Sequential measurements, disturbance and property testing

Aram W. Harrow*, Cedric Yen-Yu Lin† and Ashley Montanaro‡

October 4, 2016

Abstract

We describe two procedures which, given access to one copy of a quantum state and a sequence of two-outcome measurements, can distinguish between the case that at least one of the measurements accepts the state with high probability, and the case that all of the measurements have low probability of acceptance. The measurements cannot simply be tried in sequence, because early measurements may disturb the state being tested. One procedure is based on a variant of Marriott-Watrous amplification. The other procedure is based on the use of a test for this disturbance, which is applied with low probability. We find a number of applications:

- Quantum query complexity separations in the property testing model for testing isomorphism of functions under group actions. We give quantum algorithms for testing isomorphism, linear isomorphism and affine isomorphism of boolean functions which use exponentially fewer queries than is possible classically, and a quantum algorithm for testing graph isomorphism which uses polynomially fewer queries than the best algorithm known.
- Testing properties of quantum states and operations. We show that any finite property of quantum states can be tested using a number of copies of the state which is logarithmic in the size of the property, and give a test for genuine multipartite entanglement of states of n qubits that uses $O(n)$ copies of the state. We also show that equivalence of two unitary operations under conjugation by a unitary picked from a fixed set can be tested efficiently. This is a natural quantum generalisation of testing isomorphism of boolean functions.
- Correcting an error in a result of Aaronson on de-Merlinizing quantum protocols. This result claimed that, in any one-way quantum communication protocol where two parties are assisted by an all-powerful but untrusted third party, the third party can be removed with only a modest increase in the communication cost. We give a corrected proof of a key technical lemma required for Aaronson's result.

1 Introduction

In quantum mechanics, measuring the state of a system generally disturbs it. This effect can be harnessed for practical purposes – such as in the field of quantum cryptography [17] – but is often an annoyance. In particular, the following question is nontrivial: Given one copy of a state $|\psi\rangle$ and a set of n measurements, determine whether at least one of the measurements would accept $|\psi\rangle$ with high probability. If one simply tries the measurements one after the other, early measurements may disturb $|\psi\rangle$, leading later measurements to incorrectly reject.

*Center for Theoretical Physics, Massachusetts Institute of Technology, USA; aram@mit.edu

†Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, MD, USA; cedricl@umiacs.umd.edu

‡School of Mathematics, University of Bristol, UK; ashley.montanaro@bristol.ac.uk.

A vivid example of this phenomenon is the quantum anti-Zeno effect [4, 29, 22]. Just as the more familiar quantum Zeno effect [26] describes a situation where the time-evolution of a quantum system is halted by frequent measurements, the anti-Zeno effect describes a situation where time-evolution is *caused* by measurement[1]. For example, consider the sequence $M_1, \ldots, M_n$ of 2-outcome measurements of one qubit such that the k'th measurement is specified by the pair $\{|\psi_k\rangle\langle\psi_k|, I - |\psi_k\rangle\langle\psi_k|\}$, where we associate the first outcome with rejection, the second with acceptance, and

$$|\psi_k\rangle = \cos\left(\frac{\pi k}{2n}\right)|0\rangle + \sin\left(\frac{\pi k}{2n}\right)|1\rangle.$$

Then, if we have the state $|\psi_k\rangle$ and apply the measurement M_{k+1}, the probability of rejection is

$$\begin{aligned} |\langle\psi_k|\psi_{k+1}\rangle|^2 &= \left(\cos\left(\frac{\pi k}{2n}\right)\cos\left(\frac{\pi(k+1)}{2n}\right) + \sin\left(\frac{\pi k}{2n}\right)\sin\left(\frac{\pi(k+1)}{2n}\right)\right)^2 \\ &= \cos\left(\frac{\pi}{2n}\right)^2 \\ &= 1 - O(1/n^2) \end{aligned}$$

and the residual state following rejection is $|\psi_{k+1}\rangle$. Therefore, if each of the n measurements is performed in order, starting with the initial state $|0\rangle = |\psi_0\rangle$, the probability of ever seeing an "accept" measurement outcome is $O(1/n)$. However, the final state is $|\psi_n\rangle = |1\rangle$, which is orthogonal to the initial state; and indeed if the final measurement M_n were performed on $|0\rangle$, it would accept with certainty.

Here we describe quantum procedures that combat this effect:

Theorem 1. *Let $\Lambda_1, \ldots, \Lambda_n$ be a sequence of projective measurement operators, where Λ_i corresponds to the two-outcome measurement $M_i = \{\Lambda_i, I - \Lambda_i\}$. Let ρ be a state such that either there exists i such that $\operatorname{tr}\Lambda_i\rho = \Omega(1)$, or $\mathbb{E}_j[\operatorname{tr}\Lambda_j\rho] = o(1/n)$. Then there is a test that uses one copy of ρ and: in the first case, accepts with probability $\Omega(1)$; in the second case, accepts with probability $o(1)$.*

Theorem 1 is restated more precisely below as Corollary 11. A variant of this result was called a "quantum OR bound" by Aaronson [3], but the proof there was incorrect because it neglected the effects of disturbance in the "accept" case. The term "OR bound" refers to the fact that we would like to design a measurement that accepts if any one of the M_i are likely to accept.

We give two procedures for Theorem 1. One procedure performs a crude form of eigenvalue estimation on the POVM operator $\frac{1}{n}\sum_j \Lambda_j$. In the first case ρ has constant overlap with the $+1$-eigenspace of Λ_i, and therefore has constant overlap with the space of eigenvectors of $\frac{1}{n}\sum_j \Lambda_j$ with eigenvalue at least $\Omega(1/n)$. We employ a variant of the Marriott-Watrous gap amplification procedure [25] to determine when this is the case. The second procedure is based on repeatedly selecting a random measurement from the sequence to perform, but also incorporating a "disturbance test", which is performed with low probability. This allows us to detect whether the residual state of the system after some measurements is far from its initial state.

These procedures turn out to have a number of applications, which we now describe.

[1] One could also think of this effect as the standard Zeno effect in a rotating reference frame.

1.1 Property testing

The field of property testing aims to find super-efficient algorithms for determining whether a function has a certain property, given the promise that it either has the property, or is far from having the property. The goal is to achieve this with the minimal number of queries to the function. A number of properties are known where quantum testers outperform classical testers – sometimes exponentially or super-exponentially. See [27] for a review.

Here we present query-efficient quantum testers for properties which can be expressed as isomorphism of functions under the action of a permutation group. Let G be a nontrivial permutation group acting on a finite set X, and consider two functions $f, g : X \to Y$ for some finite set Y. We say that f and g are isomorphic if there exists $\sigma \in G$ such that $g(x) = f(\sigma(x))$ for all $x \in X$, and for conciseness sometimes write $g = f \circ \sigma$. On the other hand, we say that f and g are ϵ-far from isomorphic if, for all $\sigma \in G$, $|\{x \in X : g(x) \neq f(\sigma(x))\}| \geq \epsilon |X|$. We say that an algorithm is an ϵ-tester for the property of G-isomorphism if it distinguishes between these two cases with success probability at least $2/3$. This success probability can be improved to $1 - \delta$, for any $\delta > 0$, by running the algorithm $O(\log 1/\delta)$ times and taking the majority vote. Below we discuss concrete examples of this abstract problem.

The problem of testing G-isomorphism was studied by Babai and Chakraborty [6], who considered both the case where g is known in advance (which they call the "query-1" case), and the case where both f and g are unknown (the "query-2" case). Here we only consider the case where f and g are both unknown. For this case, Babai and Chakraborty gave an almost tight classical bound of $\widetilde{\Theta}(\sqrt{|X| \log |G|})$ queries for testers with 1-sided error in the case where the action of G is primitive, i.e. does not preserve a nontrivial partition of X, and a classical upper bound of $O(\sqrt{|X| \log |G|})$ queries which holds for all permutation groups G and also has only 1-sided error. The best general lower bound they give for classical testers with 2-sided error is substantially weaker: $\Omega(\epsilon \log n)$ queries for any transitive group G of order $2^{O(n^{1-\epsilon})}$. However, substantially stronger bounds can be shown for specific group actions, as discussed below.

Here we prove the following result:

Theorem 2. *For any set of permutations G, there is a quantum ϵ-tester for G-isomorphism which makes $O((\log |G|)/\epsilon)$ queries.*

Observe that there is no dependence on $|X|$, unlike the best known classical upper bounds; and that going slightly beyond the usual definition of G-isomorphism [6], Theorem 2 holds when G is an arbitrary subset of a permutation group, rather than needing to be itself a group. Theorem 2 encompasses a number of special cases:

- **Isomorphism of boolean functions.** Two boolean functions $f, g : \{0,1\}^n \to \{0,1\}$ are said to be isomorphic if there exists $\sigma \in S_n$ such that $g(x) = f(\sigma(x))$ for all $x \in \{0,1\}^n$, where $\sigma(x)_i = x_{\sigma(i)}$. It is known that testing isomorphism of two unknown boolean functions up to constant accuracy requires $\Omega(2^{n/2}/n^{1/4})$ classical queries, with a nearly-matching upper bound of $O(2^{n/2}\sqrt{n \log n})$ queries [5]. Theorem 2 implies (taking $G = S_n$, $X = \{0,1\}^n$) that the quantum query complexity of ϵ-testing isomorphism of boolean functions is $O((n \log n)/\epsilon)$, which is exponentially smaller.

