

On the Power of Nonstandard Quantum Oracles

Roozbeh Bassirian ✉

University of Chicago, IL, USA

Bill Fefferman ✉

University of Chicago, IL, USA

Kunal Marwaha ✉ 🏠 

University of Chicago, IL, USA

Abstract

We study how the choices made when designing an oracle affect the complexity of quantum property testing problems defined relative to this oracle. We encode a regular graph of even degree as an invertible function f , and present f in different oracle models. We first give a one-query QMA protocol to test if a graph encoded in f has a small disconnected subset. We then use representation theory to show that no classical witness can help a quantum verifier efficiently decide this problem relative to an in-place oracle. Perhaps surprisingly, a simple modification to the standard oracle prevents a quantum verifier from efficiently deciding this problem, even with access to an unbounded witness.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases quantum complexity, QCMA, expander graphs, representation theory

Digital Object Identifier 10.4230/LIPIcs.TQC.2023.11

Related Version *Full Version:* <https://arxiv.org/abs/2212.00098>

Funding *Roozbeh Bassirian:* AFOSR (YIP number FA9550-18-1-0148 and FA9550-21-1-0008); NSF under Grant CCF-2044923 (CAREER); U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers; DOE QuantISED grant DE-SC0020360.

Bill Fefferman: AFOSR (YIP number FA9550-18-1-0148 and FA9550-21-1-0008); NSF under Grant CCF-2044923 (CAREER); U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers; DOE QuantISED grant DE-SC0020360.

Kunal Marwaha: NSF Graduate Research Fellowship Program under Grant No. DGE-1746045

Acknowledgements Thanks to Casey Duckering, Juspreet Singh Sandhu, Peter Shor, and Justin Yirka for collaborating on early stages of this project. KM thanks many others for engaging discussions, including Adam Bouland, Antares Chen, Aram Harrow, Matt Hastings, Eric Hester, Neng Huang, Robin Kothari, Brian Lawrence, Yi-Kai Liu, Patrick Lutz, Tushant Mittal, Abhijit Mudigonda, Chinmay Nirkhe, and Aaron Potechin. Thanks to Chinmay Nirkhe for feedback on a previous version of this manuscript.

1 Introduction

Computational complexity is the study of the innate amount of resources required to complete some task. We assign *complexity classes* to sets of tasks that require similar amounts of resources; from here, the goal is to understand the relationship between complexity classes. There has been some success proving that two complexity classes are equal, for example $IP = PSPACE$ [25], the PCP theorem [7], and $MIP^* = RE$ [17]; however, proving that two complexity classes are *unequal* has been much more elusive. For example, we cannot prove $P \neq PSPACE$, let alone $P \neq NP$.



© Roozbeh Bassirian, Bill Fefferman, and Kunal Marwaha;
licensed under Creative Commons License CC-BY 4.0

18th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2023).

Editors: Omar Fawzi and Michael Walter; Article No. 11; pp. 11:1–11:25

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

One response to this difficulty is to equip a computational model with an *oracle*, which computes a fixed (but arbitrarily powerful) quantity in a single timestep. It is often easier to prove that a statement (e.g. $P \neq NP$) is true relative to an oracle; furthermore, this restricts the kinds of proof techniques that can show the statement is false without an oracle. In addition to separating complexity classes, oracles and *query complexity* naturally arise in cryptography (e.g. [19]) and learning theory (e.g. [20]).

Even with respect to an oracle, proving that some complexity classes are unequal can be surprisingly difficult. Notably, Aharonov and Naveh define QCMA, a subset of QMA where the witness is a classical bitstring [3], and ask if $QCMA \subsetneq QMA$. Aaronson and Kuperberg conjecture that an oracle separates these classes, but only prove a “quantum oracle” where this occurs [2]. Subsequent works [11, 22] remove the “quantumness” from the oracle model, but still use models with internal randomness or other nonstandard aspects.

We consider quantum property testing problems defined relative to oracles from various oracle models: encoding the edges of a graph in an invertible function f , we present f as either a *standard* oracle or *in-place* oracle, with or without internal randomness. With mild restrictions on the workspace of quantum verifiers, we find:

1. In several oracle models presenting f , a *quantum* witness can help a quantum verifier efficiently decide if the graph encoded in f has a small disconnected subset.
2. Where f is presented as a randomized in-place oracle, no *classical* witness can help a quantum verifier efficiently decide this problem.
3. Where f is presented as a randomized phase oracle, no witness of *any type or size* can help a quantum verifier efficiently decide this problem.

Our results highlight that the quantum complexity of a task defined relative to an oracle is influenced by the choice of oracle model.

1.1 Our techniques

We use a well-known fact of Petersen to encode the edges of any even-degree regular graph in an invertible function f . We then consider natural ways to install f within an oracle; we say that f is *presented* as a particular kind of oracle. For example, a standard oracle presents f through the map $|c, x\rangle \mapsto |c \oplus f(x), x\rangle$, while an in-place oracle presents f through the map $|x\rangle \mapsto |f(x)\rangle$. In general, we consider oracles that give access both to f and f^{-1} . An oracle may also have internal randomness: on every query to a *randomized* oracle, f is chosen uniformly at random from a fixed set of functions F .

Consider the Laplacian L_f of a graph encoded in f . We first provide a test such that for any input state $|\psi\rangle$, the test succeeds with probability expressible in terms of $\langle \psi | L_f | \psi \rangle$, independently of how an oracle presents f . We use this test to construct a QMA protocol verifying that the graph is not an *expander* graph. This problem is primarily motivated by the preimage-testing problem of Fefferman and Kimmel [11], which separates QMA and QCMA relative to a nonstandard oracle. They encode an invertible function π in an oracle *without efficient access to π^{-1}* , and test a property of π^{-1} ; by design, this property can be verified but not easily computed. Crucially, we view a permutation and its inverse as the edges of an *undirected graph*; properties of undirected graphs are not sensitive to the ordering of $(x, \pi(x))$. We use multiple permutations to study graphs of higher degree, and notice that detecting if a graph has a non-expanding region is hard without traversing most of the graph. Some of these ideas are related to the *component mixer* concept of Lutomirski [21], and are simultaneously and independently explored by Natarajan and Nirkhe [22].

A randomized oracle presenting a set of functions F can be seen as a quantum channel, so small changes to F cause statistically indistinguishable changes to the oracle. We use this flexibility to modify non-expansion testing to a simple permutation problem: do the functions

$f \in F$ stabilize a small set $V \subseteq [N]$, or is F the set of all permutations on $[N]$? Notice that F is a group in both cases. When an oracle presenting F preserves the group structure of F , we can use representation theory. For this problem, this is satisfied by an *in-place* oracle; the oracle is then an orthogonal projector onto one of two symmetric subspaces of matrices. After finding an orthogonal basis for each subspace, we construct a hybrid argument to prove that only witnesses with knowledge of V can help a quantum verifier efficiently decide this problem. We also use representation theory to give a QCMA protocol for an analogous permutation problem in randomized standard oracles.

We finally study the permutation problem in a randomized phase oracle. We directly analyze the effect of the oracle on any input density matrix; with high probability, the oracle decreases the magnitude of every off-diagonal term by a $\frac{1}{2^{\text{poly}(n)}}$ factor. We then construct a hybrid argument, bounding our measure of progress using an inequality relating the sizes of Schatten p -norms. When the state space is not too large, we prove that an exponential number of queries are required to distinguish most YES instances from the NO instance, *regardless of input state*. As a result, no witness can help a verifier distinguish YES from NO.

Note that our quantum verifiers are not fully general. Our lower bound techniques restrict the number of extra workspace qubits in the verifier; however, our upper bounds also work in this setting. In Section 3.2, we explain these restrictions in more detail and discuss prospects for generalizing our results.

1.2 Related work

Quantum oracle models

A fundamental constraint of quantum oracle models is that they must be unitary. We describe several nonstandard oracle models used in quantum computing:

- A *quantum* oracle is any unitary operation U in the full Hilbert space. Although the operation is unitary, the verifier doesn't necessarily have access to U^{-1} . Oracles like these are not typically classical because the unitary's action is not efficiently and classically representable.
- An *in-place* oracle maps $|x\rangle \rightarrow |\pi(x)\rangle$ for some classical invertible function π . Again, this computation is not efficiently reversible since the verifier may not have access to π^{-1} . When a standard oracle gives access to π^{-1} , an in-place oracle query can be simulated in two queries; otherwise, an exponential number of queries are required to construct one from the other [18].
- A *phase* oracle puts the output of a classical function f in the phase of a basis state. We consider the map $|x\rangle \rightarrow e^{f(x) \cdot 2\pi i/N} |x\rangle$. To contrast, note that the map $|c, x\rangle \rightarrow e^{cf(x) \cdot 2\pi i/N} |c, x\rangle$ is unitarily equivalent to the standard oracle.

All of these oracles can optionally have *internal randomness*, as considered by Harrow and Rosenbaum [16]; we call these *randomized* oracles. On every query to a randomized oracle, a unitary is chosen at random from a fixed set. This can be very powerful; for example, [16] gives examples of randomized oracles where problems *impossible* to decide with classical queries can be decided with a single quantum query.

QMA and QCMA

The *Merlin-Arthur* style of complexity classes considers a decision problem and two players. The magician (Merlin) has claimed the answer to the decision problem is YES, and gives the verifier a token (the *proof* or *witness*) to convince them. The verifier (Arthur) must then

11:4 On the Power of Nonstandard Quantum Oracles

ensure the answer is actually YES. Given a problem with size n , the verifier must accept a correct witness (i.e. when the answer is YES) with probability $1/q$ higher than a “lying” witness (i.e. when the answer is NO) for some $q = \text{poly}(n)$. The set of problems that can be decided this way in a classical setting is known as Merlin Arthur (MA). If the verifier is a quantum computer, this is QCMA; if the witness can be any quantum state, this is QMA.

■ **Table 1** Complexity classes in the style of *Merlin-Arthur*. QCMA is a subset of QMA where the witness can be efficiently written as a classical bitstring.

| | verifier is classical | verifier is quantum |
|----------------------|-----------------------|---------------------|
| witness is classical | MA | QCMA |
| witness is quantum | - | QMA |

Since any classical bitstring can be efficiently written as a quantum state, $\text{QCMA} \subseteq \text{QMA}$. But is the reverse true? Even the *oracle* version of this problem is open: at the top of a recent list of open problems, Aaronson asks for a standard oracle that separates the two classes [1]. All previous progress [2, 11, 22] relies on specifically chosen *nonstandardness* in the oracle.

Natarajan and Nirkhe [22] make progress on a standard oracle separation of QMA and QCMA by constructing an oracle with randomness. They simultaneously and independently provide a QMA protocol for testing non-expansion of a graph in an oracle. To prove their lower bound, they combine the adversary method, the polynomial method, and a reduction to a problem of Ambainis, Childs, and Liu [5]. However, their notion of randomness is different from ours and other works [16, 11, 6], and acts as follows: when an oracle is first queried, it chooses a function f from a distribution, but on subsequent queries, it uses the same function f . By contrast, our notion of randomness is memoryless: an oracle chooses f from a uniform distribution on F for *every* query. This allows one to make small changes to F without affecting the success of the QMA protocol; we use this flexibility to study a simpler permutation problem.

2 Our setup

Consider a d -regular graph on $N := 2^n$ vertices for any n and even d . We show that an invertible function can list the edges adjacent to each vertex in G .

► **Definition 2.1** (Graph-coded function). *Consider a d -regular graph G (for even d) on N vertices. A G -coded function is a function $f : [N] \times [d/2] \rightarrow [N]$, such that $f_i(x) := f(x, i)$ is a bijection for each $i \in [d/2]$, and each edge is uniquely represented by a tuple $(x, f_i(x))$.*

► **Remark 2.2** (Even-degree regular graphs have graph-coded functions). Every regular graph G of even degree has a G -coded function.

