

A Survey of Quantum Property Testing

Ashley Montanaro*

Ronald de Wolf†

February 3, 2015

Abstract

The area of property testing tries to design algorithms that can efficiently handle very large amounts of data: given a large object that either has a certain property or is somehow “far” from having that property, a tester should efficiently distinguish between these two cases. In this survey we describe recent results obtained for *quantum* property testing. This area naturally falls into three parts. First, we may consider quantum testers for properties of classical objects. We survey the main examples known where quantum testers can be much (sometimes exponentially) more efficient than classical testers. Second, we may consider classical testers of quantum objects. These arise for instance when one is trying to determine if quantum states or operations do what they are supposed to do, based only on classical input-output behavior. Finally, we may also consider quantum testers for properties of quantum objects, such as states or operations. We survey known bounds on testing various natural properties, such as whether two states are equal, whether a state is separable, whether two operations commute, etc. We also highlight connections to other areas of quantum information theory and mention a number of open questions.

Contents

1	Introduction	3
1.1	Quantum testing of classical properties	4
1.2	Classical testing of quantum properties	5
1.3	Quantum testing of quantum properties	6
2	Quantum testing of classical properties	6
2.1	Preliminaries	6
2.2	Upper bounds	7
2.2.1	Using amplitude amplification	7
2.2.2	Using the Bernstein-Vazirani algorithm	8
2.2.3	Testing juntas	9
2.2.4	Using Simon’s algorithm	11
2.2.5	Using Shor’s algorithm	14

*Department of Computer Science, University of Bristol, UK; ashley@cs.bris.ac.uk. Supported by an EPSRC Early Career Fellowship (EP/L021005/1).

†CWI and University of Amsterdam, the Netherlands; rdewolf@cwi.nl. Supported by a Vidi grant from the Netherlands Organization for Scientific Research (NWO) which ended in 2013, by ERC Consolidator Grant QPROGRESS, and by the European Commission IST STREP project Quantum Algorithms (QALGO) 600700.

2.2.6	Using quantum counting	15
2.2.7	Using Ambainis’s algorithm	17
2.2.8	Quantum speed-ups for testing group-theoretic properties	19
2.3	Lower bounds	19
2.3.1	The polynomial method	19
2.3.2	The adversary method	21
2.3.3	A communication complexity method?	22
3	Classical testing of quantum properties	23
3.1	Self-testing gates	24
3.2	Self-testing protocols	27
4	Quantum testing of quantum properties: States	32
4.1	Pure states	33
4.1.1	Equality	33
4.1.2	Productness	35
4.1.3	Arbitrary finite sets	37
4.1.4	Open questions	38
4.2	Mixed states	39
4.2.1	Testing equality to a fixed pure state	41
4.2.2	Unitarily invariant properties	42
5	Quantum testing of quantum properties: Dynamics	44
5.1	Unitary operators	44
5.1.1	Distance measures	45
5.1.2	Controlled and inverse unitaries	47
5.1.3	From properties of states to properties of unitaries	47
5.1.4	Membership of the Pauli and Clifford groups	49
5.1.5	Testing commutativity	50
5.1.6	Testing quantum juntas	51
5.1.7	Other properties of unitary matrices	52
5.2	Properties of quantum channels	52
5.2.1	A distance measure on channels	53
5.2.2	Testing quantum measurements	54
6	Quantum properties and computational complexity	56
6.1	Computational hardness of testing quantum properties	56
6.2	From quantum property testers to complexity class inclusions	57
6.3	The quantum PCP conjecture	58
7	Conclusion	59

1 Introduction

In the last two decades, the amounts of data that need to be handled have exploded: think of the massive amounts of data on the web, or the data warehouses of customer information collected by big companies. In many cases algorithms need to decide whether this data has certain properties or not, without having sufficient time to trawl through all or even most of the data. Somehow we would like to detect the presence or absence of some global property by only making a few “local” checks. The area of *property testing* aims to design algorithms that can efficiently test whether some large object has a certain property, under the assumption that the object either has the property or is somehow “far” from having that property. An assumption like the latter is necessary for efficient property testing: deciding the property for objects that are “just on the boundary” typically requires looking at all or most of the object, which is exactly what we are trying to avoid here. In general, different property testing settings can be captured by the following “meta-definition”:

Property testing

Let \mathcal{X} be a set of objects and $d : \mathcal{X} \times \mathcal{X} \rightarrow [0, 1]$ be a distance measure on \mathcal{X} . A subset $\mathcal{P} \subseteq \mathcal{X}$ is called a *property*. An object $x \in \mathcal{X}$ is ϵ -far from \mathcal{P} if $d(x, y) \geq \epsilon$ for all $y \in \mathcal{P}$; x is ϵ -close to \mathcal{P} if there is a $y \in \mathcal{P}$ such that $d(x, y) \leq \epsilon$.

An ϵ -property tester (sometimes abbreviated to ϵ -tester) for \mathcal{P} is an algorithm that receives as input either an $x \in \mathcal{P}$ or an x that is ϵ -far from \mathcal{P} . In the former case, the algorithm accepts with probability at least $2/3$; in the latter case, the algorithm rejects with probability at least $2/3$.

Observe that, if an input is accepted by the property tester with high probability, then it must be ϵ -close to \mathcal{P} . This is true for all inputs, including inputs neither in \mathcal{P} nor ϵ -far from \mathcal{P} . The value of $2/3$ for the success probability is arbitrary and can equivalently be replaced with any other constant in $(1/2, 1)$ since we can efficiently reduce the error probability by repeating the test a few times and taking the majority outcome. We say that the tester has *perfect completeness* if it accepts every state in \mathcal{P} with certainty. The distance parameter ϵ is usually taken to be some positive constant. We will often just speak of a “tester,” leaving the value of ϵ implicit.

Clearly, the above meta-definition leaves open many choices: what type of objects to consider, what property to test, what distance measure to use, what range of ϵ to allow (the larger ϵ , the easier it should be to test \mathcal{P}), and how to measure the complexity of the testing algorithm. A lot of work in classical computer science has gone into the study of efficient testers for various properties, as well as proofs that certain properties are not efficiently testable, see for instance [37, 74, 65, 141, 73]. Typically, \mathcal{X} will be the set of all strings of length N over some finite alphabet, where we think of N as being very large. The distance will usually be normalized Hamming distance $d(x, y) = |\{i : x_i \neq y_i\}|/N$, though also more sophisticated metrics such as “edit distance” have been used. The complexity of the tester is typically measured by the number of *queries* it makes to entries of its input x , and a tester is deemed efficient if its number of queries is much less than the length of the input N , say $\text{polylog}(N)$ or even some constant independent of N . This captures the goal that a tester is able to efficiently handle huge amounts of data. The distance bound ϵ is often taken to be a small fixed constant, but in some cases it is also interesting to quantify the dependence of the tester’s complexity on ϵ as well as on N . For example, a tester whose complexity is $\Theta(2^{2^{1/\epsilon}})$ might be considered to be of little use in practice.

As an initial (very simple) example, suppose our property $\mathcal{P} = \{0^N\}$ consists of only one object, the all-zero string, and we use normalized Hamming distance. Our input x will either be in \mathcal{P} (i.e.,

$x = 0^N$) or ϵ -far from \mathcal{P} (i.e., x has at least ϵN 1-bits). An obvious tester would choose k indices in the string at random, query them, and reject if and only if there is a 1 in at least one of those positions. This tester accepts $x = 0^N$ with certainty (so it has perfect completeness), and fails to reject an input that is ϵ -far from \mathcal{P} with probability $(1 - \epsilon)^k$. Choosing $k = \Theta(1/\epsilon)$ gives a tester with small constant error probability, and this number of queries can be shown to be optimal.¹

In this survey paper we will be concerned with *quantum* property testing. There are several natural ways in which one can generalize property testing to the quantum world:

- Quantum testing of properties of classical objects. In this setting we would like to achieve provable quantum speed-ups over any possible classical algorithm, or to prove limitations on property testers, even if they are allowed to be quantum. By their very nature, efficient quantum query algorithms rely on extracting global information about the input, by querying in superposition; property testing is thus a plausible place to find significant quantum speed-ups. A very simple example of such a speed-up is for the above-mentioned property $\mathcal{P} = \{0^N\}$: a tester based on Grover’s search algorithm [80] would use $O(1/\sqrt{\epsilon})$ queries, in contrast to the $\Theta(1/\epsilon)$ queries that classical testers need.
- Classical testing of properties of quantum objects. Here we imagine we are given a black-box device which is claimed to implement some quantum process, and we would like to test whether it does what is claimed. However, our access to the device is classical: all we can do is feed classical inputs to the device, and receive classical measurement outcomes.
- Quantum testing of properties of quantum objects. In this most general scenario, we are given access to a quantum state or operation as a black box, and apply a quantum procedure to it to test whether it has some property.

We will discuss each of these settings in turn. We usually concentrate on describing the intuition behind prior work, without giving detailed proofs. Some of the results we present have not appeared in the literature before; these are largely based on combining, generalizing or improving existing works. Various open questions are pointed out throughout the survey.

A vast amount of work in quantum computing can be interpreted through the lens of property testing. Indeed, taken to extremes, any efficient quantum algorithm for a decision problem could be seen as an efficient property tester, and any measurement scheme that tries to learn properties of a quantum state or channel could be seen as a quantum property tester. We therefore concentrate on covering those algorithms which can clearly be understood as distinguishing objects with some property from those “far” from that property, and we make no attempt to be completely comprehensive. Also, our focus is on the computer-science aspects of the field, rather than work which primarily takes a physics perspective, such as the study of interaction-free measurement and the flourishing field of quantum metrology. Finally, we do not attempt to cover the (now very extensive) field of classical testers for classical properties. For much more on these, see the references given earlier.

1.1 Quantum testing of classical properties

In the first part of this paper we will consider quantum testing of *classical* properties. Again, \mathcal{X} will typically be the set of all strings of length N over some finite alphabet, the distance will be

¹Note that the complexity of a property can differ much from that of its complement. For example, $\overline{\mathcal{P}} = \{0,1\}^N \setminus \{0^N\}$ is trivial to test if $\epsilon > 1/N$: no string is ϵ -far from $\overline{\mathcal{P}}$, so we might as well accept every input without querying anything.

normalized Hamming distance, and the complexity of both quantum and classical property testers will be measured by the number of queries to the input x .

One of our goals is to survey examples of quantum speed-up, i.e., describe properties where the complexity of quantum testers is substantially less than the complexity of classical testers. Most known quantum speed-ups for testing classical properties were derived from earlier improvements in query complexity: they rely on quantum algorithms such as those of (in chronological order) Bernstein and Vazirani [34], Simon [149], Shor [147], Grover [80], and Ambainis [16]. In Section 2.2 we describe these quantum property testers and the improvements they achieve over classical testers. Some of the properties considered are very natural, and some of the improvements achieved are quite significant.

In Section 2.3 we describe some *lower*-bound methods for quantum property testing, i.e., methods to show query complexity lower bounds for quantum algorithms that want to test specific properties. The main lower bounds in this area have been obtained using the *polynomial* method. We also describe the *adversary* method, which—when applied properly—proves optimal lower bounds. And we ask whether the recent classical property testing lower bounds of Blais et al. [36], based on communication complexity, can be applied to quantum testers as well.

1.2 Classical testing of quantum properties

In the second part we will consider *classical* testing of quantum properties. At first sight, this scenario might make no sense—how could a classical algorithm, without the ability to perform any quantum operations, be able to test quantum objects? But suppose someone gives us a quantum state and claims it is an EPR-pair. Or someone builds a quantum device to implement a Hadamard gate, or to measure in a specific basis. How can we test that these quantum objects conform to their specifications? These are questions often faced for instance by experimentalists who try to check that their quantum operations work as intended, or by parties who run quantum cryptographic hardware provided by an untrusted supplier. We do not want to assume here that we already have the ability to implement some other quantum operations reliably, because that would lead to an infinite regress: how did we establish that those other quantum objects are reliable? Accordingly, we somehow would like to test the given quantum object while only trusting our *classical* devices. Of course, in order to test a quantum object there has to be at least *some* interaction with the quantum object-to-be-tested. In the testers we consider, the only quantum involvement is with that object itself in a black-box fashion (whence the name “self-testing”): we can only observe its classical input-output behavior, but not its inner quantum workings.

This notion of quantum self-testing was introduced by Mayers and Yao [120, 121], who described a procedure to test photon sources that are supposed to produce EPR-pairs. Since then quite a lot of work has been done on self-testing. We focus on two areas for self-testing: in Section 3.1 we describe self-testing of universal sets of quantum *gates* and in Section 3.2 we describe the self-testing of *protocols* for two or more parties, focusing on protocols for the so-called *CHSH game*. Self-testing of protocols has found many applications in the fast-growing area of *device-independent* quantum cryptography, where parties want to run cryptographic protocols for things like key distribution or randomness generation, using quantum states or apparatuses (photon sources, measuring devices, etc.) that they do not fully trust. Self-testing the states or apparatuses makes this possible in some cases. Device-independent cryptography is quite a large area and we will not cover it in this survey; see, e.g., [27, 60, 9, 154, 155] for more about this area.

1.3 Quantum testing of quantum properties

In the final part of the paper we will consider cases where \mathcal{X} is a set of quantum objects and our tester is also quantum, which is a setting of both theoretical and experimental interest.

As experimentalists control ever-larger quantum systems in the lab, the question of how to characterize and certify these systems becomes ever more pressing. Small quantum systems can be characterized via a procedure known as *quantum state tomography* [133, 129]. However, completely determining the state of a system of n qubits necessarily requires exponentially many measurements in n . This is already a daunting task for fairly small experiments; for example, Häffner et al. [84] report tomography of a state of 8 ions requiring 656,100 experiments and a total measurement time of 10 hours. One way of reducing this complexity is to start with the assumption that the state is of a certain form (such as a low-rank mixed state [79, 67] or a matrix product state [61]), in which case the number of parameters required to be estimated can be dramatically reduced. The viewpoint of property testing suggests another approach: the direct determination of whether or not something produced in the lab has a particular property of interest, under the assumption that it either has the property or is far away from it.

One can view classical property testing algorithms in two ways: either as testing properties of data (such as graph isomorphism), or properties of functions (such as linearity). If one wishes to generalize property testing to the quantum realm, one is thus naturally led to two different generalizations: testing properties of quantum *states*, and properties of quantum *operations*. One can divide each of these further, according to whether the state is pure or mixed, and whether the operation is reversible or irreversible; this classification is illustrated in Table 1. We discuss each of these possibilities in Sections 4 and 5. Within some of these categories there are natural generalizations of properties studied classically. For example, testing properties of mixed states is analogous to the classical idea of testing properties of probability distributions. Some quantum properties, however, have no simple classical analog (such as properties relating to entanglement).

	Coherent	Incoherent
Static	Pure state (§4.1)	Mixed state (§4.2)
Dynamic	Unitary operator (§5.1)	Quantum channel (§5.2)

Table 1: The taxonomy of quantum properties

Classically, there are many connections known between property testing and computational complexity. In Section 6 we explore the link between quantum property testing and quantum computational complexity, including the use of property testers to prove results in computational complexity, and the use of computational complexity to prove limitations on property testers.

2 Quantum testing of classical properties

2.1 Preliminaries

We will use $[m]$ to denote $\{1, \dots, m\}$ and \mathbb{Z}_m to denote $\{0, \dots, m-1\}$ modulo m . When considering (quantum or classical) testers for classical objects, those classical objects are usually strings, $\mathcal{X} = [m]^N$, and the complexity of testers is measured by the number of *queries* they make to their input x . In some cases we let x correspond to a function $f : [N] \rightarrow [m]$, where $f(i) = x_i$, and i may be viewed as either an integer $i \in [N]$, or as its binary representation $i \in \{0, 1\}^{\lceil \log N \rceil}$.

Here we briefly define the quantum query model, referring to [49] for more details. We assume some basic familiarity with classical and quantum computing [129].

Informally, a query allows us to “read” x_i for any i of our choice. Mathematically, to make this correspond to a quantum operation, it is modeled by the unitary map

$$O_x : |i\rangle|b\rangle \mapsto |i\rangle|b + x_i\rangle.$$

Here the first register has dimension N and the second has dimension m . The answer x_i is added into this second register mod m . Part of the power of quantum query algorithms comes from their ability to apply a query to a *superposition* of different i s, thus globally “accessing” the input x while using only one query.

If $m = 2$, then putting the state $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ in the second register has the following effect:

$$O_x : |i\rangle|-\rangle \mapsto |i\rangle \frac{1}{\sqrt{2}}(|0 + x_i\rangle - |1 + x_i\rangle) = (-1)^{x_i} |i\rangle|-\rangle.$$

We will sometimes call this a “phase-query,” because the answer x_i to the query is inserted in the state as a phase (+1 if $x_i = 0$, and -1 if $x_i = 1$).

A T -query quantum algorithm is described by an initial state, say $|0^k\rangle$, and $T + 1$ fixed k -qubit unitaries U_0, \dots, U_T . The final state when the algorithm runs on input x is obtained by interleaving these unitaries with queries to x (O_x may be tensored with the identity operation on the remaining workspace qubits),

$$|\psi_x\rangle = U_T O_x U_{T-1} O_x \cdots O_x U_1 O_x U_0 |0^k\rangle.$$

This final state depends on x via the T queries. A measurement of the final state will determine the classical output of the algorithm.

