2}$ and $(\Gamma_k)_{1,2}$ such that

$$\Gamma_{1,2} = \sum_{k=0}^{n^{2/3}} \frac{n^{2/3} - k}{n} (\Gamma_k)_{1,2}, \quad (\Gamma_k)_{1,2} = \sum_{\eta \vdash k} (\Gamma_\eta)_{1,2}, \quad \text{and} \quad (\Gamma_\eta)_{1,2} = 2\Pi_{\bar{\eta}_{12}, \bar{\eta}_1} \Pi_{\text{id}} - \Pi_{\bar{\eta}_{12}}^{\bar{\eta}}.$$

The projector $\Pi_{\bar{\eta}_{12}, \bar{\eta}_1}$ commutes with the action of Δ_1 , therefore we can choose

$$\begin{aligned} \Delta_1 \diamond (\Gamma_\eta)_{1,2} &= 2\Pi_{\bar{\eta}_{12}, \bar{\eta}_1} (\Delta_1 \circ \Pi_{\text{id}}) - \Pi_{\bar{\eta}_{12}}^{\bar{\eta}} = \Pi_{\bar{\eta}_{12}, \bar{\eta}_1} V_{(12)} - \Pi_{\bar{\eta}_{12}}^{\bar{\eta}} \\ &= \sum_{\zeta \supset \eta} \Pi_{\bar{\eta}_{12}, \bar{\eta}_1}^{\bar{\zeta}} V_{(12)} = \sum_{\zeta \supset \eta} \left(-\frac{1}{d_{\bar{\zeta}, \bar{\eta}_{12}}} \Pi_{\bar{\eta}_{12}, \bar{\eta}_1}^{\bar{\zeta}} + \frac{\sqrt{d_{\bar{\zeta}, \bar{\eta}_{12}}^2 - 1}}{d_{\bar{\zeta}, \bar{\eta}_{12}}} \Pi_{\bar{\eta}_{12}, \bar{\eta}_1 \leftarrow \bar{\eta}_{12}, \bar{\zeta}_1}^{\bar{\zeta}} \right), \end{aligned}$$

where the third equality is due to the branching rule and both $\Pi_{\bar{\eta}_{12}}^{\bar{\eta}} = \Pi_{\bar{\eta}_{12}}^{\bar{\eta}} \Pi_{\text{id}}$ and $\Pi_{\text{id}} V_{(12)} = \Pi_{\text{id}}$, and the last equality comes from (5.13). To estimate the norm of $\Delta_1 \diamond \Gamma'$ via (5.9), we have

$$\begin{aligned} &\sum_{\pi \in R'} V_\pi (\Delta_1 \diamond (\Gamma_\eta)_{1,2})^* (\Delta_1 \diamond (\Gamma_\eta)_{1,2}) V_{\pi^{-1}} \\ &\preceq \sum_{\zeta \supset \eta} \sum_{\pi \in R'} V_\pi \left(\frac{1}{d_{\bar{\zeta}, \bar{\eta}_{12}}^2} \Pi_{\bar{\eta}_{12}, \bar{\eta}_1}^{\bar{\zeta}} + \Pi_{\bar{\eta}_{12}, \bar{\zeta}_1}^{\bar{\zeta}} - \frac{\sqrt{d_{\bar{\zeta}, \bar{\eta}_{12}}^2 - 1}}{d_{\bar{\zeta}, \bar{\eta}_{12}}^2} \Pi_{\bar{\eta}_{12}, \bar{\eta}_1 \leftrightarrow \bar{\eta}_{12}, \bar{\zeta}_1}^{\bar{\zeta}} \right) V_{\pi^{-1}} \\ &= (n-1) \sum_{\zeta \supset \eta} \left(\frac{1}{d_{\bar{\zeta}, \bar{\eta}_{12}}^2} \frac{\dim \bar{\eta}_{12}}{\dim \bar{\eta}_1} \Pi_{\bar{\eta}_1}^{\bar{\zeta}} + \frac{\dim \bar{\eta}_{12}}{\dim \bar{\zeta}_1} \Pi_{\bar{\zeta}_1}^{\bar{\zeta}} \right) \\ &\preceq \frac{1}{n - o(n)} \sum_{\zeta \supset \eta} \Pi_{\bar{\eta}_1}^{\bar{\zeta}} + (n-1) \sum_{\zeta \supset \eta} \frac{\dim \bar{\eta}_{12}}{\dim \bar{\zeta}_1} \Pi_{\bar{\zeta}_1}^{\bar{\zeta}}, \end{aligned} \tag{5.19}$$

where the equality in the middle comes from (5.18) and the last inequality is due to $\dim \bar{\eta}_{12} \leq \dim \bar{\eta}_1$ and $d_{\bar{\zeta}, \bar{\eta}_{12}} \geq n - 2k - 1$.

Claim 5.6. *Let $\zeta \vdash k$. Then $1 - \dim \bar{\zeta}_1 / \dim \bar{\zeta} \leq 2k/n$.*

Proof. Recall the hook-length formula (1.8). As ζ has $\zeta(1) \leq k$ columns, define $\zeta^\top(j) = 0$ for all $j \in [\zeta(1) + 1..k]$. We have

$$\dim \bar{\zeta} = \frac{n!}{h((n-k, \zeta))} = \frac{n!/(n-2k)!}{h(\zeta) \prod_{j=1}^k (n-k+1-j+\zeta^\top(j))}, \quad (5.20)$$

and therefore

$$1 - \frac{\dim \bar{\zeta}_1}{\dim \bar{\zeta}} = 1 - \frac{(n-1)!/(n-2k-1)!}{n!/(n-2k)!} \prod_{j=1}^k \frac{n-k+1-j+\zeta^\top(j)}{n-k-j+\zeta^\top(j)} < 1 - \frac{n-2k}{n} = \frac{2k}{n}.$$

□

For $\eta' \neq \eta$, we have $(\Delta_1 \diamond (\Gamma_{\eta'})_{1,2})^* (\Delta_1 \diamond (\Gamma_\eta)_{1,2}) = 0$, therefore, by summing (5.19) over all $\eta \vdash k$, we get

$$\begin{aligned} & \sum_{\pi \in R'} V_\pi (\Delta_1 \diamond (\Gamma_k)_{1,2})^* (\Delta_1 \diamond (\Gamma_k)_{1,2}) V_{\pi^{-1}} \\ & \preceq \frac{1}{n - o(n)} \sum_{\eta \vdash k} \sum_{\zeta \supset \eta} \Pi_{\bar{\eta}_1}^{\bar{\zeta}} + (n-1) \sum_{\zeta \vdash k+1} \sum_{\eta \subset \zeta} \frac{\dim \bar{\eta}_{12}}{\dim \bar{\zeta}_1} \Pi_{\bar{\zeta}_1}^{\bar{\zeta}} \\ & \preceq \frac{1}{n - o(n)} \sum_{\eta \vdash k} \sum_{\zeta \supset \eta} \Pi_{\bar{\eta}_1}^{\bar{\zeta}} + 2(k+1) \sum_{\zeta \vdash k+1} \Pi_{\bar{\zeta}_1}^{\bar{\zeta}}, \end{aligned} \quad (5.21)$$

where the first inequality holds because $\sum_{\eta \vdash k} \sum_{\zeta \supset \eta}$ and $\sum_{\zeta \vdash k+1} \sum_{\eta \subset \zeta}$ are sums over the same pairs of η and ζ , and the second inequality holds because $\dim \bar{\zeta}_1 = \dim \bar{\zeta}_{12} + \sum_{\eta \subset \zeta} \dim \bar{\eta}_{12}$ (due to the branching rule) and Claim 5.6.

Finally, by summing (5.21) over k , we get

$$\begin{aligned} (\Delta_1 \diamond \Gamma')^* (\Delta_1 \diamond \Gamma') &= \sum_{\pi \in R'} V_\pi (\Delta_1 \diamond \Gamma_{1,2})^* (\Delta_1 \diamond \Gamma_{1,2}) V_{\pi^{-1}} \\ &\preceq \sum_{k=0}^{n^{2/3}} \frac{(n^{2/3} - k)^2}{n^2} \left(\frac{1}{n - o(n)} \sum_{\eta \vdash k} \sum_{\zeta \supset \eta} \Pi_{\bar{\eta}_1}^{\bar{\zeta}} + 2(k+1) \sum_{\zeta \vdash k+1} \Pi_{\bar{\zeta}_1}^{\bar{\zeta}} \right) \preceq \mathbb{I}/3. \end{aligned} \quad (5.22)$$

Hence, $\|\Delta_1 \diamond \Gamma'\| = O(1)$. (Note: the norm of (5.21) is $\Theta(k)$ and, in (5.22), we essentially multiply it with \mathcal{T}^2/n^2 , where \mathcal{T} is the intended lower bound. This provides an intuition for why one cannot prove a lower bound higher than $\Omega(n^{2/3})$.)

5.5.3 Bounding $\|\Delta_1 \circ \Gamma''\|$

Let us decompose the adversary matrix as $\Gamma = 2\Gamma_{\mathcal{A}} - \Gamma_{\mathcal{B}}$, where we define $\Gamma_{\mathcal{A}}$ and $\Gamma_{\mathcal{B}}$ via their restriction to the rows labeled by $x \in \mathcal{D}_{1,2}$:

$$(\Gamma_{\mathcal{A}})_{1,2} := \sum_{k=0}^{n^{2/3}} \frac{n^{2/3} - k}{n} \sum_{\eta \vdash k} \Pi_{\bar{\eta}_{12}, \bar{\eta}_1} \Pi_{\text{id}} \quad \text{and} \quad (\Gamma_{\mathcal{B}})_{1,2} := \sum_{k=0}^{n^{2/3}} \frac{n^{2/3} - k}{n} \sum_{\eta \vdash k} \Pi_{\bar{\eta}_{12}},$$

respectively. We do not claim any connection between the matrices $\Gamma''_{\mathcal{A}}$ and $\Gamma''_{\mathcal{B}}$ here and the matrices with the same name in Section 4.1 for ELEMENT DISTINCTNESS with large range. Nevertheless, we use them exactly the same way: we show that $\|\Delta_1 \circ \Gamma''_{\mathcal{A}}\| = O(1)$ and $\|\Delta_1 \circ \Gamma''_{\mathcal{B}}\| = O(1)$, which together imply $\|\Delta_1 \circ \Gamma''\| = O(1)$. The argument is very similar for both $\Gamma_{\mathcal{A}}$ and $\Gamma_{\mathcal{B}}$, and let us start by showing $\|\Delta_1 \circ \Gamma''_{\mathcal{A}}\| = O(1)$.

We are interested to see how Δ_3 acts on $(\Gamma_{\mathcal{A}})_{1,2}$. Let $\theta \subset \eta$, and we will have to consider $\Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}, \bar{\eta}_1}$. For every $\lambda \supset \bar{\eta}_1$, note that $V_{(23)}$ and $\Pi_{\bar{\theta}_{123}, \bar{\eta}_1}^\lambda$ commute. So, similarly to (5.13), we have

$$V_{(23)} \Pi_{\bar{\theta}_{123}, \bar{\eta}_1} = \frac{1}{d_{\bar{\eta}_1, \bar{\theta}_{123}}} \sum_{\lambda \supset \bar{\eta}_1} \left(\Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}, \bar{\eta}_1}^\lambda - \Pi_{\bar{\theta}_{123}, \bar{\theta}_{12}, \bar{\eta}_1}^\lambda + \sqrt{d_{\bar{\eta}_1, \bar{\theta}_{123}}^2 - 1} \Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}, \bar{\eta}_1 \leftrightarrow \bar{\theta}_{123}, \bar{\theta}_{12}, \bar{\eta}_1}^\lambda \right).$$

Hence

$$\frac{\text{Tr} \left(\Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}, \bar{\eta}_1}^\lambda \Pi_{\bar{\theta}_{123}, \bar{\eta}_{13}, \bar{\eta}_1}^\lambda \right)}{\dim \bar{\theta}_{123} \dim \lambda} = \frac{\text{Tr} \left(\Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}, \bar{\eta}_1}^\lambda V_{(23)} \Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}, \bar{\eta}_1}^\lambda V_{(23)} \right)}{\dim \bar{\theta}_{123} \dim \lambda} = \frac{1}{d_{\bar{\eta}_1, \bar{\theta}_{123}}^2},$$

and therefore, similarly to (5.17), we have

$$\Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}, \bar{\eta}_1} = \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}, \bar{\eta}_1} + \frac{1}{d_{\bar{\eta}_1, \bar{\theta}_{123}}^2} \left(\Pi_{\bar{\theta}_{123}, \bar{\eta}_{13}, \bar{\eta}_1} - \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}, \bar{\eta}_1} \right) + \frac{\sqrt{d_{\bar{\eta}_1, \bar{\theta}_{123}}^2 - 1}}{d_{\bar{\eta}_1, \bar{\theta}_{123}}^2} \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}, \bar{\eta}_1 \leftrightarrow \bar{\theta}_{123}, \bar{\eta}_{13}, \bar{\eta}_1}, \quad (5.23)$$

where

$$\Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}, \bar{\eta}_1 \leftrightarrow \bar{\theta}_{123}, \bar{\eta}_{13}, \bar{\eta}_1} := \sum_{\lambda \supset \bar{\eta}_1} \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}, \bar{\eta}_1 \leftrightarrow \bar{\theta}_{123}, \bar{\eta}_{13}, \bar{\eta}_1}^\lambda$$

for short.

Without loss of generality, let us assume $n^{2/3}$ to be an integer. Then, by using the branching rule and simple derivations, one can see that

$$\sum_{k=0}^{n^{2/3}-1} \frac{n^{2/3} - k}{n} \sum_{\eta \vdash k} \left(\Pi_{\bar{\eta}_{123}, \bar{\eta}_1} + \sum_{\theta \subset \eta} \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}, \bar{\eta}_1} \right) = \sum_{k=0}^{n^{2/3}-1} \left(\frac{1}{n} \sum_{\eta \vdash k} \Pi_{\bar{\eta}_{123}, \bar{\eta}_1} + \frac{n^{2/3} - k}{n} \sum_{\theta \vdash k-1} \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}} \right). \quad (5.24)$$

Therefore we have

$$\begin{aligned}
(\Gamma_{\mathcal{A}})_{1,2} &= \sum_{k=0}^{n^{2/3}-1} \frac{n^{2/3}-k}{n} \sum_{\eta \vdash k} \left(\Pi_{\bar{\eta}_{123}, \bar{\eta}_1} + \sum_{\theta \subset \eta} \Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}, \bar{\eta}_1} \right) \Pi_{\text{id}} \\
&= \sum_{k=0}^{n^{2/3}-1} \left(\frac{1}{n} \sum_{\eta \vdash k} \Pi_{\bar{\eta}_{123}, \bar{\eta}_1} + \frac{n^{2/3}-k}{n} \sum_{\eta \vdash k} \sum_{\theta \subset \eta} \left(\frac{1}{d_{\bar{\eta}_1, \bar{\theta}_{123}}^2} (\Pi_{\bar{\theta}_{123}, \bar{\eta}_{13}, \bar{\eta}_1} - \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}, \bar{\eta}_1}) \right. \right. \\
&\quad \left. \left. + \frac{\sqrt{d_{\bar{\eta}_1, \bar{\theta}_{123}}^2} - 1}{d_{\bar{\eta}_1, \bar{\theta}_{123}}^2} \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}, \bar{\eta}_1 \leftrightarrow \bar{\theta}_{123}, \bar{\eta}_{13}, \bar{\eta}_1} \right) + \frac{n^{2/3}-k}{n} \sum_{\theta \vdash k-1} \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}} \right) \Pi_{\text{id}},
\end{aligned}$$

where the first equality comes from the branching rule and the fact that we can ignore $k = n^{2/3}$, and the second equality comes from subsequent applications of (5.23) and (5.24).

Recall that the action of Δ_3 commutes with Π_{id} and $\Delta_3 \circ \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}} = 0$. Therefore we can choose

$$\begin{aligned}
\Delta_3 \diamond (\Gamma_{\mathcal{A}})_{1,2} &= \sum_{k=0}^{n^{2/3}-1} \left(\frac{1}{n} \sum_{\eta \vdash k} \Pi_{\bar{\eta}_{123}, \bar{\eta}_1} + \frac{n^{2/3}-k}{n} \sum_{\eta \vdash k} \sum_{\theta \subset \eta} \left(\frac{1}{d_{\bar{\eta}_1, \bar{\theta}_{123}}^2} (\Pi_{\bar{\theta}_{123}, \bar{\eta}_{13}, \bar{\eta}_1} - \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}, \bar{\eta}_1}) \right. \right. \\
&\quad \left. \left. + \frac{\sqrt{d_{\bar{\eta}_1, \bar{\theta}_{123}}^2} - 1}{d_{\bar{\eta}_1, \bar{\theta}_{123}}^2} \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}, \bar{\eta}_1 \leftrightarrow \bar{\theta}_{123}, \bar{\eta}_{13}, \bar{\eta}_1} \right) \right) \Pi_{\text{id}},
\end{aligned}$$

and we have

$$\begin{aligned}
&(\Delta_3 \diamond (\Gamma_{\mathcal{A}})_{1,2})^* (\Delta_3 \diamond (\Gamma_{\mathcal{A}})_{1,2}) \\
&= \sum_{k=0}^{n^{2/3}-1} \Pi_{\text{id}} \left(\frac{1}{n^2} \sum_{\eta \vdash k} \Pi_{\bar{\eta}_{123}, \bar{\eta}_1} + \frac{(n^{2/3}-k)^2}{n^2} \sum_{\eta \vdash k} \sum_{\theta \subset \eta} \frac{1}{d_{\bar{\eta}_1, \bar{\theta}_{123}}^2} (\Pi_{\bar{\theta}_{123}, \bar{\eta}_{13}, \bar{\eta}_1} + \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}, \bar{\eta}_1}) \right) \Pi_{\text{id}}, \\
&\preceq \frac{1}{n^2} \sum_{k=0}^{n^{2/3}-1} \Pi_{\text{id}} \left(\sum_{\eta \vdash k} \Pi_{\bar{\eta}_{123}, \bar{\eta}_1} + o(1) \cdot \sum_{\eta \vdash k} \sum_{\theta \subset \eta} (\Pi_{\bar{\theta}_{123}, \bar{\eta}_{13}, \bar{\eta}_1} + \Pi_{\bar{\theta}_{123}, \bar{\theta}_{13}, \bar{\eta}_1}) \right) \Pi_{\text{id}} \preceq \frac{1}{n^2} \mathbb{I}.
\end{aligned}$$

Finally, (5.10) tells us that

$$\|\Delta_1 \diamond \Gamma_{\mathcal{A}}''\|^2 = \left\| \sum_{\pi \in R''} V_{\pi} (\Delta_3 \diamond (\Gamma_{\mathcal{A}})_{1,2})^* (\Delta_3 \diamond (\Gamma_{\mathcal{A}})_{1,2}) V_{\pi^{-1}} \right\| \leq \left\| \sum_{\pi \in R''} \frac{1}{n^2} \mathbb{I} \right\| \leq 1/2,$$

and, hence, $\|\Delta_1 \diamond \Gamma_{\mathcal{A}}''\| = O(1)$.

We show that $\|\Delta_1 \diamond \Gamma_{\mathcal{B}}''\| = O(1)$ in essentially the same way, except now, instead of the decomposition (5.23) of $\Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}, \bar{\eta}_1}$ we consider the decomposition (5.17) of $\Pi_{\bar{\theta}_{123}, \bar{\eta}_{12}}$. This concludes the proof that $\|\Delta_1 \diamond \Gamma''\| = O(1)$, which, in turn, concludes the proof of Theorem 5.5.

Chapter 6

Lower bound for the Enhanced Find-Two problem

In this chapter we introduce the ENHANCED FIND-TWO problem and we study its query complexity.

Definition 6.1. The input alphabet of the ENHANCED FIND-TWO problem is $\Sigma := \{0, 1\}$, and we say that an index $i \in [n]$ is *marked* if and only if $x_i = 1$. We are promised that exactly k indices of the input are marked. The ENHANCED FIND-TWO problem is to find two distinct marked indices using the following resources:

1. one copy of the uniform superposition over all marked indices,

$$\chi(x) := \frac{1}{\sqrt{k}} \sum_{i: x_i=1} i,$$

2. an oracle that reflects across this superposition,

$$\mathcal{O}_{(D)}(x) := \mathbb{I} - 2\chi(x)(\chi(x))^*,$$

3. and an oracle that tests if an index is marked,

$$\mathcal{O}_{(S)}(x) := \mathbb{I} - 2 \sum_{i: x_i=1} ii^*.$$

Hence, the domain of the ENHANCED FIND-TWO problem is the set of all $\binom{n}{k}$ n -bit strings of Hamming weight k ,

$$\mathcal{D} = \{x \in \{0, 1\}^n : |x| = k\},$$

and the codomain R is the set of all $\binom{n}{2}$ size-two subsets of $[n]$. In this chapter we prove that

Theorem 6.2. *The bounded-error quantum query complexity of the ENHANCED FIND-TWO problem is $\Theta(\min\{\sqrt{n/k}, \sqrt{k}\})$.*

Notice that, for $i \in [n]$, $\mathcal{O}_{(S)}(x)$ maps \mathbf{i} to $(-1)^{x_i} \mathbf{i}$. A very similar oracle, the one that maps $\mathbf{i} \otimes \mathbf{b}$ to $(-1)^{b x_i} \mathbf{i} \otimes \mathbf{b}$, where $b \in \{0, 1\}$, is obtained from the standard oracle by “sandwiching” it between two copies of the x -independent unitary $\mathbb{I}_{Q'} \otimes H_{Q''}$ where H is the *Hadamard operator*. $\mathcal{O}_{(S)}$ is known as the *membership oracle*. The operator $\mathcal{O}_{(D)}(x)$, reduced to the space $\text{span}\{\mathbf{i} : x_i = 1\}$, is essentially what is known as the *Grover diffusion operator*. Also note that both $\mathcal{O}_{(S)}(x)$ and $\mathcal{O}_{(D)}(x)$ are self inverses.