- **Graph isomorphism.** Two graphs G and H on n vertices are said to be isomorphic if there exists a bijection σ between the vertices such that vertices u and v are adjacent in G if and only if $\sigma(u)$ and $\sigma(v)$ are adjacent in H. The best classical query complexity known

for testing graph isomorphism is $\widetilde{O}(n^{5/4})$ [14], though the current best classical lower bound is only $\Omega(n)$ [14]. A quantum algorithm for testing graph isomorphism which makes $\widetilde{O}(n^{7/6})$ queries was given by Chakraborty et al. [11], as well as a quantum lower bound of $\Omega(n^{1/3})$ queries. Theorem 2 implies (taking $G = S_n$, $X = [n] \times [n]$ and $Y = \{0,1\}$) that the quantum query complexity of ϵ-testing graph isomorphism is $O((n \log n)/\epsilon)$.

- **Linear and affine isomorphism of boolean functions.** Two boolean functions $f, g : \{0,1\}^n \to \{0,1\}$ are said to be linear-isomorphic if there exists a non-singular linear transformation $A \in GL_n(\mathbb{F}_2)$ such that $g(x) = f(Ax)$ for all $x \in \{0,1\}^n$. Testing linear isomorphism was studied by Grigorescu, Wimmer and Xie [18], who gave an $\Omega(n^2)$ lower bound for testing linear isormorphism to a fixed function. It is not difficult to prove a substantially stronger $\Omega(2^{n/2})$ classical lower bound for testing linear isomorphism between two unknown functions; the argument is closely related to the lower bound for Simon's problem (qv) and we include it in Appendix B. Theorem 2 implies (taking $G = GL_n(\mathbb{F}_2)$, $X = \{0,1\}^n$) that the quantum query complexity of ϵ-testing linear isomorphism of boolean functions is $O(n^2/\epsilon)$, which is exponentially smaller.

  A similar $O(n^2/\epsilon)$ quantum upper bound holds for testing affine isomorphism, i.e. the existence of a pair $A \in GL_n(\mathbb{F}_2)$, $b \in \mathbb{F}_2^n$ such that $g(x) = f(Ax + b)$ for all x. Here an $\Omega(2^{n/2})$ classical lower bound is immediate from the known $\Omega(2^{n/2})$ lower bound on the classical query complexity of the property-testing variant of Simon's problem [10].

- **Hidden subgroup problems.** Imagine we have access to a function $f : G \to Y$ for some group G and some set Y, and are promised that either f is constant on cosets of a nontrivial subgroup $H \leq G$, or f is far from any such function. This is a property-testing version of the well-studied hidden subgroup problem [24] for the group G. Take X to be G, on which G acts in the natural way. Then the pair of functions (f, f) are $(G \setminus \{e\})$-isomorphic if and only if f is constant on cosets of a nontrivial subgroup $H \leq G$. Theorem 2 implies that the query complexity of this problem is $O((\log |G|)/\epsilon)$. A query-efficient tester for this property with the same complexity parameters was previously described by Friedl et al. [15] in the case where H is promised to be a *normal* subgroup of G; in addition, their tester is time-efficient when G is abelian. Note that the promise here is a bit different to the standard HSP. There, we are promised that f is constant on cosets of a subgroup $H \leq G$, and distinct on each coset. A quantum tester for this variant where $G = \mathbb{Z}_N$, corresponding to a periodicity determination problem, was given in [11]. In this case, the query complexity can be reduced to $O(1)$ if the additional constraint is imposed that the period of f is $O(\sqrt{N})$.

  It is known that the quantum query complexity of the property-testing version of Simon's problem, a special case of the hidden subgroup problem for $G = \mathbb{Z}_2^n$, obeys an $\Omega(\log |G|)$ lower bound [23] (see also the discussion in [27]). This implies that the dependence on $|G|$ of Theorem 2 is optimal.

Following the completion of this work, Alexander Belov (personal communication) has notified us that Theorem 2 can also be proven by a direct construction of a solution to the quantum adversary bound semidefinite program.

1.2 Testing quantum properties

We can also apply the ideas underlying our tester to testing properties of quantum states and operations. The notion of testing properties of classical data can naturally be extended to pure

quantum states (unit vectors in $\mathbb{C}^d$) as follows. Let $\mathcal{P}$ be a subset of $\mathbb{C}^d$. We are given the ability to produce copies of an initially unknown state $|\psi\rangle$, and asked to distinguish between the following two cases with success probability 2/3: either $|\psi\rangle \in \mathcal{P}$, or $\inf_{|\phi\rangle\in\mathcal{P}} \|\psi - \phi\|_{\mathrm{tr}} \geq \epsilon$. Here $\|\cdot\|_{\mathrm{tr}}$ is the trace distance,

$$\||\psi\rangle\langle\psi| - |\phi\rangle\langle\phi|\|_{\mathrm{tr}} = \frac{1}{2}\,\mathrm{tr}\,||\psi\rangle\langle\psi| - |\phi\rangle\langle\phi|| = \sqrt{1 - |\langle\psi|\phi\rangle|^2}, \tag{1}$$

and we use the notation $\psi = |\psi\rangle\langle\psi|$. We say that an algorithm which achieves these bounds is a quantum ϵ-tester for $\mathcal{P}$. Some interesting properties of quantum states (such as productness and permutation-invariance) can be tested in this framework [27].

Wang [31] has given a tester which tests membership of a state $|\psi\rangle$ in an *arbitrary* finite subset $\mathcal{P}$ of quantum states using $O(\log|\mathcal{P}|)$ copies of $|\psi\rangle$, assuming that all the states in $\mathcal{P}$ are far apart from each other. Formally, Wang's result is as follows:

Theorem 3 (Wang [31])**.** *Let $\mathcal{P}$ be a finite subset of the unit sphere in $\mathbb{C}^d$ such that $\min_{|\phi\rangle,|\phi'\rangle\in\mathcal{P}} \|\phi - \phi'\|_{\mathrm{tr}} = \zeta$. Then, for any $\epsilon > 0$, there is a test which accepts every state in $\mathcal{P}$ with certainty, rejects every state $|\psi\rangle$ such that $\min_{|\phi\rangle\in\mathcal{P}} \|\psi - \phi\|_{\mathrm{tr}} \geq \epsilon$ with probability at least 2/3, and uses $O((\log|\mathcal{P}|)\max\{\epsilon^{-2}, \zeta^{-2}\})$ copies of the input state.*

Wang applied this result to testing finite subsets of the unitary group, and in particular to testing permutations [31]. Here we can improve Theorem 3 by removing the dependence on the minimum distance ζ, at the expense of introducing two-sided error:

Theorem 4. *Let $\mathcal{P}$ be a finite subset of the unit sphere in $\mathbb{C}^d$. Then, for any $\epsilon > 0$, there is a quantum ϵ-tester for $\mathcal{P}$ which uses $O((\log|\mathcal{P}|)/\epsilon^2)$ copies of the input state.*

This resolves Question 6 in [27]. We remark that a similar bound to Wang's (albeit with two-sided error) can be obtained simply by reducing the problem of testing membership of $|\psi\rangle$ in $\mathcal{P}$ to *identifying* $|\psi\rangle$, given the promise that it is contained in $\mathcal{P}$ [21]. Theorem 4 goes beyond this idea, and efficiently tests membership in sets of states $\mathcal{P}$ for which identifying a member of $\mathcal{P}$ is more challenging, because there exist states in $\mathcal{P}$ which are close to each other.

The ability to test properties of states naturally extends to testing properties of operations, i.e. unitary operators or quantum circuits [31, 27]. Indeed, using a standard correspondence between states and operations known as the Choi-Jamiołkowski isomorphism, any tester for a property of quantum states gives a tester for a corresponding property of unitaries. Now the relevant distance measure can be defined in terms of the Hilbert-Schmidt inner product $\langle A, B\rangle = \mathrm{tr}\,A^\dagger B/d$ for d-dimensional operators A, B, as $D(A,B) = \sqrt{1 - |\langle A, B\rangle|^2}$ [27] (compare (1)). This is an average-case measure of distance, as for unitary operators U, V,

$$D(U,V)^2 \approx \int \|U|\psi\rangle\langle\psi|U^\dagger - V|\psi\rangle\langle\psi|V^\dagger\|_{\mathrm{tr}}^2\, d\psi,$$

where the integral is taken over the Haar (uniform) measure on the set of pure quantum states $|\psi\rangle$ [27]; it is also closely related to the distance between A and B in the Schatten 2-norm [27]. Let $\mathcal{P}$ be a subset of $U(d)$, the set of d-dimensional unitary operators. We are given access to an initially unknown operator $U \in U(d)$, and asked to distinguish between the following two cases with success probability 2/3: either $U \in \mathcal{P}$, or $\inf_{V\in\mathcal{P}} D(U,V) \geq \epsilon$. As in the case of properties of quantum states, we say that an algorithm which achieves these bounds is a quantum ϵ-tester for $\mathcal{P}$.

Using Theorem 4, we obtain the following corollary:

Corollary 5. *Let $\mathcal{P}$ be a finite subset of $U(d)$. Then, for any $\epsilon > 0$, there is a quantum ϵ-tester for $\mathcal{P}$ which makes $O((\log |\mathcal{P}|)/\epsilon^2)$ uses of the input unitary operator.*

Another consequence of these ideas is an efficient algorithm for a generalisation of the notion of G-isomorphism of functions to isomorphism of unitary operators[2]. Imagine we have a known set $S = U_1, \ldots, U_n$ of unitary operators, and two unknown unitary operators V and W. We would like to determine whether there exists $U \in S$ such that $UVU^\dagger = W$. More precisely, our task is to distinguish between two cases: either there exists $U \in S$ such that $UVU^\dagger = W$, or for all $U \in S$, $D(UVU^\dagger, W) \geq \epsilon$. We call this property unitary S-isomorphism. For example, V and W could be quantum circuits on n qubits, S could be the set of all permutations of n qubits, and we might like to determine whether there exists $\sigma \in S$ such that $\sigma V \sigma^{-1} = W$. This is one natural quantum generalisation of the property of isomorphism of boolean functions.