2.2 Upper bounds

In this section we survey the main speed-ups that have been obtained using *quantum* testers for *classical* properties. Typically these apply pre-existing quantum algorithms to problems in property testing. Our distance measure will be normalized Hamming distance $d(x, y) = |\{i : x_i \neq y_i\}|/N$, unless explicitly stated otherwise.

2.2.1 Using amplitude amplification

A simple but very general way that quantum algorithms can speed up many classical property testers is via the powerful primitive of *amplitude amplification*, which was introduced by Brassard et al. [40] and can be seen as a generalization of Grover’s quantum search algorithm [80]. We assume we are given query access to some function f (treated as a black box), and have a quantum algorithm which, with probability p , outputs w such that $f(w) = 1$. Then the result of Brassard et al. is that, for any $p > 0$, we can find a w such that $f(w) = 1$ with $O(1/\sqrt{p})$ queries to f , with success probability at least $2/3$.

Amplitude amplification can be immediately applied to speed up classical property testers which have perfect completeness. Here we think of w as the internal randomness of the algorithm, and $f(w)$ as the test which is applied to the unknown object, based on the random bits w . We let $f(w) = 0$ if the test accepts, and $f(w) = 1$ if the test rejects. Assuming that the test has perfect completeness, finding w such that $f(w) = 1$ is equivalent to determining whether we should reject.

Given that the original test used q queries to find such a w with probability $p > 0$, we therefore obtain a test which uses $O(q/\sqrt{p})$ queries, still has perfect completeness, and rejects with constant probability.

For example, consider the well-studied classical property of **Linearity** [37]. A function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is said to be *linear* if $f(x \oplus y) = f(x) \oplus f(y)$, and *affine* if $f(x \oplus y) = f(x) \oplus f(y) \oplus 1$, where \oplus is addition modulo 2. (Linearity is equivalent to the condition $f(x) = \bigoplus_{i \in S} x_i$ for some $S \subseteq [n]$.) A simple and natural test for linearity is to pick $x, y \in \{0, 1\}^n$ uniformly at random and accept if and only if $f(x) \oplus f(y) = f(x \oplus y)$. This test uses only 3 queries, has perfect completeness, and can be shown [32] to reject functions f which are ϵ -far from linear with probability at least ϵ . Applying amplitude amplification to this tester, we immediately get a quantum ϵ -tester for **Linearity** which uses $O(1/\sqrt{\epsilon})$ queries. Another simple example is **Symmetry**, where $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is said to be symmetric if $f(x)$ depends only on $|\{i : x_i = 1\}|$. A classical tester for this property has been given by Majewski and Pippenger [119]. The tester uses 2 queries, has perfect completeness and rejects functions which are ϵ -far from symmetric with probability at least ϵ . Therefore, we again obtain a quantum ϵ -tester which uses $O(1/\sqrt{\epsilon})$ queries.

Hillery and Andersson [92] gave different quantum testers for these two properties (though also based on amplitude amplification), each of which uses $O(\epsilon^{-2/3})$ queries, which is worse. More recently, Chakraborty and Maitra [50] described a quantum algorithm for the closely related problem of testing whether a Boolean function is affine. Their algorithm also uses $O(1/\sqrt{\epsilon})$ queries and, although presented slightly differently, is also based on amplitude amplification.

2.2.2 Using the Bernstein-Vazirani algorithm

One of the first quantum algorithms was the Bernstein-Vazirani algorithm [34]. It efficiently decodes a given *Hadamard codeword*. Let $N = 2^n$, and identify $[N]$ with $\{0, 1\}^n$ so we can use the n -bit strings to index the numbers $1, \dots, N$.² Let $h : \{0, 1\}^n \rightarrow \{0, 1\}^N$ be the Hadamard encoding, defined by $h(s)_i = s \cdot i \bmod 2$; this is nothing more than identifying s with the linear function $h(s)(i) = s \cdot i \bmod 2$ and writing out its truth table. Note that two distinct Hadamard codewords $h(s)$ and $h(s')$ are at normalized Hamming distance exactly $1/2$. Given input $h(s)$, the Bernstein-Vazirani algorithm recovers s with probability 1 using only one quantum query. In contrast, any classical algorithm needs $\Omega(\log N)$ queries for this. The quantum algorithm works as follows:

1. Start with $|0^n\rangle$ and apply Hadamard gates to each qubit to form the uniform superposition $\frac{1}{\sqrt{N}} \sum_{i \in \{0, 1\}^n} |i\rangle$
2. Apply a phase-query to obtain $\frac{1}{\sqrt{N}} \sum_{i \in \{0, 1\}^n} (-1)^{x_i} |i\rangle$
3. Apply Hadamard transforms to each qubit and measure.

If $x_i = s \cdot i$ for all $i \in \{0, 1\}^n$, then it is easy to see that the measurement yields s with probability 1.

Buhrman et al. [47] showed this algorithm can be used to obtain an unbounded quantum speed-up for *testing* most subsets of Hadamard codewords.

Bernstein-Vazirani property for $A \subseteq \{0, 1\}^n$:
 $\mathcal{P}_{BV}^A = \{x \in \{0, 1\}^N : \exists s \in A \text{ such that } x = h(s)\}$

²In many presentations of the Bernstein-Vazirani, Simon, and Grover algorithms, the input is taken to be a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ rather than a string $x \in \{0, 1\}^N$. With $N = 2^n$, these two views are of course just notational variants of one another.

Theorem 1 (Buhrman et al. [47]). *For every $A \subseteq \{0,1\}^n$ there is an $O(1/\sqrt{\epsilon})$ -query quantum ϵ -tester for \mathcal{P}_{BV}^A ; in contrast, for a $1 - o(1)$ fraction of all sets A , every classical $1/2$ -tester for \mathcal{P}_{BV}^A needs $\Omega(\log N)$ queries.*

Proof. Quantum upper bound. We run the Bernstein-Vazirani algorithm on input x , which takes one quantum query. The algorithm will output some s , and if x equals some $h(s) \in \mathcal{P}_{BV}^A$ then this will be the corresponding s with certainty. Hence if $s \notin A$ we can reject immediately. If $s \in A$ then choose $i \in [N]$ at random, query x_i , and test whether indeed $x_i = s \cdot i$. If x is ϵ -far from \mathcal{P}_{BV}^A then this test will fail with probability ϵ . Using amplitude amplification, we can detect any x that is ϵ -far from \mathcal{P}_{BV}^A with success probability at least $2/3$ using $O(1/\sqrt{\epsilon})$ queries.

Classical lower bound. Choose the set $A \subseteq \{0,1\}^n$ uniformly at random. Consider the uniform input distribution over the set H of all N Hadamard codewords. Note that the Hadamard codewords that are not in \mathcal{P}_{BV}^A are $1/2$ -far from \mathcal{P}_{BV}^A , because any two distinct Hadamard codewords have normalized Hamming distance exactly $1/2$. Hence if \mathcal{P}_{BV}^A can be $1/2$ -tested with T queries, then there exists a decision tree (i.e., a deterministic query algorithm) that is correct on at least $2/3$ of the $x \in H$. Fix a deterministic decision tree \mathcal{T} of depth T . For each $x \in H$, the probability (over the choice of A) that $x \in \mathcal{P}_{BV}^A$ is $1/2$, irrespective of the output that \mathcal{T} gives on x , so the probability that \mathcal{T} correctly decides x is $1/2$. Then the probability that \mathcal{T} correctly decides at least $2/3$ of the $x \in H$ is $2^{-\Omega(N)}$ by a Chernoff bound. The total number of deterministic decision trees of depth T is at most $2^{2^T} N^{2^T - 1}$, because for each of the (at most) $2^T - 1$ internal nodes we have to choose an index to query, and for each of the (at most) 2^T leaves we have to choose a binary output value. Hence by the union bound, the probability (over the choice of A) that there exists a depth- T decision tree that correctly decides at least $2/3$ of the $x \in H$ is at most

$$2^{-\Omega(N)} \cdot 2^{2^T} N^{2^T - 1}.$$

For $T = (\log N)/2$ this quantity is negligibly small. This shows that a $1 - o(1)$ fraction of all possible sets A , there is no classical tester for \mathcal{P}_{BV}^A with $(\log N)/2$ or fewer queries. \square

As Buhrman et al. [47] noted, the above classical lower bound is essentially optimal because for every property $\mathcal{P} \subseteq \{0,1\}^N$ there exists an ϵ -tester with $T = \lceil \ln(3|\mathcal{P}|)/\epsilon \rceil$ queries, as follows. We just query the input $x \in \{0,1\}^N$ at T uniformly randomly chosen positions, and accept if and only if there is still at least one element $y \in \mathcal{P}$ that is consistent with all query outcomes. Clearly, if the input is in \mathcal{P} this test will accept, so it has perfect completeness. If the input is ϵ -far from \mathcal{P} , then the probability for a specific $y \in \mathcal{P}$ to “survive” T queries is at most $(1 - \epsilon)^T$. Hence by the union bound the probability that there is a $y \in \mathcal{P}$ surviving all T queries is at most $|\mathcal{P}| \cdot (1 - \epsilon)^T \leq |\mathcal{P}| \cdot e^{-\epsilon T} \leq 1/3$.

2.2.3 Testing juntas

Let $f : \{0,1\}^n \rightarrow \{+1, -1\}$ be a Boolean function (such an f can also be viewed as a string x of $N = 2^n$ bits, with $x_i = f(i)$), and $J \subseteq [n]$ be the set of (indices of) variables on which f depends. If $|J| \leq k$ then f is called a k -junta.

k -junta property: $\mathcal{P}_{k\text{-junta}} = \{f : \{0,1\}^n \rightarrow \{+1, -1\} : f \text{ depends on at most } k \text{ variables}\}$

The best known classical tester, due to Blais, uses $O(k \log k + k/\epsilon)$ queries [35], and the best known classical lower bound is $\Omega(k)$ [55] (for fixed ϵ).

Atıç and Servedio [20] gave an elegant quantum ϵ -property tester for $\mathcal{P}_{k\text{-junta}}$ using $O(k/\epsilon)$ quantum queries, slightly better than Blais’s classical tester.³

Theorem 2 (essentially Atıç and Servedio [20]). *There is a quantum tester for k -juntas that uses $O(k/\sqrt{\epsilon})$ queries.*

“Essentially” in the attribution of the above theorem refers to the fact that [20] proves an $O(k/\epsilon)$ bound. We observe here that the dependence on ϵ can easily be improved by a square root using amplitude amplification.

Proof. The basic quantum subroutine is the same as the Bernstein-Vazirani algorithm in Section 2.2.2:

1. Start with $|0^n\rangle$ and apply Hadamard gates to each qubit to form the uniform superposition $\frac{1}{\sqrt{N}} \sum_{i \in \{0,1\}^n} |i\rangle$
2. Apply a phase-query to obtain $\frac{1}{\sqrt{N}} \sum_{i \in \{0,1\}^n} f(i)|i\rangle$
3. Apply Hadamard transforms to each qubit and measure.

Let us analyze this subroutine by means of some Fourier analysis on the Boolean cube (see [130, 160] for background). For every $s \in \{0, 1\}^n$, let

$$\hat{f}(s) = \frac{1}{2^n} \sum_{i \in \{0,1\}^n} f(i)(-1)^{i \cdot s}$$

be the corresponding Fourier coefficient. Going through the steps of the quantum subroutine, it is easy to see that the final state before the measurement is

$$\sum_{s \in \{0,1\}^n} \hat{f}(s)|s\rangle.$$

Accordingly, the final measurement will sample an $s \in \{0, 1\}^n$ from the distribution given by the squared Fourier coefficients $\hat{f}(s)^2$. This procedure is known as *Fourier Sampling* [34]. It uses one query to f .

Let J be the set of variables on which the input f depends. The goal of the tester is to decide whether $|J| \leq k$ or not. Identifying sets $s \subseteq [n]$ with their characteristic vectors $s \in \{0, 1\}^n$, note that $\hat{f}(s) \neq 0$ only if the support of s lies within J , so each Fourier Sample gives us a subset of J . The tester will keep track of the union W of the supports seen so far. We will always have $W \subseteq J$, so if f is a k -junta then W will never have more than k elements. On the other hand, below we show that if f is ϵ -far from any k -junta, then with high probability after $O(k/\sqrt{\epsilon})$ queries W will end up having more than k elements.

For a subset $W \subseteq [n]$ of size at most k , define $g_W(i) = \sum_{s \subseteq W} \hat{f}(s)(-1)^{i \cdot s}$. This function g_W need not be a Boolean function, but we can consider the Boolean function h_W that is the sign of g_W . This h_W only depends on the variables in W , so it is a k -junta and hence ϵ -far from f . Now

³In fact, at the time [20] was written the best classical upper bound was only $O((k \log k)^2/\epsilon)$ [66].

we have

$$\begin{aligned}
\epsilon &\leq \frac{1}{2^n} \sum_{i:f(i) \neq h_W(i)} 1 \\
&\leq \frac{1}{2^n} \sum_{i:f(i) \neq h_W(i)} (f(i) - g_W(i))^2 \\
&\leq \mathbb{E}_{i \in \{0,1\}^n} [(f(i) - g_W(i))^2] \\
&= \sum_s (\hat{f}(s) - \hat{g}_W(s))^2 \\
&= \sum_{s \notin W} \hat{f}(s)^2,
\end{aligned}$$

where the first equality is Parseval's identity. But this means that with probability at least ϵ , Fourier Sampling will output an s that is not fully contained in W . Now we use amplitude amplification to find such an s , using an expected number of $O(1/\sqrt{\epsilon})$ queries, and set $W := W \cup s$ (so W 's size grows by at least one). Repeating this at most $k + 1$ times, after an expected number of $O(k/\sqrt{\epsilon})$ queries the set W (which was initially empty) will contain more than k variables and we can reject the input. \square

Very recently Ambainis et al. [17] came up with a quantum k -junta tester that uses only $\tilde{O}(\sqrt{k/\epsilon})$ queries.⁴ Unlike the tester of Atıcı and Servedio, this actually beats the best known classical lower bound. The algorithm of [17] uses the adversary bound (see Section 2.3.2 below), building upon quantum algorithms due to Belovs [?] for *learning* the relevant variables of the junta. Their algorithm is substantially more complicated than the above, and we will not explain it here. They also give an implementation of their algorithm with *time complexity* (i.e., number of quantum gates used) $\tilde{O}(n\sqrt{k/\epsilon})$. They prove a quantum lower bound of $\Omega(k^{1/3})$ queries, leaving open the following:

Question 1. *What is the quantum query complexity of testing juntas?*

2.2.4 Using Simon's algorithm

The first exponential speed-up for quantum property testing was obtained by Buhrman et al. [47]. It is inspired by Simon's algorithm [149], which was the first algorithm to have a provable exponential speed-up over classical algorithms in the black-box model and inspired Shor's factoring algorithm [147] (which we will see in the next section). Again let $N = 2^n$ and identify $[N]$ with $\{0, 1\}^n$. Consider an input $x \in [N]^N$ for which there exists an $s \in \{0, 1\}^n \setminus \{0^n\}$ such that $x_i = x_j$ if and only if $(j = i \text{ or } j = i \oplus s)$. Simon's algorithm finds s with high probability using $O(\log N)$ queries. The core of the algorithm is the following quantum subroutine:

1. Start with $|0^n\rangle|0^n\rangle$ and apply Hadamard transforms to the first n qubits to form $\frac{1}{\sqrt{N}} \sum_{i \in \{0,1\}^n} |i\rangle|0^n\rangle$
2. Apply a query to obtain $\frac{1}{\sqrt{N}} \sum_{i \in \{0,1\}^n} |i\rangle|x_i\rangle$
3. Measure the second register. This yields some $z = x_i$ and collapses the first register to the two indices with value z : $\frac{1}{\sqrt{2}}(|i\rangle + |i \oplus s\rangle)$

⁴The $\tilde{O}(\cdot)$ notation hides logarithmic factors in k .

4. Apply Hadamard transforms to the first n qubits and measure the state, obtaining some $y \in \{0, 1\}^n$.

It is easy to calculate that the measured state is (up to phases) a uniform superposition over all 2^{n-1} strings $y \in \{0, 1\}^n$ that satisfy $s \cdot y = 0 \pmod{2}$. Each such y gives us a linear constraint $(\text{mod } 2)$ on the bits of s . Repeating this subroutine $\Theta(n)$ times gives, with high probability, $n - 1$ linearly independent $y^{(1)}, \dots, y^{(n-1)}$ all orthogonal to s . From these, s can be calculated classically by Gaussian elimination. Brassard and Høyer [39] subsequently gave an *exact* version of this algorithm, where each new y is produced by a modification of Simon’s subroutine that uses $O(1)$ queries and is *guaranteed* to be linearly independent from the previous ones (as long as such a linearly independent y exists).

This algorithm can be used to obtain a strong quantum speed-up for testing a specific property.