Proof. A d -regular graph G of even degree always has a 2-factorization [23]. This means that the edges of G can be partitioned into $d/2$ edge-disjoint subgraphs $[E_1, \dots, E_{d/2}]$ where in each E_i , all vertices have degree two (i.e. a collection of cycles). Thus, we can represent each E_i with a permutation π_i , where the edge $(x, y) \in E_i$ if and only if $\pi_i(x) = y$ or $\pi_i(y) = x$. Then $f(x, i) := \pi_i(x)$ is a G -coded function. ◀

Graph-coded functions f are bijective, and therefore invertible. We now *present* f in various oracle models. Note that we define all oracles with access both to f and f^{-1} .

► **Definition 2.3** (An oracle model *presents* a function f). *For each oracle model below (e.g. standard oracle), we say that this oracle model presents the function f .*

► **Remark 2.4.** For notational convenience, we refer to a qubit z that controls the inversion of a function f as taking on values in $\{\pm 1\}$, so that f^z is either $f^1 = f$ or f^{-1} .

► **Definition 2.5** (Standard oracle). *For any $f : [N] \rightarrow [N]$, define $U_f : \mathbb{C}^{2N^2} \rightarrow \mathbb{C}^{2N^2}$ as*

$$U_f \sum_{c,x \in [N], z \in \{\pm 1\}} \alpha_{c,x,z} |c, x, z\rangle := \sum_{c,x \in [N], z \in \{\pm 1\}} \alpha_{c,x,z} |c \oplus f^z(x), x, z\rangle. \quad (1)$$

► **Definition 2.6** (In-place oracle [18]). *For any permutation $\pi : [N] \rightarrow [N]$, define $\tilde{U}_\pi : \mathbb{C}^{2N} \rightarrow \mathbb{C}^{2N}$ as*

$$\tilde{U}_\pi \sum_{x \in [N], z \in \{\pm 1\}} \beta_{x,z} |x, z\rangle := \sum_{x \in [N], z \in \{\pm 1\}} \beta_{x,z} |\pi^z(x), z\rangle. \quad (2)$$

► **Remark 2.7** ([18]). A standard oracle U_f (with access to f^{-1}) can simulate an in-place oracle \tilde{U}_f in two queries:

$$(\mathbb{I} \otimes X) \circ U_f \circ (\text{SWAP}_{n,n} \otimes X) \circ U_f |0\rangle^{\otimes n} |x, z\rangle = |0\rangle^{\otimes n} |f^z(x), z\rangle. \quad (3)$$

► **Definition 2.8** (N^{th} root of unity). *Define the N^{th} root of unity as $\omega_N := e^{2\pi i/N}$.*

► **Definition 2.9** (Phase oracle). *For any function $f : [N] \rightarrow [N]$, define $\bar{U}_f : \mathbb{C}^{2N} \rightarrow \mathbb{C}^{2N}$ as*

$$\bar{U}_f \sum_{x \in [N], z \in \{\pm 1\}} \alpha_{x,z} |x, z\rangle := \sum_{x \in [N], z \in \{\pm 1\}} \alpha_{x,z} \omega_N^{f^z(x)} |x, z\rangle. \quad (4)$$

We describe how an oracle in our setup exhibits internal randomness. On each query, a *randomized* oracle chooses a function uniformly from a set F . We say that a randomized oracle *presents* F .

► **Remark 2.10.** Given a unitary U , we use the notation \mathcal{U} to denote an operator on density matrices; that is,

$$\mathcal{U}[\rho] := U\rho U^\dagger. \quad (5)$$

► **Definition 2.11** (Randomized oracle (e.g. [16, 11])). *For any set F of functions $f : [N] \rightarrow [N]$ corresponding to oracles $\{U_f \mid f \in F\}$, define the linear operator \mathcal{O}_F as*

$$\mathcal{O}_F := \frac{1}{|F|} \sum_{f \in F} \mathcal{U}_f. \quad (6)$$

We match the notation of randomized oracle \mathcal{O}_F with oracle U_f ; e.g. $\tilde{\mathcal{O}}_F$ is a randomized in-place oracle.

2.1 Problem statements

The problems below are not fully specified without the choice of oracle model. We prepend the names below with the choice of oracle model; for example, we denote Problem 2.12 in a standard oracle as STANDARD NON-EXPANSION(d, α, ε).

11:6 On the Power of Nonstandard Quantum Oracles

► **Problem 2.12** (NON-EXPANSION(d, α, ε)). Consider an oracle U_f presenting a G -coded function f .

1. In a YES instance, we are promised that G is a union of two disconnected d -regular graphs, and that the smaller graph has N^α vertices.
2. In a NO instance, we are promised that G is d -regular and has spectral gap at least ε (for example, an expander graph).

The problem is to decide whether U_f is a YES instance or NO instance.

We also consider a version of this problem with *randomized* oracles, where each randomized YES instance is specified by the set of vertices V of the smaller graph. On each query, an oracle chooses an graph-coded function f uniformly at random that corresponds to a graph where V and $[N]/V$ are disconnected.

► **Problem 2.13** (RANDOMIZED NON-EXPANSION(d, α, ε)). Consider a randomized oracle \mathcal{O}_F presenting a set of graph-coded functions F .

1. Each subset $V \subseteq [N]$ of size $|V| = N^\alpha$ specifies a YES instance \mathcal{O}_{F_V} . Let F_V be the set of all G -coded functions of d -regular graphs G with no edges between V and $[N]/V$.
2. There is a single NO instance $\mathcal{O}_{F_\emptyset}$. Let F_\emptyset be the set of all G -coded functions of d -regular graphs G with spectral gap at least ε .

The problem is to decide whether \mathcal{O} is a YES instance or a NO instance.

In the configuration model of a random graph, F_V contains all functions $f(x, i)$ such that $f_i(x) := f(x, i)$ is the union of a permutation on $[N]/V$ and a permutation on V . In fact, we can use the oracle's internal randomness to adjust the underlying set F , and even consider graphs that are not typically expander graphs.

► **Definition 2.14** (Subset indicator). For a set $V \subseteq [N]$, define the function $i_V : [N] \rightarrow \{V, [N]/V\}$ as

$$i_V(x) = \begin{cases} V & x \in V \\ [N]/V & x \notin V. \end{cases} \quad (7)$$

► **Definition 2.15** (Permutations that stabilize a subset). For a set $V \subseteq [N]$, define the set of permutations

$$T_V := \{\pi : [N] \rightarrow [N] : i_V(x) = i_V(\pi(x)) \forall x \in [N]\}. \quad (8)$$

We say that T_V stabilizes the subset V .

► **Problem 2.16** (RANDOMIZED HIDDEN SUBSET(α)). Consider a randomized oracle \mathcal{O}_F presenting a set of functions F .

1. Each subset $V \subseteq [N]$ of size $|V| = N^\alpha$ specifies a YES instance \mathcal{O}_{T_V} .
2. There is a single NO instance $\mathcal{O}_{T_\emptyset}$, where T_\emptyset is the set of all permutations of $[N]$.

The problem is to decide whether \mathcal{O} is a YES instance or a NO instance.

Notice that RANDOMIZED HIDDEN SUBSET(α) is exactly RANDOMIZED NON-EXPANSION($2, \alpha, 0$).

Notice that T_V is a group under function composition. One can generalize this algebraic structure to a problem distinguishing oracles presenting subgroups of T_\emptyset from an oracle presenting T_\emptyset :

► **Problem 2.17** (RANDOMIZED HIDDEN SUBGROUP($\times, \{H_i\}$)). Consider the set T_\emptyset of all permutations on $[N]$ as a group with operation \times , such that each $H_i \subsetneq T_\emptyset$ is also a group. Suppose a randomized oracle \mathcal{O} presents either T_\emptyset or any H_i .

1. Each subgroup H_i specifies a YES instance \mathcal{O}_{H_i} .
 2. There is a single NO instance $\mathcal{O}_{T_\varnothing}$, where T_\varnothing is the set of all permutations of $[N]$.
- The problem is to decide whether \mathcal{O} is a YES instance or a NO instance.

For example, RANDOMIZED HIDDEN SUBSET(α) is a special case of RANDOMIZED HIDDEN SUBGROUP($\times, \{H_i\}$) using the group operation of function composition.

3 Our results

3.1 Non-expansion and quantum witness

Our first result shows that there is a one-query QMA protocol for NON-EXPANSION(d, α, ε) in many oracle models presenting a graph-coded function.

► **Theorem 3.1.** *There is a QMA protocol for STANDARD NON-EXPANSION(d, α, ε) and IN-PLACE NON-EXPANSION(d, α, ε) at every even $d \geq 4$, all $0 < \alpha < \frac{1}{2}$, and all constant $\varepsilon > 0$.*

Graphs with good *expansion* are well-connected despite their sparsity. For any graph G , let A_G be the adjacency matrix of G , and $L_G = d\mathbb{I} - A_G$ be the *graph Laplacian* of G . The smallest eigenvalue of L_G is $\lambda_1(L_G) = 0$, and the next-smallest eigenvalue $\lambda_2(L_G)$ measures the expansion of G . In this framework, NON-EXPANSION(d, α, ε) asks if an oracle presenting a G -coded function has $\lambda_2(L_G) = 0$ (YES), or if $\lambda_2(L_G) \geq \varepsilon$ (NO).

At the heart of our protocol is the *spectral test*, which takes an input state $|\psi\rangle$ and fails with probability proportional to $\langle \psi | L_G | \psi \rangle$. We describe the spectral test for both standard oracles and in-place oracles in Appendix A.1. A state that passes the spectral test is essentially supported on a subspace according to $\lambda(L_G) = o(\frac{1}{\text{poly}(n)})$; in a NO instance, this is one-dimensional, and in a YES instance, this is at least two-dimensional. In fact, the uniform superposition over all inputs, $|+\rangle^{\otimes n}$, is always in this subspace. As a result, our protocol (Theorem 3.1) either runs the spectral test, or checks if the input state is close to $|+\rangle^{\otimes n}$.

Consider the randomized variant of NON-EXPANSION(d, α, ε). The graph of any graph-coded function presented in a YES instance is guaranteed to have a small set V (i.e. $|V| = N^\alpha$) disconnected from the rest of the graph. As a result, there is a state, defined only by the vertices of V , that is all-but-negligibly supported in the $\lambda(L_G) = 0$ subspace. This state is the *subset state* $|V\rangle$:

► **Definition 3.2** (Subset state). *For any non-empty subset $S \subseteq [N]$, define the subset state of S as*

$$|S\rangle := \frac{1}{\sqrt{|S|}} \sum_{x \in S} |x\rangle. \quad (9)$$

Since $|V\rangle$ is a good witness for *every* graph encoded in a YES instance, the QMA protocol works just as well in the randomized setting (Theorem A.4).

Randomized oracles that present a set F of graph-coded functions are stable to small changes in the set F . In fact, an oracle presenting F encoding all d -regular expander graphs is indistinguishable from an oracle presenting F encoding all d -regular graphs. The latter oracle can be simulated with $d/2$ queries to the NO instance of RANDOMIZED HIDDEN SUBSET(α); we show in Theorem A.5 that the same QMA protocol can also decide this problem.

3.2 Randomized in-place oracles and classical witness

Our second result shows that a general class of verifiers cannot decide IN-PLACE RANDOMIZED HIDDEN SUBSET(α).