Note that ENHANCED FIND-TWO corresponds to the computational problem of FIND-TWO, but it does not fit the definition of a computational problem itself (see Definition 2.1). Because of this, we have to relax the definition of the quantum query algorithm (see Definition 2.10). Nonetheless, we will do it so that all the arguments from Section 2.2 on the symmetrization of algorithms carry over.

Relaxed definition of the query algorithm. For the ENHANCED FIND-TWO, we relax the definition of the quantum query algorithm as follows. First of all, in Section 2.2.1 we defined the query register \mathcal{X}_Q as a tensor product of the query index register $\mathcal{X}_{Q'} = \mathbb{C}^{[n]}$ and the query symbol register $\mathcal{X}_{Q''} = \mathbb{C}^\Sigma$. In contrast, in this chapter we completely ignore the $\mathcal{X}_{Q''}$ register, and we define the query register as $\mathcal{X}_Q := \mathbb{C}^{[n]}$. We assume that $\chi(x)$ belongs to \mathcal{X}_Q and both $\mathcal{O}_{(D)}(x)$ and $\mathcal{O}_{(S)}(x)$ act on \mathcal{X}_Q .

Second, we assume that the initial state of the algorithm is

$$\phi_0(x) := \chi(x) \otimes \chi',$$

where $\chi' \in \mathcal{X}_{WR}$ is independent from x and has unit norm. The state χ' is part of the specification of the algorithm.

Given multiple oracles, a natural question arises: Should there be a “control” register in an algorithm that determines which oracle to query and, as considered in Remark 2.11, whether to query at all? We choose not to have such a register, and we assume that it is predetermined—independently from the input and the computation so far—at what point in the algorithm which oracle will be queried. We explain why we can do that (namely, why our lower bounds would still hold if we assumed this control register) in Remark 6.9.

Upper bound. Let us first prove the upper bound in Theorem 6.2. We start by *measuring* $\chi(x)$ in the standard basis of \mathcal{X}_Q , thus obtaining one marked index i . Then we proceed in one of the following two ways, completely ignoring the oracle $\mathcal{O}_{(D)}$ or $\mathcal{O}_{(S)}$, respectively, and aiming to prepare $\chi(x)$.

1. *We search for a marked index.* We initially prepare \mathcal{X}_Q in the state $\frac{1}{\sqrt{n}} \sum_{i \in [n]} \mathbf{i}$, and then we run Grover's search algorithm (see [Gro96, BBHT98]) using $\mathbb{I}_n - 2\mathbb{J}_n/n$ as the diffusion operator and $\mathcal{O}_{(S)}(x)$ as the membership oracle. The algorithm obtains the state $\chi(x)$ using $O(\sqrt{n/k})$ queries.
2. *We "search" for the state $\chi(x)$ within the k -dimensional subspace corresponding to the marked indices.* Since we know i , we can implement $\mathcal{O}_{(i)} := \mathbb{I} - 2\mathbf{i}\mathbf{i}^*$ without any queries. We run Grover's search algorithm in reverse, starting \mathcal{X}_Q in the state \mathbf{i} , using $\mathcal{O}_{(D)}(x)$ as the diffusion operator and $\mathcal{O}_{(i)}$ as the membership oracle. After $O(\sqrt{k})$ iterations, the algorithm restores \mathcal{X}_Q in the state $\chi(x)$.

We measure $\chi(x)$ again and with probability $1 - 1/k$ we obtain a marked index $j \neq i$.

Now let us prove the lower bound in Theorem 6.2. We first present the framework of the proof in Section 6.1. And, in Section 6.2, we conclude with a proof of a technical lemma required for this lower bound.

6.1 Framework of the lower bound

Suppose \mathcal{A} is an arbitrary algorithm for ENHANCED FIND-TWO, and from it we construct the algorithm $\bar{\mathcal{A}}^+$ by introducing both the symmetrization register \mathcal{X}_ζ and the input register \mathcal{X}_i as described in Section 2.2.2 and Section 2.2.3, respectively. For the input register, we choose $\delta_x = \binom{n}{k}^{-1/2}$ for all $x \in \mathcal{D}$. Regarding the symmetrization register, we consider $\mathbb{S}_{[n]}$ as an oracle automorphism of the problem. Let $\mathbb{S}_{[n]}$ act on the sets \mathcal{D} , $[n]$, R as, respectively,

$$\pi: (x_1, \dots, x_n) \mapsto (x_{\pi^{-1}(1)}, \dots, x_{\pi^{-1}(n)}), \quad (6.1)$$

$$\pi: i \mapsto \pi(i), \quad (6.2)$$

$$\pi: \{i, j\} \mapsto \{\pi(i), \pi(j)\}. \quad (6.3)$$

The first two actions were already defined in Section 2.2.2, and we take the third as the group action ω required in Definition 2.12 defining an oracle automorphism. Note that both conditions (2.9) and (2.10) of Definition 2.12 are satisfied by $\mathcal{O}_{(S)}(x)$ and $\mathcal{O}_{(D)}(x)$ and the group actions above. Additionally, note that

$$(U_{Q,\gamma} \otimes \mathbb{I}_{WR})\phi_0(x) = \phi_0(\gamma(x))$$

for all $x \in \mathcal{D}$ and $\pi \in \mathbb{S}_{[n]}$, where, as in Section 2.2, U_I , U_Q , and U_R are the representations of $\mathbb{S}_{[n]}$ corresponding to the group actions (6.1), (6.2), and (6.3), respectively. Because of this, all symmetrization arguments from Section 2.2 apply here.

Also as before, we denote the (inner) tensor products of representations U_I , U_Q , and U_R by concatenating the corresponding subscripts. Note that there is a natural bijection

$$x \mapsto \{i: x_i = 1\} \quad (6.4)$$

between the set of inputs and the set of labels of the rows and columns of the Johnson scheme (see Section 1.5.2). Due to this bijection, (1.18) states that \mathcal{X}_1 decomposes into irreps as

$$\mathcal{X}_1 = \bigoplus_{i=0}^k \mathcal{X}_1^{(n-i,i)}, \quad \text{where } \mathcal{X}_1^{(n-i,i)} \cong \mathcal{S}^{(n-i,i)}. \quad (6.5)$$

Let $\Pi_1^{(n-i,i)}$ denote the projector on the irrep $\mathcal{X}_1^{(n-i,i)}$. In the standard basis of $\mathcal{X}_1 = \mathbb{C}^{\mathcal{D}}$, $\Pi_1^{(n-i,i)}$ here is the same as Π_i in Section 1.5.2. Let

$$\mathcal{X}_{1,a} := \mathcal{X}_1^{(n)} \oplus \mathcal{X}_1^{(n-1,1)} \quad \text{and} \quad \mathcal{X}_{1,b} := \mathcal{X}_1 \ominus \mathcal{X}_{1,a} = \bigoplus_{i=2}^k \mathcal{X}_1^{(n-i,i)},$$

and let $\Pi_{1,a}$ and $\Pi_{1,b}$ be the projectors on these spaces, respectively.

As in Section 2.2.4, for $t \in [0..T-1]$, let $\bar{\psi}_t^+$ be the state of the algorithm $\bar{\mathcal{A}}^+$ just before the query $t+1$, let $\bar{\psi}_T^+$ be the final state of $\bar{\mathcal{A}}^+$, and let

$$\rho'_t = \text{Tr}_{\text{SWR}}(\bar{\psi}_t^+(\bar{\psi}_t^+)^*), \quad \rho'_T = \text{Tr}_{\text{SQW}}(\bar{\psi}_T^+(\bar{\psi}_T^+)^*), \quad \text{Tr}_Q(\rho'_t) = \rho_t \quad \text{and} \quad \text{Tr}_R(\rho'_T) = \rho_T.$$

Also recall the symmetries (2.19), (2.21), (2.20), (2.22) of these states, where now we have $G = \mathbb{S}_{[n]}$.

For $t \in [0..T]$, let

$$r_{a,t} := \text{Tr}(\rho_t \Pi_{1,a}) \quad \text{and} \quad r_{b,t} := 1 - r_{a,t} = \text{Tr}(\rho_t \Pi_{1,b}).$$

The probability $r_{b,t}$, in some sense, measures the entanglement between \mathcal{X}_1 and \mathcal{X}_{SA} . As we described in Section 2.3.1, the adversary bound is based the observations that a successful algorithm has to establish a certain amount of entanglement between \mathcal{X}_1 and the other registers and that a single query cannot create too much entanglement. We proceed similarly here, and the lower bound in Theorem 6.2 follows from the following three lemmas.

Lemma 6.3. *(At the very beginning of the algorithm) we have $r_{b,0} = 0$.*

Lemma 6.4. *The success probability of the algorithm is at most $\frac{2(k-1)}{n-1} + \sqrt{2r_{b,T}}$.*

Lemma 6.5. *For all $t \in \{0, \dots, T-1\}$, we have $|r_{b,t} - r_{b,t+1}| = O(\max\{\sqrt{k/n}, \sqrt{1/k}\})$.*

6.1.1 Proof of Lemma 6.3

Recall from Section 1.5.2 the operator C_1 of the Johnson scheme, and, via the bijection (6.4), we can think of C_1 being in $L(\mathcal{X}_1)$. From (1.22) we get that the support of C_1 equals $\mathcal{X}_{1,a}$, and from (1.19) and (1.20) we get that this support is spanned by states

$$\zeta_i := \frac{1}{\sqrt{\binom{n-1}{k-1}}} \sum_{\substack{x \in \mathcal{D} \\ x_i=1}} \mathbf{x}, \quad (6.6)$$

where $i \in [n]$.

Now, the initial state of $\bar{\mathcal{A}}^+$ is

$$\bar{\phi}_0^+ := \binom{n}{k}^{-1/2} \sum_{x \in \mathcal{D}} \mathbf{x}_1 \otimes (n!)^{-1/2} \sum_{\pi \in \mathbb{S}_{[n]}} \pi_S \otimes \chi(x)_Q \otimes \chi'_{WR},$$

and we have $\rho_0 = \text{Tr}_{\text{SA}}(\bar{\phi}_0^+ (\bar{\phi}_0^+)^*)$. Note that the registers \mathcal{X}_1 and \mathcal{X}_Q are not entangled with the rest of the registers. By swapping the order of summation, we have

$$\sum_{x \in \mathcal{D}} \mathbf{x}_1 \otimes \chi(x)_Q \propto \sum_{\substack{x \in \mathcal{D}, i \in [n] \\ x_i=1}} \mathbf{x}_1 \otimes \mathbf{i}_Q \propto \sum_{i \in [n]} (\zeta_i)_1 \otimes \mathbf{i}_Q.$$

Since $\Pi_{1,a} \zeta_i = \zeta_i$ for all i , we get $(\Pi_{1,a} \otimes \mathbb{I}_{\text{SA}}) \bar{\phi}_0^+ = \bar{\phi}_0^+$. This, in turn, implies that $(\Pi_{1,b} \otimes \mathbb{I}_{\text{SA}}) \bar{\phi}_0^+ = 0$, and, thus, Lemma 6.3 holds.

6.1.2 Proof of Lemma 6.4

From (2.17), we get that the success probability of $\bar{\mathcal{A}}^+$ is

$$\sum_{\substack{x \in \mathcal{D}, \{i,j\} \subset [n] \\ x_i=x_j=1}} \text{Tr}(\rho'_T(\mathbf{x}\mathbf{x}_1^* \otimes \{\mathbf{i}, \mathbf{j}\} \{\mathbf{i}, \mathbf{j}\}_R^*)) = \left\| \sum_{\substack{x \in \mathcal{D}, \{i,j\} \subset [n] \\ x_i=x_j=1}} (\mathbf{x}\mathbf{x}_1^* \otimes \{\mathbf{i}, \mathbf{j}\} \{\mathbf{i}, \mathbf{j}\}_R^* \otimes \mathbb{I}_{\text{SQW}}) \bar{\psi}_T^+ \right\|^2. \quad (6.7)$$

Let us first reduce the lemma to its special case when $r_{T,b} = 0$. This reduction was used in [Amb10] for a very similar problem. Recall that the final state of the algorithm $\bar{\psi}_T^+$ satisfies the symmetry $(U_{\text{ISR},\pi} \otimes \mathbb{I}_{\text{QW}}) \bar{\psi}_T^+ = \bar{\psi}_T^+$ for all $\pi \in \mathbb{S}_{[n]}$, and note that, for $c \in \{a, b\}$, the state

$$\bar{\psi}_{T,c}^+ := \frac{(\Pi_{1,c} \otimes \mathbb{I}_{\text{SA}}) \bar{\psi}_T^+}{\|(\Pi_{1,c} \otimes \mathbb{I}_{\text{SA}}) \bar{\psi}_T^+\|} = \frac{1}{\sqrt{r_{c,T}}} (\Pi_{1,c} \otimes \mathbb{I}_{\text{SA}}) \bar{\psi}_T^+$$

satisfies the same symmetry. We have

$$\bar{\psi}_T^+ = \sqrt{1 - r_{b,T}} \bar{\psi}_{T,a}^+ + \sqrt{r_{b,T}} \bar{\psi}_{T,b}^+.$$

Since $\bar{\psi}_{T,a}^+$ and $\bar{\psi}_{T,b}^+$ are orthogonal, we have

$$\|\bar{\psi}_T^+ - \bar{\psi}_{T,a}^+\| = \sqrt{(1 - \sqrt{1 - r_{b,T}})^2 + (\sqrt{r_{b,T}})^2} \leq \sqrt{2r_{b,T}} \quad (6.8)$$

Lemma 6.6. *Suppose \mathcal{X} is a finite Hilbert space. For any two states $\psi_0 \in \mathcal{X}$ and $\psi_1 \in \mathcal{X}$ and any projector $\Pi \in \mathbf{L}(\mathcal{X})$, $\|\Pi\psi_0\|^2 - \|\Pi\psi_1\|^2 \leq \|\psi_0 - \psi_1\|$.*

Proof. Let $\mu = \|\psi_0 - \psi_1\|/2 \leq 1$. There is an orthonormal basis $\{\xi_0, \xi_1, \dots\}$ of \mathcal{X} such that

$$\psi_0 = \sqrt{1 - \mu^2} \xi_0 + \mu \xi_1 \quad \text{and} \quad \psi_1 = \sqrt{1 - \mu^2} \xi_0 - \mu \xi_1.$$

Note that

$$\|\Pi\psi_0\|^2 - \|\Pi\psi_1\|^2 = \text{Tr}(\Pi(\psi_0\psi_0^* - \psi_1\psi_1^*)) = 2\mu\sqrt{1 - \mu^2} \text{Tr}(\Pi(\xi_0\xi_1^* + \xi_1\xi_0^*)). \quad (6.9)$$

Since $(\xi_0\xi_0^* + \xi_1\xi_1^*)\Pi(\xi_0\xi_0^* + \xi_1\xi_1^*)$ is positive semidefinite of spectral norm at most 1, $|\text{Tr}(\Pi\xi_0\xi_1^*)| \leq 1/2$. Hence, (6.9) is at most $2\mu\sqrt{1 - \mu^2} \leq 2\mu$. \square

From now on, let us assume that $r_{b,T} = 0$ and, thus, $\bar{\psi}_T^+ = \bar{\psi}_{T,a}^+$. Lemma 6.6, (6.7), and (6.8) states that this changes the success probability by at most $\sqrt{2r_{b,T}}$.

Due to the symmetry (2.22) of ρ'_T , we can rewrite the success probability (6.7) as $\text{Tr}(\hat{\Pi}\hat{\rho})$, where $\hat{\Pi}$ is the projector on the subspace of \mathcal{X}_1 spanned by all \mathbf{x} such that $x_1 = x_2 = 1$ and

$$\hat{\rho} := \binom{n}{2} (\mathbb{I}_1 \otimes \{\mathbf{1}, \mathbf{2}\}_R^*) \rho'_T (\mathbb{I}_1 \otimes \{\mathbf{1}, \mathbf{2}\}_R)$$

is a density operator on the register \mathcal{X}_1 . It is left to show that

Claim 6.7. $\text{Tr}(\hat{\Pi}\hat{\rho}) \leq 2(k-1)/(n-1)$.

Proof. Let $\hat{\mathbb{S}} := \mathbb{S}_{\{1,2\}} \times \mathbb{S}_{\{3,\dots,n\}} < \mathbb{S}_{[n]}$ be the group of all permutations $\pi \in \mathbb{S}_{[n]}$ that map $\{1, 2\}$ to itself. We have

$$\Pi_{1,a}\hat{\rho} = \hat{\rho} \quad \text{and} \quad \forall \pi \in \hat{\mathbb{S}}: U_{1,\pi}\hat{\rho}U_{1,\pi}^{-1} = \hat{\rho}. \quad (6.10)$$

We also have $U_{1,\pi}\hat{\Pi}U_{1,\pi}^{-1} = \hat{\Pi}$ for all $\pi \in \hat{\mathbb{S}}$.

We can express $\hat{\rho}$ as a mixture of its eigenvectors ξ_i , with probabilities that are equal to their eigenvalues λ_i , namely, $\hat{\rho} = \sum_i \lambda_i \xi_i \xi_i^*$. Hence we have

$$\text{Tr}(\hat{\Pi}\hat{\rho}) = \sum_i \lambda_i \text{Tr}(\hat{\Pi}\xi_i \xi_i^*) = \sum_i \lambda_i \|\hat{\Pi}\xi_i\|^2,$$

which is at most

$$\max_{\xi} (\|\hat{\Pi}\xi\|^2 / \|\xi\|^2)$$

where the maximization is over all eigenvectors of $\hat{\rho}$ with non-zero eigenvalues. Due to the symmetry (6.10), we can calculate the eigenspaces of $\hat{\rho}$ by inspecting the reduction of U_1 to the subspace $\mathcal{X}_{1,a}$, namely, $\hat{U}_1 := \Pi_{1,a}U_1$. Recall that $\mathcal{X}_{1,a}$ is the space spanned by all ζ_i (defined in (6.6)). Note that $\zeta_i^*\zeta_j = \frac{k-1}{n-1}$ for all $i, j: i \neq j$.

\hat{U}_1 is a representation of both $\mathbb{S}_{[n]}$ and its subgroup $\hat{\mathbb{S}}$, and, as a representation of $\mathbb{S}_{[n]}$, it consists of two irreps: one-dimensional $\mathcal{X}_1^{(n)}$ and $(n-1)$ -dimensional $\mathcal{X}_1^{(n-1,1)}$. In order to see how \hat{U}_1 decomposes into irreps of $\hat{\mathbb{S}}$, we need to restrict $\mathcal{S}^{(n)}$ and $\mathcal{S}^{(n-1,1)}$ from \mathbb{S}_n to $\mathbb{S}_2 \times \mathbb{S}_{n-2}$. The Littlewood–Richardson rule (1.10) gives us the decomposition of these restrictions:

$$\begin{aligned} \mathcal{S}^{(n)} \downarrow (\mathbb{S}_2 \times \mathbb{S}_{n-2}) &\cong (\mathcal{S}^{(2)} \times \mathcal{S}^{(n-2)}); \\ \mathcal{S}^{(n-1,1)} \downarrow (\mathbb{S}_2 \times \mathbb{S}_{n-2}) &\cong (\mathcal{S}^{(2)} \times \mathcal{S}^{(n-2)}) \oplus (\mathcal{S}^{(1,1)} \times \mathcal{S}^{(n-2)}) \oplus (\mathcal{S}^{(2)} \times \mathcal{S}^{(n-3,1)}). \end{aligned}$$

Hence, Schur’s lemma (Lemma 1.2) and (6.10) imply that the eigenspaces of $\hat{\rho}$ are invariant under $U_{1,\pi}$ for all $\pi \in \hat{\mathbb{S}}$, and they have one of the following forms:

1. one-dimensional subspace spanned by $\zeta_{(\alpha,\beta)} := \alpha(\zeta_1 + \zeta_2) + \beta \sum_{i=3}^n \zeta_i$ for some coefficients α, β ;
2. one-dimensional subspace spanned by $\zeta_1 - \zeta_2$;
3. $(n-3)$ -dimensional subspace consisting of all $\sum_{i=3}^n \alpha_i \zeta_i$ with $\sum_i \alpha_i = 0$ (spanned by all $\zeta_i - \zeta_j, i, j \in [3..n]$);
4. a direct sum of subspaces of the above form.