Simon property:

$$\mathcal{P}_{\text{Simon}} = \{x \in [N]^N : \exists s \in \{0, 1\}^n \setminus \{0^n\} \text{ such that } x_i = x_j \text{ if } j = i \oplus s\}$$

Note that, compared with Simon’s original problem, the ‘if and only if’ has been replaced with an ‘if.’ Hence x_i and x_j can be equal even for distinct i, j for which $j \neq i \oplus s$. However, also for such more general inputs, Simon’s quantum subroutine (and the Brassard-Høyer version thereof) only produces y such that $s \cdot y = 0 \pmod{2}$. The speed-up is as follows; for simplicity we state it for fixed $\epsilon = 1/4$ rather than making the dependence on ϵ explicit:

Theorem 3 (essentially Buhrman et al. [47]). *There is a quantum $1/4$ -property tester for the Simon property using $O(\log N)$ queries, while every classical $1/4$ -property tester needs $\Omega(\sqrt{N})$ queries.*

“Essentially” in the attribution of the above theorem refers to the fact that Buhrman et al. [47] devised a property of *binary* strings of length N . In our presentation it will be more convenient to consider a property consisting of strings over alphabet $[N]$. As remarked by Aaronson and Ambainis [3], Theorem 3 has an interesting consequence regarding the question of when we can hope to achieve exponential quantum speed-ups. In order to obtain a super-polynomial quantum speed-up for computing some function f in the query complexity model, it is known that there has to be a *promise* on the input, i.e., f has to be a partial function [31]. The Simon property indeed involves a promise on the input, namely that it is either in or far from $\mathcal{P}_{\text{Simon}}$; however, this promise is in some sense very weak, as the algorithm has to output the right answer on a $1 - o(1)$ fraction of $[N]^N$.

Proof. Quantum upper bound (sketch). We run the Brassard-Høyer version of Simon’s subroutine $n - 1$ times. We then classically compute a non-zero string s that is orthogonal to all the $n - 1$ strings y produced by these runs (there may be several such s , in which case we just pick any). We then randomly choose $i \in [N]$, query x_i and $x_{i \oplus s}$, and check if these two values are equal. If $x \in \mathcal{P}_{\text{Simon}}$ then s will have the property that $x_i = x_{i \oplus s}$ for all i . On the other hand, if x is $1/4$ -far from $\mathcal{P}_{\text{Simon}}$, then there exist at least $N/4$ $(i, i \oplus s)$ -pairs such that $x_i \neq x_{i \oplus s}$ (for otherwise we could put x into $\mathcal{P}_{\text{Simon}}$ by changing one value for each such pair, making fewer than $N/4$ changes in total). Hence in this case we reject with constant probability. Testing a few different $(i, i \oplus s)$ -pairs reduces the error probability to below $1/3$.

Classical lower bound. Consider three distributions: \mathcal{D}_1 is uniform over $\mathcal{P}_{\text{Simon}}$, \mathcal{D}_0 is uniform over all $x \in [N]^N$ that are $1/4$ -far from $\mathcal{P}_{\text{Simon}}$, and \mathcal{U} is uniform over $[N]^N$. We first show \mathcal{D}_0 and \mathcal{U} are very close.

Claim 4. The total variation distance between \mathcal{D}_0 and \mathcal{U} is $o(1)$.

Proof. Let $S = \{y : y \text{ is not } 1/4\text{-far from } \mathcal{P}_{\text{Simon}}\}$ be the elements that are not in the support of \mathcal{D}_0 . We will upper bound the size of S . Each element of $\mathcal{P}_{\text{Simon}}$ can be specified by giving an $s \in \{0, 1\}^n \setminus \{0^n\}$ and giving for each of the $N/2$ $(i, i \oplus s)$ -pairs the value $x_i = x_{i \oplus s}$. Hence

$$|\mathcal{P}_{\text{Simon}}| \leq (N-1)N^{N/2}$$

For each x , the number of y that are $1/4$ -close to x is at most $\binom{N}{N/4}N^{N/4}$. Hence the total number of elements $1/4$ -close to $\mathcal{P}_{\text{Simon}}$ is

$$|S| \leq (N-1)N^{N/2} \binom{N}{N/4} N^{N/4} = o(N^N).$$

Since \mathcal{U} is uniform over $[N]^N$ and \mathcal{D}_0 is uniform over $[N]^N \setminus S$, the total variation distance between these two distributions is $O(|S|/N^N) = o(1)$. \square

To finish the proof, below we slightly adapt the proof in [149] to show that a T -query classical algorithm distinguishing distributions \mathcal{D}_1 and \mathcal{U} has advantage of only $O(T^2/N)$ over random guessing.⁵ Since \mathcal{D}_0 and \mathcal{U} are $o(1)$ -close, a T -query classical algorithm distinguishing distributions \mathcal{D}_1 and \mathcal{D}_0 has advantage $O(T^2/N) + o(1)$ over random guessing. A classical tester for the Simon property can distinguish \mathcal{D}_1 and \mathcal{D}_0 with success probability at least $2/3$, so it needs $T = \Omega(\sqrt{N})$ queries. It remains to prove:

Claim 5. A T -query classical algorithm for distinguishing distributions \mathcal{D}_1 and \mathcal{U} has advantage $O(T^2/N)$ over random guessing.

Proof. By the well-known Yao principle [163], it suffices to prove the claim for an arbitrary *deterministic* T -query algorithm. The proof will show that both under \mathcal{D}_1 and \mathcal{U} the T queries are likely to yield a uniformly random sequence of T distinct values. Suppose the algorithm queries the indices i_1, \dots, i_T (this sequence may be adaptive, i.e., depend on x) and gets outputs x_{i_1}, \dots, x_{i_T} . Call a sequence of queries i_1, \dots, i_T *good* (for input x) if it shows a collision, i.e., $x_{i_k} = x_{i_\ell}$ for some $k \neq \ell$. Call the sequence *bad* (for x) otherwise. We will now show that the probability of a bad sequence is $O(T^2/N)$, both under input distribution \mathcal{U} and under \mathcal{D}_1 .

First, suppose the input x is distributed according to \mathcal{U} . Then each output x_{i_k} is uniformly distributed over $[N]$, independent of the other entries of x . The probability that i_k and i_ℓ form a collision is exactly $1/N$, so the expected number of collisions among the T queries is $\binom{T}{2}/N = O(T^2/N)$. Hence, by Markov's inequality, the probability that i_1, \dots, i_T form a good sequence is $O(T^2/N)$.

Second, suppose the input x is distributed according to \mathcal{D}_1 . Then there exists a nonzero $s \in \{0, 1\}^n$, unknown to the algorithm, such that $x_i = x_j$ whenever $j = i \oplus s$. Initially, all such s are equally likely under \mathcal{D}_1 (the probability that there are two distinct such s for x is negligibly small, and we will ignore this here). If i_1, \dots, i_{k-1} is bad, then we have excluded $\binom{k-1}{2}$ of the $N-1$ possible values of s , and all other values of s are equally likely. Let i_k be the next query and $S = \{i_k \oplus i_j : j < k\}$. This set S has at most $k-1$ members, so the probability (under \mathcal{D}_1) that S happens to contain the string s is at most $\frac{k-1}{N-1-\binom{k-1}{2}}$. If S does not contain s , then the only way to make the sequence good is if the uniformly random value x_{i_k} equals one of the $k-1$ earlier

⁵The ‘‘advantage’’ of the algorithm is the difference between success and failure probabilities.

values, which has probability $(k-1)/N$. Hence the probability that the bad sequence i_1, \dots, i_{k-1} remains bad, after query i_k is made, is very close to 1. More precisely:

$$\begin{aligned} \Pr[i_1, \dots, i_T \text{ is bad}] &= \prod_{k=2}^T \Pr[i_1, \dots, i_k \text{ is bad} : i_1, \dots, i_{k-1} \text{ is bad}] \\ &\geq \prod_{k=2}^T \left(1 - \frac{k-1}{N-1-\binom{k-1}{2}} - \frac{k-1}{N} \right) \\ &\geq 1 - \sum_{k=2}^T \left(\frac{k-1}{N-1-\binom{k-1}{2}} + \frac{k-1}{N} \right). \end{aligned}$$

Here we used the fact that $(1-a)(1-b) \geq 1-(a+b)$ if $a, b \geq 0$. The latter sum over k is $O(T^2/N)$, so the probability (under \mathcal{D}_1) that i_1, \dots, i_T form a good sequence is $O(T^2/N)$.

In both cases (\mathcal{U} and \mathcal{D}_1), conditioned on seeing a bad sequence, the sequence of outputs is a uniformly random sequence of T distinct values. Accordingly, the advantage (over random guessing) of the algorithm trying to distinguish these two distributions is upper bounded by the probability of seeing a good sequence, which is $O(T^2/N)$ in both cases. \square

\square

2.2.5 Using Shor's algorithm

Probably the most famous quantum algorithm to date is Shor's polynomial-time algorithm for factoring integers [147]. Its quantum core is an algorithm that can find the *period* of a periodic sequence. Chakraborty et al. [51] used this to show that *testing* periodicity exhibits a constant-versus-polynomial quantum-classical separation. Note that the Bernstein-Vazirani property (Section 2.2.2) exhibits a constant-versus-logarithmic separation, while the Simon property (Section 2.2.4) exhibits a logarithmic-versus-polynomial separation. Periodicity-testing thus exhibits a separation that is in some ways stronger than either of those.

Periodicity: let p be a prime number and m an integer such that $m \geq p$. A string $x \in [m]^N$ is *1-1- p -periodic* if it satisfies that $x_i = x_j$ if and only if $i = j \pmod p$ (equivalently, the elements in the sequence x_0, \dots, x_{p-1} are all unique, and after that the sequence repeats itself). For integers q and r such that $q \leq r \leq \sqrt{N}/2$, define the property $\mathcal{P}_{\text{period}}^{q,r} = \{x \in [m]^N : x \text{ is } 1-1-p\text{-periodic for some } p \in \{q, \dots, r\}\}$

Note that for a given p it is easy to test whether x is p -periodic or far from it: choose an $i \in [N]$ uniformly at random, and test whether $x_i = x_{i+kp}$ for a random positive integer k . If x is p -periodic then these values will be the same, but if x is far from p -periodic then we will detect this with good probability. However, $r - q + 1$ different values of p are possible in $\mathcal{P}_{\text{period}}^{q,r}$. Below we will set $q = r/2$ so $r/2 + 1$ different values for the period are possible. This makes the property hard to test for classical testers. On the other hand, in the *quantum* case the property can be tested efficiently.

Theorem 6 (Chakraborty et al. [51]). *For every even integer $r \in [2, \sqrt{N})$ and constant distance ϵ , there is a quantum property tester for $\mathcal{P}_{\text{period}}^{r/2,r}$ using $O(1)$ queries, while every classical property tester for $\mathcal{P}_{\text{period}}^{r/2,r}$ makes $\Omega(\sqrt{r/\log r \log N})$ queries. In particular, for $r = \sqrt{N}$ testing can be done with $O(1)$ quantum queries but requires $\Omega(N^{1/4}/\log N)$ classical queries.*

The quantum upper bound is obtained by a small modification of Shor’s algorithm: use Shor to find the period p of input x (if there is such a period) and then test this purported period with another $O(1)$ queries.⁶ The classical lower bound is based on modifying proofs from Lachish and Newman [113], who showed classical testing lower bounds for more general (and hence harder) periodicity-testing problems.

This quantum-classical separation is of the form $O(1)$ quantum queries vs $N^{\Omega(1)}$ classical queries, for a problem over a polynomial-sized alphabet (so each “entry” of the input takes only $O(\log N)$ bits). How large can we make this separation? This was already asked by Buhrman et al. [47] in the following way:

Question 2. *Is there a property of strings of length N (over a moderately-sized alphabet) that can be tested with $O(1)$ quantum queries but needs $\Omega(N)$ classical queries?*

A very recent result of Aaronson and Ambainis [4] is relevant here: they showed that if a (total or partial) function on $x \in \{0, 1\}^N$ can be computed with bounded error probability using k quantum queries, then the same function can be computed by a classical randomized algorithm using $O(N^{1-1/2k})$ queries. They also show that for $k = 1$ this upper bound is tight up to a logarithmic factor, for a testing problem called “Forrelation.” In that problem, roughly, the input consists of two Boolean functions f and g , each on ℓ -bit inputs so the total input length is $N = 2 \cdot 2^\ell$ bits, such that g is either strongly or weakly correlated with the Fourier transform of f (i.e., $g(x) = \text{sign}(\widehat{f}(x))$ either for most x or for roughly half of the x). They show that this problem can be tested with one quantum query, whereas classical testers need $\Omega(\sqrt{N}/\log N)$ queries.⁷

Hence for binary alphabets the answer to the above question is negative: everything that can be tested with $k = O(1)$ quantum queries can be tested with $O(N^{1-1/2k}) = o(N)$ classical queries. This classical upper bound can be extended to small alphabets, but the question remains open for instance when the alphabet size is N .

2.2.6 Using quantum counting

Grover’s quantum search algorithm [80] can be used to find the index i of a 1-bit in $x \in \{0, 1\}^N$ (i.e., $x_i = 1$) with high probability, using $O(\sqrt{N})$ queries. We will not describe the algorithm here, but just note that it can be modified to also *estimate*, for given $S \subseteq [m]$, the number of occurrences of elements from S in a string $x \in [m]^N$, using a number of queries that is much less than would be needed classically. More precisely, we have the following “quantum approximate counting” lemma, which follows from the work of Brassard et al. [40, Theorem 13]:

Lemma 7. *There exists a constant C such that for every set $S \subseteq [m]$ and every positive integer T , there is a quantum algorithm that makes T queries to input $x \in [m]^N$ and, with probability at least $2/3$, outputs an estimate p' to $p = |\{i : x_i \in S\}|/N$ such that $|p' - p| \leq C(\sqrt{p}/T + 1/T^2)$.*

We now describe an application of quantum counting to property testing, namely to testing whether two probability distributions are equal or ϵ -far from each other in total variation distance.

Equal distributions property: $\mathcal{P}_{\text{distribution}} = \{(p, p) : p \text{ is a distribution on } [m]\}$.

⁶These ingredients are already present in work of Hales and Hallgren [86], and in Hales’s PhD thesis [85]. However, their results are not stated in the context of property testing, and no classical lower bounds are proved there.

⁷The lower bound improves an earlier $N^{1/4}$ bound of Aaronson [2], which constituted the first $O(1)$ vs $N^{\Omega(1)}$ separation for quantum vs classical property testing.

Our distance measure on the set of pairs of distributions will be the sum of the total variation distances: $d((p, q), (p', q')) = \|p - p'\|_{tvd} + \|q - q'\|_{tvd}$, where the *total variation distance* between two distributions is $\|p - p'\|_{tvd} := \frac{1}{2} \sum_j |p(j) - p'(j)|$. Note that a pair of distributions (p, q) will be ϵ -far from $\mathcal{P}_{distribution}$ if and only if $\|p - q\|_{tvd} \geq \epsilon$.

There are different ways in which the distributions could be “given” to the tester, but in this section each distribution will be given as an input $x \in [m]^N$. This naturally induces a probability distribution \mathcal{D}_x on $[m]$ according to the relative frequencies of the different elements:

$$\mathcal{D}_x(j) = \frac{|\{i : x_i = j\}|}{N}.$$

We can obtain a sample according to \mathcal{D}_x by just querying x on a uniformly random index i . Assuming the distribution is given in this way is quite natural in the setting of property testing, where our input is usually a very long string x , much too long to inspect each of its elements. Note that \mathcal{D}_x does not change if we permute the elements of x ; it just depends on the relative frequencies. Also note that Lemma 7 can be used to estimate the probability of $S \subseteq [m]$ under \mathcal{D}_x .

Suppose we are given two distributions \mathcal{D}_x and \mathcal{D}_y on $[m]$ (the distributions are given in the form of two inputs $x, y \in [m]^N$), and we want to test whether these two distributions are equal or ϵ -far in total variation distance. Batu et al. [29] exhibited classical testers for this using $O((m/\epsilon)^{2/3} \log m)$ queries⁸, and Valiant [153] proved an almost matching lower bound of $\Omega(m^{2/3})$ for constant ϵ . These bounds have both recently been improved by Chan et al. [52] to $\Theta(m^{2/3}/\epsilon^{4/3})$, which is tight for all $\epsilon \geq m^{-1/4}$. Bravyi et al. [42] showed that quantum testers can do better in terms of their dependence on m :

Theorem 8 (Bravyi et al. [42]). *There is a quantum tester to test if two given distributions on $[m]$ are equal or ϵ -far using $O(\sqrt{m}/\epsilon^8)$ queries.*

The dependence on ϵ can probably be improved with more technical effort.