► **Theorem 3.3** (informal). *No quantum verifier with $O(\log(n))$ additional workspace qubits can efficiently decide IN-PLACE RANDOMIZED HIDDEN SUBSET(α) given a polynomial sized classical witness.*

Recall that in this problem, a verifier has access to a quantum channel and a polynomial-sized classical witness, and must distinguish whether the oracle presents a uniformly random permutation or a permutation that stabilizes a hidden subset V . Let \mathcal{Y} be the set of all YES instances; note that each instance is uniquely defined by a subset V .

Suppose there exists a QCMA algorithm for this problem. Since there are at most $O(2^{\text{poly}(n)})$ different classical witnesses, there exists a set of YES instances \mathcal{Y}' that share the same witness, such that $|\mathcal{Y}'|/|\mathcal{Y}| = \Omega(2^{-\text{poly}(n)})$. We can refute the existence of such an algorithm by proving that the same verification “strategy” cannot distinguish all instances of \mathcal{Y}' from the NO case with non-negligible probability. A “strategy” is exactly a quantum algorithm: a series of unitaries and oracle queries, followed by a POVM. Without loss of generality, a T -query algorithm alternates between unitaries and oracle queries on $\mathcal{H}_O \otimes \mathcal{H}_W$ followed by a measurement¹, where \mathcal{H}_O is the Hilbert space of the “oracle” qubits and \mathcal{H}_W is the extra workspace:

$$\mathcal{E}_O[\rho_0] = (\mathcal{O} \otimes \mathbb{I}) \circ \mathcal{U}_T \circ \dots \circ \mathcal{U}_2 \circ (\mathcal{O} \otimes \mathbb{I}) \circ \mathcal{U}_1[\rho_0]. \quad (10)$$

One may try to use the hybrid argument of Bennett, Bernstein, Brassard, and Vazirani [8] and Ambainis [4] to prove that the diamond norm $\left| \mathcal{E}_{\tilde{\mathcal{O}}_{T_V}} - \mathcal{E}_{\tilde{\mathcal{O}}_{T_\emptyset}} \right|_\diamond$ is small in expectation over the choice of $\tilde{\mathcal{O}}_{T_V} \in \mathcal{Y}'$. This would imply that the verifier cannot distinguish all instances of \mathcal{Y}' with the same strategy. We can consider the optimal distinguishing probability in terms of $\left| \mathcal{E}_{\tilde{\mathcal{O}}_{T_V}}[\rho_0] - \mathcal{E}_{\tilde{\mathcal{O}}_{T_\emptyset}}[\rho_0] \right|_1$ for some fixed $\rho_0 \in \mathcal{H}_O \otimes \mathcal{H}_W$.

However, this statistical argument does not hold for some choices of \mathcal{Y}' . Consider the following simple example: \mathcal{Y}' contains all V such that $1 \in V$. First, \mathcal{Y}' satisfies the size implied by the pigeonhole principle. Second, for $\rho_0 = |1\rangle\langle 1| \otimes \mathbb{I}$, $\left| \tilde{\mathcal{O}}_{T_V}[\rho_0] - \tilde{\mathcal{O}}_{T_\emptyset}[\rho_0] \right|_1$ is large for *all* instances in \mathcal{Y}' , since $|1\rangle\langle 1|$ mixes only within a small subset. Note that this only implies the existence of an *instance-specific* POVM distinguishing each YES instance in \mathcal{Y}' from the NO instance. By contrast, a verification strategy has a *fixed* POVM $\{E, \mathbb{I} - E\}$. This allows us to prove that the following value is small on average over the choice of V :

$$\left| \text{Tr} \left[E \mathcal{E}_{\tilde{\mathcal{O}}_{T_V}}[\rho_0] \right] - \text{Tr} \left[E \mathcal{E}_{\tilde{\mathcal{O}}_{T_\emptyset}}[\rho_0] \right] \right| \quad (11)$$

We must bound this value for arbitrary choices of E , ρ_0 and \mathcal{U}_i fixed in the algorithm. In order to do this, we leverage tools from representation theory; this allows us to see randomized oracles in our problem as *orthogonal projectors* into a subspace of matrices with low dimension, and prove that the density matrix of good distinguishers is characterized

¹ Note that the last operation does not have to be a unitary – one can simply replace a unitary followed by a POVM with another equivalent POVM.

by the hidden subset. One caveat of our technique is that the verifier is only allowed to have $O(\log(n))$ extra workspace qubits. This restriction is necessary to reduce the subspace dimension to regimes we can handle.

Representation theory has been previously used to study symmetric operators on variables (in probability) or qubits (in quantum computing) using the language of de Finetti theorems (e.g. [15]); these operators project into subspaces of permutation-invariant sequences or quantum states. By contrast, we notice that some randomized oracles are symmetric operators on *density matrices*. This allows us to explicitly find an orthogonal basis for the associated symmetric subspaces. We match oracle models with problems with the same group structure: RANDOMIZED HIDDEN SUBSET(α) for in-place oracles in Appendix B.1, and an analogous special case of RANDOMIZED HIDDEN SUBGROUP($\times, \{H_i\}$) for standard oracles in Appendix B.2.

3.3 Randomized phase oracles: no witness can help

Our third result shows that deciding RANDOMIZED HIDDEN SUBSET(α) in a phase oracle is much harder than other oracle models we consider. A random phase has *zero* expectation. We use this fact to show that queries to most YES instances and the NO instance reduce the magnitude of each off-diagonal of the density matrix by an exponential factor, regardless of the input state. We bound the Frobenius norm of the difference of query outputs to show that these instances are statistically indistinguishable when the state space is not too large. As a result, no untrustworthy witness can help decide this problem.

► **Theorem 3.4** (informal). *No quantum verifier with $o(n)$ additional workspace qubits can efficiently decide PHASE RANDOMIZED HIDDEN SUBSET(α) given any unbounded witness. Moreover, these verifiers require an exponential number of queries to statistically distinguish a YES instance from the NO instance, for each of asymptotically all YES instances.*

We defer the formal proof to Appendix C. Note that the query lower bound here is *statistical*. In the NO instance, a witness is designed to fool the verifier; in order to overcome this, the verifier must use the witness in tandem with the oracle. But this cannot be done efficiently: distinguishing the NO instance from nearly any YES instance requires an exponential number of queries, *regardless of input state*.

In fact, Theorem 3.4 holds for any oracle that sends $|c, x, z\rangle \rightarrow \omega_N^{c \cdot f^z(x)} |c, x, z\rangle$, where the c register has $k = o(n)$ qubits, while at $k = n$ qubits, the oracle is unitarily equivalent to a standard oracle, and thus has a QMA protocol for RANDOMIZED HIDDEN SUBSET(α) by Theorem 3.1.

References

- 1 Scott Aaronson. Open problems related to quantum query complexity, 2021. [arXiv:2109.06917](#).
- 2 Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. *Theory Comput.*, 3(1):129–157, 2007. [arXiv:quant-ph/0604056](#).
- 3 Dorit Aharonov and Tomer Naveh. Quantum np - a survey, 2002. [arXiv:quant-ph/0210077](#).
- 4 Andris Ambainis. Quantum lower bounds by quantum arguments. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing - STOC '00*. ACM Press, 2000. [arXiv:quant-ph/0002066](#).
- 5 Andris Ambainis, Andrew M. Childs, and Yi-Kai Liu. Quantum property testing for bounded-degree graphs. *Lecture Notes in Computer Science*, pages 365–376, 2011. doi:10.1007/978-3-642-22935-0_31.

- 6 Atul Singh Arora, Alexandru Gheorghiu, and Uttam Singh. Oracle separations of hybrid quantum-classical circuits, 2022. [arXiv:2201.01904](#).
- 7 Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM (JACM)*, 45(3):501–555, 1998. doi:10.1145/1236457.1236459.
- 8 Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26(5):1510–1523, October 1997. doi:10.1137/s0097539796300933.
- 9 Andrew Childs. Lecture Notes on Quantum Algorithms, 2022. URL: <https://www.cs.umd.edu/~amchilds/qa/qa.pdf>.
- 10 Paul Erdos. On the central limit theorem for samples from a finite population. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 4:49–61, 1959. URL: https://old.renyi.hu/~p_erdos/1959-13.pdf.
- 11 Bill Fefferman and Shelby Kimmel. Quantum vs classical proofs and subset verification, 2018. [arXiv:1510.06750](#).
- 12 Joel Friedman. A proof of Alon’s second eigenvalue conjecture and related problems. *CoRR*, 2004. [arXiv:cs/0405020](#).
- 13 Chris Godsil and Hanmeng Zhan. Discrete-time quantum walks and graph structures. *Journal of Combinatorial Theory, Series A*, 167:181–212, 2019. doi:10.1016/j.jcta.2019.05.003.
- 14 Alex Bredariol Grilo, Iordanis Kerenidis, and Jamie Sikora. QMA with subset state witnesses. In *Mathematical Foundations of Computer Science 2015 - 40th International Symposium*. Springer, 2015. doi:10.1007/978-3-662-48054-0_14.
- 15 Aram W. Harrow. The church of the symmetric subspace, 2013. [arXiv:1308.6595](#).
- 16 Aram W. Harrow and David J. Rosenbaum. Uselessness for an oracle model with internal randomness, 2011. [arXiv:1111.1462](#).
- 17 Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. $\text{Mip}^* = \text{re}$. *Communications of the ACM*, 64(11):131–138, 2021. [arXiv:2001.04383](#).
- 18 Elham Kashefi, Adrian Kent, Vlatko Vedral, and Konrad Banaszek. Comparison of quantum oracles. *Physical Review A*, 65(5), May 2002. doi:10.1103/physreva.65.050304.
- 19 Jonathan Katz and Yehuda Lindell. *Introduction to modern cryptography*. CRC press, 2020. URL: <http://www.cs.umd.edu/~jkatz/imc.html>.
- 20 Eyal Kushilevitz and Yishay Mansour. Learning decision trees using the fourier spectrum. *SIAM J. Comput.*, 22(6):1331–1348, 1993. doi:10.1137/0222080.
- 21 Andrew Lutomirski. Component mixers and a hardness result for counterfeiting quantum money, 2011. [arXiv:1107.0321](#).
- 22 Anand Natarajan and Chinmay Nirkhe. A classical oracle separation between qma and qcma, 2022. [arXiv:2210.15380](#).
- 23 Julius Petersen. Die Theorie der regulären graphs. *Acta Mathematica*, 15:193–220, 1900. doi:10.1007/BF02392606.
- 24 Jason D. M. Rennie. Relating the trace and frobenius matrix norms, August 2005. URL: <http://people.csail.mit.edu/jrennie/writing/traceFrobenius.pdf>.
- 25 Adi Shamir. $\text{IP} = \text{PSPACE}$. *Journal of the ACM (JACM)*, 39(4):869–877, 1992. doi:10.1145/146585.146609.
- 26 Luca Trevisan. Max cut and the smallest eigenvalue, 2008. [arXiv:0806.1978](#).
- 27 John Watrous. Quantum simulations of classical random walks and undirected graph connectivity, 1998. [arXiv:cs/9812012](#).