In the first case,

$$\hat{\Pi}\zeta_{(\alpha,\beta)} = \frac{2\alpha + (k-2)\beta}{\sqrt{\binom{n-1}{k-1}}} \sum_{\substack{x_3, \dots, x_n \in \{0,1\} \\ x_3 + \dots + x_n = k-2}} (\mathbf{1}, \mathbf{1}, \mathbf{x}_3, \dots, \mathbf{x}_n).$$

Therefore,

$$\|\hat{\Pi}\zeta_{(\alpha,\beta)}\|^2 = \frac{\binom{n-2}{k-2}}{\binom{n-1}{k-1}} |2\alpha + (k-2)\beta|^2 = \frac{k-1}{n-1} |2\alpha + (k-2)\beta|^2.$$

We also have

$$\begin{aligned} \|\zeta_{(\alpha,\beta)}\|^2 &= \zeta_{(\alpha,\beta)}^* \zeta_{(\alpha,\beta)} \\ &= 2 \left(1 + \frac{k-1}{n-1}\right) |\alpha|^2 + (n-2) \left(1 + (n-3) \frac{k-1}{n-1}\right) |\beta|^2 + 2(n-2) \frac{k-1}{n-1} (\alpha\bar{\beta} + \beta\bar{\alpha}) \\ &\geq \frac{|2\alpha + (k-2)\beta|^2}{2}. \end{aligned} \quad (6.11)$$

If $\alpha\bar{\beta} \geq 0$, the inequality in (6.11) follows by showing that coefficients of $|\alpha|^2$, $|\beta|^2$, and $\alpha\bar{\beta}$ on the left hand side are all larger than corresponding coefficients on the right hand side. Otherwise, without loss of generality, we can assume that $\alpha = 1$ and $\beta < 0$, and the inequality follows by inspecting the extreme point of the quadratic polynomial (in β) that is obtained by subtracting the right hand side from the left hand side. Therefore,

$$\frac{\|\hat{\Pi}\zeta_{(\alpha,\beta)}\|^2}{\|\zeta_{(\alpha,\beta)}\|^2} \leq \frac{2(k-1)}{n-1}.$$

In the second case, $\hat{\Pi}(\zeta_1 - \zeta_2) = 0$ because basis states $(\mathbf{1}, \mathbf{1}, \mathbf{x}_3, \dots, \mathbf{x}_n)$ have the same amplitude in ζ_1 and ζ_2 .

In the third case, it suffices to consider a state of the form $\zeta_3 - \zeta_4$, because $\{U_{1,\pi}(\zeta_3 - \zeta_4) : \pi \in \hat{\mathbb{S}}\}$ spans the whole eigenspace and $\hat{\Pi}$ and $U_{1,\pi}$ commute. Then,

$$\hat{\Pi}(\zeta_3 - \zeta_4) = \frac{1}{\sqrt{\binom{n-1}{k-1}}} \sum_{\substack{x_5, \dots, x_n \in \{0,1\} \\ x_5 + \dots + x_n = k-3}} ((\mathbf{1}, \mathbf{1}, \mathbf{1}, \mathbf{0}, \mathbf{x}_5, \dots, \mathbf{x}_n) - (\mathbf{1}, \mathbf{1}, \mathbf{0}, \mathbf{1}, \mathbf{x}_5, \dots, \mathbf{x}_n))$$

and

$$\|\hat{\Pi}(\zeta_3 - \zeta_4)\|^2 = 2 \frac{\binom{n-4}{k-3}}{\binom{n-1}{k-1}} = 2 \frac{(k-1)(k-2)(n-k)}{(n-1)(n-2)(n-3)}.$$

We also have

$$\|\zeta_3 - \zeta_4\|^2 = 2 - 2\zeta_3^* \zeta_4 = 2 - 2 \frac{k-1}{n-1} = 2 \frac{n-k}{n-1}.$$

Hence,

$$\frac{\|\hat{\Pi}(\zeta_3 - \zeta_4)\|^2}{\|\zeta_3 - \zeta_4\|^2} = \frac{(k-2)(k-3)}{(n-2)(n-3)} = \mathcal{O}\left(\frac{k^2}{n^2}\right).$$

□

6.2 Proof of Lemma 6.5

Let $\Pi_{\mathcal{I}\mathcal{Q},a}$ and $\Pi_{\mathcal{I}\mathcal{Q},b}$ be the projectors on $\mathcal{X}_{\mathcal{I}\mathcal{Q},a} := \mathcal{X}_{\mathcal{I},a} \otimes \mathcal{X}_{\mathcal{Q}}$ and $\mathcal{X}_{\mathcal{I}\mathcal{Q},b} := \mathcal{X}_{\mathcal{I},b} \otimes \mathcal{X}_{\mathcal{Q}}$, respectively. Suppose $\mathcal{O} = \mathcal{O}_{(S)}$ or $\mathcal{O} = \mathcal{O}_{(D)}$. Recall from Section 2.2.3 that, when we introduce the input register $\mathcal{X}_{\mathcal{I}}$, we replace the oracle $\mathcal{O}(x)$ acting on $\mathcal{X}_{\mathcal{Q}}$ by \mathcal{O}^+ acting on $\mathcal{X}_{\mathcal{I}\mathcal{Q}}$ as in (2.14). The state of the $\mathcal{X}_{\mathcal{I}}$ register can be affected only by oracle queries, therefore we have $\rho_{t+1} = \text{Tr}_{\mathcal{Q}}(\mathcal{O}^+ \rho'_t \mathcal{O}^+)$ and

$$r_{b,t} - r_{b,t+1} = \text{Tr}(\Pi_{\mathcal{I}\mathcal{Q},b}(\rho'_t - \mathcal{O}^+ \rho'_t \mathcal{O}^+)) \quad (6.12)$$

for all $t \in [0..T - 1]$. Since $(\pi(x))_{\pi(i)} = 1$ if and only if $x_i = 1$, we have $U_{\mathbb{I}\mathbb{Q},\pi}\mathcal{O}^+U_{\mathbb{I}\mathbb{Q},\pi}^{-1} = \mathcal{O}^+$ for all $\pi \in \mathbb{S}_{[n]}$, and recall that the same symmetry holds for ρ'_t , namely, (2.20).

So it suffices to prove that

$$\left| \text{Tr}(\Pi_{\mathbb{I}\mathbb{Q},b}(\rho' - \mathcal{O}^+\rho'\mathcal{O}^+)) \right| \leq O(\max\{\sqrt{k/n}, \sqrt{1/k}\})$$

for every density operator ρ' on $\mathcal{X}_{\mathbb{I}\mathbb{Q}}$ that satisfies $U_{\mathbb{I}\mathbb{Q},\pi}\rho'U_{\mathbb{I}\mathbb{Q},\pi}^{-1} = \rho'$ for all $\pi \in \mathbb{S}_{[n]}$ and both oracles $\mathcal{O} = \mathcal{O}_{(S)}$ and $\mathcal{O} = \mathcal{O}_{(D)}$. Due to Schur's lemma, there is a spectral decomposition

$$\rho' = \sum_{\mu} \lambda_{\mu} \frac{\Pi_{\mu}}{\dim \mu},$$

where $\sum_{\mu} \lambda_{\mu} = 1$, every μ is an irrep of $\mathbb{S}_{[n]}$, and $\Pi_{\mu} \in \mathbb{L}(\mathcal{X}_{\mathbb{I}\mathbb{Q}})$ is the projector on μ . Because of the linearity, it suffices to show the following.

Lemma 6.8. *For every irrep $\mu \subset \mathcal{X}_{\mathbb{I}\mathbb{Q}}$ and for μ' being the subspace that μ is mapped to by $\mathcal{O}_{(S)}^+$ or $\mathcal{O}_{(D)}^+$, we have*

$$\frac{1}{\dim \mu} \left| \text{Tr}(\Pi_{\mathbb{I}\mathbb{Q},b}(\Pi_{\mu} - \Pi_{\mu'})) \right| \leq O(\max\{\sqrt{k/n}, \sqrt{1/k}\}). \quad (6.13)$$

Remark 6.9. The lower bound on the query complexity of the ENHANCED FIND-TWO problem still holds if we allow an algorithm to have a register $\mathcal{X}_{\mathbb{C}} := \mathbb{C}^{\{S,D,0\}}$ that controls if to query the oracle $\mathcal{O}_{(S)}$ (the register is in the state \mathbf{S}), the oracle $\mathcal{O}_{(D)}$ (the state \mathbf{D}), or not to query at all (the state $\mathbf{0}$). With such a register, instead of $\mathcal{O}_{(S)}^+$ and $\mathcal{O}_{(D)}^+$ acting on $\mathcal{X}_{\mathbb{I}\mathbb{Q}}$, we have to consider

$$\mathcal{O}^{++} := \sum_{\gamma \in \{S,D\}} \mathcal{O}_{(\gamma)}^+ \otimes (\gamma\gamma^*)_{\mathbb{C}} + \mathbb{I}_{\mathbb{I}\mathbb{Q}} \otimes (\mathbf{0}\mathbf{0}^*)_{\mathbb{C}}$$

acting on $\mathcal{X}_{\mathbb{I}\mathbb{Q}\mathbb{C}}$. We have $\rho'_t = \text{Tr}_{\mathbb{C}}(\sigma_t)$, where σ_t is the state of the $\mathcal{X}_{\mathbb{I}\mathbb{Q}\mathbb{C}}$ registers of the algorithm just before the query $t + 1$. Using the same argument as in Section 2.2.4, we can write

$$\sigma_t = \sum_{\gamma_1, \gamma_2 \in \{S,D,0\}} (\sigma_{\gamma_1, \gamma_2})_{\mathbb{I}\mathbb{Q}} \otimes (\gamma_1\gamma_2^*)_{\mathbb{C}},$$

where each $\sigma_{\gamma_1, \gamma_2}$ satisfies $U_{\mathbb{I}\mathbb{Q},\pi}\sigma_{\gamma_1, \gamma_2}U_{\mathbb{I}\mathbb{Q},\pi}^{-1} = \sigma_{\gamma_1, \gamma_2}$ for all $\pi \in \mathbb{S}_{[n]}$. Hence, (6.12) equals

$$\text{Tr}((\Pi_{\mathbb{I}\mathbb{Q},b} \otimes \mathbb{I}_{\mathbb{C}})(\sigma_t - \mathcal{O}^{++}\sigma_t\mathcal{O}^{++})) = \sum_{\gamma \in \{S,D\}} \text{Tr}(\Pi_{\mathbb{I}\mathbb{Q},b}(\sigma_{\gamma, \gamma} - \mathcal{O}_{(\gamma)}^+\sigma_{\gamma, \gamma}\mathcal{O}_{(\gamma)}^+)).$$

Since $\text{Tr}(\sigma_{S,S} + \sigma_{D,D}) \leq 1$, this reduces the argument of the lower bound to the case when we do not have the $\mathcal{X}_{\mathbb{C}}$ register.

6.2.1 Decomposition of $\mathcal{X}_{\mathbb{Q}}$ into irreps

In order to prove Lemma 6.8, we need to inspect the representation $\mathcal{X}_{\mathbb{Q}}$ in more detail. Let us consider two approaches how to decompose $\mathcal{X}_{\mathbb{Q}}$ into irreps, more precisely, how to decompose each isotypical subspace of $\mathcal{X}_{\mathbb{Q}}$ into irreps. For an irrep \mathcal{S}^θ present in $\mathcal{X}_{\mathbb{Q}}$, where $\theta \vdash n$, let $\hat{\Pi}_\theta$ be the projector on the \mathcal{S}^θ -isotypical subspace of $\mathcal{X}_{\mathbb{Q}}$. In the standard basis of $\mathcal{X}_{\mathbb{Q}}$, all entries of $\hat{\Pi}_\theta$ are real (see Claim 2.18).

Approach 1: via the tensor product of irreps. In (6.5) we already considered how $\mathcal{X}_{\mathbb{I}}$ decomposes into irreps. By the same argument, the natural representation $\mathcal{X}_{\mathbb{Q}}$ decomposes into irreps as

$$\mathcal{X}_{\mathbb{Q}} = \mathcal{X}_{\mathbb{Q}}^{(n)} \oplus \mathcal{X}_{\mathbb{Q}}^{(n-1,1)}, \quad \text{where } \mathcal{X}_{\mathbb{Q}}^{(n)} \cong \mathcal{S}^{(n)}, \mathcal{X}_{\mathbb{Q}}^{(n-1,1)} \cong \mathcal{S}^{(n-1,1)}. \quad (6.14)$$

Let $\Pi_{\mathbb{Q}}^{(n)}$ and $\Pi_{\mathbb{Q}}^{(n-1,1)}$ denote, respectively, the projectors on $\mathcal{X}_{\mathbb{Q}}^{(n)}$ and $\mathcal{X}_{\mathbb{Q}}^{(n-1,1)}$.

By taking these decompositions of $\mathcal{X}_{\mathbb{I}}$ and $\mathcal{X}_{\mathbb{Q}}$, we can decompose $\mathcal{X}_{\mathbb{I}\mathbb{Q}}$ into irreps by decomposing $\mathcal{X}_{\mathbb{I}}^{(n-j,j)} \otimes \mathcal{X}_{\mathbb{Q}}^{(n)}$ and $\mathcal{X}_{\mathbb{I}}^{(n-j,j)} \otimes \mathcal{X}_{\mathbb{Q}}^{(n-1,1)}$ into irreps for all $j \in [0..k]$. We have $\mathcal{X}_{\mathbb{I}}^{(n-j,j)} \otimes \mathcal{X}_{\mathbb{Q}}^{(n)} \cong \mathcal{S}^{(n-j,j)}$ and $\mathcal{X}_{\mathbb{I}}^{(n)} \otimes \mathcal{X}_{\mathbb{Q}}^{(n-1,1)} \cong \mathcal{S}^{(n-1,1)}$, as $\mathcal{S}^{(n)}$ is the trivial representation, and Corollary 1.13 in Section 1.4.5 states that, for $j \geq 1$, we have

$$\mathcal{X}_{\mathbb{I}}^{(n-j,j)} \otimes \mathcal{X}_{\mathbb{Q}}^{(n-1,1)} \cong \mathcal{S}^{(n-j+1,j-1)} \oplus \mathcal{S}^{(n-j,j)} \oplus (\mathcal{S}^{(n-j,j-1,1)} \oplus \mathcal{S}^{(n-j-1,j+1)} \oplus \mathcal{S}^{(n-j-1,j,1)}),$$

where we omit the term $\mathcal{S}^{(n-j,j-1,1)}$ when $j = 1$.

We can see that, for every $j \in [0..k]$ and $\ell \in \{0, 1\}$, the representation $\mathcal{X}_{\mathbb{I}}^{(n-j,j)} \otimes \mathcal{X}_{\mathbb{Q}}^{(n-\ell,\ell)}$ is multiplicity-free. For an irrep \mathcal{S}^θ present in $\mathcal{X}_{\mathbb{I}}^{(n-j,j)} \otimes \mathcal{X}_{\mathbb{Q}}^{(n-\ell,\ell)}$, let

$$\Pi_\theta^{(n-j,j) \otimes (n-\ell,\ell)} := \hat{\Pi}_\theta \left(\Pi_{\mathbb{I}}^{(n-j,j)} \otimes \Pi_{\mathbb{Q}}^{(n-\ell,\ell)} \right),$$

which is the projector on the unique instance of \mathcal{S}^θ in $\mathcal{X}_{\mathbb{I}}^{(n-j,j)} \otimes \mathcal{X}_{\mathbb{Q}}^{(n-\ell,\ell)}$. For example, for $\theta = (n-1, 1)$, we have projectors

$$\Pi_{(n-1,1)}^{(n-1,1) \otimes (n)}, \quad \Pi_{(n-1,1)}^{(n) \otimes (n-1,1)}, \quad \Pi_{(n-1,1)}^{(n-1,1) \otimes (n-1,1)}, \quad \text{and} \quad \Pi_{(n-1,1)}^{(n-2,2) \otimes (n-1,1)}.$$

Approach 2: via spaces invariant under the oracles $\mathcal{O}_{(S)}^+$ and $\mathcal{O}_{(D)}^+$. Let us decompose $\mathcal{X}_{\mathbb{I}\mathbb{Q}}$ as the direct sum of four subspaces, each invariant under the action of $U_{\mathbb{I}\mathbb{Q}}$, $\mathcal{O}_{(S)}^+$, and $\mathcal{O}_{(D)}^+$. First, let $\mathcal{X}_{\mathbb{I}\mathbb{Q}} = \mathcal{X}^{(0)} \oplus \mathcal{X}^{(1)}$, where

$$\mathcal{X}^{(0)} := \text{span}\{(\mathbf{x}, \mathbf{i}) \in \mathcal{X}_{\mathbb{I}\mathbb{Q}} : x_i = 0\} \quad \text{and} \quad \mathcal{X}^{(1)} := \text{span}\{(\mathbf{x}, \mathbf{i}) \in \mathcal{X}_{\mathbb{I}\mathbb{Q}} : x_i = 1\}.$$

Let us further decompose $\mathcal{X}^{(0)}$ and $\mathcal{X}^{(1)}$ as

$$\mathcal{X}^{(0)} = \mathcal{X}^{(0,s)} \oplus \mathcal{X}^{(0,t)} \quad \text{and} \quad \mathcal{X}^{(1)} = \mathcal{X}^{(1,s)} \oplus \mathcal{X}^{(1,t)},$$

where

$$\mathcal{X}^{(0,s)} := \text{span}\left\{ \sum_{i: x_i=0} (\mathbf{x}, \mathbf{i}) : x \in \mathcal{D} \right\} \quad \text{and} \quad \mathcal{X}^{(1,s)} := \text{span}\left\{ \sum_{i: x_i=1} (\mathbf{x}, \mathbf{i}) : x \in \mathcal{D} \right\},$$

and $\mathcal{X}^{(0,t)} := \mathcal{X}^{(0)} \ominus \mathcal{X}^{(0,s)}$ and $\mathcal{X}^{(1,t)} := \mathcal{X}^{(1)} \ominus \mathcal{X}^{(1,s)}$.

Note that $\mathcal{X}^{(1,s)} = \text{span}\{\mathbf{x} \otimes \chi(x) : x \in \mathcal{D}\}$. Therefore, the oracle $\mathcal{O}_{(D)}^+$ acts on $\mathcal{X}^{(1,s)}$ as the minus identity and on $\mathcal{X}^{(0)} \oplus \mathcal{X}^{(1,t)}$ as the identity. Meanwhile, $\mathcal{O}_{(S)}^+$ acts on $\mathcal{X}^{(1)}$ as the minus identity and on $\mathcal{X}^{(0)}$ as the identity.

For every superscript $\text{sup} \in \{(0), (1), (0, s), (0, t), (1, s), (1, t)\}$, let Π^{sup} be the projector on the space \mathcal{X}^{sup} , and let $U^{\text{sup}} := U|_{\mathcal{X}^{\text{sup}}}$ be the reduction of U to \mathcal{X}^{sup} . Let V_{π}^{sup} be U_{π}^{sup} restricted to $\pi \in \mathbb{S}_{[1..k]} \times \mathbb{S}_{[k+1..n]}$ and the space

$$\tilde{\mathcal{X}}^{\text{sup}} := \mathcal{X}^{\text{sup}} \cap ((\mathbf{1}^k \mathbf{0}^{n-k})_1 \otimes \mathcal{X}_{\mathbb{Q}}).$$

V^{sup} is a representation of $\mathbb{S}_{[1..k]} \times \mathbb{S}_{[k+1..n]}$. One can see that

$$|\mathbb{S}_n| / |\mathbb{S}_k \times \mathbb{S}_{n-k}| = \dim \mathcal{X}^{\text{sup}} / \dim \tilde{\mathcal{X}}^{\text{sup}},$$

so we have $U^{\text{sup}} = V^{\text{sup}} \uparrow \mathbb{S}_{[n]}$. In order to see how U^{sup} decomposes into irreps, we need to see how V^{sup} decomposes into irreps, and then apply the Littlewood–Richardson rule.