Theorem 6. *For any S, there is a quantum ϵ-tester for unitary S-isomorphism which makes $O((\log |S|)/\epsilon^2)$ uses of the input unitaries V and W.*

1.3 Testing genuine multipartite entanglement

One important property of quantum states is that of being entangled. A pure state $|\psi\rangle$ on n qudits is said to be product (resp. entangled) across the cut $S : S^c$ (for S a nontrivial subset of $[n]$) if $|\psi\rangle = |\alpha\rangle_S \otimes |\beta\rangle_{S^c}$ for some states $|\alpha\rangle, |\beta\rangle$ (resp. $|\psi\rangle$ cannot be written in this way). If such an S exists then we say that $|\psi\rangle$ is product across some cut. If no such S exists then we say that $|\psi\rangle$ has "genuine multipartite entanglement" or "genuine n-partite entanglement"; this term reflects the fact that an n-partite state might be describable entirely in terms of entangled states on smaller subsets of systems.

In [19, 20] two of us analyzed a property tester for a related question, of whether $|\psi\rangle$ is equal to an n-partite product state $|\alpha_1\rangle \otimes \cdots \otimes |\alpha_n\rangle$ or far from any such state. This property tester – known as the "product test" – can also be used to test whether $|\psi\rangle$ is entangled across any fixed cut $S : S^c$. With Theorem 1 we can extend this to test whether entanglement exists across all cuts simultaneously.

Theorem 7. *There is a quantum ϵ-tester for the property of an n-partite state being product across some cut. The tester uses $O(n/\epsilon^2)$ copies of the state.*

We note that this result does not follow from Theorem 4 because the set of product states is infinite, and even defining an ϵ-net over the set of product states would cause the sample complexity to depend on the dimensions of the subsystems.

1.4 De-Merlinizing quantum protocols

Our final application, following work of Aaronson [3], is to a problem in quantum communication complexity [8]. Let $f : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$ be a (possibly partial) boolean function. We imagine that the input to f is divided into two parts, the first of which is given to Alice, the second to Bob. Their goal is to output $f(x, y)$ for given x and y with the minimum amount of communication. Let $Q^1(f)$ denote the bounded-error one-way quantum communication complexity of f: the minimal number of qubits Alice needs to send to Bob in order to achieve worst-case success probability of 2/3.

[2] We would like to thank Noah Linden for suggesting this application.

A more general model [3] is to allow Alice and Bob to be assisted by Merlin, who provides a quantum witness $|\phi\rangle$ to Bob, and aims to convince him that $f(x,y) = 1$. Let $|\psi_x\rangle$ be the state which Alice sends to Bob on input x. We demand that:

- If $f(x,y) = 1$, then there exists $|\phi\rangle$ such that Bob accepts the triple $(y, |\psi_x\rangle, |\phi\rangle)$ with probability at least $2/3$;
- If $f(x,y) = 0$, then for all $|\phi\rangle$, Bob accepts $(y, |\psi_x\rangle, |\phi\rangle)$ with probability at most $1/3$.

Let $\mathrm{QMA}^1_w(f)$ denote the minimal a such that there exists a protocol of this form where Alice sends a qubits to Bob, and Merlin sends w qubits to Bob. Then it was claimed in [3] that:

Theorem 8. *For all partial or total boolean functions f, and all $w \geq 2$,*

$$Q^1(f) = O(\mathrm{QMA}^1_w(f) \cdot w \log^2 w).$$

Several applications of this result were also given in [3], to random access coding, one-way communication and complexity theory. The basic idea behind the proof of Theorem 8 is that first Alice and Bob's protocol can be amplified such that the soundness error (the probability of incorrectly accepting) can be made exponentially small in w, while keeping the completeness error (the probability of incorrectly rejecting) at most $1/3$. Then Bob can replace his use of the witness $|\phi\rangle$ from Merlin with simply looping over all possible computational basis states of Merlin's witness register and trying each such state as a witness in turn. In a "yes" case ($f(x,y) = 1$), at least one of these will have a sufficiently high probability of acceptance that this can be distinguished from a "no" case. This procedure requires $O(2^w)$ measurements to be applied to $|\psi_x\rangle$ (actually slightly more because of the amplification step).

Unfortunately, the proof of a key lemma required for Theorem 8 in [3] (Lemma 14, the "Quantum OR bound") does not appear to be correct. One step of the proof claims the following: given a sequence of t 2-outcome measurements performed on an initial state ρ to leave a final state ρ_t, if $\|\rho_t - \rho\|_{\mathrm{tr}} > \sqrt{\alpha}$ for some α, then the probability that at least one measurement yields outcome 1 is at least α. However, this is false, as shown by the "quantum anti-Zeno" sequence of measurements discussed in the Introduction.

Here we can use our testing procedures to deal with this issue and give a corrected proof of Theorem 8.

The corrected Quantum OR bound from [3] could also be useful for other applications in quantum information. For example, it can be used to give a proof that the security of certain proposed quantum money schemes [7] must rest on computational, rather than information-theoretic, assumptions [1] (Scott Aaronson, personal communication).

Remark: It is only the proof in [3] that is incorrect, but the protocol there may well work. This is because the measurements are performed in a random order and we do not know of a set of measurements which has the anti-Zeno property for most choices of ordering.

1.5 Disturbance and G-isomorphism

To provide some intuition for our results, we now outline how applying sequential measurements to quantum states connects to testing G-isomorphism. The G-isomorphism tester is based on a simple idea: testing whether $g = f \circ \sigma$ for each permutation $\sigma \in G$. (A drawback of this technique is that the testers produced are not time-efficient, but just query-efficient.) We construct k copies

of a state $|\psi\rangle$ corresponding to evaluating f and g on every possible input in superposition, for some k to be determined. Given one copy of $|\psi\rangle$, for any σ we can distinguish between the cases that $g = f \circ \sigma$, and g is far from $f \circ \sigma$, with success probability lower bounded by a constant. Given k copies of $|\psi\rangle$, we can distinguish between these two cases with exponentially small probability of failure in k. We would like to reuse the state $|\psi\rangle^{\otimes k}$ for each test, to avoid making any further queries.

In the "far from isomorphic" case, as the probability that the measurement incorrectly says "isomorphic" is exponentially small in k, $|\psi\rangle^{\otimes k}$ will only be disturbed by an exponentially small amount [35, 28, 3]. However, in the "isomorphic" case, it could be the case that, as well as having $g = f \circ \sigma$ for some $\sigma \in G$, we have $g \approx f \circ \tau$ for some $\tau \neq \sigma$. If the test for τ has a fairly large probability of success, and yet still outputs "not isomorphic", the resulting state may be substantially disturbed, implying that the future test for isomorphism under σ may incorrectly reject. One way to address this problem is to introduce an additional test for disturbance of the state $|\psi\rangle^{\otimes k}$ (see Appendix A): if the state is substantially disturbed at any point during the algorithm, we know that the answer should be "isomorphic". Alternatively, we could repeatedly perform a random test and attempt to reset the system to its initial state after each measurement, similarly to Marriott-Watrous gap amplification (see Section 2). In either case it turns out to be sufficient to take $k = O((\log |G|)/\epsilon)$ to distinguish between the cases that f and g are G-isomorphic, or ϵ-far from G-isomorphic.

The strategy of testing each σ in turn is similar to the technique used by Ettinger, Høyer and Knill [12] to give a polynomial-query quantum algorithm for the nonabelian hidden subgroup problem (HSP) by testing each subgroup in turn. However, the algorithm of [12] did not have the issue with disturbing the state that we need to address here. This was a consequence of the hidden subgroup promise in the standard HSP. In the HSP, we are given access to a function f which is promised to be constant on cosets of some subgroup H, and distinct on each coset. The second part of this promise implies that, if we test f for being constant on cosets of any subgroup $H' \neq H$, the test is likely to fail, so the state is left almost undisturbed. In the property-testing scenarios we consider here, we do not have this promise.

It is well known that testing graph isomorphism of rigid graphs (without a promise that "no" instances are far from isomorphic) reduces to the HSP for the symmetric group. Our results on G-isomorphism, however, do not use a reduction to the property-testing variant of the HSP for G. The standard reduction would produce a function $F(\sigma) = f \circ \sigma$, where $\sigma \in G$ and $f \circ \sigma$ represents the table of all values $f(\sigma(x))$; so evaluating F for any given σ requires $|X|$ queries. In the case of boolean function isomorphism, for example, evaluating F would require 2^n queries, and would destroy any exponential speedup.

1.6 Organisation

We present two procedures for determining whether one of a sequence of n measurements accepts a state with high probability: one based on Marriott-Watrous gap amplification, and one based on testing disturbance. The procedures both have similar parameters, and either of them would suffice to prove our main results. The procedure based on testing disturbance[3] has somewhat worse constants and a less elegant proof of correctness, so we relegate a description of this to Appendix A and focus on the modified Marriott-Watrous procedure, which we now describe.

[3]The first version of this paper only included this procedure.

1. Create the state $\rho \otimes |0\rangle\langle 0|^{\otimes m}$.
2. Repeat N times or until the algorithm accepts:
   (a) Perform the projective measurement $\{\Pi, I - \Pi\}$. If the first result is returned, accept.
   (b) Perform the projective measurement $\{\Delta, I - \Delta\}$. If the second result is returned, accept.
3. Reject.