Proof sketch. Bravyi et al. [42] actually showed something stronger, namely that the total variation distance between two distributions can be estimated up to small additive error ϵ using $O(\sqrt{m}/\epsilon^8)$ quantum queries; this clearly suffices for testing. We sketch their idea here. Consider the following random process:

$$\begin{array}{l} \text{Sample } j \in [m] \text{ according to } \mathcal{D} = \frac{1}{2}(\mathcal{D}_x + \mathcal{D}_y). \\ \text{Output } \frac{|\mathcal{D}_x(j) - \mathcal{D}_y(j)|}{\mathcal{D}_x(j) + \mathcal{D}_y(j)} \end{array}$$

It is easy to see that the expected value of the output of this process is exactly the total variation distance between \mathcal{D}_x and \mathcal{D}_y , so it suffices to approximate that expected value. We sample j according to \mathcal{D} (which costs just one query), use the quantum algorithm of Lemma 7 with $S = \{j\}$ and $T = O(\sqrt{m}/\epsilon^6)$ queries to approximate both $\mathcal{D}_x(j)$ and $\mathcal{D}_y(j)$, and output the absolute difference between these two approximations divided by their sum. Bravyi et al. [42] show that repeating this $O(1/\epsilon^2)$ times and taking the average gives, with probability at least $2/3$, an ϵ -approximation of the expected value $\mathcal{D}_x - \mathcal{D}_y$ of the above random process. \square

A second problem is where we fix one of the two distributions, say to the uniform distribution on $[m]$ (assume m divides N so we can properly “fit” this distribution in $x \in [m]^N$). Goldreich

⁸All these classical bounds are stated in terms of number of samples rather than number of queries, but it is not hard to see that these two complexity measures are equivalent here.

and Ron [75] showed a classical testing lower bound of $\Omega(\sqrt{m})$ queries for this, and Batu et al. [28] proved a nearly tight upper bound of $\tilde{O}(\sqrt{m})$ queries. Bravyi et al. [42], and independently also Chakraborty et al. [51], showed that testing can be done more efficiently in the quantum case:

Theorem 9 (Bravyi et al. [42], Chakraborty et al. [51]). *There is a quantum tester to test if a given distribution on $[m]$ equals or is ϵ -far from the uniform distribution on $[m]$, using $O(m^{1/3}/\epsilon^2)$ quantum queries.*

Proof sketch. Pick a set $R \subseteq [N]$ of $r = m^{1/3}$ indices uniformly at random, and query its elements. If \mathcal{D}_x is uniform then it is very likely that all values $\{x_i\}_{i \in R}$ are distinct, so if there is some collision then we can reject immediately. Otherwise, let $S = \{x_i : i \in R\}$ be the r distinct results, and define $p = |\{i : x_i \in S\}|/N$. If \mathcal{D}_x is uniform then $p = r/m = 1/m^{2/3}$, but [51, Lemma 13] shows that if \mathcal{D}_x is ϵ -far from uniform then p will be noticeably higher: there is a constant $c > 0$ such that with high probability $p \geq (1 + c\epsilon^2)r/m$.

Now we use the quantum algorithm of Lemma 7 with $T = 4Cm^{1/3}/c\epsilon^2$ queries to obtain (with high probability) an estimate p' of p within additive error $|p' - p| \leq C(\sqrt{p}/T + 1/T^2)$. We accept if $p' \leq (1 + c\epsilon^2/2)r/m$, and reject otherwise. If $p = r/m = 1/m^{2/3}$ then the additive error is at most $C(c\epsilon^2/4Cm^{2/3} + c^2\epsilon^4/16C^2m^{2/3}) \leq \frac{c\epsilon^2}{2} \cdot \frac{r}{m}$, so then we will accept correctly. If $p \geq (1 + c\epsilon^2)r/m$ then it is easy to show that $p' \geq (1 + c\epsilon^2/2)r/m$, so then we will reject correctly. \square

Both Bravyi et al. [42] and Chakraborty et al. [51] showed that $\Omega(m^{1/3})$ quantum queries are also necessary, so the above result is essentially tight; the lower bound follows from a reduction from the collision problem [7]. Bravyi et al. [42] also exhibited a quantum tester for whether two distributions are equal or of disjoint support (i.e., orthogonal), using $O(m^{1/3})$ quantum queries. Chakraborty et al. [51] extended Theorem 9 to testing equality to *any* fixed distribution (not just the uniform one), at the expense of a polylog factor in the number of queries. They in turn used equality-testing to obtain better quantum testers for graph isomorphism.

2.2.7 Using Ambainis’s algorithm

Ambainis’s element distinctness algorithm [16] acts on an input $x \in [m]^N$, and finds (with high probability) a pair of distinct indices such that $x_i = x_j$ if such a pair exist, and reports “no collision” otherwise. It uses $O(N^{2/3})$ queries, which is optimal [7]. This algorithm spawned a large class of algorithms based on *quantum walks* (see [146] for a survey).

Ambainis et al. [18] use the element distinctness algorithm to give better quantum testers for certain *graph properties*. Graph properties have some amount of symmetry: they are invariant under relabelling of vertices. Problems with “too much” symmetry are known not to admit exponential quantum speed-up in the query complexity model [3], and the symmetry inherent to graph properties makes them an interesting test case for the question of how symmetric the problems can be for which we do obtain a significant quantum advantage. Ambainis et al. [18] use the element distinctness algorithm to give $\tilde{O}(N^{1/3})$ -query quantum testers for the properties of bipartiteness and being an expander in bounded-degree graphs. It is known that for classical testers, $\tilde{\Theta}(\sqrt{N})$ queries are necessary and sufficient to test these properties [75, 76]. Together with the graph isomorphism tester mentioned briefly at the end of Section 2.2.6, these are the only quantum results we are aware of for testing graph properties. In contrast, graph properties have been one of the main areas of focus in classical property testing.

Let us describe the results of [18] a bit more precisely. The object to be tested is an N -vertex graph G of degree d , so each vertex has at most d neighbors. We think of d as a constant and

will absorb the dependence of the bounds on d into the constant factor. The input is given as an *adjacency list*. Formally, it corresponds to an $x \in ([N] \cup \{*\})^{N \times d}$. The entries of x are indexed by a pair of a vertex $v \in [N]$ and a number $i \in [d]$, and $x_{v,i}$ is the i th neighbor of vertex v ; $x_{v,i} = *$ in case v has fewer than i neighbors. The *distance* between two graphs given as adjacency lists is defined to be the minimal number of edges one most change in the first graph to obtain the second.

A graph is **Bipartite** if its set of vertices can be partitioned into two disconnected sets, and is an **Expander** if there is a constant $c > 0$ such that every set $S \subseteq [N]$ of at most $N/2$ vertices has at least $c|S|$ neighbors outside of S .⁹

Theorem 10 (Ambainis et al. [18]). *There exist quantum testers for **Bipartite** and **Expander** using $\tilde{O}(N^{1/3})$ queries.*

Proof sketch. At a high level, the optimal *classical* testers for both properties look for collisions in a set of roughly \sqrt{N} elements. Using Ambainis’s algorithm, this can be done in roughly $N^{1/3}$ queries. Let us see how this works for the two properties.

A bipartite graph has no odd cycles. In contrast, for a graph that is far from bipartite one can show the following. Among roughly \sqrt{N} short ($O(\log N)$ -step) random walks from the same starting vertex v , there is likely to be a pair that “collides” in the sense that one walk reaches a vertex w after an even number of steps and the other reaches the same vertex w after an odd number of steps. These two paths between v and w now form an odd cycle. Hence, fixing the randomness of such a classical tester, it suffices to detect such collisions in a string $x \in [m]^{c\sqrt{N}}$, for some constant $c > 0$, where the alphabet $[m]$ corresponds to short walks in the graph. A variant of Ambainis’s algorithm can detect this in $O((c\sqrt{N})^{2/3}) = O(N^{1/3})$ queries to x . Each query to x corresponds to an $O(\log N)$ -walk through the graph, so we use $O(N^{1/3} \log N)$ queries to the input graph in total.

In the case of expanders, a short random walk will quickly converge to the uniform distribution. In contrast, for a graph that is far from any expander, such a walk will typically not be very close to uniform. If we sample k times from the uniform distribution over some s -element set, the expected number of collisions is $\binom{k}{2}/s$. In particular, for $k \approx \sqrt{2s}$ we expect to see one collision. In contrast, k samples from a non-uniform distribution give a higher expected number of collisions. Hence if we do $c\sqrt{N}$ short random walks, for some constant c , then the expected number of collisions among the $c\sqrt{N}$ endpoints is likely to be significantly smaller for an expander than for a graph that is far from every expander. Again we use a variation of Ambainis’s algorithm, this time to approximately *count* the number of collisions in an input $x \in [m]^{c\sqrt{N}}$, consisting of the endpoints of the $c\sqrt{N}$ random walks. If this number is too high, we reject. This uses $\tilde{O}(N^{1/3})$ queries to the graph. The technical details are non-trivial, but we will skip them here. \square

Ambainis et al. also proved an $\tilde{\Omega}(N^{1/4})$ quantum lower bound for testing expanders, using the polynomial lower bound method (see Section 2.3.1). They were not able to show $N^{\Omega(1)}$ lower bounds for testing bipartiteness. This all leaves the following very interesting question open:

Question 3. *Is there any graph property which admits an exponential quantum speed-up?*

⁹Equivalently, if there is a constant gap between the first and second eigenvalue of G ’s normalized adjacency matrix. A crucial property of an expander is that the endpoint of a short ($O(\log N)$ -step) random walk starting from any vertex is close to uniformly distributed over $[N]$. We refer to [94] for much more background on expander graphs and their many applications.

2.2.8 Quantum speed-ups for testing group-theoretic properties

Finally, a number of authors have considered quantum testers for properties of groups; we list these here without explaining them in detail.

- Friedl et al. [70] give efficient quantum testers for the property of periodic functions on groups (the testers are even *time-efficient* for Abelian groups), as well as a few other group-theoretic properties. The testers are based on the use of the (Abelian and non-Abelian) quantum Fourier transform.
- Friedl et al. [69] exhibit an efficient (poly(log N , $1/\epsilon$)-query) *classical* tester for the property of $N \times N$ multiplication tables corresponding to N -element Abelian groups, which is based on “dequantizing” a quantum tester. The distance used is the so-called “edit distance.”
- Inui and Le Gall [97], extending [69], exhibit an efficient (poly(log N , $1/\epsilon$)-query) quantum tester for the property of $N \times N$ multiplication tables corresponding to N -element *solvable* groups. In this case, no efficient classical tester is known.
- Le Gall and Yoshida [114] give classical lower bounds on various group testing problems, which in particular demonstrate an exponential separation between the classical and quantum complexities of testing whether the input is an Abelian group generated by k elements (where k is fixed).

2.3 Lower bounds

Here we describe the main methods for obtaining *lower bounds* on the number of queries that quantum property testers need. Most such lower bounds have been obtained using the so-called *polynomial method*, but in principle the stronger *adversary method* can give tight bounds for any property. At the end of this section we also describe an elegant approach for deriving *classical* testing lower bounds from communication complexity, leaving its generalization to lower bounds on quantum testers as an open question.

2.3.1 The polynomial method

The first lower bounds for quantum property testing were proven by Buhrman et al. [47]. They were based on the polynomial method [31], which we now briefly explain. The key property is:

The acceptance probability of a T -query quantum algorithm on input $x \in \{0, 1\}^N$ can be written as an N -variate multilinear polynomial $p(x)$ of degree $\leq 2T$.

This property can be generalized to non-Boolean inputs x , but for simplicity we will assume $x \in \{0, 1\}^N$ in our presentation.

Note that if we have a T -query quantum tester for some property $\mathcal{P} \subseteq \{0, 1\}^N$, then its acceptance probability p is a degree- $2T$ polynomial p such that $p(x) \in [2/3, 1]$ if $x \in \mathcal{P}$; $p(x) \in [0, 1/3]$ if x is far from \mathcal{P} ; and $p(x) \in [0, 1]$ for all other x . The polynomial method derives lower bounds on the query complexity T from lower bounds on the minimal degree of such polynomials.

Our first application of this method is a result which is essentially from [47]. Informally, the result says the following: if we have a property \mathcal{P} such that a (not necessarily uniform) random $x \in \mathcal{P}$ is indistinguishable from a random N -bit string if we only look at up to k bits, then the quantum query complexity of testing \mathcal{P} is $\Omega(k)$.

Theorem 11 (Buhrman et al. [47]). *Let $\mathcal{P} \subseteq \{0, 1\}^N$ be a property such that the number of elements ϵ -close to \mathcal{P} is $< 2^{N-1}$. Let $\mathcal{D} = (p_z)$ be a distribution on $\{0, 1\}^N$ such that $p_z = 0$ for $z \notin \mathcal{P}$, and $\mathbb{E}_{\mathcal{D}}[z_{i_1} \dots z_{i_\ell}] = 2^{-\ell}$ for all choices of $\ell \leq k$ distinct indices $i_1, \dots, i_\ell \in [N]$. Then every quantum ϵ -property tester for \mathcal{P} must make at least $(k + 1)/2$ queries.*

Proof. Suppose there is a quantum algorithm which tests \mathcal{P} using T queries, where $T < (k + 1)/2$. Then by the polynomial method, its acceptance probability is given by a polynomial $p(z)$ of degree at most $2T \leq k$. Intuitively, the reason the theorem holds is that such a degree- k polynomial cannot be correlated with a k -wise independent distribution. To make this precise, assume towards a contradiction that the algorithm has success probability at least $2/3$ on every input z that is in or ϵ -far from \mathcal{P} . Then

$$\mathbb{E}_{z \sim \mathcal{D}}[p(z)] \geq \frac{2}{3}$$

and, letting \mathcal{P}_{close} be the set of z that are ϵ -close to \mathcal{P} , and \mathcal{U} the uniform distribution over $\{0, 1\}^N$, we have

$$\mathbb{E}_{z \sim \mathcal{U}}[p(z)] \leq \frac{|\mathcal{P}_{close}|}{2^N} + \frac{1}{3} \left(1 - \frac{|\mathcal{P}_{close}|}{2^N} \right) < \frac{2}{3},$$

Write $p(z) = \sum_{S \subseteq [N]} \alpha_S m_S(z)$, where m_S is the monomial $\prod_{i \in S} z_i$. We have

$$\mathbb{E}_{z \sim \mathcal{D}}[p(z)] = \sum_{S \subseteq [N]} \alpha_S \mathbb{E}_{z \sim \mathcal{D}}[m_S(z)] = \sum_{S \subseteq [N]} \alpha_S 2^{-|S|} = \sum_{S \subseteq [N]} \alpha_S \mathbb{E}_{z \sim \mathcal{U}}[m_S(z)] = \mathbb{E}_{z \sim \mathcal{U}}[p(z)].$$

We have obtained a contradiction, which completes the proof. \square

A variant of Theorem 11, which generalizes the claim to an underlying set $[m]^N$ ($m > 2$) but does not consider the property testing promise, was independently shown by Kane and Kutin [104]. It is apparently quite hard to satisfy the uniformity constraint of Theorem 11; however, it can sometimes be achieved. For example, consider any property which can be expressed as membership of a linear code $\mathcal{C} \subseteq \mathbb{F}_2^N$. Such a linear code is described as the set $\{Mz : z \in \{0, 1\}^\ell\}$ for some $N \times \ell$ matrix M . A code has *dual distance* d if every codeword c' in the dual code $\mathcal{C}^\perp := \{z : z \cdot c = 0, \forall c \in \mathcal{C}\}$ satisfies $|c'| \geq d$. As Alon et al. [14] observe, it is well-known in coding theory that if \mathcal{C} has dual distance d , then any subset of at most $d - 1$ of the bits of \mathcal{C} are uniformly distributed. As the (easy) proof does not seem easy to find in the recent literature, we include it here.

Theorem 12. [117, Chapter 1, Theorem 10] *Let $\mathcal{C} \subseteq \{0, 1\}^N$ be a code with dual distance d . Then every $k < d$ bits of codewords in \mathcal{C} are uniformly distributed.*

Proof. Dual distance d implies that every set of $k \leq d - 1$ rows in the matrix M are linearly independent (otherwise such a linear combination would imply the existence of a Hamming weight $k < d$ vector z such that $Mz = 0^N$). So for each submatrix M' formed by choosing k rows from M , all the rows of M' are linearly independent, hence the output $M'z$ is uniformly distributed over $\{0, 1\}^k$. \square

Thus, if \mathcal{C} has dual distance d , taking \mathcal{D} to be uniform over \mathcal{C} in Theorem 11 gives an $\Omega(d)$ lower bound on the quantum query complexity of testing membership in \mathcal{C} . A natural example for which this result gives a tight lower bound is the Reed-Muller code $R(d, \ell)$. Each codeword of this code is a binary string of length $N = 2^\ell$ obtained by evaluating a function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$, which can be written as a degree- d polynomial in ℓ variables over \mathbb{F}_2 , at every element $z \in \{0, 1\}^\ell$. $R(d, \ell)$ is known to have dual distance 2^{d+1} [117, Chapter 13], so Theorem 11 implies that any quantum

algorithm testing the set of degree- d polynomials in ℓ variables over \mathbb{F}_2 must make $\Omega(2^d)$ queries. In particular, this means that quantum algorithms obtain no asymptotic speed-up, in terms of their dependence on d , over the best classical algorithm for testing this property [14]. One can generalize this whole argument to derive quantum lower bounds for testing membership of various interesting properties corresponding to codes over \mathbb{F}_q , for $q > 2$; we omit the details. One example of this approach outside of the property-testing setting is [104], which proves bounds on the complexity of quantum interpolation of polynomials. Here the relevant code is the Reed-Solomon code.

Buhrman et al. also applied the polynomial method to show, by a counting argument, that *most* properties do not have an efficient quantum property tester. Informally speaking, there are too many properties, and too few low-degree polynomials.

Theorem 13 (Buhrman et al. [47]). *Let $\mathcal{P} \subset \{0, 1\}^N$ be chosen at random subject to $|\mathcal{P}| = 2^{N/20}$, and fix ϵ to be a small constant. Then, except with probability exponentially small in N , any quantum ϵ -property tester for \mathcal{P} must make $\Omega(N)$ queries.*

A more involved application of the polynomial method is the tight $\Omega(\log N)$ lower bound that Koiran et al. [110] proved for the quantum query complexity of Simon’s problem. With a bit of work, their proof also works to show that the property tester presented in Section 2.2.4 is essentially optimal.