We have $\dim \tilde{\mathcal{X}}^{(0,s)} = \dim \tilde{\mathcal{X}}^{(1,s)} = 1$, and it is easy to see that $V^{(0,s)}$ and $V^{(1,s)}$ act trivially on $\tilde{\mathcal{X}}^{(0,s)}$ and $\tilde{\mathcal{X}}^{(1,s)}$, respectively. That is, $V^{(0,s)} \cong V^{(1,s)} \cong \mathcal{S}^{(k)} \times \mathcal{S}^{(n-k)}$. Now, note that

$$\tilde{\mathcal{X}}^{(0)} = \text{span}\{\mathbf{1}^k \mathbf{0}^{n-k} \otimes \mathbf{i} : i \in [k+1..n]\}.$$

The group $\mathbb{S}_{[1..k]}$ acts trivially on $\tilde{\mathcal{X}}^{(0)}$, while and the action of $\mathbb{S}_{[k+1..n]}$ on $\tilde{\mathcal{X}}^{(0)}$ defines the natural representation of $\mathbb{S}_{[k+1..n]}$. Hence, $V^{(0)} \cong \mathcal{S}^{(k)} \times (\mathcal{S}^{(n-k)} \oplus \mathcal{S}^{(n-k-1,1)})$, and $V^{(0)} = V^{(0,s)} \oplus V^{(0,t)}$ gives us $V^{(0,t)} \cong \mathcal{S}^{(k)} \times \mathcal{S}^{(n-k-1,1)}$. Analogously we obtain $V^{(1,t)} \cong \mathcal{S}^{(k-1,1)} \times \mathcal{S}^{(n-k)}$. As in (1.18), $U^{(0,s)} = V^{(0,s)} \uparrow \mathbb{S}_{[n]}$ and $U^{(1,s)} = V^{(1,s)} \uparrow \mathbb{S}_{[n]}$ are both isomorphic to $\bigoplus_{j=0}^k \mathcal{S}^{(n-j,j)}$. For $U^{(0,t)} = V^{(0,t)} \uparrow \mathbb{S}_{[n]}$ and $U^{(1,t)} = V^{(1,t)} \uparrow \mathbb{S}_{[n]}$, the Littlewood–Richardson rule (1.12) gives us, respectively,

$$\begin{aligned} (\mathcal{S}^{(k)} \times \mathcal{S}^{(n-k-1,1)}) \uparrow \mathbb{S}_{[n]} &\cong \mathcal{S}^{(n-1,1)} \oplus \mathcal{S}^{(n-2,2)} \oplus \mathcal{S}^{(n-2,1,1)} \\ &\oplus \mathcal{S}^{(n-3,3)} \oplus \mathcal{S}^{(n-3,2,1)} \oplus \mathcal{S}^{(n-4,4)} \oplus \mathcal{S}^{(n-4,3,1)} \oplus \dots \\ &\oplus \mathcal{S}^{(n-k,k)} \oplus \mathcal{S}^{(n-k,k-1,1)} \oplus \mathcal{S}^{(n-k-1,k+1)} \oplus \mathcal{S}^{(n-k-1,k,1)} \end{aligned}$$

and

$$\begin{aligned}
& (\mathcal{S}^{(k-1,1)} \times \mathcal{S}^{(n-k)}) \uparrow \mathbb{S}_{[n]} \cong \mathcal{S}^{(n-1,1)} \oplus \mathcal{S}^{(n-2,2)} \oplus \mathcal{S}^{(n-2,1,1)} \\
& \oplus \mathcal{S}^{(n-3,3)} \oplus \mathcal{S}^{(n-3,2,1)} \oplus \mathcal{S}^{(n-4,4)} \oplus \mathcal{S}^{(n-4,3,1)} \\
& \oplus \dots \oplus \mathcal{S}^{(n-k+1,k-1)} \oplus \mathcal{S}^{(n-k+1,k-2,1)} \oplus \mathcal{S}^{(n-k,k-1,1)}.
\end{aligned}$$

Note that all $U^{(0,s)}$, $U^{(0,t)}$, $U^{(1,s)}$, and $U^{(1,t)}$ are multiplicity-free. For a superscript $\text{sup} \in \{(0,s), (0,t), (1,s), (1,t)\}$ and an irrep \mathcal{S}^θ present in U^{sup} , let $\Pi_\theta^{\text{sup}} := \hat{\Pi}_\theta \Pi^{\text{sup}}$, which is the projector on the unique instance of \mathcal{S}^θ in U^{sup} . For example, for $\theta = (n-1,1)$, we have all the projectors

$$\Pi_{(n-1,1)}^{(0,s)}, \quad \Pi_{(n-1,1)}^{(0,t)}, \quad \Pi_{(n-1,1)}^{(1,s)}, \quad \text{and} \quad \Pi_{(n-1,1)}^{(1,t)}.$$

Significant irreps. Note that, since $\mathcal{O}_{(D)}^+$ acts on $\mathcal{X}^{(1,s)}$ as the minus identity and on $\mathcal{X}^{(0)} \oplus \mathcal{X}^{(1,t)}$ as the identity and $\mathcal{O}_{(S)}^+$ acts on $\mathcal{X}^{(1)}$ as the minus identity and on $\mathcal{X}^{(0)}$ as the identity, if μ is a subspace of one of the spaces $\mathcal{X}^{(0)}$, $\mathcal{X}^{(1,s)}$, or $\mathcal{X}^{(1,t)}$, then $\mu' = \mu$. And, even if that is not the case, we still have that μ and μ' are isomorphic irreps (due to Corollary 1.3).

Also note that

$$\left| \text{Tr}(\Pi_{\mathbb{I}\mathbb{Q},b}(\Pi_\mu - \Pi_{\mu'})) \right| = \left| \text{Tr}(\Pi_{\mathbb{I}\mathbb{Q},a}(\Pi_\mu - \Pi_{\mu'})) \right|. \quad (6.15)$$

Hence we need to consider only μ that are isomorphic to irreps present in both

$$(\mathcal{X}_1^{(n)} \oplus \mathcal{X}_1^{(n-1,1)}) \otimes (\mathcal{X}_{\mathbb{Q}}^{(n)} \oplus \mathcal{X}_{\mathbb{Q}}^{(n-1,1)}) \quad \text{and} \quad \bigoplus_{j=2}^k \mathcal{X}_1^{(n-j,j)} \otimes (\mathcal{X}_{\mathbb{Q}}^{(n)} \oplus \mathcal{X}_{\mathbb{Q}}^{(n-1,1)}),$$

as otherwise the left hand side of (6.13) equals 0. As one can see from Approach 1 above, the only such irreps are $\mathcal{S}^{(n-1,1)}$, $\mathcal{S}^{(n-2,2)}$, and $\mathcal{S}^{(n-2,1,1)}$.

The representation $U_{\mathbb{I}\mathbb{Q}}$ contains four instances of the irrep $\mathcal{S}^{(n-1,1)}$, four of $\mathcal{S}^{(n-2,2)}$, and two of $\mathcal{S}^{(n-2,1,1)}$. Projectors on them, according to Approach 1, are

$$\begin{aligned}
& \Pi_{(n-1,1)}^{(n-1,1) \otimes (n)}, \quad \Pi_{(n-1,1)}^{(n) \otimes (n-1,1)}, \quad \Pi_{(n-1,1)}^{(n-1,1) \otimes (n-1,1)}, \quad \Pi_{(n-1,1)}^{(n-2,2) \otimes (n-1,1)}, \\
& \Pi_{(n-2,2)}^{(n-2,2) \otimes (n)}, \quad \Pi_{(n-2,2)}^{(n-1,1) \otimes (n-1,1)}, \quad \Pi_{(n-2,2)}^{(n-2,2) \otimes (n-1,1)}, \quad \Pi_{(n-2,2)}^{(n-3,3) \otimes (n-1,1)}, \\
& \Pi_{(n-2,1,1)}^{(n-1,1) \otimes (n-1,1)}, \quad \Pi_{(n-2,1,1)}^{(n-2,2) \otimes (n-1,1)},
\end{aligned} \quad (6.16)$$

or, according to Approach 2, are

$$\begin{aligned}
& \Pi_{(n-1,1)}^{(0,s)}, \quad \Pi_{(n-1,1)}^{(0,t)}, \quad \Pi_{(n-1,1)}^{(1,s)}, \quad \Pi_{(n-1,1)}^{(1,t)}, \\
& \Pi_{(n-2,2)}^{(0,s)}, \quad \Pi_{(n-2,2)}^{(0,t)}, \quad \Pi_{(n-2,2)}^{(1,s)}, \quad \Pi_{(n-2,2)}^{(1,t)}, \\
& \Pi_{(n-2,1,1)}^{(0,t)}, \quad \Pi_{(n-2,1,1)}^{(1,t)}.
\end{aligned}$$

From this we can see right away that, if $\mu \cong \mathcal{S}^{(n-2,1,1)}$, then $\mu \subset \mathcal{X}^{(0)} \oplus \mathcal{X}^{(1,t)}$, so the application of the oracle $\mathcal{O}_{(D)}^+$ fixes μ , and the expression (6.13) equals 0.

6.2.2 Necessary and sufficient conditions for the irrep $\mathcal{S}^{(n-1,1)}$

We would like to know what are necessary and sufficient conditions for inequality (6.13) to hold. First, let us consider the irrep $\mathcal{S}^{(n-1,1)}$; later, the argument for the other two irreps will be very similar.

Basis of transporters. For $\ell_1, \ell_2 \in \{0, 1\}$ and $m_1, m_2 \in \{s, t\}$, let $\Xi_{(n-1,1)}^{(\ell_1, m_1) \leftarrow (\ell_2, m_2)}$ be a transporter from $\mathcal{X}_{(n-1,1)}^{(\ell_2, m_2)}$ to $\mathcal{X}_{(n-1,1)}^{(\ell_1, m_1)}$. We choose global phases of these transporters so that $\{\Xi_{(n-1,1)}^{(\ell_1, m_1) \leftarrow (\ell_2, m_2)}\}$ is a basis of transporters (see Section 1.3.4). As in (1.4), for a vector $\gamma = (\gamma_{0,s}, \gamma_{0,t}, \gamma_{1,s}, \gamma_{1,t})$, let

$$\Pi_{(\gamma)} := \sum_{\substack{\ell_1, \ell_2 \in \{0,1\} \\ m_1, m_2 \in \{s,t\}}} \gamma_{\ell_1, m_1} \bar{\gamma}_{\ell_2, m_2} \Xi_{(n-1,1)}^{(\ell_1, m_1) \leftarrow (\ell_2, m_2)}.$$

Claim 1.8 states that

Fact 6.10. *Let $\mu \subset \mathcal{X}_{\mathbb{Q}}$ be isomorphic to $\mathcal{S}^{(n-1,1)}$. There exists, up to a global phase, a unique vector $\gamma = (\gamma_{0,s}, \gamma_{0,t}, \gamma_{1,s}, \gamma_{1,t})$ such that $\Pi_{\mu} = \Pi_{(\gamma)}$. The norm of the vector γ is 1. The converse also holds: for any unit vector γ , $\Pi_{(\gamma)}$ is a projector to an irrep isomorphic to $\mathcal{S}^{(n-1,1)}$.*

From now on, let us work in this basis of transporters, because in this basis, the oracles $\mathcal{O}_{(S)}^+$ and $\mathcal{O}_{(D)}^+$ restricted to $\hat{\Pi}_{(n-1,1)}$ are, respectively,

$$\mathcal{O}_{(S)}^+|_{(n-1,1)} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \quad \text{and} \quad \mathcal{O}_{(D)}^+|_{(n-1,1)} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Necessary and sufficient condition for the oracle $\mathcal{O}_{(S)}^+$. In the basis of transporters we have

$$\Pi_{\mu} = \begin{pmatrix} |\gamma_{0,s}|^2 & \bar{\gamma}_{0,s} \gamma_{0,t} & \bar{\gamma}_{0,s} \gamma_{1,s} & \bar{\gamma}_{0,s} \gamma_{1,t} \\ \bar{\gamma}_{0,t} \gamma_{0,s} & |\gamma_{0,t}|^2 & \bar{\gamma}_{0,t} \gamma_{1,s} & \bar{\gamma}_{0,t} \gamma_{1,t} \\ \bar{\gamma}_{1,s} \gamma_{0,s} & \bar{\gamma}_{1,s} \gamma_{0,t} & |\gamma_{1,s}|^2 & \bar{\gamma}_{1,s} \gamma_{1,t} \\ \bar{\gamma}_{1,t} \gamma_{0,s} & \bar{\gamma}_{1,t} \gamma_{0,t} & \bar{\gamma}_{1,t} \gamma_{1,s} & |\gamma_{1,t}|^2 \end{pmatrix}, \quad (6.17)$$

and note that

$$|\gamma_{\ell,m}|^2 = \text{Tr}(\Pi_\mu \Pi_{(n-1,1)}^{(\ell,m)}) / \dim(n-1,1).$$

From (6.16), one can see that

$$\hat{\Pi}_{(n-1,1)} \Pi_{\mathbf{Q},b} = \Pi_{(n-1,1)}^{(n-2,2) \otimes (n-1,1)}.$$

Hence, for the space μ , the desired inequality (6.13) becomes

$$\frac{1}{\dim(n-1,1)} \left| \text{Tr}(\Pi_{(n-1,1)}^{(n-2,2) \otimes (n-1,1)} (\Pi_\mu - \Pi_{\mu'})) \right| \leq O(\max\{\sqrt{k/n}, \sqrt{1/k}\}). \quad (6.18)$$

Let us first obtain a necessary condition if we want this to hold for all $\mu \cong \mathcal{S}^{(n-1,1)}$.

In the same basis of transporters, let

$$\Pi_{(n-1,1)}^{(n-2,2) \otimes (n-1,1)} = \begin{pmatrix} |\beta_{0,s}|^2 & \bar{\beta}_{0,s}\beta_{0,t} & \bar{\beta}_{0,s}\beta_{1,s} & \bar{\beta}_{0,s}\beta_{1,t} \\ \bar{\beta}_{0,t}\beta_{0,s} & |\beta_{0,t}|^2 & \bar{\beta}_{0,t}\beta_{1,s} & \bar{\beta}_{0,t}\beta_{1,t} \\ \bar{\beta}_{1,s}\beta_{0,s} & \bar{\beta}_{1,s}\beta_{0,t} & |\beta_{1,s}|^2 & \beta_{1,s}\beta_{1,t} \\ \bar{\beta}_{1,t}\beta_{0,s} & \bar{\beta}_{1,t}\beta_{0,t} & \bar{\beta}_{1,t}\beta_{1,s} & |\beta_{1,t}|^2 \end{pmatrix}. \quad (6.19)$$

For $m_1, m_2 \in \{s, t\}$ and $\eta \in \mathbb{R}$, define the space $\xi_{m_1, m_2, \eta} \cong \mathcal{S}^{(n-1,1)}$ via the projector on it:

$$\Pi_{\xi_{m_1, m_2, \eta}} := \frac{1}{2} \left(\Pi_{(n-1,1)}^{(0, m_1)} + e^{i\eta} \Pi_{(n-1,1)}^{(0, m_1) \leftarrow (1, m_2)} + e^{-i\eta} \Pi_{(n-1,1)}^{(1, m_2) \leftarrow (0, m_1)} + \Pi_{(n-1,1)}^{(1, m_2)} \right).$$

We have

$$\Pi_{\xi_{m_1, m_2, \eta}} - \mathcal{O}_{(S)}^+ \Pi_{\xi_{m_1, m_2, \phi}} \mathcal{O}_{(S)}^+ = e^{i\eta} \Pi_{(n-1,1)}^{(0, m_1) \leftarrow (1, m_2)} + e^{-i\eta} \Pi_{(n-1,1)}^{(1, m_2) \leftarrow (0, m_1)},$$

so, for this space, the inequality (6.18) becomes

$$\left| e^{i\eta} \bar{\beta}_{1, m_2} \beta_{0, m_1} + e^{-i\eta} \bar{\beta}_{0, m_1} \beta_{1, m_2} \right| \leq O(\max\{\sqrt{k/n}, \sqrt{1/k}\}).$$

Since this has to hold for all m_1, m_2 , and η (in particular, consider m_1 and m_2 that maximize $|\bar{\beta}_{1, m_2} \beta_{0, m_1}|$), we must have either

$$|\beta_{1,s}|^2 + |\beta_{1,t}|^2 \leq O(\max\{k/n, 1/k\}) \quad \text{or} \quad |\beta_{1,s}|^2 + |\beta_{1,t}|^2 \geq 1 - O(\max\{k/n, 1/k\}), \quad (6.20)$$

and note that

$$|\beta_{1,s}|^2 + |\beta_{1,t}|^2 = \text{Tr}(\Pi_{(n-1,1)}^{(n-2,2) \otimes (n-1,1)} \cdot \Pi^{(1)}) / \dim(n-1,1).$$

The condition (6.20) is necessary, but it is also sufficient. Because, if it holds, then

$$|\bar{\beta}_{1, m_2} \beta_{0, m_1}| \leq O(\max\{\sqrt{k/n}, \sqrt{1/k}\})$$

for all $m_1, m_2 \in \{s, t\}$ and, clearly, $|\bar{\gamma}_{1, m_2} \gamma_{0, m_1}| = O(1)$ for all unit vectors γ . Therefore, if we plug (6.17) and (6.19) into (6.18), the inequality is satisfied.

Necessary and sufficient condition for the oracle $\mathcal{O}_{(D)}^+$. Almost identical analysis shows that, in order for the main conjecture hold when μ is isomorphic to $\mathcal{S}^{(n-1,1)}$ and when one applies $\mathcal{O}_{(D)}^+$, it is necessary and sufficient that

$$|\beta_{1,s}|^2 \leq O(\max\{k/n, 1/k\}) \quad \text{or} \quad |\beta_{1,s}|^2 \geq 1 - O(\max\{k/n, 1/k\}). \quad (6.21)$$

Note that

$$|\beta_{1,s}|^2 = \text{Tr}\left(\Pi_{(n-1,1)}^{(n-2,2) \otimes (n-1,1)} \cdot \Pi_{(n-1,1)}^{(1,s)}\right) / \dim(n-1,1).$$

6.2.3 Solution for the irreps $\mathcal{S}^{(n-2,2)}$ and $\mathcal{S}^{(n-2,1,1)}$

For irreps $\mathcal{S}^{(n-2,2)}$ and $\mathcal{S}^{(n-2,1,1)}$, let us exploit equation (6.15). We do that because the space $\mathcal{X}_{\mathbb{Q},b}$ contains three instances of the irrep $\mathcal{S}^{(n-2,2)}$, while $\mathcal{X}_{\mathbb{Q},a}$ contains only one. From (6.16) we get

$$\hat{\Pi}_{(n-2,2)} \Pi_{\mathbb{Q},a} = \Pi_{(n-2,2)}^{(n-1,1) \otimes (n-1,1)} \quad \text{and} \quad \hat{\Pi}_{(n-2,1,1)} \Pi_{\mathbb{Q},a} = \Pi_{(n-2,1,1)}^{(n-1,1) \otimes (n-1,1)}.$$

Oracle $\mathcal{O}_{(S)}^+$. An analysis analogous to that of the irrep $\mathcal{S}^{(n-1,1)}$ shows that, in order for the desired inequality (6.13) to hold for the oracle $\mathcal{O}_{(S)}^+$ and the irreps $\mathcal{S}^{(n-2,2)}$ and $\mathcal{S}^{(n-2,1,1)}$, it is sufficient to have

$$\frac{\text{Tr}\left(\Pi_{(n-2,2)}^{(n-1,1) \otimes (n-1,1)} \cdot \Pi^{(1)}\right)}{\dim(n-2,2)} \leq O(k/n) \quad \text{and} \quad \frac{\text{Tr}\left(\Pi_{(n-2,1,1)}^{(n-1,1) \otimes (n-1,1)} \cdot \Pi^{(1)}\right)}{\dim(n-2,1,1)} \leq O(k/n).$$

Let us prove this. Consider the irrep $\mathcal{S}^{(n-2,2)}$ and the hook-length formula (1.8) gives us $\dim(n-2,2) = n(n-3)/2$. We have

$$\text{Tr}\left(\Pi_{(n-2,2)}^{(n-1,1) \otimes (n-1,1)} \cdot \Pi^{(1)}\right) \leq \text{Tr}\left(\left(\Pi_{\mathbb{I}}^{(n-1,1)} \otimes \mathbb{I}_{\mathbb{Q}}\right) \cdot \Pi^{(1)}\right),$$

and we can evaluate the right hand side of this exactly. $\Pi^{(1)}$ is diagonal (in the standard basis), and, on the diagonal, it has $(n-k)\binom{n}{k}$ zeros and $k\binom{n}{k}$ ones. The diagonal entries of $\Pi_{\mathbb{I}}^{(n-1,1)}$ are all the same because $\Pi_{\mathbb{I}}^{(n-1,1)}$ projects to an eigenspace of the Johnson scheme. More precisely, we have $\text{Tr}(\Pi_{\mathbb{I}}^{(n-1,1)}) = \dim(n-1,1) = n-1$, therefore the diagonal entries of both $\Pi_{\mathbb{I}}^{(n-1,1)}$ and $\Pi_{\mathbb{I}}^{(n-1,1)} \otimes \mathbb{I}_{\mathbb{Q}}$ are $(n-1)/\binom{n}{k}$. Hence,

$$\text{Tr}\left(\left(\Pi_{\mathbb{I}}^{(n-1,1)} \otimes \mathbb{I}_{\mathbb{Q}}\right) \Pi^{(1)}\right) = k(n-1)$$

and, in turn,

$$\frac{\text{Tr}\left(\Pi_{(n-2,2)}^{(n-1,1) \otimes (n-1,1)} \Pi^{(1)}\right)}{\dim(n-2,2)} \leq \frac{2k(n-1)}{n(n-3)} = O(k/n)$$

as required. The same argument works for the irrep $\mathcal{S}^{(n-2,1,1)}$ as, by the hook-length formula, $\dim(n-2,1,1) = (n-1)(n-2)/2 = \dim(n-2,2) + 1$.