A Verifying non-expansion with a quantum witness

A.1 The spectral test

We give a test that takes an input state $|\psi\rangle = \sum_{x \in [N]} a_x |x\rangle$ on n qubits, and fails with probability proportional to $\langle \psi | L_G | \psi \rangle$. This relies on a curious fact:

► **Lemma A.1.** *Consider a d -regular graph G (for even d) on 2^n vertices and a G -coded function f . Suppose we have a normalized quantum state $|\psi\rangle = \sum_{x \in [N]} a_x |x\rangle$ on n qubits. Then*

$$\sum_{i \in [d/2]} \sum_{x \in [N]} \|a_x \pm a_{f(x,i)}\|^2 = d \pm \langle \psi | A_G | \psi \rangle. \quad (12)$$

Proof.

$$\sum_{i \in [d/2]} \sum_{x \in [N]} \|a_x \pm a_{f(x,i)}\|^2 = \sum_{i \in [d/2]} \sum_{x \in [N]} \|a_x\|^2 + \|a_{f(x,i)}\|^2 \pm (a_x a_{f(x,i)}^* + a_x^* a_{f(x,i)}) \quad (13)$$

$$= d \pm \sum_{i \in [d/2]} \sum_{x \in [N]} (a_x a_{f(x,i)}^* + a_x^* a_{f(x,i)}) \quad (14)$$

$$= d \pm \langle \psi | A_G | \psi \rangle. \quad (15)$$

◀

We construct the spectral test with one query either to a standard oracle or in-place oracle presenting a graph-coded function f . The former (Procedure A.2) is a SWAP test but with an oracle query in the middle. The latter (Procedure A.3) relies on controlled access to the in-place oracle.

► **Procedure A.2** (Spectral test with a standard oracle). *Consider a d -regular graph G on $N = 2^n$ vertices where d is even, and normalized state $|\psi\rangle = \sum_{x \in [N]} a_x |x\rangle \in \mathbb{C}^N$. We assume access to a standard oracle $U_f : \mathbb{C}^{k \times k}$ for $k = N^2 2^{\lceil \log_2 d \rceil}$, which acts on a basis vector as*

$$U_f |c, x, i, z\rangle = |c \oplus f^z(x, i), x, i, z\rangle, \quad (16)$$

for $c, x \in [N]$, $i \in 2^{\lceil \log_2 d \rceil - 1}$, and $z \in \{\pm 1\}$.

1. Pick $i \in [d/2]$ uniformly at random, and prepare the state $|i\rangle \in \mathbb{C}^{2^{\lceil \log_2 d \rceil - 1}}$.
2. Prepare a qubit in the state $|+\rangle = \frac{|1\rangle + |-1\rangle}{\sqrt{2}} \in \mathbb{C}^2$. (Recall that we label the values of this register in $\{\pm 1\}$.)
3. Combine n registers $|0\rangle^{\otimes n}$, the input state $|\psi\rangle$, and $|i\rangle$ and $|+\rangle$ to create $|0\rangle^{\otimes n} |\psi\rangle |i\rangle |+\rangle$.
4. Apply the oracle U_f , which creates the state

$$\frac{1}{\sqrt{2}} \sum_{x \in [N]} a_x (|f(x, i)\rangle |x\rangle |i\rangle |1\rangle + |f^{-1}(x, i)\rangle |x\rangle |i\rangle |-1\rangle). \quad (17)$$

5. Swap the first two sets of n qubits, controlled by the last qubit. This creates the state

$$\frac{1}{\sqrt{2}} \sum_{x \in [N]} a_x (|f(x, i)\rangle |x\rangle |i\rangle |1\rangle + |x\rangle |f^{-1}(x, i)\rangle |i\rangle |-1\rangle) \quad (18)$$

$$= \frac{1}{\sqrt{2}} \sum_{x \in [N]} a_x |f(x, i)\rangle |x\rangle |i\rangle |1\rangle + \frac{1}{\sqrt{2}} \sum_{x \in [N]} a_{f(x,i)} |f(x, i)\rangle |x\rangle |i\rangle |-1\rangle. \quad (19)$$

11:12 On the Power of Nonstandard Quantum Oracles

6. Apply a Hadamard on the last qubit, which creates the state

$$\frac{1}{2} \sum_{x \in [N]} (a_x + a_{f(x,i)}) |f(x,i)\rangle |x\rangle |i\rangle |1\rangle + \frac{1}{2} \sum_{x \in [N]} (a_x - a_{f(x,i)}) |f(x,i)\rangle |x\rangle |i\rangle |-1\rangle . \quad (20)$$

7. Measure the last qubit and accept if it is 1.

Moreover, by Lemma A.1, this procedure fails with probability

$$\frac{1}{d/2} \sum_{i \in [d/2]} \sum_{x \in [N]} \frac{\|a_x - a_{f(x,i)}\|^2}{4} = \frac{\langle \psi | L_G | \psi \rangle}{2d} . \quad (21)$$

► **Procedure A.3** (Spectral test with an in-place oracle). Consider a d -regular graph G on $N = 2^n$ vertices where d is even, and normalized state $|\psi\rangle = \sum_{x \in [N]} a_x |x\rangle \in \mathbb{C}^N$. We assume controlled access to an in-place oracle $\tilde{U}_f : \mathbb{C}^{k \times k}$ for $k = N2^{\lceil \log_2 d \rceil + 1}$, which acts on a basis vector as

$$\tilde{U}_f |a, x, i, z\rangle = |a, f^{a \cdot z}(x, i), i, z\rangle . \quad (22)$$

for control qubit $a \in \{0, 1\}$, $x \in [N]$, $i \in 2^{\lceil \log_2 d \rceil - 1}$, and $z \in \{\pm 1\}$.²

1. Pick $i \in [d/2]$ uniformly at random, and prepare the state $|i\rangle \in \mathbb{C}^{2^{\lceil \log_2 d \rceil - 1}}$.
2. Prepare a qubit in the state $|+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}} \in \mathbb{C}^2$.
3. Combine $|+\rangle$, the input state $|\psi\rangle$, $|i\rangle$, and a register $|1\rangle$ to create $|+\rangle |\psi\rangle |i\rangle |1\rangle$.
4. Apply the oracle \tilde{U}_f , which creates the state

$$\frac{1}{\sqrt{2}} \sum_{x \in [N]} a_x (|0\rangle |x\rangle + |1\rangle |f(x,i)\rangle) |i\rangle |1\rangle \quad (23)$$

$$= \frac{1}{\sqrt{2}} \sum_{x \in [N]} (a_x |0\rangle + a_{f^{-1}(x,i)} |1\rangle) |x\rangle |i\rangle |1\rangle . \quad (24)$$

5. Apply a Hadamard on the first qubit, which creates the state

$$\frac{1}{2} \sum_{x \in [N]} ((a_x + a_{f^{-1}(x,i)}) |0\rangle + (a_x - a_{f^{-1}(x,i)}) |1\rangle) |x\rangle |i\rangle |1\rangle . \quad (25)$$

6. Measure the first qubit and accept if it is 0.

Moreover, by Lemma A.1, this procedure fails with probability

$$\frac{1}{d/2} \sum_{i \in [d/2]} \sum_{x \in [N]} \frac{\|a_x - a_{f^{-1}(x,i)}\|^2}{4} = \frac{\langle \psi | L_G | \psi \rangle}{2d} . \quad (26)$$

A.2 A one-query protocol

► **Theorem 3.1.** There is a QMA protocol for STANDARD NON-EXPANSION(d, α, ε) and IN-PLACE NON-EXPANSION(d, α, ε) at every even $d \geq 4$, all $0 < \alpha < \frac{1}{2}$, and all constant $\varepsilon > 0$.

² Note that this procedure does not actually need access to the inverse of f to conduct the spectral test.

Proof. Suppose the oracle presents a G -coded function. Let $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N$ be the eigenvalues of the graph Laplacian L_G . Note that the smallest eigenvalue of a regular graph G is $\lambda_1 = 0$. We always choose the eigenvector associated with λ_1 as a uniform superposition over vertices of the graph (i.e. $|\lambda_1\rangle := |[N]\rangle = |+\rangle^{\otimes n}$).

Suppose Arthur receives a state $|\psi\rangle = \sum_{i \in [N]} \alpha_i |\lambda_i\rangle$ from Merlin. Consider the following strategy:

- With probability $\frac{1}{2}$, measure $|\psi\rangle$ in the Hadamard basis. Fail if it is in the basis state according to $|+\rangle^{\otimes n}$, and pass otherwise.
- With probability $\frac{1}{2}$, use the spectral test (Procedure A.2 or Procedure A.3, respectively). The probability of failure FAIL is

$$\text{FAIL} = \frac{1}{2} \|\langle \psi | |+\rangle^{\otimes n}\|^2 + \frac{1}{2} \frac{\langle \psi | L_G | \psi \rangle}{2d} \quad (27)$$

$$= \frac{1}{2} \left(\|\alpha_1\|^2 + \frac{1}{2d} \sum_{i=1}^N \lambda_i \|\alpha_i\|^2 \right). \quad (28)$$

In a NO instance, $\lambda_k = \Omega(1)$ for all $k > 1$. So the probability of failure is always a positive constant:

$$\text{FAIL}_{\text{NO}} = \frac{1}{2} \|\alpha_1\|^2 + \frac{1}{2d} \sum_{i=2}^N \Omega(\|\alpha_i\|^2) = \Omega\left(\sum_{i=1}^N \|\alpha_i\|^2\right) = \Omega(1). \quad (29)$$

In a YES instance, the spectrum of L_G is the combined spectrum of the two disconnected graphs. This means $\lambda_1 = \lambda_2 = 0$, and the associated eigenvectors are linear combinations of $|V\rangle$ and $|[N]/V\rangle$. Recall that $|\lambda_1\rangle := |+\rangle^{\otimes n}$. We find the orthogonal eigenvector $|\lambda_2\rangle$ in this subspace by inspection:

$$|\lambda_2\rangle = \sqrt{\frac{N-|V|}{N}} |V\rangle + \sqrt{\frac{|V|}{N}} |[N]/V\rangle. \quad (30)$$

Note that any vector with $\|\alpha_2\|^2 = 1 - o(\frac{1}{\text{poly}(n)})$ has negligible probability of failure:

$$\text{FAIL}_{\text{YES}} = \frac{1}{2} \|\alpha_1\|^2 + \frac{1}{2d} \sum_{i=1}^N \lambda_i \|\alpha_i\|^2 = O(\|\alpha_1\|^2 + \sum_{i=3}^N \|\alpha_i\|^2) = O(1 - \|\alpha_2\|^2). \quad (31)$$

Suppose Merlin sends the subset state $|V\rangle$. Since $\|\langle V | \lambda_2 \rangle\|^2 = 1 - \frac{|V|}{N} = 1 - O(\frac{1}{2^{\text{poly}(n)}})$, the strategy has probability of failure $O(\frac{1}{2^{\text{poly}(n)}})$. ◀

In general, the spectral test can be used in a QMA protocol to test the magnitude of the second-smallest or largest eigenvalue of a graph Laplacian to inverse polynomial precision. The former is a measure of the quality of a graph's expansion, and the latter is related to a measure of a graph's bipartiteness named the *bipartiteness ratio* [26].

Because this QMA protocol requires only one query of either a standard oracle or an in-place oracle, it works even when these oracles are randomized.

► **Theorem A.4.** *There is a QMA protocol for STANDARD RANDOMIZED NON-EXPANSION(d, α, ε) and IN-PLACE RANDOMIZED NON-EXPANSION(d, α, ε) at every even $d \geq 4$, all $0 < \alpha < \frac{1}{2}$, and all constant $\varepsilon > 0$.*

11:14 On the Power of Nonstandard Quantum Oracles

Proof. The strategy in Theorem 3.1 also works here. Consider any G -coded function presented in a YES instance; the same vertices V are exactly the vertices of the smaller component of G . So the witness $|V\rangle$ is close to the second eigenvector $|\lambda_2\rangle$, and the failure probability is negligible. Now consider any G -coded function presented in a NO instance. By definition, G is an expander graph, so the failure probability is always a positive constant. ◀

Because a randomized oracle chooses a function uniformly from a set F , it is statistically indistinguishable from an oracle with exponentially small changes to F . We use this fact to simplify the NO instance in RANDOMIZED NON-EXPANSION(d, α, ε). Suppose the NO instance instead presents graph-coded functions of *all* d -regular graphs. Since $1 - O(\frac{1}{\text{poly}(N)}) = 1 - O(\frac{1}{2^{\text{poly}(n)}})$ graphs have a constant spectral gap [12] when $d > 2$, the failure probability in the QMA protocol changes by at most $O(\frac{1}{2^{\text{poly}(n)}})$.