Another highly non-trivial application of the polynomial method is the $\tilde{\Omega}(N^{1/4})$ lower bound of Ambainis et al. [18] for testing the property of a bounded-degree graph being an **Expander** (see Section 2.2.7). Their lower bound is inspired by the one for the collision problem [7], and at a high level works as follows. They give an input distribution \mathcal{D}_ℓ over N -vertex d -regular graphs with ℓ components, obtained from M -vertex graphs that consist of ℓ equal-sized random parts (M is slightly bigger than N and divisible by ℓ ; its role in the proof is rather technical). They then show that the acceptance probability of a T -query quantum tester can be written as an $O(T \log T)$ -degree bivariate polynomial $p(\ell, M)$ in ℓ and M . A random graph of $\ell = 1$ components is very likely to be an expander, so $p(1, M) \approx 1$; on the other hand, every graph with $\ell > 1$ components will be far from an expander, so $p(\ell, M) \approx 0$ for integers $\ell > 1$. They then use results about polynomial approximation to show that such polynomials need degree $\Omega(N^{1/4})$.

2.3.2 The adversary method

The two main lower bound methods that we know for quantum query complexity are the above polynomial method, and the so-called *adversary method*, introduced by Ambainis [15]. For a long time this adversary method faced the so-called “property testing barrier” [96]: for every N -bit partial Boolean function where all 0-inputs are at Hamming distance $\Omega(N)$ from all 1-inputs, the method can prove only a *constant* lower bound on the query complexity. Note that all testing problems for classical properties with respect to Hamming distance fall in this regime, since the 0-inputs are required to be far from all 1-inputs (i.e., elements of the property).

However, Høyer et al. [96] generalized Ambainis’s method to something substantially stronger, which can prove optimal bounds for quantum property testing. We now describe their “negative weights” adversary bound. Let $F : D \rightarrow \{0, 1\}$, with $D \subseteq [m]^N$, be a Boolean function. An *adversary matrix* Γ for F is a real-valued matrix whose rows and columns are indexed by all $x \in D$, satisfying that $\Gamma_{xy} = 0$ whenever $f(x) = f(y)$. Let Δ_j be the Boolean matrix whose rows and columns are indexed by all $x \in D$, such that $\Delta_j[x, y] = 1$ if $x_j \neq y_j$, and $\Delta_j[x, y] = 0$ otherwise.

The (negative-weights) adversary bound for F is given by the following expression:

$$\text{ADV}^\pm(F) = \max_{\Gamma} \frac{\|\Gamma\|}{\max_{j \in [N]} \|\Gamma \circ \Delta_j\|},$$

where Γ ranges over all adversary matrices for F , ‘ \circ ’ denotes entry-wise product of two matrices, and ‘ $\|\cdot\|$ ’ denotes operator norm (largest singular value) of the matrix.¹⁰

Høyer et al. [96] showed that this quantity is indeed a valid lower bound: every quantum algorithm that computes F with error probability $\leq \epsilon$ needs to make at least $\frac{1}{2}(1 - \sqrt{\epsilon(1 - \epsilon)})\text{ADV}^\pm(F)$ queries. Subsequently, Reichardt et al. [139, 115] showed this lower bound is actually essentially tight: for every Boolean function F there is a quantum algorithm computing it with error $\leq 1/3$ using $O(\text{ADV}^\pm(F))$ queries. Since property testing is just a special case of this (the 1-inputs of F are all $x \in \mathcal{P}$, and the 0-inputs are all x that are far from \mathcal{P}), in principle the adversary bound characterizes the quantum complexity of testing classical properties. However, in practice it is often hard to actually calculate the value of $\text{ADV}^\pm(F)$, and we are not aware of good quantum property testing lower bounds that have been obtained using this method.

2.3.3 A communication complexity method?

Recently, a very elegant lower bound method for classical property testing was developed by Blais et al. [36], based on *communication complexity*. In the basic setting of communication complexity [164, 112], two parties (Alice with input x and Bob with input y) try to compute a function $F(x, y)$ that depends on both of their inputs, using as little communication as possible. This is a very well-studied model with many applications, particularly for deriving lower bounds in other areas, such as circuits, data structures, streaming algorithms, and many others (for which see [112]).

Blais et al. [36] showed for the first time how to derive property testing lower bounds from communication complexity. Their idea is to convert a T -query property tester for some property \mathcal{P} into a protocol for some related communication problem F , by showing that 1-inputs (x, y) for F somehow correspond to elements of \mathcal{P} , while 0-inputs (x, y) for F correspond to elements that are far from \mathcal{P} . The more efficient the tester, the less communication the protocol needs. Communication complexity lower bounds for F then imply lower bounds on the complexity T of the tester.

This is best explained by means of an example. A k -linear function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is a linear function that depends on exactly k of its input bits: there exists a weight- k $x \in \{0, 1\}^n$ such that $f(i) = i \cdot x \bmod 2$ for all $i \in \{0, 1\}^n$. Let \mathcal{P} be the set of k -linear functions, and assume k is even. Suppose we have a randomized T -query tester \mathcal{T} for \mathcal{P} . We will show how such a tester induces an efficient communication protocol for the communication complexity problem of deciding whether weight- $k/2$ strings $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^n$ are *disjoint* or not (i.e., whether $x \wedge y = 0^n$). Alice, who received input x , forms the function $f(i) = i \cdot x$ and Bob forms the function $g(i) = i \cdot y$. Consider the function $h(i) = i \cdot (x \oplus y)$. Since $|x \oplus y| = |x| + |y| - 2|x \wedge y|$ and $|x| + |y| = k$, the function h is a $(k - 2|x \wedge y|)$ -linear function. In particular, h is a k -linear function if x and y are disjoint, and $1/2$ -far from any k -linear function if x and y intersect. Now Alice and Bob use a shared random coin to jointly sample one of the deterministic testers that make up the property tester \mathcal{T} . Note that they can simulate a query i to h by 2 bits of communication: Alice sends $i \cdot x$ to Bob and Bob sends $i \cdot y$ to Alice. Hence a T -query tester for \mathcal{P} implies a $2T$ -bit communication protocol for disjointness on weight- $k/2$ inputs x and y . Plugging in the known communication lower bound [103, 138] of

¹⁰Crucially, the adversary matrix Γ may have negative entries. Restricting it to non-negative entries gives one of the many equivalent formulations of Ambainis’s earlier adversary method [150].

$\Omega(k)$ bits for multi-round disjointness on weight- $k/2$ inputs implies that every classical tester for k -linear functions needs $\Omega(k)$ queries, which is nearly tight (the best upper bound is $O(k \log k)$ due to Blais [35]). Plugging in a better $\Omega(k \log k)$ lower bound for *one-way* communication complexity gives $T = \Omega(k \log k)$ for *non-adaptive* classical testers (i.e., testers where the next index to query is independent of the outcomes of the earlier queries), which is tight [63, 48].

Can we use the same idea to prove lower bounds on *quantum* testers? In principle we can, but notice that the overhead when converting a *quantum* tester into a communication protocol is much worse than in the classical case. In the classical case, thanks to the fact that Alice and Bob can use shared randomness to fix a deterministic tester, they both know at each point in the protocol which query i will be made next. Hence they only need to communicate the constant number of bits corresponding to the answer to that query, so the overall communication is $O(T)$. In the quantum case, the queries can be made in superposition, so the conversion will have an overhead of $O(n)$ qubits of communication: each query will be “simulated” by an n -qubit message from Alice to Bob, and another such message from Bob to Alice. More precisely, suppose we let Alice run the T -query quantum tester for \mathcal{P} . Whenever the tester wants to make a query to the function h , its state will be in a superposition $\sum_{i \in \{0,1\}^n} \alpha_i |i\rangle |\phi_i\rangle$ over all indices i , possibly entangled with another register. To perform a phase-query to h , Alice unitarily maps $|i\rangle \mapsto (-1)^{i \cdot x}$, sends the first n qubits of the state to Bob, who unitarily maps $|i\rangle \mapsto (-1)^{i \cdot y}$ and sends back the n qubits. This correctly implements a phase-query to h

$$|i\rangle \mapsto (-1)^{i \cdot x + i \cdot y} = (-1)^{h(i)},$$

on Alice’s state, at the expense of $2n$ qubits of communication. Thus a T -query quantum tester induces a quantum protocol for disjointness that uses $2nT$ qubits of communication. But the best communication lower bound one can hope for on communication complexity problems with n -bit inputs is $\Omega(n)$, which gives only a trivial $T = \Omega(1)$ lower bound! This, however, is not due to a suboptimal reduction: for example, testing k -linear functions can be done with $O(1)$ quantum queries using the Bernstein-Vazirani algorithm, as in Section 2.2.2.

Question 4. *Can some modification of the ideas of Blais et al. [36] be used to obtain non-trivial lower bounds on quantum testers?*

3 Classical testing of quantum properties

In this section we will survey what is known about classical testing of two kinds of quantum objects: implementations of basic unitary operations, and implementations of quantum protocols that win certain two-player games (most famously the “CHSH game”) with high probability. Even though they are testing properties of *quantum* objects, our testers will be classical in the sense that they will base their decision solely on classical data, in particular classical measurement outcomes from feeding classical inputs into the quantum objects.

Before we go there, let us mention that there is another way in which one can consider classical testing of quantum properties: by imagining that we are given classical access to a quantum object which is too large for an efficient classical description. For example, we might be given access to an unknown pure state $|\psi\rangle$ of n qubits by being allowed to query arbitrary amplitudes in the computational basis at unit cost. This then becomes an entirely classical property testing problem. Some natural properties of quantum states in this context have indeed been studied classically; one example is the *Schmidt rank*. A bipartite state $|\psi\rangle$ is said to have Schmidt rank r if it can be

written as $|\psi\rangle = \sum_{i=1}^r \sqrt{\lambda_i} |v_i\rangle |w_i\rangle$ for orthonormal sets of states $\{|v_i\rangle\}$, $\{|w_i\rangle\}$ and non-negative λ_i ; this is known as the Schmidt decomposition of $|\psi\rangle$. A tester for this property follows from work of Krauthgamer and Sasson [111], who have given an efficient tester for low-rank matrices. Their algorithm distinguishes between the case that a $d \times d$ matrix M is rank at most r , and the case that at least an ϵ -fraction of the entries in M must be changed to reduce its rank to r . Their algorithm queries only $O((r/\epsilon)^2)$ elements of the matrix. If we think of M as the amplitudes of a bipartite pure quantum state $|\psi\rangle \in (\mathbb{C}^d)^{\otimes 2}$ (i.e., $M_{ij} = \langle i | \langle j | \psi \rangle$), this is equivalent to a tester for the property of $|\psi\rangle$ having Schmidt rank at most r .

3.1 Self-testing gates

When experimentalists try to implement a quantum computer in the usual circuit model, they will have to faithfully implement a number of basic quantum operations, called elementary *gates*. Suppose we can implement some superoperator¹¹ \mathbf{G} . How can we test whether it indeed implements the gate it is supposed to implement? We are dealing here with the situation of classical testing of quantum properties, which means we can only “trust” classical states; we cannot assume that we have trusted machinery to faithfully prepare specific quantum states. What we *can* do is faithfully prepare an initial computational basis state (i.e., a classical state), apply \mathbf{G} to it a number of times, measure the resulting state in the computational basis, and look at the classical outcomes.

For example, say \mathbf{G} is supposed to implement (conjugation by) the Hadamard gate $H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$. If we prepare $|0\rangle$, apply \mathbf{G} once and measure in the computational basis, the probability to see a 0 should be 1/2. Similarly, if we prepare $|0\rangle$, apply \mathbf{G} *twice* and measure, the probability to see 0 should be 1. These are examples of so-called *experimental equations*. In general, an experimental equation specifies the probability of obtaining a certain outcome from an experiment that starts from a specific classical state and applies a specific sequence of the available superoperators. A *self-tester* for a set of gates repeatedly performs the experiments corresponding to a specific set of experimental equations, in order to verify that the probabilities of the specified outcomes are indeed (close to) what the equations claim. A good self-tester will test experimental equations which (when approximately satisfied by \mathbf{G}), “essentially” tell us what \mathbf{G} is, in a sense made precise below.

It should be noted that such experimental equations cannot *fully* pin down a gate. For example, if \mathbf{G} is the Hadamard gate in a basis where $|1\rangle$ is replaced with $e^{i\phi}|1\rangle$, then no experiment as described above can detect this: H and its cousin satisfy exactly the same experimental equations, and no self-tester is able to distinguish the two. Still, van Dam et al. [62] showed that such experimental equations are surprisingly powerful and can essentially characterize many gate sets, including some universal sets.¹² For concreteness we will focus below on a specific universal set, namely the one consisting of the Hadamard gate H , the $\pi/4$ -phase gate $T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$, and the controlled-NOT operation. This set has the added benefit that it supports fault-tolerant quantum computing: implementing these gates up to small error suffices for universal quantum computing.

Let us first define experimental equations a bit more precisely. Following van Dam et al. [62], we

¹¹Completely positive trace-preserving linear map, a.k.a. “quantum channel.” See Section 5.2 for more on these.

¹²A finite set of gates is *universal* if every n -qubit unitary can be approximated arbitrarily well (in the operator norm) by means of a circuit consisting of these gates. We cannot hope to represent all unitaries *exactly*, because the set of circuits over a finite (or even countable) set of elementary gates is only countable, hence much smaller than the uncountable set of all unitaries.

use $\Pr^c[\rho]$ to denote the probability that measuring the (pure or mixed) state ρ in the computational basis gives outcome c . Then an experimental equation in one superoperator variable \mathbf{G} is of the form

$$\Pr^c[\mathbf{G}^k(|b\rangle\langle b|)] = r,$$

for $b, c \in \{0, 1\}$, positive integer k , and $r \in [0, 1]$. Note that we assume here that we can apply exactly the same superoperator \mathbf{G} more than once. An experimental equation in two variables \mathbf{F} and \mathbf{G} is of the form

$$\Pr^c[\mathbf{F}^{k_1} \mathbf{G}^{\ell_1} \dots \mathbf{F}^{k_t} \mathbf{G}^{\ell_t}(|b\rangle\langle b|)] = r,$$

for $b, c \in \{0, 1\}$, integers $k_1, \dots, k_t, \ell_1, \dots, \ell_t$, and $r \in [0, 1]$ (concatenation of superoperators here denotes composition). We can similarly write experimental equations in more than two operators, and on systems of more than one qubit. Such experimental equations are all the things a self-tester can test.

Suppose one-qubit operators \mathbf{H} and \mathbf{T} are intended to be the Hadamard gate H and the $\pi/4$ -phase gate T , respectively, and two-qubit operator \mathbf{C} is supposed to be CNOT (with slight abuse of notation we identify unitary gates with the corresponding superoperators here). Let us see to what extent we can test this. To start, the following experimental equations are clearly *necessary* for \mathbf{H} :

$$\Pr^0[\mathbf{H}(|0\rangle\langle 0|)] = 1/2$$

$$\Pr^0[\mathbf{H}^2(|0\rangle\langle 0|)] = 1$$

$$\Pr^1[\mathbf{H}^2(|1\rangle\langle 1|)] = 1$$

Van Dam et al. [62, Theorem 4.2] showed that these equations *characterize* the Hadamard gate up to the one remaining degree of freedom that we already mentioned, in the following sense: \mathbf{H} satisfies the above three equations if and only if there exists $\phi \in [0, 2\pi)$ such that \mathbf{H} equals (the superoperator corresponding to) H_ϕ , which is the Hadamard gate where $|1\rangle$ is replaced with $e^{i\phi}|1\rangle$:

$$H_\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & e^{-i\phi} \\ e^{i\phi} & -1 \end{pmatrix}.$$

The unknown phase ϕ cannot be ignored, because it might interact with the effects of other gates.

The following two experimental equations are clearly necessary for \mathbf{T} :

$$\Pr^0[\mathbf{T}(|0\rangle\langle 0|)] = 1$$

$$\Pr^1[\mathbf{T}(|1\rangle\langle 1|)] = 1$$

These two equations are far from sufficient for characterizing the T gate; for example, every diagonal unitary will satisfy these two equations, as would the superoperator that fully decoheres a qubit in the computational basis. However, by introducing some additional equations involving both \mathbf{H} and \mathbf{T} we can do better:

$$\Pr^0[\mathbf{HT}^8\mathbf{H}(|0\rangle\langle 0|)] = 1$$

$$\Pr^0[\mathbf{HTH}(|0\rangle\langle 0|)] = \frac{1}{2}(1 + \cos(\pi/4))$$

Note that if $\mathbf{H} = H$, then both $\mathbf{T} = T$ and its inverse $\mathbf{T} = T^{-1}$ would satisfy the above equations; this is unfortunate, but will turn out below not to matter. Van Dam et al. [62, Theorem 4.4] showed

that a pair of superoperators \mathbf{H} and \mathbf{T} satisfy the above set of 7 equations if and only if there exists $\phi \in [0, 2\pi)$ such that $\mathbf{H} = H_\phi$, and \mathbf{T} corresponds to either T or T^{-1} .