Oracle $\mathcal{O}_{(D)}^+$. As we mentioned in the very end of Section 6.2.1, $\mathcal{O}_{(D)}^+$ affects no space μ isomorphic to the irrep $\mathcal{S}^{(n-2,1,1)}$. Nevertheless, the following argument for the irrep $\mathcal{S}^{(n-2,2)}$ works for $\mathcal{S}^{(n-2,1,1)}$ as well. We have

$$\frac{\mathrm{Tr}\left(\Pi_{(n-2,2)}^{(n-1,1)\otimes(n-1,1)}\Pi^{(1,s)}\right)}{\dim(n-2,2)} \leq \frac{\mathrm{Tr}\left(\Pi_{(n-2,2)}^{(n-1,1)\otimes(n-1,1)}\Pi^{(1)}\right)}{\dim(n-2,2)} \leq \mathcal{O}(k/n),$$

which, similarly to the condition (6.21) for the irrep $\mathcal{S}^{(n-1,1)}$, is sufficient to show that Lemma 6.8 holds for the irrep $\mathcal{S}^{(n-2,2)}$ and the oracle $\mathcal{O}_{(D)}^+$.

6.2.4 Solution for the irrep $\mathcal{S}^{(n-1,1)}$

Recall that the conditions (6.20) and (6.21) are sufficient for Lemma 6.8 to hold for the oracles $\mathcal{O}_{(S)}^+$ and $\mathcal{O}_{(D)}^+$, respectively. Hence, it suffices for us to show that

$$\begin{aligned} \frac{\mathrm{Tr}\left(\Pi_{(n-1,1)}^{(n-2,2)\otimes(n-1,1)}\cdot\Pi^{(1)}\right)}{\dim(n-1,1)} &\geq \frac{\mathrm{Tr}\left(\Pi_{(n-1,1)}^{(n-2,2)\otimes(n-1,1)}\cdot\Pi_{(n-1,1)}^{(1,s)}\right)}{\dim(n-1,1)} = \\ &= \frac{k-1}{k} \cdot \frac{n(n-k-1)}{(n-1)(n-2)} \geq 1 - \mathcal{O}(\max\{k/n, 1/k\}). \end{aligned} \quad (6.22)$$

It is easy to see that both inequalities in this expression hold, and we need to concern ourselves only with the equality in the middle. Note that

$$\Pi_{(n-1,1)}^{(n-2,2)\otimes(n-1,1)}\cdot\Pi_{(n-1,1)}^{(1,s)} = (\Pi_{\mathbb{1}}^{(n-2,2)} \otimes \mathbb{I}_{\mathbb{Q}})\cdot\Pi_{(n-1,1)}^{(1,s)},$$

and let us evaluate the trace of the latter.

Johnson scheme on $\mathcal{X}_{\mathbb{1}}$. Recall that we defined $\Pi_{\mathbb{1}}^{(n-2,2)}$, via the bijection (6.4), to be the projector on an eigenspace of the Johnson scheme. Hence, given $x, x' \in \mathcal{D}$ such that the Hamming distance $|x - x'|$ between them is $2i$, (1.25) states that $(\Pi_{\mathbb{1}}^{(n-2,2)} \otimes \mathbb{I}_{\mathbb{Q}})[(x, j), (x', j')]$ equals

$$\left(\binom{k-i}{2} - \frac{(k-1)^2}{n-2}(k-i) + \frac{k^2(k-1)^2}{2(n-1)(n-2)} \right) / \binom{n-4}{k-2} \quad (6.23)$$

if $j = j'$, and 0 if $j \neq j'$.

Johnson scheme on $\mathcal{X}^{(1,s)}$. Recall that, for $x \in \mathcal{D}$, we have $\chi(x) := \sum_{j: x_j=1} \mathbf{j}/\sqrt{k}$, and let us define

$$A_i^{(1,s)} := \sum_{\substack{x, x' \in \mathcal{D} \\ |x-x'|=2i}} (\mathbf{x} \otimes \chi(x))(\mathbf{x}' \otimes \chi(x'))^* \in \mathbf{L}(\mathcal{X}^{(1,s)})$$

for all $i \in [0..k]$. The matrices A_i of the Johnson scheme (see (1.17)) and $A_i^{(1,s)}$ here have the same eigenvalues corresponding to the same irreps. Hence, given $x, x' \in \mathcal{D}$ such that $|x-x'| = 2i$, (1.24) gives us that

$$\Pi_{(n-1,1)}^{(1,s)} \llbracket (x, j), (x', j) \rrbracket = \frac{1}{k} \cdot \left((k-i) - \frac{k^2}{n} \right) / \binom{n-2}{k-1} \quad (6.24)$$

if $x_j = x'_j = 1$, and 0 otherwise. Note that there are exactly $k-i$ indices j such that $x_j = x'_j = 1$.

Both Johnson schemes together. There are $\binom{n}{k}$ inputs $x \in \mathcal{D}$, and, for every x , there are $\binom{k}{i} \binom{n-k}{i}$ inputs $x' \in \mathcal{D}$ such that $|x-x'| = 2i$. From (6.23) and (6.24), we get

$$\begin{aligned} & \text{Tr} \left((\Pi_1^{(n-2,2)} \otimes \mathbb{I}_{\mathbb{Q}}) \cdot \Pi_{(n-1,1)}^{(1,s)} \right) = \\ &= \binom{n}{k} \sum_{i=0}^k \binom{k}{i} \binom{n-k}{i} \frac{\left(\frac{(k-i)(k-i-1)}{2} - \frac{(k-1)^2}{n-2} (k-i) + \frac{k^2(k-1)^2}{2(n-1)(n-2)} \right) \left((k-i) - \frac{k^2}{n} \right) k - i}{\binom{n-4}{k-2} \binom{n-2}{k-1} k}. \end{aligned} \quad (6.25)$$

It is straightforward to rewrite this expression as a linear combination of

$$\sum_{i=0}^k \binom{k}{i} \binom{n-k}{i} \frac{(k-i)!}{(k-i-\ell)!} = \frac{k!}{(k-\ell)!} \binom{n-\ell}{n-k} \quad (6.26)$$

where $\ell \in \{0, 1, 2, 3, 4\}$ and the coefficients of the linear combination do not depend on i . The equality (6.26) is essentially the same as

$$\sum_{i=0}^k \binom{k-\ell}{i} \binom{n-k}{i} = \binom{n-\ell}{k-\ell}.$$

By using the equality (6.26), one can show that (6.25) equals to

$$\frac{k-1}{k} \cdot \frac{n(n-k-1)}{(n-2)}.$$

We get the desired equality in (6.22) by dividing this by $\dim(n-1, 1) = n-1$.

Conclusion

In this thesis, we have studied applications of the adversary bound, lower bounds on the learning graph complexity, and connections between the two. We have also studied the query complexity of a problem that cannot be directly addressed by the adversary method, that is, the ENHANCED FIND-TWO problem. This work has answered some of the questions in the area of quantum query complexity, yet there are many interesting and important questions left. We conclude by mentioning some of those questions.

We proved that the $O(n^{9/7})$ non-adaptive learning graph for TRIANGLE in Ref. [LMS13] is almost optimal. Aside from the triangle subgraph, it would be interesting to prove lower bounds on the non-adaptive learning graph complexity of other subgraph-finding problems. For example, Ref. [LMS13] also gives an $O(n^{10/7})$ non-adaptive learning graph for the ASSOCIATIVITY TESTING problem, which essentially looks for a path of length four in a graph. A natural question arises: Can we construct a tight $\Omega(n^{10/7})$ lower bound proving the optimality of this learning graph?

Now that the power of non-adaptive learning graphs is better characterized, it would be interesting to do the same for adaptive learning graphs. One concrete question to consider would be: Is there a decision problem with small 1-certificate complexity for which there is a gap between its quantum query and adaptive learning graph complexities? In fact, k -DISTINCTNESS might be such a problem, as an $O(n^{1-2^{k-2}/(2^k-1)})$ adaptive learning graph for this problem is known only under strong additional promises [BL11]. As we already mentioned in the introduction, another problem to consider is the following: Is there a general technique to construct adversary bounds from adaptive dual learning graphs, and what class of problems would such bounds address?

Our adversary construction for ELEMENT DISTINCTNESS with small input alphabet in Chapter 5 is rather technical. Claims 5.2 and 5.4 suggest that the adversary matrix that we consider in Theorem 5.5 is a natural choice. While any other optimal adversary matrix probably cannot look too different (in terms of the singular value decomposition), it does not mean that it cannot have a simpler specification. Such a simpler specification might facilitate the construction of adversary bounds for other problems.

For example, it might help to narrow the gap between the best known lower bound and upper bound for k -DISTINCTNESS, $\Omega(n^{2/3})$ and $O(n^{1-2^{k-2}/(2^k-1)})$, respectively. Or, it might help to

reduce the required alphabet size in the $\Omega(n^{k/(k+1)})$ lower bound for k -SUM. As pointed out in Ref. [BŠ13], the quantum query complexity of k -SUM becomes $O(\sqrt{n})$ for alphabets of constant size. Therefore it would be interesting to find tradeoffs between the quantum query complexity and the size of the alphabet. These tradeoffs might be relatively smooth, unlike the jump in the query complexity of ELEMENT DISTINCTNESS between alphabet sizes $n - 1$ and n .

Jeffery, Magniez, and de Wolf recently studied the model of parallel quantum query algorithms, which can make p queries in parallel in each time step [JMdW13]. They show that such algorithms have to make $\Theta((n/p)^{2/3})$ p -parallel quantum queries to solve ELEMENT DISTINCTNESS. For the lower bound, they generalize the adversary bound given in [BR13a] and therefore require that the alphabet size is at least $\Omega(n^2)$. Can the representation-theoretic techniques from Chapter 5 remove this requirement?

In addition to allowing negative weights in the adversary matrix, the adversary method has been generalized in multiple directions. From function evaluation, it has been generalized to *quantum state generation* [AMRR11] and further to *quantum state conversion* [LMR⁺11]. The *multiplicative adversary method* was introduced by [Špa08], which, unlike the *additive* method presented here, allows to obtain useful bounds even for large error probabilities. Can the adversary method be generalized in yet another direction, namely, for relational problems, when the correct solution is not unique? And, secondly, can it be generalized to lower bound the *distributional quantum query complexity* of a problem, where one only requires that the algorithm succeeds with a high probability on a probabilistic distribution over inputs? (A tight $\Omega(n^{1/3})$ lower bound for finding a collision in a random input was recently given by Zhandry using the polynomial method [Zha13].)

References

- [AA11] Scott Aaronson and Andris Ambainis. The need for structure in quantum speedups. In *Proc. of 2nd ACM Conference on Innovations in Theoretical Computer Science*, pages 338–352, 2011.
- [Aar02] Scott Aaronson. Quantum lower bound for the collision problem. In *Proc. of 34th ACM Symposium on Theory of Computing*, pages 635–642, 2002.
- [Amb02] Andris Ambainis. Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64(4):750–767, 2002.
- [Amb03] Andris Ambainis. Polynomial degree vs. quantum query complexity. In *Proc. of 44th IEEE Symposium on Foundations of Computer Science*, pages 230–239, 2003.
- [Amb05] Andris Ambainis. Polynomial degree and lower bounds in quantum complexity: Collision and element distinctness with small range. *Theory of Computing*, 1:37–46, 2005.
- [Amb07] Andris Ambainis. Quantum walk algorithm for element distinctness. *SIAM Journal on Computing*, 37(1):210–239, 2007.
- [Amb10] Andris Ambainis. A new quantum lower bound method, with an application to a strong direct product theorem for quantum search. *Theory of Computing*, 6(1):1–25, 2010.
- [AMRR11] Andris Ambainis, Loïck Magnin, Martin Rötteler, and Jérémie Roland. Symmetry-assisted adversaries for quantum state generation. In *Proc. of 26th IEEE Conference on Computational Complexity*, pages 167–177, 2011.
- [ARU14] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. Quantum attacks on classical proof systems — the hardness of quantum rewinding. 2014. Available at [arXiv:1404.6898](https://arxiv.org/abs/1404.6898).

- [AS04] Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM*, 51(4):595–605, 2004.
- [AŠdW09] Andris Ambainis, Robert Špalek, and Ronald de Wolf. A new quantum lower bound method, with applications to direct product theorems and time-space tradeoffs. *Algorithmica*, 55(3):422–461, 2009.
- [Aud06] Koenraad M.R. Audenaert. A digest on representation theory of the symmetric group. Manuscript, 2006. Available at http://personal.rhul.ac.uk/usah/080/qitnotes_files/irreps_v06.pdf
- [BBC⁺01] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *Journal of the ACM*, 48(4):778–797, 2001.
- [BBHT98] Michel Boyer, Gilles Brassard, Peter Høyer, and Alain Tapp. Tight bounds on quantum searching. *Fortschritte der Physik*, 46(4-5):493–505, 1998.
- [BDH⁺05] Harry Buhrman, Christoph Dürr, Mark Heiligman, Peter Høyer, Frédéric Magniez, Miklos Santha, and Ronald de Wolf. Quantum algorithms for element distinctness. *SIAM Journal on Computing*, 34(6):1324–1330, 2005.
- [Bel12a] Aleksandrs Belovs. Personal communication, 2012.
- [Bel12b] Aleksandrs Belovs. Adversary lower bound for element distinctness. 2012. Available at [arXiv:1204.5074](https://arxiv.org/abs/1204.5074).
- [Bel12c] Aleksandrs Belovs. Learning-graph-based quantum algorithm for k -distinctness. In *Proc. of 53rd IEEE Symposium on Foundations of Computer Science*, pages 207–216, 2012.
- [Bel12d] Aleksandrs Belovs. Span programs for functions with constant-sized 1-certificates. In *Proc. of 44th ACM Symposium on Theory of Computing*, pages 77–84, 2012.
- [Bel13] Aleksandrs Belovs. *Applications of the Adversary Method in Quantum Query Algorithms*. PhD thesis, University of Latvia, Faculty of Computing, 2013. Available at [arXiv:1402.3858](https://arxiv.org/abs/1402.3858).
- [Bel14] Aleksandrs Belovs. Quantum algorithms for learning symmetric juntas via adversary bound. 2014. Available at [arXiv:1311.6777](https://arxiv.org/abs/1311.6777).
- [Bha97] Rajendra Bhatia. *Matrix Analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer, 1997.

- [BHT98] Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum cryptanalysis of hash and claw-free functions. In *Proc. of 3rd Latin American Theoretical Informatics Symposium*, volume 1380 of *Lecture Notes in Computer Science*, pages 163–169. Springer, 1998.
- [BL11] Aleksandrs Belovs and Troy Lee. Quantum algorithm for k -distinctness with prior knowledge on the input. 2011. Available at [arXiv:1108.3022](https://arxiv.org/abs/1108.3022).
- [Boe63] Hermann Boerner. *Representation of groups with special consideration for the needs of modern physics*. North-Holland Publishing Co., 1963.
- [BR13a] Aleksandrs Belovs and Ansis Rosmanis. On the power of non-adaptive learning graphs. In *Proc. of 28th IEEE Conference on Computational Complexity*, pages 44–55, 2013.
- [BR13b] Aleksandrs Belovs and Ansis Rosmanis. On adversary lower bounds for the collision and the set equality problems (Version 1). 2013. Available at [arXiv:1310.5185v1](https://arxiv.org/abs/1310.5185v1).
- [BR14] Aleksandrs Belovs and Ansis Rosmanis. Adversary lower bounds for the collision and the set equality problems. 2014. Available at [arXiv:1310.5185](https://arxiv.org/abs/1310.5185).
- [BŠ13] Aleksandrs Belovs and Robert Špalek. Adversary lower bound for the k -sum problem. In *Proc. of 4th ACM Conference on Innovations in Theoretical Computer Science*, pages 323–328, 2013.
- [EHK04] 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.
- [Gil10] Bryan Gillespie. On randomness of subsets of \mathbb{Z}_N , as described by uniformity of Fourier coefficients. Manuscript, 2010. Available at <http://math.berkeley.edu/~bgillesp/files/doc/on-randomness.pdf>
- [God05] Chris Godsil. Association schemes. Lecture Notes, 2005. Available at <http://quoll.uwaterloo.ca/pdfs/assoc1.pdf>
- [Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proc. of 28th ACM Symposium on Theory of Computing*, pages 212–219, 1996.
- [HLŠ07] Peter Høyer, Troy Lee, and Robert Špalek. Negative weights make adversaries stronger. In *Proc. of 39th ACM Symposium on Theory of Computing*, pages 526–535, 2007.
- [HSS99] A. S. Hedayat, N. J. A. Sloane, and John Stufken. *Orthogonal arrays: theory and applications*. Springer, 1999.

- [JK81] Gordon James and Adalbert Kerber. *The Representation Theory of the Symmetric Group*, volume 16 of *Encyclopedia of Mathematics and its Applications*. Addison-Wesley, 1981.
- [JKM13] Stacey Jeffery, Robin Kothari, and Frédéric Magniez. Nested quantum walks with quantum data structures. In *Proc. of 24th ACM-SIAM Symposium on Discrete Algorithms*, pages 1474–1485, 2013.
- [JMdW13] Stacey Jeffery, Frédéric Magniez, and Ronald de Wolf. Optimal parallel quantum query algorithms. 2013. Available at [arXiv:1309.6116](https://arxiv.org/abs/1309.6116).
- [KL07] Jonathan Katz and Yehuda Lindell. *Introduction to Modern Cryptography*. Taylor & Francis, 2007.
- [Kut05] Samuel Kutin. Quantum lower bound for the collision problem with small range. *Theory of Computing*, 1(1):29–36, 2005.
- [LG14] François Le Gall. Improved quantum algorithm for triangle finding via combinatorial arguments. 2014. Available at [arXiv:1407.0085](https://arxiv.org/abs/1407.0085).
- [LMR⁺11] Troy Lee, Rajat Mittal, Ben W. Reichardt, Robert Špalek, and Mario Szegedy. Quantum query complexity of the state conversion problem. In *Proc. of 52nd IEEE Symposium on Foundations of Computer Science*, pages 344–353, 2011.
- [LMS12] Troy Lee, Frédéric Magniez, and Miklos Santha. A learning graph based quantum query algorithm for finding constant-size subgraphs. *Chicago Journal of Theoretical Computer Science*, (10):1–21, 2012.
- [LMS13] Troy Lee, Frédéric Magniez, and Miklos Santha. Improved quantum query algorithms for triangle finding and associativity testing. In *Proc. of 24th ACM-SIAM Symposium on Discrete Algorithms*, pages 1486–1502, 2013.
- [Lov95] László Lovász. Semidefinite programs and combinatorial optimization. Lecture Notes, 1995. Available at <http://www.cs.elte.hu/~lovasz/semidef.ps>
- [Mid04] Gatis Midrijānis. A polynomial quantum query lower bound for the set equality problem. In *Proc. of 31st International Colloquium on Automata, Languages, and Programming*, volume 3142 of *Lecture Notes in Computer Science*, pages 996–1005. Springer, 2004.
- [NC00] Michael A. Nielsen and Isaac L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, 2000.

- [Rei09] Ben W. Reichardt. Span programs and quantum query complexity: The general adversary bound is nearly tight for every boolean function. 2009. Available at [arXiv:0904.2759](https://arxiv.org/abs/0904.2759).
- [Rei11] Ben W. Reichardt. Reflections for quantum query algorithms. In *Proc. of 22nd ACM-SIAM Symposium on Discrete Algorithms*, pages 560–569, 2011.
- [Ros14] Ansis Rosmanis. Adversary lower bound for element distinctness with small range. 2014. Available at [arXiv:1401.3826](https://arxiv.org/abs/1401.3826).
- [Rot95] Joseph J. Rotman. *An Introduction to the Theory of Groups*, volume 148 of *Graduate Texts in Mathematics*. Springer, 1995.
- [Sag01] Bruce E. Sagan. *The symmetric group: representations, combinatorial algorithms, and symmetric functions*, volume 203 of *Graduate Texts in Mathematics*. Springer, 2001.
- [Ser77] Jean-Pierre Serre. *Linear Representations of Finite Groups*, volume 42 of *Graduate Texts in Mathematics*. Springer, 1977.
- [Shi02] Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. In *Proc. of 43th IEEE Symposium on Foundations of Computer Science*, pages 513–519, 2002.
- [Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.
- [Sim97] Daniel R. Simon. On the power of quantum computation. *SIAM Journal on Computing*, 26(5):1474–1483, 1997.
- [Špa08] Robert Špalek. The multiplicative quantum adversary. In *Proc. of 23rd IEEE Conference on Computational Complexity*, pages 237–248, 2008.
- [Špa13] Robert Špalek. Adversary lower bound for the orthogonal array problem. 2013. Available at [arXiv:1304.0845](https://arxiv.org/abs/1304.0845).
- [ŠS06] Robert Špalek and Mario Szegedy. All quantum adversary methods are equivalent. *Theory of Computing*, 2:1–18, 2006.
- [TV06] Terence Tao and Vav H. Vu. *Additive combinatorics*, volume 105 of *Cambridge Studies in Advanced Mathematics*. 2006.
- [Wat11] John Watrous. Theory of quantum information. Lecture Notes, 2011. Available at <https://cs.uwaterloo.ca/~watrous/LectureNotes.html>

- [WSV00] Henry Wolkowicz, Romesh Saigal, and Lieven Vandenbergh, editors. *Handbook of Semidefinite Programming: Theory, Algorithms, and Applications*. Kluwer Academic Publishers, 2000.
- [Zha05] Shengyu Zhang. On the power of Ambainis lower bounds. *Theoretical Computer Science*, 339(2):241–256, 2005.
- [Zha12] Mark Zhandry. How to construct quantum random functions. In *Proc. of 53rd IEEE Symposium on Foundations of Computer Science*, pages 679–687, 2012.
- [Zha13] Mark Zhandry. A note on the quantum collision and set equality problems. 2013. Available at [arXiv:1312.1027](https://arxiv.org/abs/1312.1027).
- [Zhu12] Yechao Zhu. Quantum query complexity of subgraph containment with constant-sized certificates. *International Journal of Quantum Information*, 10(3), 2012.