Notice that with this modification, the oracles are exactly $d/2$ copies of the oracles in RANDOMIZED HIDDEN SUBSET(α). One way to interpret this is that the randomization offers a substitute for expander graphs. An expander graph is sparse but well-mixing; a randomized oracle query instantaneously mixes across a graph's connected component. As a result, we can distinguish degree-2 graphs with this QMA protocol, even though they are not typically expander graphs:

► **Theorem A.5.** *There is a QMA protocol for STANDARD RANDOMIZED HIDDEN SUBSET(α) and IN-PLACE RANDOMIZED HIDDEN SUBSET(α) for all $0 < \alpha < \frac{1}{2}$.*

Proof. Perhaps surprisingly, the strategy in Theorem 3.1 also works here:

- Consider the graph G of any G -coded function presented in a YES instance. By definition, the vertices V are disconnected from all vertices in $[N]/V$. So the witness $|V\rangle$ is close to the second eigenvector $|\lambda_2\rangle$, and the failure probability is negligible.
- Consider the NO instance. Then f is chosen uniformly from the set T_\emptyset of all permutations of $[N]$. Then the spectral test fails with probability

$$\mathbb{E}_{\pi \in T_\emptyset} \left[\frac{d - \langle \psi | A_\pi | \psi \rangle}{2d} \right] \Bigg|_{d=2} = \frac{1}{2} - \frac{1}{4} \mathbb{E}_{\pi \in T_\emptyset} [\langle \psi | A_\pi | \psi \rangle] \quad (32)$$

$$= \frac{1}{2} - \frac{1}{4} \left(\frac{1}{N!} \sum_{\pi \in T_\emptyset} \langle \psi | A_\pi | \psi \rangle \right) \quad (33)$$

$$= \frac{1}{2} - \frac{1}{8} \left(\frac{1}{(N!)^2} \sum_{\pi_1, \pi_2 \in T_\emptyset} \langle \psi | (A_{\pi_1} + A_{\pi_2}) | \psi \rangle \right). \quad (34)$$

The matrix $A_{\pi_1} + A_{\pi_2}$ determines the adjacency matrix of a random 4-regular graph in the configuration model; as a result,

$$\mathbb{E}_{\pi \in T_\emptyset} \left[\frac{d - \langle \psi | A_\pi | \psi \rangle}{2d} \right] \Bigg|_{d=2} = \mathbb{E}_{\pi_1, \pi_2 \in T_\emptyset} \left[\frac{d - \langle \psi | A_{\pi_1, \pi_2} | \psi \rangle}{2d} \right] \Bigg|_{d=4}. \quad (35)$$

Since a random 4-regular graph has constant spectral gap with probability $1 - O(\frac{1}{\text{poly}(N)}) = 1 - O(\frac{1}{2^{\text{poly}(n)}})$ [12], the failure probability is at least FAIL_{YES} from Theorem 3.1, less $O(\frac{1}{2^{\text{poly}(n)}})$. So the failure probability is $\Omega(1)$, just as before. ◀

B Randomized oracles and symmetric subspaces

We first formalize how randomized oracles are orthogonal projectors. We include the proofs in the arXiv version.

► **Definition B.1** (Representation of a group). *Consider a group G and a vector space V . A representation of G is a map R that sends each $g \in G$ to a linear operator $R(g) : V \rightarrow V$ such that $R(g_1 g_2) = R(g_1) \circ R(g_2)$ for all $g_1, g_2 \in G$.*

► **Theorem B.2** (Projecting onto the symmetric subspace [15, Proposition 2]). *Consider a finite group G , a vector space V , and a representation $R : G \rightarrow L(V)$. Then the operator*

$$\Pi_R := \frac{1}{|G|} \sum_{g \in G} R(g) \quad (36)$$

is an orthogonal projector onto $V^G \subseteq V$, where

$$V^G := \{v \in V : R(g)[v] = v \forall g \in G\}. \quad (37)$$

► **Theorem B.3** (Oracles on density matrices form a representation). *Consider a group G of functions $f : [N] \rightarrow [N]$ with bitwise \oplus as the group operation. Then the map $f \mapsto \mathcal{U}_f$ is a representation over the vector space of $2N^2 \times 2N^2$ complex matrices.*

Similarly, consider a group \tilde{G} of permutations $\pi : [N] \rightarrow [N]$ with composition as the group operation. Then the map $\pi \mapsto \tilde{\mathcal{U}}_\pi$ is a representation over the vector space of $2N \times 2N$ complex matrices.

► **Theorem B.4** (Some randomized oracles are orthogonal projectors). *Consider a group G of functions $f : [N] \rightarrow [N]$ with bitwise \oplus as the group operation. Then \mathcal{O}_G is an orthogonal projector, under the Frobenius inner product $\langle x|y \rangle = \text{Tr}[x^\dagger y]$ for $x, y \in \mathbb{C}^{2N^2 \times 2N^2}$, onto*

$$V_G := \{\rho \in \mathbb{C}^{2N^2 \times 2N^2} : \mathcal{U}_f[\rho] = \rho \forall f \in G\}. \quad (38)$$

Similarly, consider a group \tilde{G} of permutations $\pi : [N] \rightarrow [N]$ with composition as the group operation. Then $\tilde{\mathcal{O}}_{\tilde{G}}$ is an orthogonal projector, under the Frobenius inner product $\langle x|y \rangle = \text{Tr}[x^\dagger y]$ for $x, y \in \mathbb{C}^{2N \times 2N}$, onto

$$\tilde{V}_{\tilde{G}} := \{\rho \in \mathbb{C}^{2N \times 2N} : \tilde{\mathcal{U}}_\pi[\rho] = \rho \forall \pi \in \tilde{G}\}. \quad (39)$$

In IN-PLACE RANDOMIZED HIDDEN SUBSET(α), a quantum verifier is either given $\tilde{\mathcal{O}}_{T_\emptyset}$ (NO) or $\tilde{\mathcal{O}}_{T_V}$ for some $V \subseteq [N]$ where $|V| = N^\alpha$ (YES). Since $T_V \subseteq T_\emptyset$, the symmetric subspace according to T_\emptyset is a subspace of that according to T_V , i.e. $\tilde{V}_{T_\emptyset} \subseteq \tilde{V}_{T_V}$. So we can exactly find the basis of the symmetric subspaces \tilde{V}_{T_\emptyset} and \tilde{V}_{T_V} (see the arXiv version for details). This key property is used throughout Appendix B.1.

B.1 In-place oracles: when classical witnesses are not enough

We interpret IN-PLACE RANDOMIZED HIDDEN SUBSET(α) as distinguishing the set of all permutations from a subgroup that stabilizes a small subset $V \subseteq [N]$. In Theorem 3.3, we prove that classical witnesses designed for the verifier to choose YES cannot help a quantum verifier efficiently decide this problem. This requires three main lemmas. First, we show in Lemma B.6 that input states distinguishing a YES instance or NO instance must have knowledge of the hidden subset V (either as a subset state $|V\rangle$ or a mixed state \mathbb{I}_V). However,

11:16 On the Power of Nonstandard Quantum Oracles

no density matrix can be close to too many subset states $|V\rangle$ (Lemma B.7), and no POVM can choose the right answer for too many mixed states \mathbb{I}_V (Lemma B.8). We combine these facts in a hybrid argument; note that we must fix an algorithm by its unitaries *and* its POVM. We formally state the lemmas (proofs are in the arXiv version), and then prove Theorem 3.3.

We use the following measure of “progress” for the hybrid argument:

► **Definition B.5** (Difference of oracle queries). *For any ρ , let $d_{V,\rho}$ be the difference of the two oracle queries*

$$d_{V,\rho} := \tilde{\mathcal{O}}_{T_V}[\rho] - \tilde{\mathcal{O}}_{T_\emptyset}[\rho]. \quad (40)$$

If the nuclear norm of $d_{V,\rho}$ is non-negligible, we say that ρ is a good distinguisher of $\tilde{\mathcal{O}}_{T_V}$ and $\tilde{\mathcal{O}}_{T_\emptyset}$. We show that every good distinguisher ρ has a certain form; we include the proof in the arXiv version.

► **Lemma B.6** (Good distinguishers have a certain form). *Consider a density matrix ρ and up to $O(\log(n))$ extra workspace qubits. Suppose $\|d_{V,\rho}\|_1 = \Omega(\frac{1}{\text{poly}(n)})$. Then among the quantities*

$$\langle V, z | \rho | V, z \rangle, \quad (41)$$

$$\text{Tr} \left[\rho \left(\mathbb{I}_{V,z} - \frac{|V|}{N} \mathbb{I}_{[N],z} \right) \right], \quad (42)$$

for any $z \in \{\pm 1\}$, at least one has magnitude $\Omega(\frac{1}{\text{poly}(n)})$.

We now state two lemmas about subsets and subset states. These help us prove that no quantum state can be a good distinguisher of too many YES instances. We include the proofs in the arXiv version.

► **Lemma B.7** (Can’t approximate too many subset states). *Consider a Hermitian $N \times N$ matrix ρ that is positive semidefinite and has trace at most 1. Consider the set of all subsets $V \subseteq [N]$, where $|V| = N^\alpha$ for a fixed $0 < \alpha < \frac{1}{2}$. Then the fraction of subsets V such that $\langle V | \rho | V \rangle = \Omega(\frac{1}{\text{poly}(n)})$ decreases faster than any exponential in $\text{poly}(n)$.*

► **Lemma B.8** (Not too many subsets can have elevated mean). *Consider any $N \times N$ POVM $\{E, \mathbb{I} - E\}$, and the set of all subsets $V \subseteq [N]$, where $|V| = N^\alpha$ for a fixed $0 < \alpha < \frac{1}{2}$. Then the fraction of subsets V where*

$$|f(V)| := \left| \frac{1}{|V|} \text{Tr}[\mathbb{I}_V E] - \frac{1}{N} \text{Tr}[E] \right| = \Omega\left(\frac{1}{\text{poly}(n)}\right), \quad (43)$$

decreases faster than any exponential in $\text{poly}(n)$.

Intuitively, Lemma B.7 and Lemma B.8 hold because subset states can approximate *any* quantum state well. Grilo, Kerenidis, and Sikora [14] show that for any n -qubit quantum state $|\psi\rangle$, there exists a subset state $|S\rangle$ such that $|\langle S | \psi \rangle| \geq \frac{1}{8\sqrt{n+3}}$.

We now prove the main statement:

► **Theorem 3.3** (formal). *No quantum verifier that entangles oracle queries with at most $O(\log(n))$ additional qubits can efficiently decide IN PLACE RANDOMIZED HIDDEN SUBSET(α) for any $0 < \alpha < \frac{1}{2}$, even with a polynomial-length classical witness designed for the verifier to choose YES.*

Proof. Let the set of YES instances be \mathcal{Y} ; note that each YES instance corresponds to a set $V \subseteq [N]$ where $|V| = N^\alpha$, for some fixed $0 < \alpha < \frac{1}{2}$.

Suppose for contradiction that there is a protocol for this problem at some $\alpha < \frac{1}{2}$. Then the verifier can distinguish $\tilde{\mathcal{O}}_{T_\emptyset}$ from any $\tilde{\mathcal{O}}_{T_V}$ in a polynomial number of queries using a classical witness of size $O(\text{poly}(n))$. By the pigeonhole principle, there must exist a set of YES instances \mathcal{Y}' such that $|\mathcal{Y}'|/|\mathcal{Y}| = \Omega(\frac{1}{2^{\text{poly}(n)}})$, where the verifier can use the *same algorithm* to distinguish $\tilde{\mathcal{O}}_{T_\emptyset}$ from *every* YES instance in \mathcal{Y}' .