To complete our self-test, consider the superoperator \mathbf{C} . The following experimental equations are clearly necessary for \mathbf{C} to equal CNOT:

$$\begin{aligned}\Pr^{00}[\mathbf{C}(|00\rangle\langle 00|)] &= 1 \\ \Pr^{01}[\mathbf{C}(|01\rangle\langle 01|)] &= 1 \\ \Pr^{11}[\mathbf{C}(|10\rangle\langle 10|)] &= 1 \\ \Pr^{10}[\mathbf{C}(|11\rangle\langle 11|)] &= 1\end{aligned}$$

These equations ensure that \mathbf{C} implements the same permutation of basis states as the CNOT gate. This is still far from sufficient. We add the following experimental equations, which describe the desired interaction between CNOT and H :

$$\begin{aligned}\Pr^{00}[(\mathbf{I} \otimes \mathbf{H})\mathbf{C}(\mathbf{I} \otimes \mathbf{H})(|00\rangle\langle 00|)] &= 1 \\ \Pr^{10}[(\mathbf{I} \otimes \mathbf{H})\mathbf{C}(\mathbf{I} \otimes \mathbf{H})(|10\rangle\langle 10|)] &= 1 \\ \Pr^{00}[(\mathbf{H} \otimes \mathbf{I})\mathbf{C}^2(\mathbf{H} \otimes \mathbf{I})(|00\rangle\langle 00|)] &= 1 \\ \Pr^{01}[(\mathbf{H} \otimes \mathbf{I})\mathbf{C}^2(\mathbf{H} \otimes \mathbf{I})(|01\rangle\langle 01|)] &= 1 \\ \Pr^{00}[(\mathbf{H} \otimes \mathbf{H})\mathbf{C}(\mathbf{H} \otimes \mathbf{H})(|00\rangle\langle 00|)] &= 1\end{aligned}$$

Van Dam et al. [62, Theorem 4.5] showed that if superoperators \mathbf{H} , \mathbf{T} , \mathbf{C} satisfy the above 16 experimental equations, then there exists $\phi \in [0, 2\pi)$ such that:

$$\mathbf{H} = H_\phi; \mathbf{T} = T \text{ or } \mathbf{T} = T^{-1}; \mathbf{C} = C_\phi,$$

where C_ϕ denotes (the superoperator corresponding to the) controlled-NOT gate with $|1\rangle$ replaced with $e^{i\phi}|1\rangle$.

Because our apparatuses are never perfect, we cannot hope to implement the elementary gates exactly. Fortunately, thanks to quantum fault-tolerant computing it suffices if we can implement them up to small error (in fact different applications of the same superoperator can have different errors and need not all be identical). Hence we also cannot expect the gates that we are testing to *exactly* satisfy all of the above experimental equations. Furthermore, even if they did satisfy these equations exactly, we would never be able to perfectly test this with a finite number of experiments. Accordingly, we would like the test consisting of these experimental equations to be *robust*, in the sense that if \mathbf{H} , \mathbf{T} , and \mathbf{C} *approximately* satisfy these equations, then they will be *close* to the gates they purport to be. We say that superoperators ϵ -*satisfy* a set of experimental equations if for each of the equations the left- and right-hand sides differ by at most ϵ . We measure closeness between superoperators in the norm induced by the trace norm:¹³

$$\|G\|_\infty = \sup\{\|G(V)\|_1 : \|V\|_1 = 1\},$$

where the trace norm (Schatten 1-norm) is defined as $\|M\|_1 := \text{tr}(|M|)$.

Van Dam et al. [62, Theorem 6.5, last item] indeed showed that the above equations constitute a robust self-test:

¹³This norm $\|G\|_\infty$ is different from (and weaker than) the diamond norm defined later in Eq. (7), which is also often used to measure distance between superoperators.

Theorem 14 (van Dam et al. [62]). *There exists a constant c such that for all $\epsilon > 0$ the following holds. If superoperators \mathbf{H} , \mathbf{T} , \mathbf{C} ϵ -satisfy the above 16 experimental equations, then there exists $\phi \in [0, 2\pi)$ such that:*

$$\| \mathbf{H} - H_\phi \|_\infty \leq c\sqrt{\epsilon}; \| \mathbf{T} - T \|_\infty \leq c\sqrt{\epsilon} \text{ or } \| \mathbf{T} - T^{-1} \|_\infty \leq c\sqrt{\epsilon}; \| \mathbf{C} - C_\phi \|_\infty \leq c\sqrt{\epsilon}.$$

Let us mention explicitly how this testing of sets of gates fits in the framework outlined in the introduction. The universe now consists of all triples of superoperators $(\mathbf{H}, \mathbf{T}, \mathbf{C})$. The property \mathcal{P} consists of all triples for which there is a ϕ such that $\mathbf{H} = H_\phi$, $\mathbf{T} = T$ or $\mathbf{T} = T^{-1}$, and $\mathbf{C} = C_\phi$. The distance measure would be

$$d((\mathbf{H}, \mathbf{T}, \mathbf{C}), (\mathbf{H}', \mathbf{T}', \mathbf{C}')) = \max (\| \mathbf{H} - \mathbf{H}' \|_\infty, \| \mathbf{T} - \mathbf{T}' \|_\infty, \| \mathbf{C} - \mathbf{C}' \|_\infty).$$

One can derive a tester from Theorem 14 by running the experiments for each experimental equation $O(1/\epsilon)$ times, estimating the probabilities in their right-hand side up to additive error $c\sqrt{\epsilon}$, and accepting if and only if for each of the 16 equations, the estimate is $c\sqrt{\epsilon}$ -close to what it should be. This will accept (with high probability) every triple in \mathcal{P} , and reject (with high probability) every triple that is $2c\sqrt{\epsilon}$ -far from \mathcal{P} .

Each triple $(\mathbf{H}, \mathbf{T}, \mathbf{C})$ that passes the test is a universal (and fault-tolerant) set of elementary gates, so can in principle be used to realize any quantum circuit. The fact that we do not know ϕ is not important when implementing a circuit using this triple of gates: since ϕ cannot be detected by any experimental equations, it cannot affect the classical input-output behavior of a quantum circuit built from these superoperators. We also do not know whether \mathbf{T} approximately equals T or its inverse T^{-1} . Using Hadamard and CNOTs cannot help distinguish these two cases, because they only differ in a minus sign for the imaginary unit (something gates with real entries cannot pick up). However, precisely because such a change is undetectable experimentally, we can just build our circuit assuming \mathbf{T} is close to T ; if it is close to T^{-1} instead, that will incur no observable differences in the input-output behavior of our circuit, so for all intents and purposes we may just assume \mathbf{T} is close to T .

In addition to the above result, van Dam et al. [62] also showed a number of other families of gates to be robustly self-testable, and proved more general robustness results. In follow-up work, Magniez et al. [118] study self-testing of quantum circuits together with measurement apparatuses and sources of EPR-pairs, introducing notions of simulation and equivalence.

3.2 Self-testing protocols

In addition to quantum gates and circuits, a large area of application of quantum self-testing is in multi-party quantum *protocols*. Here typically two or more parties share an entangled state on which they operate locally. In the two-party case these are often EPR-pairs—or at least *should* be EPR-pairs. Experimentalists often need to test that their apparatuses actually produce the required entangled state, or at least something close to it, and that the local operations and measurements act as required. Unless we somehow already have some other trusted quantum objects available, we are in the self-testing regime: like in the previous section, we can only trust preparations of classical states and measurements in the computational basis. We would like to test a quantum object by classically interacting with it, without making assumptions about the measurement apparatuses, the states used, or even the dimension of the Hilbert spaces that are involved.

Again, for concreteness we will focus on testing protocols for one specific example in the two-party setting¹⁴, namely the famous CHSH game [59]. This is defined as follows.

CHSH game. Alice and Bob receive uniformly distributed inputs $x, y \in \{0, 1\}$, respectively. They output $a, b \in \{0, 1\}$, respectively. The players (equivalently, the protocol) *win* the game if and only if the XOR of the outputs equals the AND of the inputs: $a \oplus b = xy$.

Alice and Bob want to coordinate to maximize their probability¹⁵ of winning this game, without communication between them. It is known that classical protocols can win with probability 0.75, but not more, even when they use shared randomness. In contrast, the following quantum protocol P^* wins the game with probability $\cos(\pi/8)^2 \approx 0.854$.¹⁶ It is defined in terms of the four single-qubit *Pauli matrices*, which are

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Standard protocol for CHSH. P^* uses one EPR-pair $|\phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ as starting state. Depending on their inputs, Alice and Bob apply the following specific ± 1 -valued observables¹⁷: Alice measures X if $x = 0$, or Z if $x = 1$. She outputs 0 if her measurement yields 1, and she outputs 1 if it yields -1 . Bob measures the observable $(X + Z)/\sqrt{2}$ if $y = 0$ and $(X - Z)/\sqrt{2}$ if $y = 1$, and outputs 0 or 1 accordingly.

Note that for ± 1 -valued observables A and B , $\langle \phi^+ | A \otimes B | \phi^+ \rangle = \text{tr}(AB^T)/2$ is the difference between the probability that the two output bits are equal and the probability that the outputs are different. If $xy = 0$ a protocol tries to get this difference close to 1, and if $x = y = 1$ it tries to get the difference close to -1 . In the above protocol P^* , the difference is $1/\sqrt{2}$ if $xy = 0$, and $-1/\sqrt{2}$ if $x = y = 1$, so the sum of these 4 terms (negating the last one) equals $2\sqrt{2}$. Tsirelson famously proved that this value of $2\sqrt{2}$ is optimal among all possible protocols [58], no matter how much entanglement they use; hence the corresponding winning probability $\frac{1}{2} + \frac{1}{2\sqrt{2}} = \cos(\pi/8)^2$ is optimal as well.

Theorem 15 (Tsirelson [58]). *Suppose Alice and Bob run a protocol for CHSH that starts with a shared pure state $|\psi\rangle$, where Alice applies ± 1 -valued observables A_0 or A_1 depending on her input x , and Bob applies ± 1 -valued observables B_0 or B_1 depending on y . Then $|\langle \psi | (A_0B_0 + A_0B_1 + A_1B_0 - A_1B_1) | \psi \rangle| \leq 2\sqrt{2}$.*

For simplicity we abbreviate $A \otimes B$ to AB in the above statement as well as the rest of this section (and $A \otimes I$ to just A). The assumption that the starting state is pure and that Alice and Bob apply ± 1 -valued observables is without loss of generality, so Tsirelson’s bound covers all possible quantum protocols.

¹⁴In the three-party setting, the most famous game is the GHZ game [78]. Colbeck [60] seems to have been the first to give a self-testing result for this.

¹⁵This probability is taken over the input distribution as well as over the internal randomness of the protocol.

¹⁶This “Bell inequality violation” has been confirmed by many experiments, albeit with a few remaining experimental “loopholes,” suggesting that Nature does not behave according to classical physics. See the recent survey by Brunner et al. [44] for much more on such “nonlocal” behavior, where two spatially separated entangled players are correlated in ways that are impossible for classical players.

¹⁷A ± 1 -valued observable A can be written as the difference $A = P_+ - P_-$ of two orthogonal projections that satisfy $P_+ + P_- = I$. It corresponds to a projective measurement in a natural way, with outcome $+1$ corresponding to P_+ and outcome -1 corresponding to P_- . Note that such an A is both Hermitian and unitary, and hence $A^2 = I$.

Proof. Define $C = A_0B_0 + A_0B_1 + A_1B_0 - A_1B_1$. Using that $A_x^2 = B_y^2 = I$, the square of C works out to

$$C^2 = 4I + [A_0, A_1] \otimes [B_1, B_0],$$

where $[A, B] = AB - BA$ denotes the *commutator* of two operators. Note that if $\|A\|, \|B\| \leq 1$ then $\|[A, B]\| \leq 2$. Hence, using Cauchy-Schwarz we get:

$$\begin{aligned} |\langle \psi | C | \psi \rangle|^2 &\leq \langle \psi | C^2 | \psi \rangle = 4 + \langle \psi | [A_0, A_1] \otimes [B_1, B_0] | \psi \rangle \\ &= 4 + \langle \psi | ([A_0, A_1] \otimes I) \cdot (I \otimes [B_1, B_0]) | \psi \rangle \\ &\leq 4 + \|[A_0, A_1]\| \cdot \|[B_1, B_0]\| \leq 4 + 2 \cdot 2 = 8, \end{aligned}$$

which implies $|\langle \psi | C | \psi \rangle| \leq 2\sqrt{2}$. □

There are many different protocols that achieve the optimal value $2\sqrt{2}$ or something close to it. For example, applying a local basis change to P^* results in a different protocol that still achieves the maximal value. How much freedom do we have in such optimal or near-optimal protocols for the CHSH game? Surprisingly, this freedom is essentially limited to local basis transformations. Popescu and Rohrlich [136] and Braunstein et al. [41] independently showed that any protocol that wins CHSH with maximal probability needs to start with an EPR-pair, or something that can be turned into an EPR-pair (possibly in tensor product with another state shared between Alice and Bob) using local isometries.¹⁸

However, as in the previous section, *robustness* is important: we expect that if a protocol wins the CHSH game with *close-to-maximal* probability, then its entangled state must be *close* to an EPR-pair, and its measurement operators must be in some sense close to those of the standard protocol. Such a robust result was proved independently in [123, 124]¹⁹:

Theorem 16 ([123, 124]). *Suppose Alice and Bob run a protocol for CHSH that starts with a shared pure state $|\psi\rangle$, where Alice applies ± 1 -valued observables A_0 or A_1 depending on her input x , and Bob applies ± 1 -valued observables B_0 or B_1 depending on y . Suppose the protocol wins CHSH with probability at least $\cos(\pi/8)^2 - \epsilon$. Define new operators for Alice and Bob, respectively:*

$$\begin{aligned} X'_A &= A_0, & Z'_A &= A_1, \\ X'_B &= \frac{B_0 + B_1}{\sqrt{2}}, & Z'_B &= \frac{B_0 - B_1}{\sqrt{2}}. \end{aligned}$$

Then there exists a local isometry $\Phi = \Phi_A \otimes \Phi_B$ and a pure state $|junk\rangle$ shared between Alice and Bob, such that for all $M, N \in \{I, X, Z\}$ we have

$$\|\Phi(M'_A N'_B |\psi\rangle) - |junk\rangle \otimes M_A N_B |\phi^+\rangle\| = O(\sqrt{\epsilon}),$$

where, e.g., if $M = X$ the notation M'_A denotes the operator X'_A .

In words, up to a local basis change and small errors depending on ϵ , $|\psi\rangle$ behaves like an EPR-pair and X'_A, Z'_A, X'_B, Z'_B behave like the standard Pauli operators X and Z for Alice and Bob, respectively, applied to that EPR-pair. Note that this also implies that A_0, A_1, B_0, B_1 behave like the observables of the standard protocol P^* . We give the proof of [123] here for the special case where $\epsilon = 0$. This allows us to describe the main ideas, without going into the technical but straightforward details needed to keep track of the errors and approximations.

¹⁸The correct attribution of this result is not completely clear, see also the work of Summers and Werner [151] and Tsirelson [152, p. 11].

¹⁹The earlier work of Mayers and Yao [120, 121] that started the area of self-testing of quantum states also had a protocol for robustly self-testing EPR-pairs, albeit based on more than the CHSH game.

Proof for $\epsilon = 0$. Consider the proof of Tsirelson’s bound (Theorem 15). If a protocol achieves the maximum value $2\sqrt{2}$, then the inequalities in the proof must be equalities. This implies $|\langle \psi|[A_0, A_1]|\psi\rangle| = 2$ and hence A_0 and A_1 anti-commute on $|\psi\rangle$:

$$A_0 A_1 |\psi\rangle = -A_1 A_0 |\psi\rangle.$$

Similarly $|\langle \psi|[B_1, B_0]|\psi\rangle| = 2$, and hence B_0 and B_1 anti-commute on $|\psi\rangle$ as well:

$$B_0 B_1 |\psi\rangle = -B_1 B_0 |\psi\rangle.$$

We list some properties of the operators X'_A, Z'_A, X'_B, Z'_B that were defined in the statement of the theorem. All are clearly Hermitian. On Alice’s side, X'_A and Z'_A are unitary because A_0 and A_1 are. They anti-commute on $|\psi\rangle$ because A_0 and A_1 do. On Bob’s side, X'_B and Z'_B anti-commute. We cannot assume X'_B and Z'_B are unitary. However, since $(X'_B)^2 = I + (B_0 B_1 + B_1 B_0)/2$ and B_0 and B_1 anti-commute on $|\psi\rangle$, we have $(X'_B)^2 |\psi\rangle = |\psi\rangle$. Hence $\|X'_B |\psi\rangle\|^2 = \langle \psi|(X'_B)^2 |\psi\rangle = 1$, so X'_B preserves the norm of $|\psi\rangle$. Similarly, Z'_B preserves the norm of $|\psi\rangle$, as does $X'_B Z'_B$.

We now want to show that $X'_A X'_B |\psi\rangle = |\psi\rangle$. First,

$$\langle \psi|A_0(B_0 + B_1)|\psi\rangle + \langle \psi|A_1(B_0 - B_1)|\psi\rangle = \langle \psi|C|\psi\rangle = 2\sqrt{2}. \quad (1)$$

Second, by squaring the operator $A_0(B_0 + B_1)$ and using anti-commutativity of B_0 and B_1 on $|\psi\rangle$ we can show $\langle \psi|A_0(B_0 + B_1)|\psi\rangle \leq \sqrt{2}$, and similarly $\langle \psi|A_1(B_0 - B_1)|\psi\rangle \leq \sqrt{2}$. Combining with Eq. 1, it follows that both terms *equal* $\sqrt{2}$. Then we have

$$\langle \psi|X'_A X'_B |\psi\rangle = \frac{1}{\sqrt{2}} \langle \psi|A_0(B_0 + B_1)|\psi\rangle = 1,$$

hence $X'_A X'_B |\psi\rangle = |\psi\rangle$. Since X'_A is unitary and Hermitian it is self-inverse, which implies $X'_A |\psi\rangle = X'_B |\psi\rangle$. A similar argument shows $Z'_A Z'_B |\psi\rangle = |\psi\rangle$ and $Z'_A |\psi\rangle = Z'_B |\psi\rangle$.