APPENDICES

Appendix A

Proofs of lemmas in Section 4.2

Here we prove the technical lemmas that are used in proving adversary bounds for the COLLISION and SET EQUALITY problems in Section 4.2. All the proofs here use the representation theory of the symmetric group introduced in Section 1.4. The proofs in Appendices A.2 and A.3 also use the representation theory of the unitary group.

A.1 Proof of Lemma 4.4

Fix the value of k . Assume $m \geq 2k$ is some integer, and consider the symmetric group \mathbb{S}_m . Let κ be an element of the group algebra $\mathbb{C}\mathbb{S}_m$ defined by

$$\kappa := \frac{1}{2^k} (\varepsilon - (\mathbf{a}_1, \mathbf{b}_1))(\varepsilon - (\mathbf{a}_2, \mathbf{b}_2)) \cdots (\varepsilon - (\mathbf{a}_k, \mathbf{b}_k)), \quad (\text{A.1})$$

where $a_1, b_1, \dots, a_k, b_k$ 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 .

Fix an e -basis $\{e_i\}$ of \mathcal{H} . Recall that the e -basis of $\mathcal{H}^{\otimes n}$ and $\mathcal{H}^{\otimes 2n}$ consists of tensor products of the vectors in $\{e_i\}$. The vector e_0 in such a tensor product is called the zero component, and the weight of the corresponding basis vector is the number of non-zero components in the product. The spaces $\mathcal{H}_k^{(n)} \subset \mathcal{H}^{\otimes n}$ and $\mathcal{H}_k^{(2n)} \subset \mathcal{H}^{\otimes 2n}$ are spanned by all the e -basis vectors of weight k .

Lemma A.1. *For any vector $v \in \mathcal{H}_k^{(2n)}$ satisfying $\kappa v = v$ 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 (4.17) and (4.19).*

Proof. We prove the result for W_k^μ , the proofs for X_k^μ and Y_k^μ being similar. All vectors in the proof are considered in the e -basis. We say that a basis vector is *used* in v , if it has a non-zero coefficient. Let $A_i = \{a_i, b_i\}$ be the pairs from the definition of κ . Note that

$$\text{each basis vector used in } v \text{ has exactly one non-zero component positioned in each } A_i. \quad (\text{A.2})$$

Split the basis vectors into equivalence classes by assigning two vectors to the same equivalence class if and only if they can be obtained from one another by a permutation used in κ , i.e., by permuting elements inside A_i . In v , the coefficients of the basis vectors in one equivalence class are all equal up to a sign.

If μ contains a pair A_i for some i , then $W_k^\mu v = 0$, so assume it is not the case. We construct the following graph G that depends on κ and μ . Its vertex set is formed by the pairs A_1, \dots, A_k and the singletons $\{j\}$ for $j \in [2n] \setminus \bigcup_i A_i$. For each pair in μ , connect the sets containing the elements of the pair by an edge. The graph G does not contain loops, but it may have parallel edges. The graph G has maximal degree 2, so it is a collection of paths and cycles. Let c denote the number of cycles in G .

The operator W_k^μ maps basis vectors of $\mathcal{H}^{\otimes 2n}$ into basis vectors of $\mathcal{H}^{\otimes n}$ (that correspond to the labelling of the edges of G) or the zero vector. Let v' be the vector v with all terms that are mapped to 0 by W_k^μ removed. We claim that $\|v'\| \leq \|v\|/\sqrt{2^c}$. Indeed, in any equivalence class, for each cycle, at least half of the vectors are mapped to 0 (for an edge matches two non-zero components in them).

Next, we claim that each basis vector in the range of W_k^μ has exactly 0 or 2^c preimages among the basis vectors of the domain that satisfy (A.2). Indeed, consider a labelling of the edges of G , and our task is to count the number of basis vectors in $\mathcal{H}^{\otimes 2n}$ such that each A_i contains exactly one non-zero component, and each edge matches e_0 and its label (which is either e_0 as well or a non-zero component). Assume there is at least one way to satisfy these requirements. Then, for each path in G , all edges but one are labeled by a non-zero component and there is a unique way to satisfy it. For each cycle, there are two possibilities.

Since distinct basis vectors in the range of W_k^μ have no common preimage, we have $\|W_k^\mu v\| = \|W_k^\mu v'\| \leq \sqrt{2^c} \|v'\| \leq \|v\|$. \square

Now we are ready to prove Lemma 4.4. 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 irrep of \mathbb{S}'_Q 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 irrep is isomorphic to either $\mathcal{S}^{(2n-1-k,\lambda)}$ with $\lambda \vdash k$ for COLLISION, or $\mathcal{S}^{(n-1-\ell,\lambda)} \otimes \mathcal{S}^{(n-k+\ell,\lambda')}$ with $\lambda \vdash \ell$ and $\lambda' \vdash k - \ell$ for SET EQUALITY. By Lemma 1.11,

in both cases, there exists a non-zero vector v in the space of the irrep satisfying $\kappa v = v$ for some choice of a_1, \dots, b_k (in the case of SET EQUALITY, one has to take the tensor product of two vectors obtained by two applications of Lemma 1.11). By Lemma A.1, $\|X_k^\mu v\| \leq \|v\|$, hence, $\|X_{Q,k} v\| \leq \|v\|$, and $\|\bar{X}_{Q,k}\| \leq 1$.

Consider case (b) now. Similarly as for (a), we get a right singular vector v of singular value $\|\bar{Y}_{Q,k}\|$ such that $\kappa v = v$. Note that, if μ matches 1 with an element outside $\{a_1, b_1, \dots, a_k, b_k\}$, then, because of (A.2), $Y_{Q,k}^\mu v = 0$. Otherwise, we still get $\|Y_{Q,k}^\mu v\| \leq \|v\|$ by Lemma A.1. The latter case only holds for an $O(k/n)$ fraction of all matchings, hence $\|\bar{Y}_{Q,k}\| = O(\sqrt{k/n})$.

Now, let us prove (c). From Lemma 4.3, we know that $\|\Phi_k^{(m)}\| = O(1/\sqrt{m})$ and $W_{Q,k} \Phi_{Q,k} = W_{Q,k} \bar{\Pi}_k \bar{\Phi}_{Q,k}$. So, it suffices to prove that $\|W_{Q,k} \bar{\Pi}_k\| = O(1)$. This time consider

$$\kappa' := \frac{1}{2^{k-1}} (\varepsilon - (\mathbf{a}_1, \mathbf{b}_1))(\varepsilon - (\mathbf{a}_2, \mathbf{b}_2)) \cdots (\varepsilon - (\mathbf{a}_{k-1}, \mathbf{b}_{k-1})).$$

By an argument similar to (a), we get that $W_{Q,k}$ has a right singular vector $v \in \bar{\mathcal{H}}_{Q,k}$ of singular value $\|W_{Q,k} \bar{\Pi}_k\|$ that satisfies $\kappa' v = v$. Now we proceed by modifying the proof of Lemma A.1. Again, we define $A_i = \{a_i, b_i\}$ for $i \in [k-1]$, and the equivalence classes as before. Each of A_i has to contain one non-zero component. One of them may contain two non-zero components, or there can be one non-zero component in a singleton.

It suffices to prove that $\|W_k^\mu v\| \leq \sqrt{3}\|v\|$ for any matching μ . Again, if μ has one of A_i as a pair, then $W_k^\mu v = 0$, so we may assume it is not the case, and define G as before. Consider a labelling of the edges of G (that is, an e -basis vector of $\mathcal{H}_k^{(n)}$), and let us count the number of preimages of this labelling. Each cycle must have all its edges labelled by non-zero components, and there are two ways to satisfy it. All paths, except one, have exactly one edge labelled by the zero component e_0 . They can be satisfied in a unique way. One path has all its edges labelled by non-zero components. We call it special. Let T be the length of the special path. Then, it can be satisfied in $T+1$ ways, hence, the edge labelling has $(T+1)2^c$ preimages.

For basis vectors in $\mathcal{H}^{\otimes 2n}$, we call a path special if the total number of non-zero components at the endpoints of its edges is equal to its length. All vectors of an equivalence class have the same special path. Let us label all paths in G with numbers from 1 to ℓ , and let T_i be the length of the i th path. Let us decompose $v = v_1 + \dots + v_\ell$ where v_i only uses basis vectors with the i th special path. Similarly to the proof of Lemma A.1, let v'_i be the vector v_i with terms mapped to 0 by W_k^μ removed. It is not hard to check that $\|v'_i\| \leq \|v_i\|/\sqrt{2^{c+\max\{T_i-2,0\}}}$. Hence, $\|W_k^\mu v_i\| \leq \sqrt{(T_i+1)/2^{\max\{T_i-2,0\}}}\|v_i\| \leq \sqrt{3}\|v_i\|$. As $\{W_k^\mu v_i\}$ are orthogonal, we get that $\|W_k^\mu v\| \leq \sqrt{3}\|v\|$.

A.2 Proof of Lemma 4.3

Consider a unitary transformation $U \in \mathbf{U}(\mathcal{H}_1)$. We embed U into $\mathbf{U}(\mathcal{H})$ using the assumption $Ue_0 = e_0$. Note that all permutation matrices satisfy $Ue_0 = e_0$, so, in some sense, U corresponds to permuting symbols within the input alphabet $[q]$. Nevertheless, as $\mathbf{U}(\mathcal{H}_1)$ is a larger group than $\mathbb{S}_{[q]}$, we have to deal with less irreps when considering $\mathbf{U}(\mathcal{H}_1)$ rather than $\mathbb{S}_{[q]}$. Therefore we choose to consider $\mathcal{H}^{\otimes m}$ as a representation of the direct product $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_{[m]}$. Let \mathbb{S}_m stand for $\mathbb{S}_{[m]}$.

Decomposition of $\mathcal{H}^{\otimes m}$ into irreps. The unitary group $\mathbf{U}(\mathcal{H}_1)$ acts on $\mathcal{H}^{\otimes m}$ by simultaneous matrix multiplication, as in (1.14), and the symmetric group \mathbb{S}_m acts on $\mathcal{H}^{\otimes m}$ by permuting the tensor factors, as in (1.15). These actions of $\mathbf{U}(\mathcal{H}_1)$ and \mathbb{S}_m commute, so they provide a legitimate representation of $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_m$.

It is not hard to see that the subspace $\mathcal{H}_k^{(m)}$ is stable under the action of this group. Thus, it remains to show how $\mathcal{H}_k^{(m)}$ decomposes into irreps. First, consider the subspace $\mathcal{H}_1^{\otimes k} \otimes \mathcal{H}_0^{\otimes(m-k)} \subseteq \mathcal{H}_k^{(m)}$, which is stable under the action of $\mathbf{U}(\mathcal{H}_1) \times (\mathbb{S}_{[1..k]} \times \mathbb{S}_{[k+1..m]})$ and therefore defines a representation of this group. Note that $\mathbf{U}(\mathcal{H}_1)$ and $\mathbb{S}_{[k+1..m]}$ act trivially on the last $m - k$ multipliers in the tensor product. The largest value of m we care about will be $2n$, thus, we have $\dim \mathcal{H}_1 = q - 1 \geq m$. Hence, the Schur–Weyl duality (Theorem 1.14) says that this representation decomposes as a direct sum of the irreps $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^\lambda \times \mathcal{S}^{(m-k)}$, where the sum is over all $\lambda \vdash k$. Now we induce this representation of $\mathbf{U}(\mathcal{H}_1) \times (\mathbb{S}_{[1..k]} \times \mathbb{S}_{[k+1..m]})$ on $\mathcal{H}_1^{\otimes k} \otimes \mathcal{H}_0^{\otimes(m-k)}$ to a representation of $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_m$ on $\mathcal{H}_k^{(m)}$,¹ and the Littlewood–Richardson rule (1.12) gives us the following:

Lemma A.2. *The subspace $\mathcal{H}_k^{(m)}$ can be decomposed as the direct sum of subspaces $\mathcal{H}_{(m-|\tilde{\lambda}|, \tilde{\lambda})}^\lambda$ corresponding to irreps $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-|\tilde{\lambda}|, \tilde{\lambda})}$ of $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_m$, where the sum is taken over all $\lambda \vdash k$ and $\tilde{\lambda} \in \Lambda(\lambda)$ such that $m - |\tilde{\lambda}| > \tilde{\lambda}_1$.*

In particular, $\mathcal{H}_k^{(m)}$ as a representation of $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_m$ is multiplicity-free.

Let \mathbb{S}_{m-1} stand for $\mathbb{S}_{[2..m]}$. For $\sigma \vdash m$, define $\mathcal{H}_{k,\sigma}^{(m)}$ to be the subspace of $\mathcal{H}_k^{(m)}$ spanned by the irreps of \mathbb{S}_m isomorphic to \mathcal{S}^σ . From Lemma A.2, we get that $\mathcal{H}_{k,(m-k,\lambda)}^{(m)} = \mathcal{H}_{(m-k,\lambda)}^\lambda$, that is, it is the subspace corresponding to the unique copy of the irrep $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k,\lambda)}$ of $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_m$ appearing in $\mathcal{H}_k^{(m)}$. We want to see how the subspace $\mathcal{H}_k^{(m)}$ divides into the irreps

¹Technically, we defined the induction only for finite groups, while neither $\mathbf{U}(\mathcal{H}_1) \times (\mathbb{S}_{[1..k]} \times \mathbb{S}_{[k+1..m]})$ nor $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_m$ is finite. However, in this case, we are essentially inducing from $\mathbb{S}_{[1..k]} \times \mathbb{S}_{[k+1..m]}$ to \mathbb{S}_m , and one can formally verify that Lemma A.2 resulting from this induction is correct.

of $U(\mathcal{H}_1) \times \mathbb{S}_{m-1}$. Unlike in the case of $U(\mathcal{H}_1) \times \mathbb{S}_m$, the representation is not multiplicity-free any longer. Therefore, when there are multiple copies of the same irrep in the representation, we would like to have a way to address a single copy. Let us consider two ways of doing that.

Way 1: restricting an irrep of $U(\mathcal{H}_1) \times \mathbb{S}_m$. We can take an irrep $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^\sigma$ of $U(\mathcal{H}_1) \times \mathbb{S}_m$, corresponding to the subspace $\mathcal{H}_\sigma^\lambda$, and restrict it to $U(\mathcal{H}_1) \times \mathbb{S}_{m-1}$. The restricted representation will be a direct sum over all irreps $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{\sigma^-}$ such that $\sigma^- \subset \sigma$, and we denote the space corresponding to such an irrep by $\mathcal{H}_{\sigma^-; \sigma}^\lambda$. Thus, we may obtain multiple copies of the same irrep, but all of them are uniquely labeled.

The specific irreps we are interested in appear in $\mathcal{H}_k^{(m)}$ only once or twice. Irreps of the form $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k-1, \lambda)}$ appear only once as they are only present in the restriction of $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda)}$. For $\lambda^- \subset \lambda$, irreps of the form $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda^-)}$ appear twice as they are present in the restrictions of $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k+1, \lambda^-)}$ and $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda)}$. For brevity, let us use denotations

$$\begin{aligned}\hat{\mathcal{H}}_{(m-k-1, \lambda)}^\lambda &:= \mathcal{H}_{(m-k-1, \lambda); (m-k, \lambda)}^\lambda, \\ \check{\mathcal{H}}_{(m-k, \lambda^-)}^\lambda &:= \mathcal{H}_{(m-k, \lambda^-); (m-k+1, \lambda^-)}^\lambda, \\ \hat{\mathcal{H}}_{(m-k, \lambda^-)}^\lambda &:= \mathcal{H}_{(m-k, \lambda^-); (m-k, \lambda)}^\lambda.\end{aligned}$$

We have

$$\mathcal{H}_{(m-k, \lambda)}^\lambda = \hat{\mathcal{H}}_{(m-k-1, \lambda)}^\lambda \oplus \bigoplus_{\lambda^- \subset \lambda} \hat{\mathcal{H}}_{(m-k, \lambda^-)}^\lambda. \quad (\text{A.3})$$

Way 2: considering $\mathcal{H}_0 \otimes \mathcal{H}_k^{(m-1)}$ and $\mathcal{H}_1 \otimes \mathcal{H}_{k-1}^{(m-1)}$ separately. Notice that, for $k \neq 0$, we have

$$\mathcal{H}_k^{(m)} = \mathcal{H}_0 \otimes \mathcal{H}_k^{(m-1)} \oplus \mathcal{H}_1 \otimes \mathcal{H}_{k-1}^{(m-1)},$$

and both $\mathcal{H}_0 \otimes \mathcal{H}_k^{(m-1)}$ and $\mathcal{H}_1 \otimes \mathcal{H}_{k-1}^{(m-1)}$ are stable under $U(\mathcal{H}_1) \times \mathbb{S}_{m-1}$, thus defining two representations of the group. Let us first consider $\mathcal{H}_0 \otimes \mathcal{H}_k^{(m-1)}$. The action of $U(\mathcal{H}_1) \times \mathbb{S}_{m-1}$ on $\mathcal{H}_0 \otimes \mathcal{H}_k^{(m-1)}$ is isomorphic to its action on $\mathcal{H}_k^{(m-1)}$. In turn, Lemma A.2 describes how $\mathcal{H}_k^{(m-1)}$ decomposes into irreps of $U(\mathcal{H}_1) \times \mathbb{S}_{m-1}$. Hence, $\mathcal{H}_0 \otimes \mathcal{H}_k^{(m-1)}$ decomposes as the direct sum of spaces $\mathcal{H}_0 \otimes \tilde{\mathcal{H}}_{(m-|\tilde{\lambda}|-1, \tilde{\lambda})}^\lambda$ corresponding to irreps $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-|\tilde{\lambda}|-1, \tilde{\lambda})}$ of $U(\mathcal{H}_1) \times \mathbb{S}_{m-1}$, where the sum is taken over all $\lambda \vdash k$ and $\tilde{\lambda} \in \Lambda(\lambda)$.

This means that the representation of $U(\mathcal{H}_1) \times \mathbb{S}_{m-1}$ defined by its action on $\mathcal{H}_0 \otimes \mathcal{H}_k^{(m-1)}$ is multiplicity-free. We denote the subspace of $\mathcal{H}_0 \otimes \mathcal{H}_k^{(m-1)}$ corresponding to the irrep $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{\sigma'}$ by $\mathcal{H}_{0, \sigma'}^\lambda$. Since $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k-1, \lambda)}$ is contained in both $\mathcal{H}_k^{(m)}$ and $\mathcal{H}_0 \otimes \mathcal{H}_k^{(m-1)}$ exactly once, we have

$$\hat{\mathcal{H}}_{(m-k-1, \lambda)}^\lambda = \mathcal{H}_{0, (m-k-1, \lambda)}^\lambda = \mathcal{H}_0 \otimes \tilde{\mathcal{H}}_{(m-k-1, \lambda)}^\lambda.$$

For the irrep $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda^-)}$, where $\lambda^- \subset \lambda$, the situation is slightly more complicated as $\mathcal{H}_k^{(m)}$ contains two copies of it: one in $\mathcal{H}_0 \otimes \mathcal{H}_k^{(m-1)}$ and one in $\mathcal{H}_1 \otimes \mathcal{H}_{k-1}^{(m-1)}$. Let us denote the copy corresponding to the latter by $\mathcal{H}_{1, (m-k, \lambda^-)}^\lambda$. We have

$$\hat{\Pi}_{(m-k, \lambda^-)}^\lambda + \check{\Pi}_{(m-k, \lambda^-)}^\lambda = \Pi_{0, (m-k, \lambda^-)}^\lambda + \Pi_{1, (m-k, \lambda^-)}^\lambda. \quad (\text{A.4})$$

Overlaps of the copies of the same irrep. We would like to calculate what is the overlap between the subspaces $\hat{\mathcal{H}}_{(m-k, \lambda^-)}^\lambda$ and $\mathcal{H}_{0, (m-k, \lambda^-)}^\lambda$. The following lemma puts an upper bound on it:

Lemma A.3. *Let $d_{\lambda^-}^\lambda(m)$ be the distance in the $(m+1)$ -box Young diagram $(m-k+1, \lambda)$ between the two boxes we have to remove in order to obtain $(m-k, \lambda^-)$. Then*

$$\text{Tr}(\hat{\Pi}_{(m-k, \lambda^-)}^\lambda \Pi_{0, (m-k, \lambda^-)}^\lambda) \leq \frac{1}{d_{\lambda^-}^\lambda(m)} \dim(\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda^-)}). \quad (\text{A.5})$$

We leave the proof of Lemma A.3 to Appendix A.3.