We then construct a *hybrid argument* in the style of Bennett, Bernstein, Brassard, and Vazirani [8] and Ambainis [4], which interpolates from queries of one oracle to queries of another oracle. For simplicity we write the proof without extra workspace qubits; however, we can have up to $O(\log(n))$ extra workspace qubits to satisfy Lemma B.6. Any polynomial query algorithm can be written as a set of unitaries $A = \{U^{(1)}, \dots, U^{(k)}\}$ for some $k = O(\text{poly}(n))$ (alternating between unitary evolutions and oracle queries), and a POVM $\{E, \mathbb{I} - E\}$. Consider the following “hybrid” algorithms:

► **Definition B.9.** *Given any set of k unitaries $A = \{U^{(1)}, \dots, U^{(k)}\}$, define the hybrid algorithm*

$$A_{V,\ell}[\rho_0] = \tilde{\mathcal{O}}_{T_V}^{(k)} \circ \mathcal{U}^{(k)} \circ \dots \circ \tilde{\mathcal{O}}_{T_V}^{(\ell+1)} \circ \mathcal{U}^{(\ell)} \circ \tilde{\mathcal{O}}_{T_\emptyset}^{(\ell)} \circ \mathcal{U}^{(\ell)} \circ \dots \circ \tilde{\mathcal{O}}_{T_\emptyset}^{(1)} \circ \mathcal{U}^{(1)}[\rho_0], \quad (44)$$

which evolves ρ_0 under the oracle $\mathcal{O}_{T_\emptyset}$ for ℓ steps and under \mathcal{O}_{T_V} for the other $k - \ell$ steps.

Then the following is true for each $\tilde{\mathcal{O}}_{T_V} \in \mathcal{Y}'$:

$$\Omega\left(\frac{1}{\text{poly}(n)}\right) = |\text{Tr}[EA_{V,k}[\rho_0]] - \text{Tr}[EA_{V,0}[\rho_0]]| \leq \sum_{i=0}^{k-1} |\text{Tr}[EA_{V,i+1}[\rho_0]] - \text{Tr}[EA_{V,i}[\rho_0]]|, \quad (45)$$

which implies

$$\Omega\left(\frac{1}{\text{poly}(n)}\right) = \sum_{i=0}^{k-1} \left| \text{Tr} \left[E \left(\tilde{\mathcal{O}}_{T_V}^{(k)} \circ \dots \circ \mathcal{U}^{(i)} \circ \left(\tilde{\mathcal{O}}_{T_V}^{(i)} - \tilde{\mathcal{O}}_{T_\emptyset}^{(i)} \right) [\rho^{(i)}] \right) \right] \right| = \sum_{i=0}^{k-1} \left| \text{Tr} [E^{V,(i)} d_{V,\rho^{(i)}}] \right|, \quad (46)$$

for the operator $E^{V,(i)}$ constructed by

$$E^{V,(i)} = \mathcal{U}^{\dagger(i)} \circ \tilde{\mathcal{O}}_{T_V}^{(i)} \circ \dots \circ \mathcal{U}^{\dagger(k)} \circ \tilde{\mathcal{O}}_{T_V}^{(k)} [E]. \quad (47)$$

By Fact D.8, the operators $E^{V,(i)}$ and $\mathbb{I} - E^{V,(i)}$ are also Hermitian and positive semidefinite, so $\{E^{V,(i)}, \mathbb{I} - E^{V,(i)}\}$ is a POVM.

Using the pigeonhole principle, there must be a step ℓ in the summation with magnitude $\Omega(\frac{1}{\text{poly}(n)})$. Each $\tilde{\mathcal{O}}_{T_V} \in \mathcal{Y}'$ has such a step. Again by the pigeonhole principle, there is a ℓ^* and set $\mathcal{Y}^* \subseteq \mathcal{Y}'$ where

$$\left| \text{Tr} [E^{V,(\ell^*)} d_{V,\rho^{(\ell^*)}}] \right| = \Omega\left(\frac{1}{\text{poly}(n)}\right), \quad (48)$$

and $|\mathcal{Y}^*|/|\mathcal{Y}'| \geq \frac{1}{k} = \Omega(\frac{1}{\text{poly}(n)})$. Notice that this implies $|\mathcal{Y}^*|/|\mathcal{Y}| = \Omega(\frac{1}{2^{\text{poly}(n)}})$.

Since the trace of M with a POVM operator is at most $\|M\|_1$ (Fact D.5), we have for all $\tilde{\mathcal{O}}_{T_V} \in \mathcal{Y}^*$,

$$\Omega\left(\frac{1}{\text{poly}(n)}\right) = \left| \text{Tr} [E^{V,(\ell^*)} d_{V,\rho^{(\ell^*)}}] \right| \leq \|d_{V,\rho^{(\ell^*)}}\|_1. \quad (49)$$

11:18 On the Power of Nonstandard Quantum Oracles

When queries are entangled with at most $O(\log(n))$ additional qubits, the premise of Lemma B.6 holds; then one of the quantities in the theorem statement must be large. However, Lemma B.7 says that a given ρ can only satisfy either of the first two quantities for a smaller-than-exponential fraction of \mathcal{Y} . So for most choices of $\tilde{\mathcal{O}}_{T_V} \in \mathcal{Y}^*$,

$$\mathrm{Tr} \left[\rho^{(\ell^*)} \left(\mathbb{I}_{V,z} - \frac{|V|}{N} \mathbb{I}_{[N],z} \right) \right] = \Omega\left(\frac{1}{\mathrm{poly}(n)}\right). \quad (50)$$

for at least one of $z \in \{\pm 1\}$.

Inspecting the proof of Lemma B.6, this implies $d_{V,\rho^{(\ell^*)}}$ can only have $\Omega(\frac{1}{\mathrm{poly}(n)})$ weight on $C_{4,z}$ for some $z \in \{\pm 1\}$ across all matrices in \mathcal{C} . In fact, for most choices of $\tilde{\mathcal{O}}_{T_V} \in \mathcal{Y}^*$, we show that this is also true for

$$d_{V,\ell^*,j} := \tilde{\mathcal{O}}_{T_V}^{(\ell^*+j)} \circ \mathcal{U}^{(\ell^*+j)} \circ \dots \circ \tilde{\mathcal{O}}_{T_V}^{(\ell^*+1)} \circ \mathcal{U}^{(\ell^*+1)} \circ d_{V,\rho^{(\ell^*)}}, \quad (51)$$

for all $0 \leq j \leq k - \ell^*$. We show this by induction. Note that by Fact D.5 and the fact that $d_{V,\ell^*,k-\ell^*}$ is the difference of two objects with nuclear norm 1, $\|d_{V,\ell^*,k-\ell^*}\|_1 = \Omega(\frac{1}{\mathrm{poly}(n)}) = O(1)$.

Consider $d_{V,\ell^*,i}$ for some $1 \leq i \leq k - \ell^*$, which can be represented with the basis \mathcal{C} . By Fact D.9, it has Frobenius norm at most $\|d_{V,\rho^{(\ell^*)}}\|_{Fr} = O(\frac{1}{\sqrt{|V|}})$. So it must have $o(\frac{1}{\mathrm{poly}(n)})$ weight on pure states. Inspecting the basis \mathcal{C} , this means $d_{V,\rho^{(\ell^*)}}$ can only have $\Omega(\frac{1}{\mathrm{poly}(n)})$ weight on $C_{4,z}$ or $\frac{1}{N} \mathbb{I}_{[N],z}$ for $z \in \{\pm 1\}$. By Fact D.9, $d_{V,\rho^{(\ell^*)}}$ has nuclear norm at least $\|d_{V,\ell^*,k-\ell^*}\|_1 = \Omega(\frac{1}{\mathrm{poly}(n)})$, so it must have $\Omega(\frac{1}{\mathrm{poly}(n)})$ weight on at least one such matrix. Suppose for contradiction that the matrix is $\frac{1}{N} \mathbb{I}_{[N],z}$ for $z \in \{\pm 1\}$. Then

$$\Omega\left(\frac{1}{\mathrm{poly}(n)}\right) = \mathrm{Tr} \left[\mathbb{I}_{[N],z} \tilde{\mathcal{O}}_{T_V} \left[\mathcal{U}^{(\ell^*+i)} [d_{V,\ell^*,i-1}] \right] \right] = \mathrm{Tr} \left[\left(\mathcal{U}^{(\ell^*+i)\dagger} \circ \tilde{\mathcal{O}}_{T_V}^\dagger [\mathbb{I}_{[N],z}] \right) d_{V,\ell^*,i-1} \right]. \quad (52)$$

By the inductive hypothesis, $d_{V,\ell^*,i-1}$ only has $\Omega(\frac{1}{\mathrm{poly}(n)})$ weight on some $C_{4,z}$ for $z \in \{\pm 1\}$. Then for some $z' \in \{\pm 1\}$,

$$\Omega\left(\frac{1}{\mathrm{poly}(n)}\right) = \mathrm{Tr} \left[\left(\mathcal{U}^{(\ell^*+i)\dagger} \circ \tilde{\mathcal{O}}_{T_V}^\dagger [\mathbb{I}_{[N],z}] \right) \left(\frac{1}{|V|} \mathbb{I}_{V,z'} - \frac{1}{N} \mathbb{I}_{[N],z'} \right) \right]. \quad (53)$$

Notice that for any unitary U , the object $\{ \mathcal{U}^{(\ell^*+i)\dagger} \circ \tilde{\mathcal{O}}_{T_V}^\dagger [\mathbb{I}_{[N],z=+1}], \mathcal{U}^{(\ell^*+i)\dagger} \circ \tilde{\mathcal{O}}_{T_V}^\dagger [\mathbb{I}_{[N],z=-1}] \}$ forms a POVM. By Lemma B.8, this can only be satisfied at either $z \in \{\pm 1\}$ for a smaller-than-exponential fraction of choices of V . So for most choices of $\tilde{\mathcal{O}}_{T_V} \in \mathcal{Y}^*$ (i.e. a $\Omega(\frac{1}{2^{\mathrm{poly}(n)}})$ fraction of choices of V), $d_{V,\ell^*,i}$ has $\Omega(\frac{1}{\mathrm{poly}(n)})$ weight on $C_{4,z}$ for at least one of $z \in \{\pm 1\}$, and for no other matrices in \mathcal{C} .