Algorithm 1: Modified Marriott-Watrous gap-amplification procedure

2 Modified Marriott-Watrous gap amplification

Our procedure will be based on a subroutine which, roughly speaking, performs eigenvalue estimation on a POVM element Λ applied to one copy of ρ. More precisely, we can use this subroutine to determine if ρ has high support on the space of eigenvectors with large eigenvalues.

Assume that we have one copy of some quantum state ρ and a 2-outcome POVM $\{\Lambda, I - \Lambda\}$, where we are given an explicit decomposition $\Lambda \otimes |0\rangle\langle 0|^{\otimes m} = \Delta\Pi\Delta$, where m is some integer, $\Delta = I \otimes |0\rangle\langle 0|^{\otimes m}$, and Π is an orthogonal projector. This is the case in most applications: for example if we are given an explicit circuit description of a measurement corresponding to Λ that consists of appending m ancilla qubits in the state $|0\rangle$, applying a unitary U, and then making a projective measurement P, then the success probability on input ψ is

$$\langle\psi|\langle 0|^{\otimes m} U^\dagger P U |\psi\rangle|0\rangle^{\otimes m} = \langle\psi|\langle 0|^{\otimes m} \Delta\Pi\Delta |\psi\rangle|0\rangle^{\otimes m}$$

where we take $\Pi = U^\dagger P U$. In this case the projector Δ serves to check that the ancilla qubits are initialized to $|0\rangle$ properly; if they are, the success probability should be equal to $\langle\psi|\Lambda|\psi\rangle$, and therefore $\Lambda \otimes |0\rangle\langle 0|^{\otimes m} = \Delta\Pi\Delta$ as desired.

Alternatively, if the Naimark extension of Λ is given then we also have such a decomposition for Λ with $m = 1$; or if Λ is a projector then we can take a trivial decomposition with $m = 0$.

One special case that will be important for us is when we have a sequence of projectors Λ_i, where Λ_i corresponds to the two-outcome measurement $M_i = \{\Lambda_i, I - \Lambda_i\}$, and we would like to implement the POVM element $\Lambda := \frac{1}{n}\sum_{j=1}^{n} \Lambda_j$. Define the projector $\Pi = \sum_{i=0}^{n-1} \Lambda_{i+1} \otimes (Q|i\rangle\langle i|Q^{-1})$, where Q is the quantum Fourier transform on $\mathbb{Z}_n$, and also define $\Delta = I \otimes |0\rangle\langle 0|$. Given quantum circuits which implement each measurement M_i, it is easy to write down a circuit that implements the measurement $\{\Pi, I - \Pi\}$. Then

$$\Delta\Pi\Delta = \left(\frac{1}{n}\sum_{j=1}^{n} \Lambda_j\right) \otimes |0\rangle\langle 0| = \Lambda \otimes |0\rangle\langle 0|,$$

so this gives us an implementation of Λ as desired.

Our procedure is described as Algorithm 1 above. The algorithm is based on the Marriott-Watrous procedure for in-place amplification for QMA [25], but we give a simplified procedure that is similar to the "OR-type repetition procedure" in [13], but with different analysis.

Theorem 9. *Let $p_{\mathrm{acc}}(N)$ be the acceptance probability of Algorithm 1 when applied to the measurement operator Λ and state ρ, and write $P_{\geq\delta}$ for the projector onto* $\mathrm{span}\{|\phi\rangle : \Lambda|\phi\rangle = \lambda|\phi\rangle, \lambda \geq \delta\}$. *Then*

$$(1 - e^{-1}) \operatorname{tr} P_{\geq \frac{1}{2N}} \rho \;\leq\; p_{\mathrm{acc}}(N) \;\leq\; 2N \operatorname{tr} \Lambda\rho.$$

Proof. First assume ρ is pure, $\rho = |\psi\rangle\langle\psi|$. Decompose $|\psi\rangle$ into an eigenbasis of Λ:

$$|\psi\rangle = \sum_i \alpha_i |\psi_i\rangle,$$

where $\Lambda|\psi_i\rangle = \lambda_i|\psi_i\rangle$ and the ψ_i's are normalized states. Appending the ancilla qubits, we have

$$|\psi\rangle \otimes |0\rangle^{\otimes m} = \sum_i \alpha_i |\widetilde{\psi}_i\rangle$$

where the states $|\widetilde{\psi}_i\rangle := |\psi_i\rangle \otimes |0\rangle^{\otimes m}$ are eigenvectors of $\Delta\Pi\Delta$, and moreover $\Delta|\widetilde{\psi}_i\rangle = |\widetilde{\psi}_i\rangle$. Note that

$$\Delta(I - \Pi)|\widetilde{\psi}_i\rangle = \Delta|\widetilde{\psi}_i\rangle - \Delta\Pi\Delta|\widetilde{\psi}_i\rangle = (1 - \lambda_i)|\widetilde{\psi}_i\rangle.$$

Therefore if Algorithm 1 does not accept in Step 2, the residual unnormalized state is

$$(\Delta(I - \Pi))^N |\psi\rangle \otimes |0\rangle^{\otimes m} = \sum_i \alpha_i (1 - \lambda_i)^N |\widetilde{\psi}_i\rangle$$

and the probability that Algorithm 1 accepts (doesn't reach Step 3) is

$$p_{\mathrm{acc}}(N) = 1 - \sum_i |\alpha_i|^2 (1 - \lambda_i)^{2N} = \sum_i |\alpha_i|^2 \left[1 - (1 - \lambda_i)^{2N}\right].$$

It is now easy to derive the lower and upper bounds claimed in the theorem. First,

$$p_{\mathrm{acc}}(N) \geq \sum_{i:\lambda_i \geq 1/(2N)} |\alpha_i|^2 \left[1 - (1 - \lambda_i)^{2N}\right] > (1 - e^{-1}) \operatorname{tr} P_{\geq \frac{1}{2N}} \psi,$$

where we used $(1 - a)^{1/a} < e^{-1}$ for any $a > 0$. Second,

$$p_{\mathrm{acc}}(N) \leq \sum_i |\alpha_i|^2 (2N\lambda_i) = 2N \operatorname{tr} \Lambda\psi,$$

where we used $1 - (1 - \lambda_i)^{2N} \leq 2N\lambda_i$. Finally, if ρ is mixed, the claim follows from convexity. This completes the proof. □

We now specialise Theorem 9 to cases of interest for applications.

3 Property testing

The first setting in which we would like to use Theorem 9 is where we have a sequence of measurements, and would like to determine whether one of them accepts with high probability. We will need the following lemma:

Lemma 10 (Gentle measurement / "almost as good as new" lemma [35, 28, 2, 3])**.** *Let ρ be a quantum state and let $0 \leq \Lambda \leq I$ be a measurement operator. Then*

$$\left\| \rho - \frac{\sqrt{\Lambda}\rho\sqrt{\Lambda}}{\operatorname{tr}\Lambda\rho} \right\|_{\mathrm{tr}} \leq \sqrt{\operatorname{tr}(I-\Lambda)\rho}.$$

We state this lemma somewhat differently to some previous works; a proof of this version can be found in [33, Lemma 9.4.1]. For completeness, we also provide a concise proof here.

Proof. Let $F(\rho,\sigma) := \operatorname{tr}\sqrt{\rho^{1/2}\sigma\rho^{1/2}}$ be the fidelity between quantum states ρ and σ. We have

$$F\left(\rho, \frac{\sqrt{\Lambda}\rho\sqrt{\Lambda}}{\operatorname{tr}\Lambda\rho}\right) = \frac{\operatorname{tr}\sqrt{\sqrt{\rho}\sqrt{\Lambda}\rho\sqrt{\Lambda}\sqrt{\rho}}}{\sqrt{\operatorname{tr}\Lambda\rho}} = \frac{\operatorname{tr}\sqrt{\rho}\sqrt{\Lambda}\sqrt{\rho}}{\sqrt{\operatorname{tr}\Lambda\rho}} = \frac{\operatorname{tr}\sqrt{\Lambda}\rho}{\sqrt{\operatorname{tr}\Lambda\rho}} \geq \frac{\operatorname{tr}\Lambda\rho}{\sqrt{\operatorname{tr}\Lambda\rho}} = \sqrt{\operatorname{tr}\Lambda\rho},$$

where the second equality follows because $\sqrt{\rho}\sqrt{\Lambda}\sqrt{\rho}$ is positive semidefinite, the third is cyclicity of the trace, and the inequality is $\sqrt{\Lambda} \geq \Lambda$ for $0 \leq \Lambda \leq I$. Using the inequality $\|\rho-\sigma\|_{\mathrm{tr}} \leq \sqrt{1-F(\rho,\sigma)^2}$, we obtain the claimed result. □

Corollary 11. *Let $\Lambda_1, \ldots, \Lambda_n$ be a sequence of projectors, where Λ_i corresponds to the two-outcome measurement $M_i = \{\Lambda_i, I-\Lambda_i\}$, and fix parameters $\epsilon \leq 1/2$, δ. Let ρ be a state such that either there exists i such that $\operatorname{tr}\Lambda_i\rho \geq 1-\epsilon$ (case 1), or $\mathbb{E}_j[\operatorname{tr}\Lambda_j\rho] \leq \delta$ (case 2). Then there is a test that uses one copy of ρ and: in case 1, accepts with probability at least $(1-\epsilon)^2/7$; in case 2, accepts with probability at most $4\delta n$.*