One can see that $d_{\lambda^-}^\lambda(m) \geq m - 2k + 2$, with equality achieved by $\lambda = (k)$ and $\lambda^- = (k-1)$. Since we consider $k = o(m)$, we have

$$\frac{1}{d_{\lambda^-}^\lambda(m)} = \frac{1}{m} + o\left(\frac{1}{m}\right).$$

Let $\Xi_{1 \leftarrow 0, (m-k, \lambda^-)}^\lambda$ be the transporter from the copy of the irrep $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda^-)}$ corresponding to the subspace $\mathcal{H}_{0, (m-k, \lambda^-)}^\lambda$ to the copy corresponding to the subspace $\mathcal{H}_{1, (m-k, \lambda^-)}^\lambda$, and $\Xi_{0 \leftarrow 1, (m-k, \lambda^-)}^\lambda := (\Xi_{1 \leftarrow 0, (m-k, \lambda^-)}^\lambda)^*$. From (A.4) and (A.5) we get

$$\hat{\Pi}_{(m-k, \lambda^-)}^\lambda = \left(1 - O\left(\frac{1}{m}\right)\right) \Pi_{1, (m-k, \lambda^-)}^\lambda + O\left(\frac{1}{m}\right) \Pi_{0, (m-k, \lambda^-)}^\lambda + O\left(\frac{1}{\sqrt{m}}\right) \Xi_{(m-k, \lambda^-)}^\lambda, \quad (\text{A.6})$$

where

$$\Xi_{(m-k, \lambda^-)}^\lambda := \Xi_{1 \leftarrow 0, (m-k, \lambda^-)}^\lambda + \Xi_{0 \leftarrow 1, (m-k, \lambda^-)}^\lambda,$$

which has the norm one.

Connection between $\mathcal{H}_1 \otimes \mathcal{H}_{k-1}^{(m-1)}$ and $\mathcal{H}_0 \otimes \mathcal{H}_{k-1}^{(m-1)}$. The representation of $\text{U}(\mathcal{H}_1) \times \mathbb{S}_{m-1}$ on $\mathcal{H}_1 \otimes \mathcal{H}_{k-1}^{(m-1)}$ is not necessarily multiplicity-free. Nor, for every irrep contained in it, the corresponding subspace can be written in the form $\mathcal{H}_1 \otimes \mathcal{K}$ for some space \mathcal{K} . However, we have the following useful result.

Lemma A.4. For any $\lambda' \vdash k-1$ we have

$$\bigoplus_{\lambda \supset \lambda'} \mathcal{H}_{1,(m-k,\lambda')}^{\lambda^+} = \mathcal{H}_1 \otimes \tilde{\mathcal{H}}_{(m-k,\lambda')}^{\lambda'}.$$

Proof. Both spaces $\mathcal{H}_1 \otimes \mathcal{H}_{k-1}^{(m-1)}$ and $\mathcal{H}_0 \otimes \mathcal{H}_{k-1}^{(m-1)}$ are stable under \mathbb{S}_{m-1} , and thus are representations of this group. These representations decompose as $\bigoplus_{\sigma} \mathcal{H}_1 \otimes \mathcal{K}'_{\sigma}$ and $\bigoplus_{\sigma} \mathcal{H}_0 \otimes \mathcal{K}''_{\sigma}$, respectively, where $\sigma \vdash m-1$ and $\mathcal{H}_1 \otimes \mathcal{K}'_{\sigma}$ and $\mathcal{H}_0 \otimes \mathcal{K}''_{\sigma}$ are the \mathcal{S}^{σ} -isotypical subspaces. Since we ignore the action of $U(\mathcal{H}_1)$ here, the first space in the tensor products plays no role, and we have $\mathcal{K}'_{\sigma} = \mathcal{K}''_{\sigma} = \mathcal{H}_{k-1,\sigma}^{(m-1)}$. Finally, on the one hand side, we have

$$\mathcal{H}_1 \otimes \mathcal{K}'_{(m-k,\lambda')} = \bigoplus_{\lambda \supset \lambda'} \mathcal{H}_{1,(m-k,\lambda')}^{\lambda^+}$$

and, on the other,

$$\mathcal{H}_0 \otimes \mathcal{K}''_{(m-k,\lambda')} = \mathcal{H}_{0,(m-k,\lambda')}^{\lambda'} = \mathcal{H}_0 \otimes \tilde{\mathcal{H}}_{(m-k,\lambda')}^{\lambda'}.$$

□

Putting everything together. We have

$$\begin{aligned} \bar{\Pi}_k^{(m)} &= \sum_{\lambda \vdash k} \Pi_{(m-k,\lambda)}^{\lambda} = \sum_{\lambda \vdash k} \hat{\Pi}_{(m-k-1,\lambda)}^{\lambda} + \sum_{\lambda \vdash k} \sum_{\lambda^- \subset \lambda} \hat{\Pi}_{(m-k,\lambda^-)}^{\lambda} \\ &= \sum_{\lambda \vdash k} \hat{\Pi}_{(m-k-1,\lambda)}^{\lambda} + \sum_{\lambda^- \vdash k-1} \sum_{\lambda \supset \lambda^-} \left(\left(1 - O\left(\frac{1}{m}\right)\right) \Pi_{1,(m-k,\lambda^-)}^{\lambda} + \right. \\ &\quad \left. + O\left(\frac{1}{m}\right) \Pi_{0,(m-k,\lambda^-)}^{\lambda} + O\left(\frac{1}{\sqrt{m}}\right) \Xi_{(m-k,\lambda^-)}^{\lambda} \right) \\ &= \Pi_0 \otimes \sum_{\lambda \vdash k} \tilde{\Pi}_{(m-k-1,\lambda)}^{\lambda} + \Pi_1 \otimes \sum_{\lambda \vdash k-1} \tilde{\Pi}_{(m-1-|\lambda|,\lambda')}^{\lambda'} + \Phi_k^{(m)} \\ &= \Pi_0 \otimes \bar{\Pi}_k^{(m-1)} + \Pi_1 \otimes \bar{\Pi}_{k-1}^{(m-1)} + \Phi_k^{(m)}, \end{aligned}$$

where the second equality comes from (A.3), the third equality comes from (A.6), the fourth equality comes from Lemma A.4, and

$$\Phi_k^{(m)} = \sum_{\lambda \vdash k} \sum_{\lambda^- \subset \lambda} \left(-O\left(\frac{1}{m}\right) \Pi_{1,(m-k,\lambda^-)}^{\lambda} + O\left(\frac{1}{m}\right) \Pi_{0,(m-k,\lambda^-)}^{\lambda} + O\left(\frac{1}{\sqrt{m}}\right) \Xi_{(m-k,\lambda^-)}^{\lambda} \right). \quad (\text{A.7})$$

The norm of $\Phi_k^{(m)}$ is in $O(1/\sqrt{m})$, because the operators in the brackets of (A.7) are orthogonal for different pairs of λ and λ^- . Also, one can see that the support of $\Phi_k^{(m)}$ is contained in $\bar{\mathcal{H}}_k^{(m)}$.

A.3 Proof of Lemma A.3

We use inductive argument on m . Consider the groups

$$\mathbf{U}(\mathcal{H}_1), \quad \mathbb{S}_{m-1} := \mathbb{S}_{[m-1]}, \quad \mathbb{S}_m := \mathbb{S}_{[m]}, \quad \text{and} \quad \mathbb{S}_{m+1} := \mathbb{S}_{[m+1]}$$

and their action on $\mathcal{H}^{\otimes(m+1)}$. The representation of $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_{m-1}$ on $\mathcal{H}^{\otimes(m+1)}$ contains three copies of the irrep $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda^-)}$. Using notations similar to Appendix A.2, they are

$$\begin{aligned} \mathcal{H}_a &:= \mathcal{H}_{(m-k, \lambda^-); (m-k+1, \lambda^-); (m-k+2, \lambda^-)}^\lambda, \\ \mathcal{H}_b &:= \mathcal{H}_{(m-k, \lambda^-); (m-k+1, \lambda^-); (m-k+1, \lambda)}^\lambda, \\ \mathcal{H}_c &:= \mathcal{H}_{(m-k, \lambda^-); (m-k, \lambda); (m-k+1, \lambda)}^\lambda, \end{aligned}$$

where we first restrict an irrep of $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_{m+1}$ to $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_m$, and then restrict the irreps of this restriction further to $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_{m-1}$. Let us consider few other ways how to address copies of $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda^-)}$. The representation of $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_m$ on $\mathcal{H}^{\otimes m} \otimes \mathcal{H}_0$ contains both irreps $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k+1, \lambda^-)}$ and $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda)}$ exactly once, and their restrictions to $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_{m-1}$ each contains a unique copy of $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda^-)}$. Let us denote these copies

$$\begin{aligned} \mathcal{H}'_d &:= \mathcal{H}_{(m-k, \lambda^-); (m-k+1, \lambda^-); m+1}^\lambda, \\ \mathcal{H}'_c &:= \mathcal{H}_{(m-k, \lambda^-); (m-k, \lambda); m+1}^\lambda, \end{aligned}$$

respectively. Since $\mathcal{H}^{\otimes(m+1)}$ also contains only one copy of $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda)}$, we have $\mathcal{H}_c = \mathcal{H}'_c$. Finally, $\mathcal{H}^{\otimes(m-1)} \otimes \mathcal{H}_0^{\otimes 2}$ contains a unique copy of $\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda^-)}$, which we denote by

$$\mathcal{H}''_e := \mathcal{H}_{(m-k, \lambda^-); m; m+1}^\lambda.$$

Now let us consider overlaps of these irreps. Let

$$\gamma_{\lambda^-}^\lambda(m) := \frac{\text{Tr}(\Pi_{(m-k, \lambda^-); (m-k, \lambda); m+1}^\lambda \Pi_{(m-k, \lambda^-); m; m+1}^\lambda)}{\dim(\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda^-)})} = \frac{\text{Tr}(\Pi'_c \Pi''_e)}{\dim(\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda^-)})} \in [0, 1]$$

and, consistently, let

$$\gamma_{\lambda^-}^\lambda(m+1) := \frac{\text{Tr}(\Pi_{(m-k+1, \lambda^-); (m-k+1, \lambda)}^\lambda \Pi_{(m-k+1, \lambda^-); m+1}^\lambda)}{\dim(\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k+1, \lambda^-)})} = \frac{\text{Tr}(\Pi_b \Pi'_d)}{\dim(\mathcal{W}_{q-1}^\lambda \times \mathcal{S}^{(m-k, \lambda^-)})} \in [0, 1],$$

where the last equality comes from the fact that, if two copies of the same irrep have a certain overlap, then the unique irreps in their restrictions also have the same overlap. Let $\gamma_m := \gamma_{\lambda^-}^\lambda(m)$ and $\gamma_{m+1} := \gamma_{\lambda^-}^\lambda(m+1)$ for short. Thus, due to the orthogonality of Π'_d and Π_c , we have

$$\Pi'_d = \begin{pmatrix} (1 - \gamma_{m+1})\Pi_a & \sqrt{\gamma_{m+1}(1 - \gamma_{m+1})}\Xi_{a \leftarrow b} \\ \sqrt{\gamma_{m+1}(1 - \gamma_{m+1})}\Xi_{b \leftarrow a} & \gamma_{m+1}\Pi_b \end{pmatrix},$$

which implies that Π_e'' is equal to

$$\begin{pmatrix} (1-\gamma_m)(1-\gamma_{m+1})\Pi_a & (1-\gamma_m)\sqrt{\gamma_{m+1}(1-\gamma_{m+1})}\Xi_{a\leftarrow b} & \sqrt{\gamma_m(1-\gamma_m)(1-\gamma_{m+1})}\Xi_{a\leftarrow c} \\ (1-\gamma_m)\sqrt{\gamma_{m+1}(1-\gamma_{m+1})}\Xi_{b\leftarrow a} & (1-\gamma_m)\gamma_{m+1}\Pi_b & \sqrt{\gamma_m(1-\gamma_m)\gamma_{m+1}}\Xi_{b\leftarrow c} \\ \sqrt{\gamma_m(1-\gamma_m)(1-\gamma_{m+1})}\Xi_{c\leftarrow a} & \sqrt{\gamma_m(1-\gamma_m)\gamma_{m+1}}\Xi_{c\leftarrow b} & \gamma_m\Pi_c \end{pmatrix}$$

(here, $\Xi_{a\leftarrow b}, \Xi_{a\leftarrow c}, \dots, \Xi_{c\leftarrow b}$ are transporters between different copies of the irrep).

Let d_m be equal to $d_{\lambda^-}^\lambda(m)$ in Lemma A.3, that is, let d_m be the distance in $(m-k+1, \lambda)$ between the two boxes we have to remove in order to obtain $(m-k, \lambda^-)$. Consistently, let $d_{m+1} = d_m + 1$ be the distance in $(m-k+2, \lambda)$ between the two boxes we have to remove in order to obtain $(m-k+1, \lambda^-)$. We have to prove that $\gamma_m \leq 1/d_m$. If $\gamma_m = 0$, we are done, so assume $\gamma_m > 0$.

Claim A.5. *We have $\gamma_m \neq 1$ and*

$$\gamma_{m+1} = \frac{d_m - 1}{(d_m + 1)(1/\gamma_m - 1)}.$$

Proof. Let $V_{(m,m+1)}$ be the operator permuting the last two instances of the space \mathcal{H} in $\mathcal{H}^{\otimes(m+1)}$. Any space corresponding to an irrep of $\mathbf{U}(\mathcal{H}_1) \times \mathbb{S}_{m+1}$ is stable under $V_{(m,m+1)}$, and $\mathcal{H}_{(m-k+1, \lambda)}^\lambda$ is such a space. This, in turn, means that the space $\mathcal{H}_b \oplus \mathcal{H}_c$ is stable under $V_{(m,m+1)}$ as well, and the action of $V_{(m,m+1)}$ on this space is given by the orthogonal form of irrep $S^{(m-k+1, \lambda)}$ (see Section 1.4.4), namely,

$$V_{(m,m+1)}|_{\mathcal{H}_b \oplus \mathcal{H}_c} = \begin{pmatrix} -\frac{1}{d_m}\Pi_b & \sqrt{1 - \frac{1}{d_m^2}}\Xi_{b\leftarrow c} \\ \sqrt{1 - \frac{1}{d_m^2}}\Xi_{c\leftarrow b} & \frac{1}{d_m}\Pi_c \end{pmatrix}.$$

Since $V_{(m,m+1)}(\mathbb{I}_{\mathcal{H}}^{\otimes(m-1)} \otimes \Pi_0^{\otimes 2}) = \mathbb{I}_{\mathcal{H}}^{\otimes(m-1)} \otimes \Pi_0^{\otimes 2}$, we also have $V_{(m,m+1)}\Pi_e'' = \Pi_e''$. Thus, we have

$$\begin{aligned} & \begin{pmatrix} -\frac{1}{d_m} & \sqrt{1 - \frac{1}{d_m^2}} \\ \sqrt{1 - \frac{1}{d_m^2}} & \frac{1}{d_m} \end{pmatrix} \begin{pmatrix} (1-\gamma_m)\gamma_{m+1} & \sqrt{\gamma_m(1-\gamma_m)\gamma_{m+1}} \\ \sqrt{\gamma_m(1-\gamma_m)\gamma_{m+1}} & \gamma_m \end{pmatrix} = \\ & = \begin{pmatrix} (1-\gamma_m)\gamma_{m+1} & \sqrt{\gamma_m(1-\gamma_m)\gamma_{m+1}} \\ \sqrt{\gamma_m(1-\gamma_m)\gamma_{m+1}} & \gamma_m \end{pmatrix}, \end{aligned}$$

which clearly cannot hold for $\gamma_m = 1$. Simple further calculation proves the claim. \square

Suppose the contrary: $\gamma_m > 1/d_m$. Then there exists $c > 0$ such that $\gamma_m = 1/(d_m - \binom{d_m}{2}c)$. Claim A.5 gives us

$$\gamma_{m+1} = \frac{d_m - 1}{(d_m + 1)(d_m - \binom{d_m}{2}c - 1)} = \frac{1}{d_m + 1 - \frac{d_m+1}{d_m-1}\binom{d_m}{2}c} = \frac{1}{d_{m+1} - \binom{d_{m+1}}{2}c}.$$

By repeating the same argument, we get that $\gamma_{m'} = 1/(d_{m'} - \binom{d_{m'}}{2}c)$ for all $m' \geq m$. However, since $d_{m'}$ grows linearly with m' , but $\binom{d_{m'}}{2}$ quadratically, there exists m' such that $\gamma_{m'} \notin [0, 1]$, which is a contradiction.

Appendix B

Necessary conditions on the adversary matrix for Element Distinctness with small range

B.1 Action of Δ_i on Π_λ^λ and transporters

Let us consider $i \neq 2$. Recall the projectors $\hat{\Pi}_i^s$ from Section 5.4.2, and note that $V_\pi^\tau \hat{\Pi}_i^s = \hat{\Pi}_i^s V_\pi^\tau$ for all $(\pi, \tau) \in \mathbb{S}_{[n] \setminus \{i\}} \times \mathbb{S}_{\Sigma \setminus \{s\}}$. Analogously to Claim 5.1,

$$\hat{\Pi}_i^s = \sum_{\mu \vdash n-1} \hat{\Pi}_{i, \mu_i}^{s, \mu_s},$$

where $\hat{\Pi}_{i, \mu_i}^{s, \mu_s} := \hat{\Pi}_i^s \Pi_{\mu_i}^{\mu_s} = \Pi_{\mu_i}^{\mu_s} \hat{\Pi}_i^s$ projects on a single instance of the irrep $\mathcal{S}^\mu \times \mathcal{S}^\mu$ of $\mathbb{S}_{[n] \setminus \{i\}} \times \mathbb{S}_{\Sigma \setminus \{s\}}$.

Due to the symmetry, $V_\pi^\tau (\bar{\Delta}_i \circ \Pi_\lambda^\lambda) = (\bar{\Delta}_i \circ \Pi_\lambda^\lambda) V_\pi^\tau$ for all $(\pi, \tau) \in \mathbb{S}_{[n] \setminus \{i\}} \times \mathbb{S}_\Sigma$, therefore we can express

$$\bar{\Delta}_i \circ \Pi_\lambda^\lambda = \sum_{\lambda' \vdash n} \sum_{\mu \subset \lambda'} \phi_\mu^{\lambda'} \Pi_{\mu_i}^{\lambda'}.$$

We have

$$\begin{aligned} \phi_\mu^{\lambda'} &= \frac{\text{Tr}((\bar{\Delta}_i \circ \Pi_\lambda^\lambda) \Pi_{\mu_i}^{\lambda'})}{\text{Tr}(\Pi_{\mu_i}^{\lambda'})} = \frac{\text{Tr}(\sum_{s \in \Sigma} \hat{\Pi}_i^s \Pi_\lambda^\lambda \hat{\Pi}_i^s \Pi_{\mu_i}^{\lambda'})}{\text{dim } \lambda' \text{ dim } \mu} \\ &= n \frac{\text{Tr}(\hat{\Pi}_{i, \mu_i}^{s, \mu_s} \Pi_{\mu_i}^\lambda \hat{\Pi}_{i, \mu_i}^{s, \mu_s} \Pi_{\mu_i}^{\lambda'})}{\text{dim } \lambda' \text{ dim } \mu} = n \frac{\text{Tr}(\hat{\Pi}_{i, \mu_i}^{s, \mu_s} \Pi_{\mu_i}^\lambda) \cdot \text{Tr}(\hat{\Pi}_{i, \mu_i}^{s, \mu_s} \Pi_{\mu_i}^{\lambda'})}{\text{dim } \lambda' (\text{dim } \mu)^3} \\ &= n \frac{\text{Tr}(\hat{\Pi}_i^s \Pi_{\mu_i}^\lambda) \cdot \text{Tr}(\hat{\Pi}_i^s \Pi_{\mu_i}^{\lambda'})}{\text{dim } \lambda' (\text{dim } \mu)^3} = \frac{\text{Tr}(\Pi_{\mu_i}^\lambda) \cdot \text{Tr}(\Pi_{\mu_i}^{\lambda'})}{n \text{ dim } \lambda' (\text{dim } \mu)^3} = \begin{cases} \frac{\text{dim } \lambda}{n \text{ dim } \mu}, & \text{if } \mu \subset \lambda, \\ 0, & \text{if } \mu \not\subset \lambda \text{ (i.e., } \Pi_{\mu_i}^\lambda = 0), \end{cases} \end{aligned}$$

where the second equality is due to (5.11), the third and sixth equalities are due to the symmetry among all $s \in \Sigma$, and the fourth equality is from Claim 1.6. Hence

$$\Delta_i \circ \Pi_\lambda^\lambda = \Pi_\lambda^\lambda - \frac{\dim \lambda}{n} \sum_{\mu \subset \lambda} \left(\frac{1}{\dim \mu} \sum_{\lambda' \supset \mu} \Pi_{\mu_i}^{\lambda'} \right) = \Pi_\lambda^\lambda - \frac{\dim \lambda}{n} \sum_{\mu \subset \lambda} \left(\frac{1}{\dim \mu} \Pi_{\mu_i} \right). \quad (\text{B.1})$$

Now consider $j \neq i$, $\lambda \vdash n$, and $\nu \subset_{rc} \lambda$. Let $\mu, \mu' \vdash n-1$ be such that $\nu \subset \mu \subset \lambda$, $\nu \subset \mu' \subset \lambda$, and $\mu \neq \mu'$. Let us see how $\overline{\Delta}_i$ acts on the transporter $\Pi_{\nu_{ij}, \mu'_i \leftarrow \nu_{ij}, \mu_i}^\lambda$. We have

$$\hat{\Pi}_i^s \Pi_{\nu_{ij}, \mu'_i \leftarrow \nu_{ij}, \mu_i}^\lambda \hat{\Pi}_i^s = \hat{\Pi}_i^s \Pi_{\mu'_i}^{\mu'_s} \Pi_{\nu_{ij}, \mu'_i \leftarrow \nu_{ij}, \mu_i}^\lambda \Pi_{\mu_i}^{\mu_s} \hat{\Pi}_i^s = 0$$

because $\Pi_{\nu_{ij}, \mu'_i \leftarrow \nu_{ij}, \mu_i}^{\mu'_s} \Pi_{\mu_i}^{\mu_s}$ is a transporter between two instances of the irrep $\mathcal{S}^\nu \times \mathcal{S}^{\mu'}$ of $\mathbb{S}_{[n] \setminus \{i, j\}} \times \mathbb{S}_{\Sigma \setminus \{s\}}$ and, therefore, orthogonal to Π^{μ_s} . Hence,

$$\overline{\Delta}_i \circ \Pi_{\nu_{ij}, \mu'_i \leftarrow \nu_{ij}, \mu_i}^\lambda = 0 \quad \text{and} \quad \Delta_i \circ \Pi_{\nu_{ij}, \mu'_i \leftarrow \nu_{ij}, \mu_i}^\lambda = \Pi_{\nu_{ij}, \mu'_i \leftarrow \nu_{ij}, \mu_i}^\lambda. \quad (\text{B.2})$$

B.2 Necessary conditions for $\|\Delta_1 \circ \Gamma\| = O(1)$

We will use the following lemmas and corollaries in the proof of Claim 5.4. Let $\Gamma_{1,2}$ be given as in (5.7), and Γ be obtained from $\Gamma_{1,2}$ via (5.5).