We now need to show that, after a local isometry, $|\psi\rangle$ behaves like an EPR-pair (tensored with some “junk” state) and X'_A, Z'_A, X'_B, Z'_B behave like X_A, Z_A, X_B, Z_B . Consider the dimension-increasing map on states $|\phi\rangle$ (in the same space as $|\psi\rangle$) that is described by Figure 1. It adds one auxiliary qubit for Alice (at the top line of the figure) and one for Bob (at the bottom), both initially $|0\rangle$. Because all operators involved preserve norm on all states involved, this can be extended to a local isometry $\Phi = \Phi_A \otimes \Phi_B$.




Figure 1: Isometry for transforming a perfect CHSH protocol to the standard one.

For convenience we will write the two auxiliary qubits on the right of the state, the first for Alice and the second for Bob. Let $M, N \in \{I, X, Z\}$. Following the state through the different

steps of Figure 1, a straightforward calculation shows

$$\begin{aligned}
\Phi(M'_A N'_B |\psi\rangle) &= \frac{1}{4}(I + Z'_A)(I + Z'_B)M'_A N'_B |\psi\rangle |00\rangle \\
&+ \frac{1}{4}X'_B(I + Z'_A)(I - Z'_B)M'_A N'_B |\psi\rangle |01\rangle \\
&+ \frac{1}{4}X'_A(I - Z'_A)(I + Z'_B)M'_A N'_B |\psi\rangle |10\rangle \\
&+ \frac{1}{4}X'_A X'_B(I - Z'_A)(I - Z'_B)M'_A N'_B |\psi\rangle |11\rangle.
\end{aligned} \tag{2}$$

First consider the case where $M = N = I$. Then the second term vanishes, because $I|\psi\rangle = Z'_A Z'_B |\psi\rangle$ and $Z'_A |\psi\rangle = Z'_B |\psi\rangle$. Similarly the third term vanishes. The fourth term equals the first (except in the last two qubits) because $X'_A X'_B(I - Z'_A)(I - Z'_B)|\psi\rangle = (I + Z'_A)(I + Z'_B)X'_A X'_B |\psi\rangle$ by anti-commutativity, and $X'_A X'_B |\psi\rangle = |\psi\rangle$. Hence we end up with

$$\Phi(|\psi\rangle) = \left(\frac{1}{2\sqrt{2}}(I + Z'_A)(I + Z'_B)|\psi\rangle \right) \otimes \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) = |\text{junk}\rangle \otimes |\phi^+\rangle,$$

where we defined $|\text{junk}\rangle := \frac{1}{2\sqrt{2}}(I + Z'_A)(I + Z'_B)|\psi\rangle$.

If $MN = XX$ then the same proof applies, because $X'_A X'_B |\psi\rangle = |\psi\rangle$ and $X_A X_B |\phi^+\rangle = |\phi^+\rangle$. The same holds if $MN = ZZ$.

Now consider the case $MN = XZ$. Looking at Eq. (2), the first term vanishes because

$$(I + Z'_A)(I + Z'_B)X'_A Z'_B |\psi\rangle = X'_A(I - Z'_A)(I + Z'_B)|\psi\rangle = 0,$$

using the anti-commutativity of X'_A and Z'_A , and the fact that $(I + Z'_B)Z'_B |\psi\rangle = (I + Z'_B)|\psi\rangle$ (because $(Z'_B)^2 |\psi\rangle = I|\psi\rangle$). Similarly the fourth term vanishes. For the second term we use

$$\begin{aligned}
X'_B(I + Z'_A)(I - Z'_B)X'_A Z'_B |\psi\rangle &= (I + Z'_A)(I + Z'_B)X'_A X'_B Z'_B |\psi\rangle \\
&= -(I + Z'_A)(I + Z'_B)X'_A X'_B |\psi\rangle = -(I + Z'_A)(I + Z'_B)|\psi\rangle,
\end{aligned}$$

where we used $X'_B Z'_B |\psi\rangle = -Z'_B X'_B |\psi\rangle$, $X'_A X'_B |\psi\rangle = |\psi\rangle$, and $(I + Z'_B)Z'_B |\psi\rangle = (I + Z'_B)|\psi\rangle$. We similarly analyze the third term. We end up with

$$\begin{aligned}
\Phi(X'_A Z'_B |\psi\rangle) &= -\frac{1}{4}(I + Z'_A)(I + Z'_B)|\psi\rangle |01\rangle + \frac{1}{4}(I + Z'_A)(I + Z'_B)|\psi\rangle |10\rangle \\
&= \frac{1}{2\sqrt{2}}(I + Z'_A)(I + Z'_B)|\psi\rangle \otimes \frac{1}{\sqrt{2}}(|10\rangle - |01\rangle) = |\text{junk}\rangle \otimes X_A Z_B |\phi^+\rangle.
\end{aligned}$$

For the other five possible M, N pairs, a similar calculation (starting from Eq. (2) and using the known commutation and anti-commutation properties) works to establish the desired property:

$$\Phi(M'_A N'_B |\psi\rangle) = |\text{junk}\rangle \otimes M_A N_B |\phi^+\rangle.$$

□

Accordingly, we can use this robust result to test whether a given protocol behaves essentially like P^* , based only on classical-input output behavior: run it multiple times on uniformly distributed classical input bits, observe the classical output bits, and see if the winning probability is close to the optimal value $\cos(\pi/8)^2$. If so, then (up to local change of basis) the state must be

close to an EPR-pair tensored with some other “junk” state, and the behavior of the measurements must be close to the ones of the standard CHSH protocol P^* .

There has been a lot more work along these lines. McKague et al. [123] give a more general framework for bipartite robust self-testing that subsumes the CHSH inequality, the Mayers-Yao self-test (simplifying [118]), as well as others. Yang and Navascués [161] give robust self-tests for any entangled two-qubit states, not just maximally entangled ones; the noise-resistance was further improved in [162]. McKague [122] and Miller and Shi [124] give results about self-testing of states shared by more than two parties.

In some applications one needs to have many states that all behave like EPR-pairs, not just the one EPR-pair that is needed for an optimal protocol for CHSH. Recently, Reichardt et al. [140] proved a subtle robustness result for playing *many instances* of CHSH. Roughly, their result says: if a quantum protocol wins a fraction of nearly $\cos(\pi/8)^2$ of a sequence of k given instances of the CHSH game, then most blocks of $m = k^{\Omega(1)}$ instances have the property that they start “essentially” (again, up to local operations and small differences like in Theorem 16) from m EPR-pairs and run m independent instances of the standard protocol P^* . With significant additional work it is possible to use this result to devise methods that allow a classical system to “command” an untrusted quantum system, in the sense of forcing that quantum system to either use essentially the states and operations you want it to use, or be detected if it deviates too much from those states and operations. Such control enables various kinds of device-independent quantum cryptography, as well as the ability to offload general quantum computation to untrusted devices.

4 Quantum testing of quantum properties: States

In the third part of this survey we discuss quantum testers for quantum properties. The first decision we have to take in this setting is how the quantum object which we wish to test is presented to us. The two options are a quantum presentation (i.e., we are given access to the object as a black box, which can be used in a quantum algorithm), or a classical presentation (i.e., we are given an efficient classical description of the object, such as a quantum circuit). We concentrate on the former option (Sections 4–5), as this seems to be the most natural generalization of ideas from classical property testing. However, in Section 6 we also discuss the latter option, which turns out to be important in quantum computational complexity.

Our focus in this part of the survey is on quantum tests for quantum properties which generalize the idea of classical property testing. That is, tests which are designed to distinguish quantum states (or operations) with some property from those far from having that property, given access to the state (or operation) as a black box. We also mention here two related and well-studied areas elsewhere in quantum information theory. The first is quantum state discrimination, which can be seen as a quantum generalization of classical hypothesis testing. The archetypal problem in this setting is as follows: given the ability to create copies of an unknown quantum state ρ picked from a known set S of quantum states, identify ρ with minimal probability of error. Some authors use the term “quantum hypothesis testing” for this problem [53]; others reserve this term for the case $|S| = 2$, where precise results have been obtained relating the optimal error probability to the number of copies of ρ consumed, and trade-offs between different kinds of error have been determined [22]. See the surveys [26, 53] for detailed reviews of quantum state discrimination. The second area is the question of directly estimating some quantity of interest about a completely unknown quantum state ρ , given access to multiple copies of the state, without performing full tomography. Results of this form include direct estimation of the spectrum of ρ [106], estimation of

polynomials in the entries of ρ [43], and estimation of quantities related to entanglement (e.g., [81]).

We begin our discussion of quantum properties by considering properties of quantum states, first pure states and then mixed states.

4.1 Pure states

A pure state $|\psi\rangle$ of a d -dimensional quantum system is described by a d -dimensional complex unit vector (technically, a ray; that is, $e^{i\theta}|\psi\rangle$ is equivalent to $|\psi\rangle$ for all real θ). A property of d -dimensional pure quantum states is therefore a set $\mathcal{P} \subseteq \mathbb{C}^d$. One can naturally generalize this to properties of pairs of quantum states, where $\mathcal{P} \subseteq \mathbb{C}^d \times \mathbb{C}^d$, etc.

There is a natural measure of distance between quantum states $|\psi\rangle$ and $|\phi\rangle$: the trace distance

$$D(|\psi\rangle, |\phi\rangle) := \frac{1}{2} \|\ |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_1 = \sqrt{1 - |\langle\psi|\phi\rangle|^2}. \quad (3)$$

Here, as in Section 3.1, $\|\cdot\|_1$ is the trace norm (Schatten 1-norm) $\|M\|_1 := \text{tr}(|M|)$. Given a state promised to be either $|\psi\rangle$ or $|\phi\rangle$, with equal probability of each, the optimal probability of determining via a measurement which state we have is exactly $(1 + D(|\psi\rangle, |\phi\rangle))/2$ [91, 129]. We therefore say that $|\psi\rangle$ is ϵ -close to having property \mathcal{P} if

$$D(|\psi\rangle, \mathcal{P}) := \inf_{|\phi\rangle \in \mathcal{P}} D(|\psi\rangle, |\phi\rangle) \leq \epsilon,$$

and similarly that $|\psi\rangle$ is ϵ -far from having property \mathcal{P} if $D(|\psi\rangle, \mathcal{P}) \geq \epsilon$. If $|\psi\rangle$ is ϵ -close to having property \mathcal{P} , there is no hope of certifying that $|\psi\rangle \notin \mathcal{P}$ with worst-case bias larger than ϵ , given access to only one copy of $|\psi\rangle$.

The complexity of algorithms for testing pure quantum states is measured by the number of copies of the test state $|\psi\rangle$ required to distinguish between the two cases that (a) $|\psi\rangle \in \mathcal{P}$, or (b) $|\psi\rangle$ is ϵ -far away from having property \mathcal{P} . We therefore say that \mathcal{P} can be ϵ -tested with q copies if there exists a quantum algorithm which uses q copies of the input state to distinguish between these two cases, and fails with probability at most $1/3$ on any input. As with classical property testers, we say that a tester has *perfect completeness* if it accepts every state in \mathcal{P} with certainty. Crucially, we look for algorithms where the number of copies used scales only in terms of ϵ , and there is no dependence on the dimension d , making this a fair analog of the classical concept. If we cannot find such an algorithm, we attempt to minimize the dependence on d .

On the other hand, if we do not care about the dependence on d , *any* (even infinite) property $\mathcal{P} \subseteq \mathbb{C}^d$ can be tested using $O(d/\epsilon^2)$ copies of the input state $|\psi\rangle$; it suffices to obtain an estimate $|\psi'\rangle$ such that $D(|\psi'\rangle, |\psi\rangle) < \epsilon/2$, and accept if and only if $D(|\psi'\rangle, \mathcal{P}) \leq \epsilon/2$. In order to produce such an estimate one can use a procedure known as *quantum state estimation*, which needs $O(d/\epsilon^2)$ copies of $|\psi\rangle$ to achieve the required accuracy with success probability at least $2/3$ [45].

4.1.1 Equality

The first property we consider is extremely basic, but a useful building block for more complicated protocols: whether the input state is equal to some fixed state. We say that a state $|\psi\rangle$ satisfies the **Equality to $|\phi\rangle$** property if $|\psi\rangle = e^{i\theta}|\phi\rangle$ for some real θ , so $\mathcal{P} = \{e^{i\theta}|\phi\rangle : \theta \in \mathbb{R}\}$; it is necessary to allow an arbitrary phase θ in the definition of this property, as $|\psi\rangle$ cannot be distinguished from $e^{i\theta}|\psi\rangle$ by any measurement. A natural test for **Equality to $|\phi\rangle$** is simply to perform the measurement $\{|\phi\rangle\langle\phi|, I - |\phi\rangle\langle\phi|\}$ on $|\psi\rangle$, and accept if and only if the first outcome is obtained. The

probability of acceptance is precisely $|\langle\psi|\phi\rangle|^2$, so if $|\psi\rangle$ satisfies the property, the test accepts with certainty. On the other hand, if $D(|\psi\rangle, |\phi\rangle) = \epsilon$, the test rejects with probability $1 - |\langle\psi|\phi\rangle|^2 = \epsilon^2$. Via repetition, we find that for any $|\phi\rangle$, **Equality to $|\phi\rangle$** can be tested with $O(1/\epsilon^2)$ copies.

A matching lower bound follows from considering the special case where the input state is promised to be either $|\phi\rangle$ or some state $|\phi'\rangle$ such that $D(|\phi\rangle, |\phi'\rangle) = \epsilon$, with equal probability of each. Then any test which uses k copies to test whether the input is equal to $|\phi\rangle$ is equivalent to a procedure which discriminates between $|\phi\rangle^{\otimes k}$ and $|\phi'\rangle^{\otimes k}$, which has success probability upper-bounded by $(1 + D(|\psi\rangle^{\otimes k}, |\phi\rangle^{\otimes k}))/2$. Using the definition (3) of the trace distance, we require $k = \Omega(1/\epsilon^2)$ to achieve success probability $2/3$. This same argument in fact shows that *any* non-trivial property of pure states requires $\Omega(1/\epsilon^2)$ copies to be tested.

We remark that testing equality to a fixed state immediately generalizes to the problem of testing whether $|\psi\rangle \in \mathbb{C}^d$ is contained in a known subspace $S \subseteq \mathbb{C}^d$. Here the prescription is to perform the measurement $\{\Pi_S, I - \Pi_S\}$ $O(1/\epsilon^2)$ times, where Π_S is the projector onto S , and accept if and only if the first outcome is obtained every time. For example, this allows the property **Permutation Invariance** to be tested efficiently, where $|\psi\rangle \in (\mathbb{C}^d)^{\otimes n}$ satisfies the property if it is invariant under any permutation of the n subsystems. As $|\psi\rangle$ is permutation-invariant if and only if it is contained in the symmetric subspace of $(\mathbb{C}^d)^{\otimes n}$, projecting onto this subspace gives an efficient test for this property. This procedure, which is known as *symmetrization*, has been studied in the context of quantum fault-tolerance and can be performed efficiently [25]; see Section 4.2.2 below for a description of how this can be achieved via the powerful primitive of generalized phase estimation.

Another immediate generalization of **Equality to $|\phi\rangle$** is the question of testing whether two *unknown* states are the same. We say that a pair of states $|\psi\rangle, |\phi\rangle$ satisfies the **Equality** property if $|\phi\rangle = e^{i\theta}|\psi\rangle$ for some real θ , so now the property is $\mathcal{P} = \{(|\psi\rangle, e^{i\theta}|\psi\rangle) : |\phi\rangle \text{ is a pure state, } \theta \in \mathbb{R}\}$. In order to test this property, we will use a simple but important procedure known as the *swap test*. This was used by Buhrman et al. [46] to demonstrate an exponential separation between the quantum and classical models of simultaneous message passing (SMP) communication complexity, and has since become a standard tool in quantum algorithm design. In the test, we take two (possibly mixed²⁰) states ρ, σ as input and attach an ancilla qubit in state $|0\rangle$. We then apply a Hadamard gate to the ancilla, followed by a controlled-SWAP gate (controlled on the ancilla), and another Hadamard gate. We then measure the ancilla qubit and accept if the answer is 0. This procedure is illustrated by the circuit in Figure 2.




Figure 2: The swap test.

One can show [46, 109] that the swap test accepts with probability

$$\frac{1}{2} + \frac{1}{2}\text{tr}(\rho\sigma),$$

which for pure states $|\psi\rangle, |\phi\rangle$ is equal to $(1 + |\langle\psi|\phi\rangle|^2)/2 = 1 - D(|\psi\rangle, |\phi\rangle)^2/2$. In particular, if this test is applied to two pure states which satisfy the **Equality** property then the test accepts with

²⁰See Section 4.2 for more about mixed states and a formal definition.

certainty. On the other hand, if the states are ϵ -far away from equal, then by definition

$$\inf_{|\xi\rangle} D(|\psi\rangle|\phi\rangle, |\xi\rangle^{\otimes 2}) \geq \epsilon.$$

But

$$\inf_{|\xi\rangle} D(|\psi\rangle|\phi\rangle, |\xi\rangle^{\otimes 2}) = \sqrt{1 - \sup_{|\xi\rangle} |\langle\psi|\xi\rangle\langle\phi|\xi\rangle|^2} \leq \sqrt{1 - |\langle\psi|\phi\rangle|^2} = D(|\psi\rangle, |\phi\rangle),$$

where the inequality follows by taking $|\xi\rangle = |\phi\rangle$. Thus the test rejects with probability at least $\epsilon^2/2$, so $O(1/\epsilon^2)$ repetitions suffice to detect states ϵ -far away from equal with constant probability; in other words, **Equality** can be tested with $O(1/\epsilon^2)$ copies. The swap test is in fact optimal among all testers for this property which have perfect completeness and use one copy of each of the input states. To see this, observe that the swap test is precisely the operation of projecting onto the symmetric subspace of $(\mathbb{C}^d)^{\otimes 2}$. Any tester which accepts every pair of equal states $|\psi\rangle^{\otimes 2}$ must accept every state in this subspace, so the swap test is the most refined test of this type. One can generalize this to prove that the swap test is also optimal among tests which are allowed two-sided error, in the sense that it achieves the largest possible gap between the acceptance probabilities in equal and orthogonal instances [102].