Proof. We apply Algorithm 1 to $\Lambda = \frac{1}{n}\sum_j \Lambda_j$ and ρ, taking $N = \lceil n/(1-\epsilon)\rceil$. We first consider case 1. Let Q denote the projector onto $\operatorname{span}\{|\phi\rangle : \Lambda|\phi\rangle = \lambda|\phi\rangle, \lambda \geq 1/(2N)\}$, and write $Q^\perp = I - Q$. To apply Theorem 9, we need to lower-bound $\operatorname{tr} Q\rho$. By Lemma 10,

$$\operatorname{tr} Q\rho \geq \left\| \rho - \frac{Q^\perp\rho Q^\perp}{\operatorname{tr} Q^\perp\rho} \right\|_{\mathrm{tr}}^2.$$

Write $\sigma = Q^\perp\rho Q^\perp / \operatorname{tr} Q^\perp\rho$. Then, as we are in case 1, $\operatorname{tr}\Lambda_i\rho \geq 1-\epsilon$ and hence

$$\frac{1-\epsilon}{n} \leq \frac{\operatorname{tr}\Lambda_i\rho}{n} = \frac{\operatorname{tr}\Lambda_i(\rho-\sigma) + \operatorname{tr}\Lambda_i\sigma}{n} \leq \frac{\|\rho-\sigma\|_{\mathrm{tr}}}{n} + \frac{\operatorname{tr}\Lambda_i\sigma}{n} \leq \frac{\|\rho-\sigma\|_{\mathrm{tr}}}{n} + \operatorname{tr}\Lambda\sigma.$$

We have $\operatorname{tr}\Lambda\sigma < 1/(2N) \leq (1-\epsilon)/(2n)$ because σ is only supported on eigenvectors of Λ with eigenvalues less than $1/(2N)$. Rearranging, we obtain that

$$\|\rho-\sigma\|_{\mathrm{tr}} \geq \frac{1-\epsilon}{2},$$

so $\operatorname{tr} Q\rho \geq (1-\epsilon)^2/4$ and hence Algorithm 1 accepts with probability at least $(1-e^{-1})(1-\epsilon)^2/4 \geq (1-\epsilon)^2/7$. In case 2, $\operatorname{tr}\Lambda\rho \leq \delta$ by assumption. So by Theorem 9, Algorithm 1 accepts with probability at most $2\delta\lceil n/(1-\epsilon)\rceil \leq 4\delta n$, using $\epsilon \leq 1/2$. □

3.1 Application to testing G-isomorphism

We now apply Corollary 11 to testing G-isomorphism. In fact, we reduce this problem to the following more general question: testing whether a state is an eigenvector of some unitary operator picked from a known set of unitary operators.

Lemma 12. *Assume that we have access to a sequence of controlled unitaries $U_1, \ldots, U_n$, and their inverses, and the ability to produce copies of a state $|\psi\rangle$. We are promised that either there exists i such that $U_i|\psi\rangle = |\psi\rangle$, or for all i, $|\langle\psi|U_i|\psi\rangle| \leq 1 - \epsilon$. There is an algorithm which distinguishes between these two cases using $O((\log n)/\epsilon)$ copies of $|\psi\rangle$ and succeeds with probability at least $2/3$.*

Proof. Write $|\phi\rangle := (\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)|\psi\rangle)^{\otimes k}|0\rangle$ for some k and, for each i, define the two-outcome projective measurement M_i by the following sequence of steps:

1. Apply controlled-U_i on each of the first k registers (controlled on the first qubit within each).
2. Apply a Hadamard gate to each first qubit in each of the first k registers.
3. Apply a controlled-X gate to the last qubit, controlled on all of these k qubits being in the 0 state.
4. Measure the last qubit. Associate outcome 1 with acceptance, 0 with rejection.
5. Invert steps 1-3.

The state after the second step is $(\frac{1}{2}(|0\rangle(I + U_i)|\psi\rangle + |1\rangle(I - U_i)|\psi\rangle))^{\otimes k}|0\rangle$, so the probability of acceptance is

$$\left(\frac{\|(I + U_i)|\psi\rangle\|^2}{4}\right)^k = \left(\frac{1}{2} + \frac{1}{2}\mathrm{Re}\langle\psi|U_i|\psi\rangle\right)^k.$$

If $U_i|\psi\rangle = |\psi\rangle$, the measurement accepts with certainty. If $|\langle\psi|U_i|\psi\rangle| \leq 1 - \epsilon$, the measurement accepts with probability at most $(1 - \epsilon/2)^k \leq e^{-\epsilon k/2}$. It is sufficient to take $k = O((\log n)/\epsilon)$ to obtain an acceptance probability in this case which is at most c/n for an arbitrary constant $c > 0$. We can now apply Corollary 11. In the former case, we have that the test accepts with probability at least $1/7$. In the latter case, we can choose c such that, by Corollary 11, the test accepts with probability at most $1/8$. A constant number of repetitions suffices to distinguish between these two cases with probability at least $2/3$. □

Lemma 12 can be applied to testing G-isomorphism. Recall that in this problem we have a group G acting on a set X, and would like to distinguish between these two cases: a) there exists $\sigma \in G$ such that $g(x) = f(\sigma(x))$ for all $x \in X$; b) for all $\sigma \in G$, $|\{x \in X : g(x) \neq f(\sigma(x))\}| \geq \epsilon|X|$. Formally, the G-isomorphism problem is actually a special case of a quantum problem discussed below (unitary S-isomorphism); however, we state and prove it separately for clarity.

Theorem 2 (restated). *For any G, there is a quantum ϵ-tester for G-isomorphism which makes $O((\log |G|)/\epsilon)$ queries to f and g.*

Proof. Write $d(f, g) := |\{x \in X : f(x) \neq g(x)\}|/|X|$. For any function $f : X \to Y$, define the corresponding state

$$|f\rangle := \frac{1}{\sqrt{|X|}} \sum_{x \in X} |x\rangle|f(x)\rangle;$$

then $\langle f|g\rangle = 1 - d(f,g)$. Also define the unitary operator U_σ by $U_\sigma|x\rangle = |\sigma(x)\rangle$ for $\sigma \in G$. Then $U_\sigma|f\rangle = |f \circ \sigma\rangle$. Consider the state $|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle|f\rangle + |1\rangle|g\rangle)$ and the operator U'_σ which maps $|\psi\rangle$ to $\frac{1}{\sqrt{2}}(|0\rangle|g \circ \sigma^{-1}\rangle + |1\rangle|f \circ \sigma\rangle)$ for any f and g. U'_σ can easily be implemented using Pauli-X, controlled-U_σ and controlled-U_σ^{-1} operations. Then

$$\langle\psi|U'_\sigma|\psi\rangle = \langle f \circ \sigma|g\rangle = 1 - d(f \circ \sigma, g).$$

Applying Lemma 12 to the state $|\psi\rangle$ and the unitary operator U'_σ, we can distinguish between the case that there exists σ such that $g = f \circ \sigma$, and the case that $d(g, f \circ \sigma) \geq \epsilon$ for all $\sigma \in G$, with $O((\log|G|)/\epsilon)$ copies of $|\psi\rangle$. Each copy can be created with one query to f and g. This proves Theorem 2. □

3.2 Testing quantum states and operations

We next apply our results to testing properties of quantum states, and then properties of quantum operations. These are all quite straightforward corollaries of previous results in the paper.

Theorem 4 (restated). *Let $\mathcal{P}$ be a finite subset of the unit sphere in $\mathbb{C}^d$. Then, for any $\epsilon > 0$, there is a quantum ϵ-tester for $\mathcal{P}$ which uses $O((\log|\mathcal{P}|)/\epsilon^2)$ copies of the input state.*

Proof. Let k be an integer parameter to be determined. We apply Corollary 11 to the state $|\psi\rangle^{\otimes k}$ and the measurements $\Pi_{|\phi\rangle} = |\phi\rangle\langle\phi|^{\otimes k}$, $|\phi\rangle \in \mathcal{P}$. If $|\psi\rangle \in \mathcal{P}$, there exists $|\phi\rangle$ such that $\langle\psi|^{\otimes k}\Pi_{|\phi\rangle}|\psi\rangle^{\otimes k} = 1$. If $\min_{|\phi\rangle\in\mathcal{P}} \|\psi - \phi\|_{\mathrm{tr}} \geq \epsilon$, then for all $|\phi\rangle \in \mathcal{P}$,

$$\langle\psi|^{\otimes k}\Pi_{|\phi\rangle}|\psi\rangle^{\otimes k} = |\langle\psi|\phi\rangle|^{2k} = (1 - \|\psi - \phi\|_{\mathrm{tr}}^2)^k \leq (1 - \epsilon^2)^k.$$

In the former case, the algorithm of Corollary 11 accepts with probability at least 1/7. In the latter case, it accepts with probability at most $4|\mathcal{P}|(1-\epsilon^2)^k$, so it is sufficient to take $k = O((\log|\mathcal{P}|)/\epsilon^2)$ to make the acceptance probability at most 1/8. Taking the median of $O(1)$ runs is enough to distinguish these two cases except with probability at most 1/3. □

We now turn to testing properties of quantum operations. Lifting our results on testing quantum states to testing unitary operators on $\mathbb{C}^d$ is based on the following connection, known as the Choi-Jamiołkowski isomorphism. Let $|\Phi\rangle = \frac{1}{\sqrt{d}}\sum_{i=1}^d |i\rangle|i\rangle$, and for any $U \in U(d)$, define

$$|U\rangle = (U \otimes I)|\Phi\rangle = \frac{1}{\sqrt{d}}\sum_{i,j=1}^d U_{ij}|i\rangle|j\rangle.$$

Then it is easy to see that

$$\langle U, V\rangle = \langle U|V\rangle, \quad (A \otimes B)|V\rangle = |AVB^T\rangle. \tag{2}$$

As a consequence of the first equality, $D(U,V) := \sqrt{1 - |\langle U,V\rangle|^2} = \||U\rangle\langle U| - |V\rangle\langle V|\|_{\mathrm{tr}}$. A copy of $|U\rangle$ can be created with one use of U.