Lemma B.1. *Consider $\lambda \vdash n$, $\mu \subset \lambda$, $\mu' \subset \lambda$, and $\nu \subset \mu, \mu'$ (we allow $\mu = \mu'$ here). If $\|\Delta_1 \circ \Gamma'\| \leq 1$, then*

$$\|\Pi_{\nu_{12}, \mu_1}^\lambda (\Delta_1 \circ \Gamma_{1,2}) \Pi_{\nu_{12}, \mu'_1}^\lambda\| \leq \sqrt{\frac{\dim \mu'}{(n-1) \dim \nu}}.$$

Proof. For the proof, let us assume that $\nu \subset_{rc} \lambda$ and $\mu \neq \mu'$. It is easy to see that the proof works in all the other cases too. Let $\Psi_{\nu, \mu}^\lambda := \sum_{\pi \in R'} U_\pi \Pi_{\nu_{12}, \mu_1}^\lambda U_{\pi^{-1}}$, where the transversal R' was defined in Section 5.4.1. From (5.8), we have

$$\Psi_{\nu, \mu}^\lambda (\Delta_1 \circ \Gamma') = \sum_{\pi \in R'} U_\pi \Pi_{\nu_{12}, \mu_1}^\lambda (\Delta_1 \circ \Gamma_{1,2}) V_{\pi^{-1}}, \quad (\text{B.3})$$

whose norm is at most 1 because $\Psi_{\nu, \mu}^\lambda$ is a projector.

We can express

$$\Pi_{\nu_{12}, \mu_1}^\lambda (\Delta_1 \circ \Gamma_{1,2}) = \psi \Pi_{\nu_{12}, \mu_1}^\lambda + \psi' \Pi_{\nu_{12}, \mu_1 \leftarrow \nu_{12}, \mu'_1}^\lambda,$$

where

$$\psi := \|\Pi_{\nu_{12}, \mu_1}^\lambda (\Delta_1 \circ \Gamma_{1,2}) \Pi_{\nu_{12}, \mu_1}^\lambda\| \quad \text{and} \quad \psi' := \|\Pi_{\nu_{12}, \mu_1}^\lambda (\Delta_1 \circ \Gamma_{1,2}) \Pi_{\nu_{12}, \mu'_1}^\lambda\|.$$

Hence,

$$(\Delta_1 \circ \Gamma_{1,2})^* \Pi_{\nu_{12}, \mu_1}^\lambda (\Delta_1 \circ \Gamma_{1,2}) = \psi^2 \Pi_{\nu_{12}, \mu_1}^\lambda + (\psi')^2 \Pi_{\nu_{12}, \mu'_1}^\lambda + \psi \psi' \Pi_{\nu_{12}, \mu_1 \leftrightarrow \nu_{12}, \mu'_1}^\lambda. \quad (\text{B.4})$$

From (B.3), (B.4), and (5.18), we get

$$(\Delta_1 \circ \Gamma')^* \Psi_{\nu, \mu}^\lambda (\Delta_1 \circ \Gamma') = \psi^2 (n-1) \frac{\dim \nu}{\dim \mu} \Pi_\mu^\lambda + (\psi')^2 (n-1) \frac{\dim \nu}{\dim \mu'} \Pi_{\mu'}^\lambda.$$

The norm of this matrix is at most 1, which completes the proof. \square

Corollary B.2. *Let $\nu \vdash n-2$, $\mu \supset \nu$, and $\lambda, \lambda' \supset \mu$. If $\|\Delta_1 \circ \Gamma'\| \leq 1$, then*

$$\left| \frac{\text{Tr}(\Pi_{\nu_{12}, \mu_1}^\lambda \Gamma_{1,2})}{\dim \lambda \dim \nu} - \frac{\text{Tr}(\Pi_{\nu_{12}, \mu_1}^{\lambda'} \Gamma_{1,2})}{\dim \lambda' \dim \nu} \right| \leq 2 \sqrt{\frac{\dim \mu}{(n-1) \dim \nu}}.$$

Proof. From Lemma B.1, we have

$$\begin{aligned} \left\| \Pi_{\nu_{12}, \mu_1}^\lambda (\Delta_1 \circ \Gamma_{1,2}) \Pi_{\nu_{12}, \mu_1}^\lambda \right\| &= \frac{\left| \text{Tr}(\Pi_{\nu_{12}, \mu_1}^\lambda (\Delta_1 \circ \Gamma_{1,2})) \right|}{\dim \lambda \dim \nu} \\ &= \frac{\left| \text{Tr}((\Delta_1 \circ \Pi_{\nu_{12}, \mu_1}^\lambda) \Gamma_{1,2}) \right|}{\dim \lambda \dim \nu} = \frac{\left| \text{Tr}((\Pi_{\nu_{12}, \mu_1}^\lambda - \frac{\dim \lambda}{n \dim \mu} \Pi_{\nu_{12}, \mu_1}) \Gamma_{1,2}) \right|}{\dim \lambda \dim \nu} \\ &= \left| \frac{\text{Tr}(\Pi_{\nu_{12}, \mu_1}^\lambda \Gamma_{1,2})}{\dim \lambda \dim \nu} - \frac{\text{Tr}(\Pi_{\nu_{12}, \mu_1} \Gamma_{1,2})}{n \dim \mu \dim \nu} \right| \leq \sqrt{\frac{\dim \mu}{(n-1) \dim \nu}}, \end{aligned}$$

where the second and third equalities are due to (5.11) and (B.1), respectively. We obtain the same inequality with λ' instead of λ , and the result follows from the triangle inequality. \square

Corollary B.3. *Consider $\lambda \vdash n$, $\nu \subset_{rc} \lambda$, and $\mu, \mu' \vdash n-1$ such that $\nu \subset \mu \subset \lambda$, $\nu \subset \mu' \subset \lambda$, and μ appears after μ' in the lexicographical order. If $\|\Delta_1 \circ \Gamma'\| \leq 1$, then*

$$\begin{aligned} \left| \alpha_{\text{id}, \nu}^\lambda \frac{\sqrt{d_{\lambda, \nu}^2 - 1}}{2d_{\lambda, \nu}} - \alpha_{\text{sgn}, \nu}^\lambda \frac{d_{\lambda, \nu} - 1}{2d_{\lambda, \nu}} \right| &\leq \sqrt{\frac{\dim \mu}{(n-1) \dim \nu}}, \\ \left| \alpha_{\text{id}, \nu}^\lambda \frac{\sqrt{d_{\lambda, \nu}^2 - 1}}{2d_{\lambda, \nu}} + \alpha_{\text{sgn}, \nu}^\lambda \frac{d_{\lambda, \nu} + 1}{2d_{\lambda, \nu}} \right| &\leq \sqrt{\frac{\dim \mu'}{(n-1) \dim \nu}}, \end{aligned}$$

Proof. Since λ is the unique n -box Young diagram that has both μ and μ' as subdiagrams, we have

$$\Pi_{\nu_{12}, \mu'_1} \Gamma_{1,2} \Pi_{\nu_{12}, \mu_1} = \Pi_{\nu_{12}, \mu'_1}^\lambda \Gamma_{1,2} \Pi_{\nu_{12}, \mu_1}^\lambda.$$

Hence, due to (B.2) and the commutativity relations (5.12), we have

$$\Pi_{\nu_{12}, \mu'_1}^\lambda (\Delta_1 \circ \Gamma_{1,2}) \Pi_{\nu_{12}, \mu_1}^\lambda = \Pi^\lambda (\Delta_1 \circ (\Pi_{\nu_{12}, \mu'_1} \Gamma_{1,2} \Pi_{\nu_{12}, \mu_1})) \Pi^\lambda = \Pi_{\nu_{12}, \mu'_1}^\lambda \Gamma_{1,2} \Pi_{\nu_{12}, \mu_1}^\lambda.$$

The same holds with μ and μ' swapped. From (5.14) and (5.15), we get that

$$\begin{aligned} \Pi_{\nu_{12}, \mu'_1}^\lambda \Gamma_{1,2} \Pi_{\nu_{12}, \mu_1}^\lambda &= \left(\alpha_{\text{id}, \nu}^\lambda \frac{\sqrt{d_{\lambda, \nu}^2 - 1}}{2d_{\lambda, \nu}} - \alpha_{\text{sgn}, \nu}^\lambda \frac{d_{\lambda, \nu} - 1}{2d_{\lambda, \nu}} \right) \Pi_{\nu_{12}, \mu'_1 \leftarrow \nu_{12}, \mu_1}^\lambda, \\ \Pi_{\nu_{12}, \mu_1}^\lambda \Gamma_{1,2} \Pi_{\nu_{12}, \mu'_1}^\lambda &= \left(\alpha_{\text{id}, \nu}^\lambda \frac{\sqrt{d_{\lambda, \nu}^2 - 1}}{2d_{\lambda, \nu}} + \alpha_{\text{sgn}, \nu}^\lambda \frac{d_{\lambda, \nu} + 1}{2d_{\lambda, \nu}} \right) \Pi_{\nu_{12}, \mu_1 \leftarrow \nu_{12}, \mu'_1}^\lambda, \end{aligned}$$

and we apply Lemma B.1 to complete the proof. \square

Lemma B.4. *Let θ be a Young diagram having at most $n/2 - 2$ boxes and $\eta \supset \theta$. If $\|\Delta_1 \circ \Gamma''\| \leq 1$, then*

$$\left| \alpha_{\text{id}, \bar{\eta}_{12}}^{\bar{\eta}} - \alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\theta}} + \frac{2(\alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\eta}} - \alpha_{\text{id}, \bar{\eta}_{12}}^{\bar{\eta}})}{d_{\bar{\eta}, \bar{\theta}_{12}}(d_{\bar{\eta}, \bar{\theta}_{12}} - 1)} \right| \leq 2 \sqrt{\frac{\dim \bar{\theta}_3}{\binom{n-1}{2} \dim \bar{\theta}_{123}}}.$$

Proof. Note that $\Pi_{\bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}} (\Delta_3 \circ \Gamma_{1,2})$ can be expressed as a linear combination of $\Pi_{\bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}}$ and $\Pi_{\bar{\theta}_{123}, \bar{\theta}_3 \leftarrow \text{id}, \bar{\theta}_{123}, \bar{\eta}_3}^{\bar{\eta}}$, while $\Pi_{\bar{\theta}_{123}}^{\bar{\theta}} (\Delta_3 \circ \Gamma_{1,2})$ is proportional to $\Pi_{\bar{\theta}_{123}}^{\bar{\theta}}$. Similarly to Lemma B.1, we can show that

$$\left\| \Pi_{\bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}} (\Delta_3 \circ \Gamma_{1,2}) \Pi_{\bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}} \right\| \leq \sqrt{\frac{\dim \bar{\theta}_3}{\binom{n-1}{2} \dim \bar{\theta}_{123}}}, \quad \left\| \Pi_{\bar{\theta}_{123}}^{\bar{\theta}} (\Delta_3 \circ \Gamma_{1,2}) \Pi_{\bar{\theta}_{123}}^{\bar{\theta}} \right\| \leq \sqrt{\frac{\dim \bar{\theta}_3}{\binom{n-1}{2} \dim \bar{\theta}_{123}}},$$

where, instead of (5.18), we have to use (analogously proven)

$$\sum_{\pi \in R''} V_\pi \Pi_{\bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}} V_{\pi^{-1}} = \binom{n-1}{2} \frac{\dim \bar{\theta}_{123}}{\dim \bar{\theta}_3} \Pi_{\bar{\theta}_3}^{\bar{\eta}} \quad \text{and} \quad \sum_{\pi \in R''} V_\pi \Pi_{\text{id}, \bar{\theta}_{123}, \bar{\eta}_3 \leftrightarrow \bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}} V_{\pi^{-1}} = 0.$$

Then, similarly to Corollary B.2, we get

$$\left| \frac{\text{Tr}(\Pi_{\bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}} \Gamma_{1,2})}{\dim \bar{\eta} \dim \bar{\theta}_{123}} - \frac{\text{Tr}(\Pi_{\bar{\theta}_{123}}^{\bar{\theta}} \Gamma_{1,2})}{\dim \bar{\theta} \dim \bar{\theta}_{123}} \right| \leq 2 \sqrt{\frac{\dim \bar{\theta}_3}{\binom{n-1}{2} \dim \bar{\theta}_{123}}}.$$

We conclude by noticing that

$$\Pi_{\bar{\theta}_{123}}^{\bar{\theta}} \Gamma_{1,2} = \Pi_{\bar{\theta}_{123}}^{\bar{\theta}} (\alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\theta}} \Pi_{\bar{\theta}_{12}}^{\bar{\theta}}) = \alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\theta}} \Pi_{\bar{\theta}_{123}}^{\bar{\theta}}$$

and, due to (5.16),

$$\begin{aligned} \Pi_{\bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}} \Gamma_{1,2} \Pi_{\bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}} &= \Pi_{\bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}} \left(\alpha_{\text{id}, \bar{\eta}_{12}}^{\bar{\eta}} \Pi_{\bar{\eta}_{12}}^{\bar{\eta}} + \alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\eta}} \Pi_{\text{id}, \bar{\theta}_{12}}^{\bar{\eta}} \right) \Pi_{\bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}} \\ &= \left(\left(1 - \frac{2}{d_{\bar{\eta}, \bar{\theta}_{12}} (d_{\bar{\eta}, \bar{\theta}_{12}} - 1)} \right) \alpha_{\text{id}, \bar{\eta}_{12}}^{\bar{\eta}} + \frac{2}{d_{\bar{\eta}, \bar{\theta}_{12}} (d_{\bar{\eta}, \bar{\theta}_{12}} - 1)} \alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\eta}} \right) \Pi_{\bar{\theta}_{123}, \bar{\theta}_3}^{\bar{\eta}}. \end{aligned}$$

□

B.3 Proof of Claim 5.4

We can assume that all the coefficients β in the expression (5.4) for Γ are at most n , as n is the trivial upper bound on the quantum query complexity of ELEMENT DISTINCTNESS. That, in turn, means that we can assume that the coefficients α in Point 1, Point 2, and Point 3 of Claim 5.4 are, respectively, at most $O(1)$, $O(\sqrt{n})$, and $O(n)$. Let us prove sequentially every point of the claim.

Point 1. Consider $k = O(1)$, $\theta \vdash k$, and $\eta \supset \theta$, so $d_{\bar{\eta}, \bar{\theta}_{12}} = n - O(1)$ and $\dim \bar{\theta}_3 / \dim \bar{\theta}_{123} = \Theta(1)$. From Lemma B.4, we get that $|\alpha_{\text{id}, \bar{\eta}_{12}}^{\bar{\eta}} - \alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\theta}}| = O(1/n)$, which proves that $\alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\theta}} = n^{-1/3} + O(1/n)$ by the induction over k , where we take $\alpha_{\text{id}, (n-2)}^{(n)} = n^{-1/3}$ as the base case.

Point 2. Consider $\theta \vdash O(1)$ and $\eta \supset \theta$, so $\dim \bar{\theta}_1 / \dim \bar{\theta}_{12} = \Theta(1)$. From the first inequality of Corollary B.3 (in which we choose $\lambda := \bar{\eta}$ and $\nu := \bar{\theta}_{12}$, forcing $\mu = \bar{\theta}_1$), we get that $|\alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\eta}} - \alpha_{\text{sgn}, \bar{\theta}_{12}}^{\bar{\eta}}| = O(1/\sqrt{n})$. From Corollary B.2 (in which we choose $\nu := \bar{\theta}_{12}$, $\mu := \bar{\theta}_1$, $\lambda := \bar{\theta}$, and $\lambda' := \bar{\eta}$), we get

$$\left| \frac{\text{Tr}(\Pi_{\bar{\theta}_{12}}^{\bar{\theta}} \Gamma_{1,2})}{\dim \bar{\theta} \dim \bar{\theta}_{12}} - \frac{\text{Tr}(\Pi_{\bar{\theta}_{12}, \bar{\theta}_1}^{\bar{\eta}} \Gamma_{1,2})}{\dim \bar{\eta} \dim \bar{\theta}_{12}} \right| = O(1/\sqrt{n}),$$

where we have

$$\Pi_{\bar{\theta}_{12}}^{\bar{\theta}} \Gamma_{1,2} = \alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\theta}} \Pi_{\bar{\theta}_{12}}^{\bar{\theta}}, \quad \Pi_{\bar{\theta}_{12}, \bar{\theta}_1}^{\bar{\eta}} \Gamma_{1,2} \Pi_{\bar{\theta}_{12}, \bar{\theta}_1}^{\bar{\eta}} = \left(\alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\eta}} \frac{d_{\bar{\eta}, \bar{\theta}_{12}} - 1}{2d_{\bar{\eta}, \bar{\theta}_{12}}} + \alpha_{\text{sgn}, \bar{\theta}_{12}}^{\bar{\eta}} \frac{\sqrt{d_{\bar{\eta}, \bar{\theta}_{12}}^2 - 1}}{2d_{\bar{\eta}, \bar{\theta}_{12}}} \right) \Pi_{\bar{\theta}_{12}, \bar{\theta}_1}^{\bar{\eta}}$$

from (5.14) and (5.15). Therefore, $|\alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\theta}} - (\alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\eta}} + \alpha_{\text{sgn}, \bar{\theta}_{12}}^{\bar{\eta}})/2| = O(1/\sqrt{n})$, which together with previously proven $\alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\theta}} = n^{-1/3} + O(1/n)$ and $|\alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\eta}} - \alpha_{\text{sgn}, \bar{\theta}_{12}}^{\bar{\eta}}| = O(1/\sqrt{n})$ imply $\alpha_{\text{id}, \bar{\theta}_{12}}^{\bar{\eta}} = n^{-1/3} + O(1/\sqrt{n})$ and $\alpha_{\text{sgn}, \bar{\theta}_{12}}^{\bar{\eta}} = n^{-1/3} + O(1/\sqrt{n})$.

Point 3. Consider $\lambda \vdash n$ and $\nu \subset_c \lambda$ that is obtained from λ by removing two boxes in different columns below the first row. Let us consider two cases.

Case 1: $\nu \subset_{rc} \lambda$. Let $\mu, \mu' \vdash n - 1$ be such that $\nu \subset \mu \subset \lambda$, $\nu \subset \mu' \subset \lambda$, and $\mu \neq \mu'$. Since $d_{\lambda, \nu} \geq 2$, $\dim \mu / \dim \nu = \Theta(n)$, and $\dim \mu' / \dim \nu = \Theta(n)$, both inequalities of Corollary B.3 together imply $\alpha_{\text{id}, \nu}^\lambda = O(1)$ and $\alpha_{\text{sgn}, \nu}^\lambda = O(1)$.

Case 2: $\nu \subset_c \lambda$ and $\nu \not\subset_r \lambda$ (i.e., ν is obtained from λ by removing two boxes in the same, but not the first, row). Let $\mu \vdash n - 1$ be the unique Young diagram that satisfies $\nu \subset \mu \subset \lambda$, and let λ' be obtained from μ by adding a box in the first row. For Point 2 we already have shown that $\alpha_{\text{id}, \nu}^{\lambda'} = o(1)$ and $\alpha_{\text{sgn}, \nu}^{\lambda'} = o(1)$, so, from Corollary B.2 and $\dim \mu / \dim \nu = \Theta(n)$, we get that $\alpha_{\text{id}, \nu}^\lambda = O(1)$.