Since $\Omega(\frac{1}{\mathrm{poly}(n)}) = |\mathrm{Tr}[E d_{V,\ell^*,k-\ell^*}]|$, our supposition then implies that for one of $z \in \{\pm 1\}$,

$$\Omega\left(\frac{1}{\mathrm{poly}(n)}\right) = |\mathrm{Tr}[E C_{4,z}]| = \left| \mathrm{Tr} \left[E \left(\frac{1}{|V|} \mathbb{I}_{V,z} - \frac{1}{N} \mathbb{I}_{[N],z} \right) \right] \right| \pm O\left(\frac{1}{2^{\mathrm{poly}(n)}}\right). \quad (54)$$

But by Lemma B.8, this can only be satisfied at either $z \in \{\pm 1\}$ for a smaller-than-exponential fraction of \mathcal{Y} . This is a contradiction. So there can be no efficient protocol for this problem. \blacktriangleleft

B.2 Standard oracles: when classical witnesses are enough

As shown in Theorem B.3, randomized standard oracles can also form a representation. But the preserved group structure is much different than for randomized in-place oracles. Consider the set T_\emptyset of permutations on $[N]$. For any $f_1, f_2 \in T_\emptyset$, the element $f_1 f_2$ in this group structure acts for all $x \in [N]$ and $z \in \{\pm 1\}$ as

$$(f_1 f_2)^z(x) = f_1^z(x) \oplus f_2^z(x). \quad (55)$$

Note that this operation is abelian; that is, $(f_1 f_2) = (f_2 f_1)$. Any finite abelian group can always be represented as the direct sum of cyclic groups. In fact, under this group operation, T_\emptyset can be decomposed by the input $x \in [N]$ and function inverter $z \in \{\pm 1\}$:

$$T_\emptyset = \bigoplus_{x \in [2^n], z \in \{\pm 1\}} \mathbb{Z}_{2^n}. \quad (56)$$

With this group operation, the only possible subgroups of T_\emptyset have the form

$$\bigoplus_{x \in [2^n], z \in \{\pm 1\}} \mathbb{Z}_{2^{k_{x,z}}}, \quad (57)$$

for $0 \leq k_{x,z} \leq n$. As a result, there is a QCMA protocol to distinguish any strict subgroup of T_\emptyset from T_\emptyset .

► **Theorem B.10.** *There is a one-query QCMA protocol for STANDARD RANDOMIZED HIDDEN SUBGROUP($\times, \{H_i\}$) when the group operation \times is bitwise XOR, for any valid $\{H_i\}$.*

Proof. Suppose the classical witness is a bitstring of length at least $n + 1$. The verifier can then:

1. Use the first n bits to construct x and the next bit to construct z .
2. Prepare the state $|0\rangle^{\otimes n} |x, z\rangle$.
3. Apply \mathcal{O}_H , creating the state $|f^z(x)\rangle |x, z\rangle$ for some $f \in H$.
4. Measure the first n qubits, and accept if the result is even.³

Consider a YES instance associated with a subgroup $H \subsetneq T_\emptyset$. Then H will have some $x \in [N], z \in \{\pm 1\}$ such that $k_{x,z} < n$. A witness can store x and z ; since $k_{x,z} < n$, $f^z(x)$ will be even with probability 1.

In the NO instance, $H = T_\emptyset$. Then $f^z(x)$ is even with probability 0.5 for every $x \in [N], z \in \{\pm 1\}$. ◀

Note that Theorem B.10 holds even if the randomized standard oracle \mathcal{O}_F does not have access to the function inverse.

C No witness is enough for phase oracles

► **Theorem 3.4 (formal).** *No quantum verifier that entangles oracle queries with at most $o(n)$ additional qubits can efficiently decide PHASE RANDOMIZED HIDDEN SUBSET(α) for any $0 < \alpha < \frac{1}{2}$, even with any witness designed for the verifier to choose YES. Moreover, these verifiers require an exponential number of queries to statistically distinguish a YES instance from the NO instance, for each of asymptotically all YES instances.*

³ Depending on the encoding, one can simply measure the n^{th} qubit, and accept if the result is 0.

11:20 On the Power of Nonstandard Quantum Oracles

Proof. We now prove the query lower bound. Let k be the number of queries required to distinguish $\overline{\mathcal{O}}_{T_V}$ from $\overline{\mathcal{O}}_{T_\emptyset}$. Consider any algorithm that distinguishes the two instances, defined by a starting state ρ_0 , k unitaries, k oracle queries, and a POVM $\{E, \mathbb{I} - E\}$. In the framework of hybrid algorithms (Definition B.9),

$$\Omega\left(\frac{1}{\text{poly}(n)}\right) = |\text{Tr}[EA_{V,k}[\rho_0]] - \text{Tr}[EA_{V,0}[\rho_0]]| \quad (58)$$

$$\leq \|A_{V,k}[\rho_0] - A_{V,0}[\rho_0]\|_1 \quad (59)$$

$$\leq k \max_{i \in \{0, \dots, k-1\}} \|A_{V,i+1}[\rho_0] - A_{V,i}[\rho_0]\|_1 \quad (60)$$

$$\leq k \max_{i \in \{0, \dots, k-1\}} \left\| \overline{\mathcal{O}}_{T_V}[\rho^{(i)}] - \overline{\mathcal{O}}_{T_\emptyset}[\rho^{(i)}] \right\|_1, \quad (61)$$

where the last line follows because randomized oracles do not increase the nuclear norm (Fact D.9).

We now bound $\|\overline{\mathcal{O}}_{T_V}[\rho] - \overline{\mathcal{O}}_{T_\emptyset}[\rho]\|_1$ for *any* ρ . Recall that a phase oracle $\overline{\mathcal{O}}_F$ acts as

$$\overline{\mathcal{O}}_F[|x_1, z_1\rangle\langle x_2, z_2|] = \frac{1}{|F|} \sum_{f \in F} \omega_N^{f^{z_1}(x_1) - f^{z_2}(x_2)} |x_1, z_1\rangle\langle x_2, z_2|, \quad (62)$$

for any $x_1, x_2 \in [N]$ and $z_1, z_2 \in \{\pm 1\}$. So every basis vector $|x_1, z_1\rangle\langle x_2, z_2|$ acquires a coefficient c_{x_1, z_1, x_2, z_2} .

We start with $\overline{\mathcal{O}}_{T_\emptyset}$ (the NO instance). When $(x_1, z_1) = (x_2, z_2)$, the coefficient is 1. When $x_1 \neq x_2$, $f^z(x_1)$ and $f^{-z}(x_2)$ are uniformly likely to be any value, so the coefficient is

$$\frac{1}{N^2} \sum_{a \in [N], b \in [N]} \omega_N^{a-b} = \frac{1}{N^2} \left\| \sum_{a \in [N]} \omega_N^a \right\|^2 = 0. \quad (63)$$

Similarly, when $x_1 \neq x_2$, $f^z(x_1)$ and $f^z(x_2)$ are uniformly likely to be any unequal values; the coefficient is

$$\frac{1}{N(N-1)} \sum_{a \in [N], b \in [N], a \neq b} \omega_N^{a-b} = \frac{1}{N(N-1)} \left[\sum_{a \in [N], b \in [N]} \omega_N^{a-b} - \sum_{a \in [N]} \omega_N^{a-a} \right] = -\frac{1}{N-1}. \quad (64)$$

We now consider $\overline{\mathcal{O}}_{T_V}$ (a YES instance). When $(x_1, z_1) = (x_2, z_2)$, the coefficient is again 1. When $x_1 \neq x_2$, the values of $f^z(x_1)$ and $f^{-z}(x_2)$ are uniformly likely to be any value in $i_V(x_1)$ and $i_V(x_2)$, respectively, so the coefficient is

$$\frac{1}{|i_V(x_1)| \times |i_V(x_2)|} \sum_{a \in i_V(x_1), b \in i_V(x_2)} \omega_N^{a-b}. \quad (65)$$

Similarly, when $x_1 \neq x_2$, $f^z(x_1)$ and $f^z(x_2)$ are uniformly likely to be any unequal values in $i_V(x_1)$ and $i_V(x_2)$, respectively, so the coefficient is

$$\frac{1}{|i_V(x_1)| \times |i_V(x_2)|} \sum_{a \in i_V(x_1), b \in i_V(x_2), a \neq b} \omega_N^{a-b} = \frac{\left(\sum_{a \in i_V(x_1), b \in i_V(x_2)} \omega_N^{a-b} \right) - \delta |i_V(x_1)|}{|i_V(x_1)| \times |i_V(x_2)| - \delta}, \quad (66)$$

where δ is 1 if $i_V(x_1) = i_V(x_2)$ and 0 otherwise.

Consider the object $\overline{\mathcal{O}}_{T_V}[\rho] - \overline{\mathcal{O}}_{T_\emptyset}[\rho]$ as the sum of two matrices $A_V + B_V$. Let A_V contain the $(V, z) \times (V, z)$ submatrix for both $z \in \{\pm 1\}$, and B_V contain the rest of the entries. When the oracle query is entangled with $o(n)$ additional qubits, A_V has rank $O(|V|)$, and B_V has rank $O(N)$.

Since the roots of unity sum to zero, $\sum_{a \in V} \omega_N^a = -\sum_{a \in [N]/V} \omega_N^a$ for any $V \subseteq [N]$. Because of this,

$$\left\| \sum_{a, b \in i_V(x)} \omega_N^{a-b} \right\| \leq \left\| \sum_{a \in i_V(x)} \omega_N^a \right\|^2 \leq \left\| \sum_{a \in V} \omega_N^a \right\|^2 = O(|V|^2). \quad (67)$$

As a result, all coefficients in B_V are $O(\frac{1}{N^{1-\alpha}})$.

In Lemma C.1, we show that for asymptotically all choices of V , all coefficients in A_V are $O(\frac{1}{N^{3\alpha/4}})$. This argument uses a Chernoff bound and a central limit argument on samples without replacement.

We bound the nuclear norm of $\overline{\mathcal{O}}_{T_V}[\rho] - \overline{\mathcal{O}}_{T_\emptyset}[\rho]$ with the rank and Frobenius norm of A_V and B_V (Fact D.3):

$$\|\overline{\mathcal{O}}_{T_V}[\rho] - \overline{\mathcal{O}}_{T_\emptyset}[\rho]\|_1 \leq \|A_V\|_1 + \|B_V\|_1 \quad (68)$$

$$= O(\sqrt{V})\|A_V\|_{Fr} + O(\sqrt{N})\|B_V\|_{Fr} \quad (69)$$

$$\leq \left(O(\sqrt{V})O(N^{-3\alpha/4}) + O(\sqrt{N})O(N^{\alpha-1}) \right) \|\rho\|_{Fr} \quad (70)$$

$$= O(N^{-\alpha/4} + N^{\alpha-1/2}). \quad (71)$$

Thus, for most choices of V , distinguishing $\overline{\mathcal{O}}_{T_V}$ and $\overline{\mathcal{O}}_{T_\emptyset}$ requires $k = \Omega(\min(N^{\alpha/4}, N^{1/2-\alpha}))$ queries. \blacktriangleleft

We now prove the Chernoff bound:

► **Lemma C.1.** *Fix any $0 < \alpha < \frac{1}{2}$, and consider all subsets $V \subseteq [N]$ such that $|V| = N^\alpha$. Then for all but a doubly exponentially small fraction of choices of V ,*

$$\left\| \frac{1}{N^{2\alpha}} \sum_{a, b \in V} \omega_N^{a-b} \right\| = O\left(\frac{1}{N^{3\alpha/4}}\right). \quad (72)$$

Proof. Consider the distribution $X = \{\omega_N^k\}$ where k is chosen uniformly from N . Both $Re(X)$ and $Im(X)$ have mean zero and variance at most 1.