Corollary 5 (restated). *Let $\mathcal{P}$ be a finite subset of $U(d)$. Then, for any $\epsilon > 0$, there is a quantum ϵ-tester for $\mathcal{P}$ which makes $O((\log|\mathcal{P}|)/\epsilon^2)$ uses of the input unitary operator.*

Proof. Apply Theorem 4 to test membership of $|U\rangle$ in the set $\mathcal{P}' = \{|V\rangle : V \in \mathcal{P}\}$. The test uses $O((\log|\mathcal{P}|)/\epsilon^2)$ copies of $|U\rangle$, each of which can be constructed with one use of U. □

Next we consider the property of unitary S-isomorphism. Recall that in this problem we have a set $S = U_1, \ldots, U_n$ of unitary operators, and two unitary operators V and W. We would like to distinguish between two cases: either there exists $U \in S$ such that $UVU^\dagger = W$, or for all $U \in S$, $D(UVU^\dagger, W) \geq \epsilon$.

Theorem 6 (restated). *For any S and any $\epsilon > 0$, there is a quantum ϵ-tester for unitary S-isomorphism which makes $O((\log |S|)/\epsilon^2)$ uses of the input unitaries V and W.*

Proof. The argument is similar to the proof of Theorem 2. We can produce a copy of the state $|\psi\rangle = |V\rangle|W\rangle$ with a single use of each of V and W. Similarly, by (2), for any U we can implement an operation U' mapping $|\psi\rangle$ to $|\psi'\rangle = |U^\dagger W U\rangle|UVU^\dagger\rangle$ by applying $U \otimes U^*$ to the first register, and $U^\dagger \otimes U^T$ to the second register; and then swapping the two registers. If $UVU^\dagger = W$, then $|\psi'\rangle = |\psi\rangle$. On the other hand, if there exists $U \in S$ such that $D(UVU^\dagger, W) = \epsilon$, then

$$|\langle\psi|U'|\psi\rangle| = |\langle UVU^\dagger, W\rangle|^2 = 1 - \epsilon^2.$$

By Lemma 12, these two cases can be distinguished with $O((\log |S|)/\epsilon^2)$ uses of V and W. □

The apparently worse scaling with ϵ of this result compared with Theorem 2 is an artifact of the distance measure used being defined differently.

3.3 Testing genuinely multipartite entanglement

Our final quantum tester is for the property of not possessing genuine multipartite entanglement – i.e. a quantum state being product across some partition of the qubits into two parts.

Theorem 7 (restated). *There is a quantum ϵ-tester for the property of an n-partite state being product across some cut. The tester uses $O(n/\epsilon^2)$ copies of the state.*

Proof. Suppose we are given $|\psi\rangle^{\otimes k}$ for some n-partite state $|\psi\rangle$. For any fixed proper nonempty $S \subset [n]$ let $D(S)$ denote the minimum distance of $|\psi\rangle$ to a state of the form $|\alpha\rangle_S \otimes |\beta\rangle_{S^c}$. According to Lemma 20 of [20] (see also [32]), there is a test using two copies of $|\psi\rangle$ that will accept with probability $1 - \Theta(D(S)^2)$. In particular if $|\psi\rangle$ is product across $S : S^c$ then it will accept with certainty, and if it is ϵ-far from product across $S : S^c$ then it will accept with probability $1 - \Omega(\epsilon^2)$. Taking $k = O(n/\epsilon^2)$ we can reduce this acceptance probability to $2^{-\Omega(n)}$. Since there are $2^{n-1} - 1$ ways to partition $[n]$ into two pieces, our result follows from Corollary 11. □

4 De-Merlinizing quantum protocols

We finally apply Algorithm 1 to prove Theorem 8. The key technical ingredient can be stated as follows. Let Λ be a measurement operator corresponding to the 2-outcome POVM measurement $\{\Lambda, I - \Lambda\}$ on a bipartite Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$. Let $\{|1\rangle, \ldots, |d\rangle\}$ be an orthonormal basis for $\mathcal{H}_B$. For all $j \in \{1, \ldots, d\}$, let Λ_j be the measurement operator on $\mathcal{H}_A$ corresponding to the 2-outcome POVM $\{\Lambda_j, I - \Lambda_j\}$ induced by applying Λ to $\mathcal{H}_A \otimes |j\rangle$.

Fix a pure state $|\psi\rangle \in \mathcal{H}_A$. We would like to distinguish between the following two cases:

1. There exists σ in $\mathcal{H}_B$ such that Λ accepts with probability at least $\eta > 0$ when applied to $\psi \otimes \sigma$.

2. For all states σ in $\mathcal{H}_B$, Λ accepts with probability at most ζ when applied to $\psi \otimes \sigma$.

To do so, we apply Algorithm 1 to $|\psi\rangle$ and the operator $\Lambda = \frac{1}{d}\sum_{j=1}^{d} \Lambda_j$, taking $N = \lceil d/\eta \rceil$.

Corollary 13. *In case 1, Algorithm 1 accepts with probability at least $\eta^2/7$. In case 2, Algorithm 2 accepts with probability at most $2\zeta\lceil d/\eta \rceil$.*

Proof. The proof is similar to that of Corollary 11. We first observe that, for an arbitrary state ϕ,

$$\begin{aligned} \operatorname{tr} \Lambda\phi &= \mathbb{E}_j \left[\operatorname{tr} \Lambda(\phi \otimes |j\rangle\langle j|)\right] = \operatorname{tr} \Lambda \left(\phi \otimes \frac{I}{d}\right) \geq \frac{\operatorname{tr} \Lambda(\phi \otimes \sigma)}{d} \\ &\geq \frac{\operatorname{tr} \Lambda(\psi \otimes \sigma) - \|\phi - \psi\|_{\mathrm{tr}}}{d} \geq \frac{\eta - \|\phi - \psi\|_{\mathrm{tr}}}{d}, \end{aligned}$$

implying that for any state $|\phi\rangle$ such that $\operatorname{tr} \Lambda\phi \leq \eta/(2d)$, $\|\phi - \psi\|_{\mathrm{tr}} \geq \eta/2$. Let Q denote the projector onto $\operatorname{span}\{|\phi\rangle : \Lambda|\phi\rangle = \lambda|\phi\rangle, \lambda \geq 1/(2N)\}$, set $Q^{\perp} = I - Q$, and take $|\phi\rangle = Q^{\perp}|\psi\rangle / \|Q^{\perp}|\psi\rangle\|$. As $N = \lceil d/\eta \rceil$, $\operatorname{tr} Q^{\perp}\phi \leq \eta/(2d)$. By the same "gentle measurement" argument as used in the proof of Corollary 11, the fact that $\|\phi - \psi\|_{\mathrm{tr}} \geq \eta/2$ implies that $\operatorname{tr} Q\psi \geq \eta^2/4$. So by Theorem 9, Algorithm 1 accepts with probability at least $(1 - e^{-1})\eta^2/4 \geq \eta^2/7$. In case 2, $\operatorname{tr} \Lambda\psi = \operatorname{tr} \Lambda(\psi \otimes I/d) \leq \zeta$, so it is immediate from Theorem 9 that Algorithm 1 accepts with probability at most $2\zeta\lceil d/\eta \rceil$. □

It is now straightforward to give a corrected proof of Theorem 8 by using Corollary 13 within the framework of Aaronson [3]. We will use the following lemma from [3]:

Lemma 14 (Aaronson [3])**.** *Suppose Bob receives an a-qubit message $|\psi\rangle$ from Alice and a w-qubit message $|\phi\rangle$ from Merlin, where $w \geq 2$. Let $A = O(aw \log^2 w)$ and $W = O(w \log w)$. Then by using A qubits from Alice and W qubits from Merlin, Bob can amplify his soundness error to 5^{-W} while keeping his completeness error $1/3$.*

Theorem 8 (restated). *For all partial or total boolean functions f, and all $w \geq 2$,*

$$Q^1(f) = O(\mathrm{QMA}^1_w(f) \cdot w \log^2 w).$$

Proof. Let Λ be the measurement corresponding to Bob's amplified algorithm from Lemma 14 and let $|\psi_x\rangle$ be the state of Alice's register. We know that, if $f(x, y) = 1$, there exists a state $|\phi\rangle$ of the witness register such that Λ accepts $\psi_x \otimes \phi$ with probability at least $2/3$. On the other hand, if $f(x, y) = 0$, then for all witness states $|\phi\rangle$, Λ accepts $\psi_x \otimes \phi$ with probability at most 5^{-W}. Inserting these parameters within Corollary 13 and using $d = 2^W$, we find that in the former case Algorithm 1 accepts with probability at least $4/63$, and in the latter case accepts with probability at most $4 \cdot 5^{-W} \cdot 2^W = o(1)$. The two cases can therefore be distinguished with $O(1)$ repetitions. □

Acknowledgements

We would like to thank Noah Linden for suggesting the application to testing whether one unitary operator is a permutation of another, Mark Wilde for pointing out references [16, 30, 34] and Scott Aaronson for helpful comments on a previous version. AM was supported by EPSRC Early Career Fellowship EP/L021005/1. AWH was funded by NSF grants CCF-1629809 and CCF-1452616. CYL is supported by the Department of Defense.

A An alternative protocol via testing disturbance

In this appendix we describe a alternative approach towards determining whether one of a sequence of n measurements accepts an input state, based around testing disturbance of the input state. We will need the following result regarding sequences of measurements:

Lemma 15 (Improved quantum union bound [16])**.** *Let ρ be a mixed state, and let $M_1, \ldots, M_T$ be a sequence of 2-outcome projective measurements. Suppose each M_t yields outcome 1 with probability at most ϵ when applied to ρ. Then if we apply $M_1, \ldots, M_T$ in sequence to ρ, the probability that at least one measurement yields outcome 1 is at most $4T\epsilon$.*

Lemma 15, which is due to Gao [16], improves previous bounds of a similar nature [35, 28, 3, 30, 34] up to quadratically.

We assume that we have one copy of some state ρ, and have access to quantum circuits which allow us to coherently implement each of a sequence of 2-outcome POVMs specified by projectors $\Lambda_1, \ldots, \Lambda_n$, where each Λ_i corresponds to the measurement $M_i = \{\Lambda_i, I - \Lambda_i\}$, and the first outcome is associated with acceptance, the second with rejection. We further assume that η, ζ are picked such that exactly one of the following two cases holds:

1. ρ is pure and, for all pure states $|\phi\rangle$ such that $\|\rho - \phi\|_{\mathrm{tr}} \leq \eta$, we have $\mathbb{E}_j[\mathrm{tr}\,\Lambda_j\phi] \geq \eta/n$ ("the average probability of acceptance is quite high for all states relatively close to ρ").
2. For all j, $\mathrm{tr}\,\Lambda_j\rho \leq \zeta$ ("the probability that any measurement accepts ρ is low").

Our task is to accept in the first case, and reject in the second. The first case may seem somewhat unintuitive, but we state it in this way so that it encompasses all our applications.

We use Algorithm 2 below to complete this task. The intuition behind this algorithm is as follows, in the case that $\rho = |\psi\rangle\langle\psi|$ and $\eta = \Theta(1)$. Throughout the algorithm, the state of the system is of the form $\alpha|0\rangle|\psi\rangle + \beta|1\rangle|\widetilde{\psi}\rangle$ for some state $|\widetilde{\psi}\rangle$. Assume we are in case 1 above. If $|\widetilde{\psi}\rangle \approx |\psi\rangle$ and β is not too small, the next random choice of measurement will accept with fairly high probability (roughly $\Omega(1/n)$). On the other hand, if $|\widetilde{\psi}\rangle$ is far from $|\psi\rangle$ or β is small, the test in step 2a would accept with high probability if it were performed. So the overall probability that the test accepts at this stage is $\Omega(1/n)$ in either case; repeating $O(n)$ times, the overall acceptance probability is $\Omega(1)$. On the other hand, if we are in case 2, we can use Lemma 15 to infer that after making $O(n)$ measurements the overall acceptance probability is $O(n\zeta)$.

We now prove the correctness of Algorithm 2 more formally.

Theorem 16. *In case 1, Algorithm 2 accepts with probability at least $\eta^2/7 - O(1/n)$. In case 2, Algorithm 2 accepts with probability at most $2\lceil 5n/\eta + 5/\eta^2\rceil\zeta$.*

Proof. First consider case 1, in which we would like the algorithm to accept. In this case we assume that ρ is pure, so write $\rho = |\psi\rangle\langle\psi|$. The algorithm accepts if and only if either the measurement in step 2a is made and the outcome is 1, or the first measurement outcome in step 2c is obtained. Call either of these a "good" measurement outcome.

The overall state of the algorithm at the start of the i'th step of the loop can be written as $\alpha_i|0\rangle|\psi\rangle + \beta_i|1\rangle|\psi_i\rangle$ for some normalised state $|\psi_i\rangle$ and some $\alpha_i, \beta_i \in \mathbb{C}$ such that $|\alpha_i|^2 + |\beta_i|^2 = 1$, with $|\psi_1\rangle = |\psi\rangle$ and $\alpha_1 = \beta_1 = 1/\sqrt{2}$. For any such state, the probability that the measurement in step 2a would return an outcome of 1, if it were made, is

$$\frac{1}{2}\|\alpha_i|\psi\rangle - \beta_i|\psi_i\rangle\|^2 = \frac{1}{2}(1 - 2\,\mathrm{Re}(\alpha_i^*\beta_i\langle\psi|\psi_i\rangle)). \tag{3}$$

1. Create the state $|+\rangle\langle+| \otimes \rho$, where $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$.
2. Repeat the following $k := \lceil 5n/\eta + 5/\eta^2 \rceil$ times:
   (a) With probability $1/(\eta n+1)$, perform a Hadamard gate on the first qubit and measure it in the computational basis. If the outcome is 0, reject; otherwise, accept.
   (b) Pick $j \in [n]$ uniformly at random.
   (c) Perform the measurement $\{|1\rangle\langle 1| \otimes \Lambda_j, I - |1\rangle\langle 1| \otimes \Lambda_j\}$. If the first outcome is returned, accept. Otherwise, retain the residual state of the two registers.
3. Reject.

Algorithm 2: Sequential measurement test

We say that $|\psi_i\rangle$ is disturbed if $\|\psi_i - \psi\|_{\mathrm{tr}} \geq \eta$, and undisturbed otherwise. First assume that $|\psi_i\rangle$ is undisturbed. Then the probability that the measurement in step 2c is made and the first measurement outcome is obtained is

$$\left(1 - \frac{1}{\eta n + 1}\right) |\beta_i|^2 \mathbb{E}_j[\mathrm{tr}\, \Lambda_j \psi_i] \geq \left(1 - \frac{1}{\eta n + 1}\right) |\beta_i|^2 \frac{\eta}{n} \tag{4}$$

as we are in case 1. By (3), the probability that the measurement in step 2a is made and an outcome of 1 is obtained is lower-bounded by

$$\frac{1}{\eta n + 1}\left(1 - 2|\beta_i|\sqrt{1 - |\beta_i|^2}|\langle\psi|\psi_i\rangle|\right) \geq \frac{1}{\eta n + 1}\left(1 - 2|\beta_i|\sqrt{1 - |\beta_i|^2}\right). \tag{5}$$

If $|\beta_i| \geq 1/\sqrt{5}$, then (4) is $\geq \eta^2/(5(\eta n + 1))$; if $|\beta_i| \leq 1/\sqrt{5}$ then (5) is $\geq 1/(5(\eta n + 1))$. As $\eta^2 \leq 1$, the first bound is always lower, so the probability that a good measurement outcome is received if $|\psi_i\rangle$ is undisturbed obeys the overall lower bound of $\eta^2/(5(\eta n + 1))$.

On the other hand, if $|\psi_i\rangle$ is disturbed, by (3) the probability that the measurement in step 2a is made and returns an outcome of 1 is lower-bounded by

$$\frac{1}{\eta n + 1} \cdot \frac{1}{2}(1 - 2|\alpha_i||\beta_i||\langle\psi|\psi_i\rangle|) \geq \frac{1 - |\langle\psi|\psi_i\rangle|}{2(\eta n + 1)} \geq \frac{1 - |\langle\psi|\psi_i\rangle|^2}{4(\eta n + 1)} \geq \frac{\eta^2}{4(\eta n + 1)}.$$

Therefore, at the i'th step of the loop, the probability of a good measurement outcome is at least $p := \eta^2/(5(\eta n + 1))$ whether or not $|\psi_i\rangle$ is disturbed. The probability that the protocol fails at any given step – by incorrectly rejecting – is at most $q := 1/(\eta n + 1)$. So the probability that the protocol terminates with a good measurement outcome occurring before failure is lower-bounded by

$$(1-q)p + (1-p)(1-q)^2 p + (1-p)^2(1-q)^3 p + \cdots + (1-p)^{k-1}(1-q)^k p = p(1-q)\left(\frac{1 - (1-p)^k(1-q)^k}{1 - (1-p)(1-q)}\right).$$

We have $(1-q)^k \leq (1-p)^k \leq e^{-pk} \leq 1/e$, so the overall probability of success is lower-bounded by

$$\left(1 - \frac{1}{e^2}\right)\frac{p - pq}{p + q - pq} = \left(1 - \frac{1}{e^2}\right)\frac{\eta^3 n}{5\eta n + \eta^3 n + 5} \geq \frac{\eta^2}{7} - O\left(\frac{1}{n}\right).$$

Now consider case 2. By assumption, Λ_j accepts ρ with probability at most ζ for all j, so the measurement operator $|1\rangle\langle 1| \otimes \Lambda_j$ accepts the starting state $|+\rangle\langle +| \otimes \rho$ with probability at most $\zeta/2$. Also, the measurement in step 2a would reject the starting state with certainty. At most k measurements are made in the algorithm. By the quantum union bound (Lemma 15), the probability that any measurement made in the algorithm leads to acceptance is upper-bounded by $2k\zeta$. □

An alternative protocol is to start with two copies of the input state, perform the measurements M_j on only the first copy, and test disturbance between the two copies using the swap test [9]. This would avoid the need for controlled measurements, but would require an additional copy of the input state.

Given Theorem 16, it is easy to show variants of Corollary 11 and Corollary 13 with slightly worse constants, which imply the rest of the results in the paper.