Take a size- N^α sample from the distribution X , *without replacement*. Denote Y as the distribution of the sample mean. Both $Re(Y)$ and $Im(Y)$ have expectation $Re(X) = Im(X) = 0$, and variance

$$\frac{\sigma_X^2}{N^\alpha} \left(1 - \frac{N^\alpha - 1}{N - 1}\right) \leq \frac{1}{N^\alpha}. \quad (73)$$

Even when sampling without replacement, Y is asymptotically normally distributed [10]. So its moment generating function is

$$\text{MGF}_Y[t] = e^{t\mu_Y + \sigma_Y^2 t^2 / 2} \leq e^{t^2 / N^\alpha}. \quad (74)$$

We use a Chernoff bound to estimate when Y has magnitude at least $N^{-3\alpha/8}$. Notice that

$$\Pr[Y \geq a] = \Pr[e^{tY} \geq e^{ta}] \leq e^{-at} \text{MGF}_Y[t] \leq e^{t^2 / N^\alpha} e^{-at}, \quad (75)$$

11:22 On the Power of Nonstandard Quantum Oracles

so

$$\Pr\left[Y \geq \frac{0.5}{N^{3\alpha/8}}\right] \leq \inf_{t \geq 0} \exp\left(\frac{t^2}{N^\alpha} - \frac{0.5t}{N^{3\alpha/8}}\right) \Big|_{t=2N^{\alpha/2}} \leq \exp(4 - N^{\alpha/8}) = O\left(\frac{1}{\exp(\exp(n))}\right). \quad (76)$$

This implies that Y has magnitude at most $N^{-3\alpha/8}$ (and Y^2 at most $N^{-3\alpha/4}$) except in a doubly exponentially small fraction of choices of V . ◀

► **Remark C.2.** Consider an oracle that sends $|c, x, z\rangle \rightarrow \omega_N^{c \cdot f^z(x)} |c, x, z\rangle$, where the c register has k qubits. Note that Theorem 3.4 applies whenever $k = o(n)$. However, there must be a phase transition, since at $k = n$, this oracle is unitarily equivalent to a standard oracle, and thus has a QMA protocol for RANDOMIZED HIDDEN SUBSET(α) in Appendix A.

D Norms and inner products

Note that we work with arbitrary matrices, not just positive semidefinite ones.

► **Definition D.1** (Nuclear norm of a matrix). *The nuclear norm of a matrix M is the sum of its singular values; that is,*

$$\|M\|_1 = \sum_i \sigma_i(M) = \text{Tr}\left[\sqrt{M^\dagger M}\right]. \quad (77)$$

► **Definition D.2** (Frobenius norm and inner product of a matrix). *The Frobenius inner product of $N \times N$ matrices A, B is*

$$(A|B) = \text{Tr}[A^\dagger B] \quad (78)$$

This induces a norm, which is the square root of the sum of squares of the singular values:

$$\|A\|_{Fr} = \sqrt{\sum_i \sigma_i(A)^2} = \sqrt{\sum_{ij \in [N]} |A_{ij}|^2}. \quad (79)$$

► **Fact D.3.** *The nuclear norm of a matrix is at most the product of its Frobenius norm and the square root of its rank.*

Proof. See Rennie [24] for a proof with explanation. ◀

► **Fact D.4** (Nuclear norm of a positive semidefinite matrix). *The nuclear norm of a positive semidefinite Hermitian matrix is simply its trace; that is, if ρ is Hermitian and positive semidefinite, then*

$$\|\rho\|_1 = \text{Tr}[\rho]. \quad (80)$$

Proof. For a Hermitian and positive semidefinite matrix, the eigenvalues are all real and nonnegative, so the singular values are exactly the eigenvalues. Alternatively, notice that $\rho = \sqrt{\rho^\dagger \rho}$ and use Definition D.1. ◀

► **Fact D.5** (POVM trace is at most the nuclear norm). *Consider any Hermitian matrix M and a POVM $\{E, \mathbb{I} - E\}$. Then*

$$\text{Tr}[EM] \leq \|M\|_1. \quad (81)$$

Proof. Consider the singular value decomposition of a Hermitian $M = UDU^\dagger$. Then

$$\mathrm{Tr}[EM] = \mathrm{Tr}[(U^\dagger EU)D] = \mathrm{Tr}[E'D] \quad (82)$$

for some matrix E' . Note that $\{E', \mathbb{I} - E'\}$ make a POVM; they have the same eigenvalues of $\{E, \mathbb{I} - E\}$, respectively, and so are both positive semidefinite. Recall that the diagonal elements of a POVM are all nonnegative and at most 1. Then

$$\mathrm{Tr}[EM] = \mathrm{Tr}[E'D] = \sum_i E'_{ii} D_i \leq \sum_i |D_i| = \|D\|_1 = \|M\|_1. \quad (83)$$

► **Fact D.6** (Trace of outer product is inner product). Consider vectors $|x\rangle, |y\rangle \in \mathbb{C}^m$ and a matrix $A \in \mathbb{C}^{m \times m}$. Then the inner product of $|y\rangle\langle x|$ and A is

$$\mathrm{Tr}[(|y\rangle\langle x|)^\dagger A] = \mathrm{Tr}[A|x\rangle\langle y|] = \langle y|A|x\rangle. \quad (84)$$

Proof.

$$\mathrm{Tr}[A|x\rangle\langle y|] = \mathrm{Tr} \left[\sum_{i,k \in [m]} \left(\sum_{j \in [m]} A_{ij} x_j \right) y_k^\dagger \right] = \sum_{k,j \in [m]} A_{kj} x_j y_k^\dagger = \langle y|A|x\rangle. \quad (85)$$

► **Remark D.7** (Orthogonal basis for an input density matrix). We can decompose ρ into a basis \mathcal{M} that is orthogonal under the Frobenius inner product ($a|b$) = $\mathrm{Tr}[a^\dagger b]$:

$$\rho = \sum_{M \in \mathcal{M}} c_M M. \quad (86)$$

Because the basis is orthogonal, for any $M \in \mathcal{M}$,

$$\mathrm{Tr}[M^\dagger \rho] = \sum_{M' \in \mathcal{M}} c_{M'} \mathrm{Tr}[M^\dagger M'] = c_M \|M\|_{Fr}^2. \quad (87)$$

Moreover, by Cauchy-Schwarz, the inner product of M and ρ is at most the product of the norm of each, so

$$\|c_M M\|_{Fr} = \frac{|\mathrm{Tr}[M^\dagger \rho]|}{\|M\|_{Fr}} \leq \|\rho\|_{Fr}. \quad (88)$$

We also state two properties that hold for *any* randomized oracle:

► **Fact D.8** (Randomized oracles preserve trace, Hermiticity, and positive semidefiniteness). Consider any randomized oracle \mathcal{O}_F corresponding to a set of functions $f \in F$. Then \mathcal{O}_F preserves the trace of its input. Moreover, if the input M is Hermitian, so is $\mathcal{O}_F[M]$; if M is also positive semidefinite, so is $\mathcal{O}_F[M]$.

Proof. Consider any input matrix M . Then

$$\mathrm{Tr}[\mathcal{O}_F[M]] = \mathrm{Tr} \left[\frac{1}{|F|} \sum_{f \in F} \mathcal{U}_f[M] \right] = \frac{1}{|F|} \sum_{f \in F} \mathrm{Tr}[\mathcal{U}_f M \mathcal{U}_f^\dagger] = \frac{1}{|F|} \sum_{f \in F} \mathrm{Tr}[M] = \mathrm{Tr}[M]. \quad (89)$$

11:24 On the Power of Nonstandard Quantum Oracles

Now suppose M is Hermitian; that is, $M^\dagger = M$. Then

$$\mathcal{O}_F[M]^\dagger = \left(\frac{1}{|F|} \sum_{f \in F} \mathcal{U}_f[M] \right)^\dagger = \frac{1}{|F|} \sum_{f \in F} (U_f M U_f^\dagger)^\dagger = \frac{1}{|F|} \sum_{f \in F} U_f M^\dagger U_f = \mathcal{O}_F[M]. \quad (90)$$

Furthermore, suppose M is positive semidefinite; that is, there is a matrix B such that $M = B^\dagger B$. Then

$$\mathcal{O}_F[M] = \frac{1}{|F|} \sum_{f \in F} \mathcal{U}_f[M] = \frac{1}{|F|} \sum_{f \in F} U_f B^\dagger B U_f^\dagger = \sum_{f \in F} (B U_f)^\dagger (B U_f), \quad (91)$$

which is a sum of positive semidefinite matrices. Thus, $\mathcal{O}_F[M]$ is positive semidefinite. ◀

► **Fact D.9** (Randomized oracles do not increase nuclear norm or Frobenius norm). *Consider any randomized oracle \mathcal{O}_F corresponding to a set of functions $f \in F$. Then \mathcal{O}_F does not increase the nuclear norm nor the Frobenius norm of its input.*

Proof. Recall that both the nuclear norm and Frobenius norm are unitarily invariant. Now consider any input matrix M . Then the nuclear norm of $\mathcal{O}_F[M]$ is

$$\|\mathcal{O}_F[M]\|_1 = \left\| \frac{1}{|F|} \sum_{f \in F} U_f M U_f^\dagger \right\|_1 \leq \frac{1}{|F|} \sum_{f \in F} \|U_f M U_f^\dagger\|_1 = \frac{1}{|F|} \sum_{f \in F} \|M\|_1 = \|M\|_1. \quad (92)$$

The Frobenius norm of $\mathcal{O}_F[M]$ follows in exactly the same way. ◀

We use one additional property of density matrices in the proof of Lemma B.6:

► **Fact D.10.** *Consider any $N \times N$ density matrix ρ and normalized states $|v\rangle, |w\rangle$. If $|\langle v | \rho | w \rangle| = \Omega(\frac{1}{\text{poly}(n)})$, then both $\langle v | \rho | v \rangle$ and $\langle w | \rho | w \rangle$ are $\Omega(\frac{1}{\text{poly}(n)})$.*

Proof. Recall that a density matrix is Hermitian and positive semidefinite, so it is diagonalizable and has real and nonnegative eigenvalues. As a result, it has a decomposition

$$\rho = S^\dagger \Lambda S = S^\dagger \sqrt{\Lambda} \sqrt{\Lambda} S = (\sqrt{\Lambda} S)^\dagger (\sqrt{\Lambda} S) = A^\dagger A, \quad (93)$$

for some diagonal Λ and $A := \sqrt{\Lambda} S$. Then by Cauchy-Schwarz,

$$|\langle v | \rho | w \rangle|^2 = |(A|v\rangle)^\dagger (A|w\rangle)|^2 \leq |(A|v\rangle)^\dagger (A|v\rangle)| \cdot |(A|w\rangle)^\dagger (A|w\rangle)| = \langle v | \rho | v \rangle \cdot \langle w | \rho | w \rangle. \quad (94)$$

Since $\text{Tr}[\rho] = 1$, $\langle \psi | \rho | \psi \rangle \leq 1$ for all normalized states $|\psi\rangle$. Thus, both $\langle v | \rho | v \rangle$ and $\langle w | \rho | w \rangle$ are at least $|\langle v | \rho | w \rangle|^2 = \Omega(\frac{1}{\text{poly}(n)})$. ◀

E Our setup contrasted with a discrete-time quantum walk

The way one stores a graph in an oracle drastically changes the difficulty of some problems. Consider a *discrete-time quantum walk* [27], which allows a vertex access to a superposition of its neighbors.⁴ Given a d -regular graph $G(V, E)$, the operator $W : \mathbb{C}^{N^2 \times N^2}$ acts as

$$W = \left(\sum_{(j,k) \in E} |j, k\rangle \langle k, j| \right) C \quad (95)$$

$$C = \sum_{j \in V} |j\rangle \langle j| \otimes (2|\partial_j\rangle \langle \partial_j| - \mathbb{I}) \quad (96)$$

$$|\partial_j\rangle = \frac{1}{\sqrt{d}} \sum_{(j,k) \in E} |k\rangle. \quad (97)$$

Using a discrete-time quantum walk, we can learn about the mixing properties of the associated graph; these are fundamentally related to the graph's spectral gap [13].

By contrast, we query each neighbor of a vertex $v \in G$ with the value of the registers encoding $i \in [d/2]$ (defined by a G -coded function). For example, [5] uses a similar oracle to show that deciding whether a graph is a single expander graph or two equal-sized disconnected expander graphs is outside of BQP. Intuitively, a lack of superposition access to neighbors of a vertex makes it harder for a quantum computer to “traverse” the graph.

⁴ [9, Chapter 17] has a good introduction to this topic.