B Classical lower bound for testing linear isomorphism

Proposition 17. *There is a universal constant $\epsilon > 0$ such that any classical ϵ-tester for linear isomorphism of two unknown boolean functions must make $\Omega(2^{n/2})$ queries.*

Proof. The proof is very similar to the lower bound on the property-testing variant of Simon's problem [10]. By the Yao principle, it is sufficient to bound the success probability of deterministic algorithms which distinguish between the following two distributions:

- $\mathcal{D}_{\text{yes}}$: $f : \{0,1\}^n \to \{0,1\}$ is picked uniformly at random, $A \in GL_n(\mathbb{F}_2)$ is picked uniformly at random, and g is defined by $g(x) = f(Ax)$.
- $\mathcal{D}_{\text{no}}$: $f, g : \{0,1\}^n \to \{0,1\}$ are each picked uniformly at random, conditioned on $g(0^n) = f(0^n)$ and
$$\forall A \in GL_n(\mathbb{F}_2) \quad |\{x : g(x) = f(Ax)\}| \leq (1-\epsilon)2^n. \tag{6}$$

To be precise, the algorithm is given an input picked from the distribution $\mathcal{D} = \frac{1}{2}(\mathcal{D}_{\text{yes}} + \mathcal{D}_{\text{no}})$, and is asked to determine whether it was picked from $\mathcal{D}_{\text{yes}}$ or $\mathcal{D}_{\text{no}}$. Since $\mathcal{D}_{\text{no}}$ is not easy to analyze, we first argue that it is close to a much simpler distribution. Define $\mathcal{D}_{\text{rand}}$ to be the uniform distribution over $f, g : \{0,1\}^n \to \{0,1\}$ subject only to the constraint that $g(0^n) = f(0^n)$. We claim that $\mathcal{D}_{\text{rand}}$ satisfies (6) with high probability, which will imply that $\mathcal{D}_{\text{rand}}$ and $\mathcal{D}_{\text{no}}$ are close in variational distance. Indeed, fix a choice of f and A. Let $f \circ A$ denote the function $x \mapsto f(Ax)$. Then the probability that a random g agrees with $f \circ A$ in a $\geq 1-\epsilon$ fraction of positions is $\approx 2^{-(1-H_2(\epsilon))2^n}$ where $H_2(\epsilon) = -\epsilon\log(\epsilon) - (1-\epsilon)\log(1-\epsilon)$. Since there are $\leq 2^{n^2}$ choices of A, the probability that a random g fails to satisfy (6) is at most

$$\exp\left(n^2 - 2^n(1 - H_2(\epsilon))\right).$$

For $\epsilon < 1/2$ and sufficiently large n this probability is negligible. We now proceed as though the input were chosen from the distribution $\frac{1}{2}(\mathcal{D}_{\text{yes}} + \mathcal{D}_{\text{rand}})$.

Now consider any deterministic decision tree which queries positions in f and g, without loss of generality querying a distinct position at each step. The values $f(0^n)$ and $g(0^n)$ give no useful information for the algorithm to distinguish between $\mathcal{D}_{\text{yes}}$ and $\mathcal{D}_{\text{rand}}$, so we can assume that they are never queried. Thus, in a "no" instance, the response to queries is always uniformly random.

In a "yes" instance, the response to a new query, say, $g(x')$ is uniformly random unless there exists x such that $f(x)$ has been previously queried and $x' = Ax$. If every query by the algorithm receives a uniformly random response, the algorithm cannot distinguish this from a "no" instance.

For any sequence of k previous queries, the probability (over the random choice of A) that the next query corresponds to a pair $x' = Ax$ of this form is at most $k/(2^n - 1)$ by a union bound. Therefore, the probability that an algorithm making k queries has found such a pair at any point in its execution is $O(k^2/2^n)$. So, to achieve success probability $2/3$, it is necessary to make $k = \Omega(2^{n/2})$ queries. □

References

[1] S. Aaronson. The security of private-key quantum money. In preparation.

[2] S. Aaronson. Limitations of quantum advice and one-way communication. *Theory of Computing*, 1:1–28, 2004. `quant-ph/0402095`.

[3] S. Aaronson. QMA/qpoly ⊆ PSPACE/poly: De-Merlinizing quantum protocols. In *Proc. 21st Annual IEEE Conf. Computational Complexity*, pages 261–273, 2006. `quant-ph/0510230`.

[4] Y. Aharonov and M. Vardi. Meaning of an individual "Feynman path". *Phys. Rev. D*, 21(8):2235–2240, 1980.

[5] N. Alon, E. Blais, S. Chakraborty, D. Garcia-Soriano, and A. Matsliah. Nearly tight bounds for testing function isomorphism. *SIAM J. Comput.*, 42(2):459–493, 2013.

[6] L. Babai and S. Chakraborty. Property testing of equivalence under a permutation group action. *ACM Transactions on Computation Theory*, to appear.

[7] C. Bennett, G. Brassard, S. Breidbart, and S. Wiesner. Quantum cryptography, or unforgeable subway tokens. In *Proc. CRYPTO'83*, pages 267–275, 1983.

[8] H. Buhrman, R. Cleve, S. Massar, and R. de Wolf. Non-locality and communication complexity. *Rev. Mod. Phys.*, 82(1):665–698, 2010. `arXiv:0907.3584`.

[9] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf. Quantum fingerprinting. *Phys. Rev. Lett.*, 87(16):167902, 2001. `quant-ph/0102001`.

[10] H. Buhrman, L. Fortnow, I. Newman, and H. Röhrig. Quantum property testing. *SIAM J. Comput.*, 37(5):1387–1400, 2008. `quant-ph/0201117`.

[11] S. Chakraborty, E. Fischer, A. Matsliah, and R. de Wolf. New results on quantum property testing. In *Proceedings of FSTTCS*, pages 145–156, 2010. `arXiv:1005.0523`.

[12] M. Ettinger, P. Høyer, and E. Knill. The quantum query complexity of the hidden subgroup problem is polynomial. *Inf. Proc. Lett.*, 91:43–48, 2004. `quant-ph/0401083`.

[13] B. Fefferman, H. Kobayashi, C. Y.-Y. Lin, T. Morimae, and H. Nishimura. Space-efficient error reduction for unitary quantum computations, 2016. `arXiv:1604.08192`.

[14] E. Fischer and A. Matsliah. Testing graph isomorphism. *SIAM J. Comput.*, 38(1):207–225, 2008.

[15] K. Friedl, F. Magniez, M. Santha, and P. Sen. Quantum testers for hidden group properties. *Fundamenta Informaticae*, 91(2):325–340, 2009. quant-ph/0208184.

[16] J. Gao. Quantum union bounds for sequential projective measurements. *Phys. Rev. A*, 92:052331, 2015. arXiv:1410.5688.

[17] N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden. Quantum cryptography. *Reviews of Modern Physics*, 74:145–195, 2002. quant-ph/0101098.

[18] E. Grigorescu, K. Wimmer, and N. Xie. Tight lower bounds for testing linear isomorphism. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2013)*, pages 559–574, 2013.

[19] A. Harrow and A. Montanaro. An efficient test for product states, with applications to quantum Merlin-Arthur games. In *Proc. 51st Annual Symp. Foundations of Computer Science*, pages 633–642, 2010. arXiv:1001.0017.

[20] A. Harrow and A. Montanaro. Testing product states, quantum Merlin-Arthur games and tensor optimization. *J. ACM*, 60(1), 2013. arXiv:1001.0017.

[21] A. Harrow and A. Winter. How many copies are needed for state discrimination? *IEEE Trans. Inform. Theory*, 58(1):1–2, 2012. quant-ph/0606131.

[22] B. Kaulakys and V. Gontis. Quantum anti-Zeno effect. *Phys. Rev. A*, 56(2):1131, 1997. quant-ph/9708024.

[23] P. Koiran, V. Nesme, and N. Portier. A quantum lower bound for the query complexity of Simon's problem. In *Proc. 32nd International Conference on Automata, Languages and Programming (ICALP'05)*, pages 1287–1298, 2005. quant-ph/0501060.

[24] C. Lomont. The hidden subgroup problem – review and open problems, 2004. quant-ph/0411037.

[25] C. Marriott and J. Watrous. Quantum Arthur-Merlin games. *Computational Complexity*, 14(2):122–152, 2005. cs/0506068.

[26] B. Misra and E. Sudarshan. The Zeno's paradox in quantum theory. *J. Math. Phys.*, 18:756, 1977.

[27] A. Montanaro and R. de Wolf. A survey of quantum property testing. *Theory of Computing Graduate Surveys*, 2016(7):1–81, 2016. arXiv:1310.2035.

[28] T. Ogawa and H. Nagaoka. A new proof of the channel coding theorem via hypothesis testing in quantum information theory. In *Proc. 2002 IEEE International Symposium on Information Theory*, page 73, 2002. quant-ph/0208139.

[29] W. Schieve, L. Horwitz, and J. Levitan. Numerical study of Zeno and anti-Zeno effects in a local potential model. *Phys. Lett. A*, 54:264, 1989.

[30] P. Sen. Achieving the Han-Kobayashi inner bound for the quantum interference channel by sequential decoding. In *Proc. IEEE International Symposium on Information Theory (ISIT'12)*, pages 736–740, 2012. arXiv:1109.0802.

[31] G. Wang. Property testing of unitary operators. *Phys. Rev. A*, 84:052328, 2011. arXiv:1110.1133.

[32] T.C. Wei and P.M. Goldbart. Geometric measure of entanglement and applications to bipartite and multipartite quantum states. *Phys. Rev. A*, 68(4):42307, 2003. quant-ph/0307219.

[33] M. Wilde. *Quantum Information Theory*. Cambridge University Press, 2013.

[34] M. Wilde. Sequential decoding of a general classical-quantum channel. *Proc. Roy. Soc. A*, 469:20130259, 2013. arXiv:1303.0808.

[35] A. Winter. Coding theorem and strong converse for quantum channels. *IEEE Trans. Inform. Theory*, 45(7):2481–2485, 1999.