The property of **Equality** can be generalized further, to the question of testing whether n pure states $|\psi_1\rangle, \dots, |\psi_n\rangle$ are all equal. The natural tester for this property, generalizing the swap test, is to project onto the symmetric subspace of $(\mathbb{C}^d)^{\otimes n}$, i.e., to perform symmetrization [25]. Kada et al. [102] have studied this procedure under the name of the *permutation test*, and show that the test accepts n -tuples where at least one pair of states is orthogonal with probability at most $1/n$, and that this is optimal among tests with perfect completeness. No explicit bounds appear to be known on this tester's parameters if the promise is relaxed, for example to specify that at least one pair of states has overlap at most ϵ . Kada et al. also study a related tester, called the *circle test*, and prove that this tester is also optimal for prime n [102]. This procedure is somewhat simpler as it only involves taking a quantum Fourier transform over \mathbb{Z}_n , rather than S_n .

4.1.2 Productness

A pure state $|\psi\rangle \in (\mathbb{C}^d)^{\otimes n}$ of n d -dimensional subsystems is said to be *product* (i.e., satisfy the **Product** property) if it can be written as a tensor product $|\psi\rangle = |\psi_1\rangle|\psi_2\rangle \dots |\psi_n\rangle$ for some local states $|\psi_1\rangle, \dots, |\psi_n\rangle \in \mathbb{C}^d$. A state which is not product is called *entangled*. Entanglement is a ubiquitous phenomenon in quantum information theory (see for example [95] for an extensive review), so the property of being a product state is an obvious target to test.

Given just one copy of $|\psi\rangle$, our ability to test whether it is product is very limited. Indeed, as every quantum state can be written as a linear combination of product states, any tester which accepts all product states with certainty must accept all states with certainty. However, if we are given two copies of $|\psi\rangle$, there are non-trivial tests we can perform. In particular, consider the following procedure, which was first discussed by Mintert et al. [125] and is called the *product test* [88]: apply the swap test across each corresponding pair of subsystems of $|\psi\rangle^{\otimes 2}$, and accept if and only if all of the tests accept. The overall procedure is illustrated in Figure 3.

If $|\psi\rangle$ is indeed product, then all of the swap tests will accept. On the other hand, if $|\psi\rangle$ is far from product, the intuition is that the entanglement in $|\psi\rangle$ will cause at least some of the tests to reject with fairly high probability. This intuition can be formalized to give the following result.

Theorem 17 (Harrow and Montanaro [88]). *If $|\psi\rangle$ is ϵ -far from product, the product test rejects with probability $\Omega(\epsilon^2)$.*




Figure 3: Schematic of the product test applied to an n -partite state $|\psi\rangle$. The swap test (vertical boxes) is applied to the n pairs of corresponding subsystems of two copies of $|\psi\rangle$ (horizontal boxes).

Thus the property of productness can be tested with $O(1/\epsilon^2)$ copies. We will not give the full, and somewhat technical, proof of Theorem 17 here, but merely sketch the proof technique; see [88] for details.

Proof sketch. Let $P_{\text{test}}(|\psi\rangle)$ denote the probability of the product test accepting when applied to two copies of $|\psi\rangle$, and let the distance of $|\psi\rangle$ from the nearest product state be ϵ . The proof is split into two parts, depending on whether ϵ is low or high. For $S \subseteq [n]$, let ψ_S be the mixed state obtained by tracing out (discarding) the qubits not in S . Then the starting point is the observation that

$$P_{\text{test}}(|\psi\rangle) = \frac{1}{2^n} \sum_{S \subseteq [n]} \text{tr}(\psi_S^2). \quad (4)$$

The quantity $\text{tr}(\psi_S^2)$ measures the purity of the reduced state ψ_S , which can be seen as a measure of the entanglement of $|\psi\rangle$ across the bipartition (S, S^c) ; if $|\psi\rangle$ were product across this bipartition, ψ_S would be pure and $\text{tr}(\psi_S^2)$ would equal 1. By (4), the probability that the test passes is equal to the average purity of the reduced state obtained by a random bipartition of the n systems. Writing $|\psi\rangle = \sqrt{1 - \epsilon^2}|0^n\rangle + \epsilon|\phi\rangle$ (without loss of generality), for some product state $|0^n\rangle$ and arbitrary orthogonal state $|\phi\rangle$, Eq. (4) allows an explicit expression for $\text{tr}(\psi_S^2)$ in terms of ϵ and $|\phi\rangle$ to be obtained. Expanding $|\phi\rangle = \sum_{x \in \{0, \dots, d-1\}^n} \alpha_x |x\rangle$ and summing over S , we get an expression containing terms of the form $\sum_{x \in \{0, \dots, d-1\}^n} |\alpha_x|^2 c^{|x|}$ for some $c < 1$, where $|x| := |\{i : x_i \neq 0\}|$. In order to obtain a non-trivial bound from this, the final step of the first part of the proof is to use the fact that $|0^n\rangle$ is the closest product state to $|\psi\rangle$ to argue that $|\phi\rangle$ cannot have any amplitude on basis states $|x\rangle$ such that $|x| \leq 1$. A bound is eventually obtained that is applicable when ϵ is small, namely that

$$P_{\text{test}}(|\psi\rangle) \leq 1 - \epsilon^2 + \epsilon^3 + \epsilon^4.$$

In the case where ϵ is large, this does not yet give a useful upper bound, so the second part of the proof finds a constant upper bound on $P_{\text{test}}(|\psi\rangle)$. This quantity can be shown to be upper bounded by the probability that a relaxed test, for being product across some partition of the n subsystems into $k \leq n$ parties, passes. If $|\psi\rangle$ is far from product across the n subsystems, the proof shows that one can find a partition into k parties (for some $k \leq n$) such that the distance from the closest product state (with respect to this partition) falls into the regime where the first part of the proof works. The eventual result is that if $\epsilon^2 \geq 11/32 > 0.343$, then $P_{\text{test}}(|\psi\rangle) \leq 501/512 < 0.979$; combining these two bounds completes the proof. \square

We mention two implications of Theorem 17. First, by the characterization (4), the content of Theorem 17 can be understood as: if a pure state of n systems is still fairly pure on average after discarding a random subset of the systems, it must in fact have been close to a product state in the

first place. In the classical property testing literature, one of the motivations for analysing tests for combinatorial properties is to obtain some insight into the structure of the property being tested; Theorem 17 can be seen as achieving something similar in a quantum setting.

Second, by allowing one to efficiently certify productness given two copies of $|\psi\rangle$, the product test can be used to show that quantum Merlin-Arthur proof systems with multiple provers can be simulated efficiently by two provers, or in complexity-theoretic terminology that $\text{QMA}(k) = \text{QMA}(2)$ [88]. Roughly speaking, to simulate a k -Merlin protocol, one can simply ask two Merlins to provide identical copies of the k -Merlin proofs, and perform the product test to ensure that they are indeed product states. Since the product test uses only two copies of the state, two Merlins suffice. Via a previous result of Aaronson et al. [5] giving a multiple-prover quantum proof system for 3-SAT, this in turn allows one to prove hardness of various tasks in quantum information theory, conditioned on the hardness of 3-SAT [88]. This is again analogous to the classical literature, where efficient property testers are used as components in hardness-of-approximation results.

Although the product test itself is natural, the detailed proof of Theorem 17 given in [88] is a lengthy case analysis which does not provide much intuition and gives suboptimal constants. For example, the lower bound obtained on the probability of the product test rejecting does not increase monotonically with ϵ , which presumably should be the case for an optimal bound. We therefore highlight the following open question.

Question 5. *Can the analysis of the product test be improved?*

4.1.3 Arbitrary finite sets

The following algorithm of Wang [156] gives a tester for *any* finite property $\mathcal{P} \subset \mathbb{C}^d$ (this is similar to the result for any finite *classical* property mentioned at the end of Section 2.2.2). The tester cannot necessarily be implemented time-efficiently in general. Given access to copies of an input state $|\psi\rangle$, the tester proceeds as follows:

1. Create the state $|\psi\rangle^{\otimes T}$, for some T to be determined.
2. Let $S = \text{span}\{|\phi\rangle^{\otimes T} : |\phi\rangle \in \mathcal{P}\}$. Perform the measurement $\{\Pi_S, I - \Pi_S\}$, where Π_S is the projector onto S , and accept if the first outcome is obtained. Otherwise, reject.

Theorem 18 (Wang [156]). *Let $\mathcal{P} \subset \mathbb{C}^d$ be such that $\min_{|\phi\rangle \neq |\phi'\rangle \in \mathcal{P}} D(|\phi\rangle, |\phi'\rangle) = \delta$. Then it suffices to take $T = O(\log |\mathcal{P}| \max\{\epsilon^{-2}, \delta^{-2}\})$ to obtain a tester which accepts every state in \mathcal{P} with certainty, and rejects every state $|\psi\rangle$ such that $D(|\psi\rangle, \mathcal{P}) \geq \epsilon$ with probability at least $2/3$.*

Observe that the dependence on $|\mathcal{P}|$ is only logarithmic. The intuition behind Theorem 18 is that, if all the states in \mathcal{P} have large pairwise distances, $\{|\phi\rangle^{\otimes T}\}$ is an approximately orthonormal basis for S , so if $|\psi\rangle$ is ϵ -far from \mathcal{P} , the probability of incorrectly accepting is

$$\langle \psi |^{\otimes T} \Pi_S | \psi \rangle^{\otimes T} \approx \sum_{|\phi\rangle \in \mathcal{P}} |\langle \psi | \phi \rangle|^{2T} \leq |\mathcal{P}|(1 - \epsilon^2)^T,$$

which is sufficiently small when $T = O((\log |\mathcal{P}|)/\epsilon^2)$. Wang describes an application of Theorem 18 to testing the set of permutations of n qubits using $O((n \log n)/\epsilon^2)$ copies [156]. However, the dependence of the theorem on δ seems to limit its applicability. It is an interesting question whether this dependence can be improved or removed, either by better analysis of the above tester or by designing a new tester.

Question 6. *Does there exist a tester for arbitrary finite properties $\mathcal{P} \subset \mathbb{C}^d$ which uses $\text{polylog} |\mathcal{P}|$ copies, and whose parameters have no dependence on $\min_{|\phi\rangle \neq |\phi'\rangle \in \mathcal{P}} D(|\phi\rangle, |\phi'\rangle)$?*

The above tester is a general algorithm for testing any property \mathcal{P} . For some properties \mathcal{P} it is possible to prove better bounds on the performance of this algorithm than Theorem 18 would give, or prove bounds with fewer preconditions. For example, the product test is a particular case of this algorithm (with $T = 2$), and Theorem 17 gives non-trivial bounds on its performance, even though it is applied to the infinite set of product states. We also remark that an alternative algorithm to the above tester would be to produce $|\psi\rangle^{\otimes T}$, and for each $|\phi\rangle \in \mathcal{P}$ in turn, perform the measurement $\{|\phi\rangle\langle\phi|^{\otimes T}, I - |\phi\rangle\langle\phi|^{\otimes T}\}$, and accept if and only if the first outcome is obtained from any measurement. This algorithm would achieve similar scaling in terms of ϵ and δ (as can be shown using the “quantum union bound” of Aaronson [1], or the “gentle measurement lemma” of Winter [159] and Ogawa and Nagaoka [132]) but would not have perfect completeness.

4.1.4 Open questions

There are a number of interesting sets of pure states for which an efficient tester is not known. One such set is the *stabilizer states*. Recall that the Pauli matrices on one qubit are defined to be the set

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

They form a basis for the space of single-qubit linear operators, and by tensoring form a basis for the space of linear operators on n qubits; for $s \in \{I, X, Y, Z\}^n$, we write σ_s for the corresponding operator on n qubits. We call each such tensor product operator a (n -qubit) Pauli matrix, and use \mathcal{P}_n to denote the set of all n -qubit Pauli matrices, together with phases $\pm 1, \pm i$, which forms a group under multiplication.

A state $|\psi\rangle$ of n qubits is said to be a stabilizer state if there exists a maximal Abelian subgroup G of \mathcal{P}_n such that $U|\psi\rangle = |\psi\rangle$ for all $U \in G$. Stabilizer states are important in the study of quantum error-correction [77] and measurement-based quantum computation [137], as well as many other areas of quantum information. It is known that, given access to copies of an unknown stabilizer state $|\psi\rangle$ of n qubits, $|\psi\rangle$ can be learned with $O(n)$ copies [6]; there is a matching $\Omega(n)$ lower bound following from an information-theoretic argument [93]. However, it might be possible to *test* whether $|\psi\rangle$ is a stabilizer state using far fewer copies.

Question 7. *Is there a tester for the property of being a stabilizer state whose parameters do not depend on the number of qubits n ?*

Other sets of pure states for which it would be interesting to have an efficient tester are matrix product states (see, e.g., [134]) and states of low Schmidt rank, or with low complexity with respect to some other entanglement measure [81]. See Section 4.2 below for evidence for a lower bound on the complexity of testing the Schmidt rank.

Another interesting, and as yet largely unexplored, direction for future research is testing properties of quantum states in a distributed setting. Here we imagine that two parties, Alice and Bob, each hold part of one copy of a large unknown state $|\psi\rangle$. Their goal is to determine whether $|\psi\rangle$ satisfies some property while exchanging only a small number of qubits; in particular, Alice cannot just send her half of the state to Bob. Our normal complexity measure “number of copies consumed” is thus replaced with “number of qubits sent.” Aharonov et al. [12] recently showed that the d -dimensional maximally entangled state $\frac{1}{\sqrt{d}} \sum_{i=1}^d |i\rangle|i\rangle$ can be tested up to accuracy ϵ

by communicating only $O(\log 1/\epsilon)$ qubits. There are many other properties where the question of existence of communication-efficient testers remains open.

4.2 Mixed states

A *mixed* state ρ is a convex combination of pure states. Mixed states are described by density matrices, which are positive semidefinite matrices with unit trace; we let $\mathcal{B}(\mathbb{C}^d)$ denote the set of d -dimensional density matrices. The concept of property testing can easily be generalized from pure states to mixed states. We retain the same, natural distance measure

$$D(\rho, \sigma) := \frac{1}{2} \|\rho - \sigma\|_1,$$

which is called the trace distance between ρ and σ . Note that for classical probability distributions (i.e., diagonal density matrices) this is just the total variation distance. As before, say that ρ is ϵ -far from having property $\mathcal{P} \subseteq \mathcal{B}(\mathbb{C}^d)$ if

$$D(\rho, \mathcal{P}) := \inf_{\sigma \in \mathcal{P}} D(\rho, \sigma) \geq \epsilon,$$

and ϵ -close to having property \mathcal{P} if $D(\rho, \mathcal{P}) \leq \epsilon$. Another important distance measure for mixed states is the *fidelity*, which is defined as $F(\rho, \sigma) := \|\sqrt{\rho}\sqrt{\sigma}\|_1$, where $\sqrt{\rho}$ denotes the positive semidefinite square root of the operator ρ . For any mixed state ρ and pure state $|\psi\rangle$, $F(\rho, |\psi\rangle\langle\psi|) = \sqrt{\langle\psi|\rho|\psi\rangle}$. The fidelity and trace distance are related by the inequalities [129, Eq. 9.110]

$$1 - F(\rho, \sigma) \leq D(\rho, \sigma) \leq \sqrt{1 - F(\rho, \sigma)^2}. \quad (5)$$

In a mixed-state property testing scenario, we are given k copies of ρ , for some unknown ρ , and asked to perform a measurement on $\rho^{\otimes k}$ to determine whether $\rho \in \mathcal{P}$, or ρ is ϵ -far away from \mathcal{P} .

Similarly to the case of pure states, *any* property $\mathcal{P} \subseteq \mathcal{B}(\mathbb{C}^d)$ can be tested with $O(d^4/\epsilon^2)$ copies. To distinguish between the two cases that $\rho \in \mathcal{P}$ or ρ is ϵ -far from \mathcal{P} , it suffices to use an estimate $\tilde{\rho}$ such that $D(\tilde{\rho}, \rho) < \epsilon/2$, and accept if and only if $D(\tilde{\rho}, \mathcal{P}) \leq \epsilon/2$. Producing such an estimate can be achieved using quantum state tomography [133, 129]; in order to achieve the required accuracy with success probability $2/3$, $O(d^4/\epsilon^2)$ copies suffice [67]. If one knows in advance that ρ is rank r , compressed sensing techniques allow this bound to be improved to $O((rd/\epsilon)^2 \log d)$ [67].

Some properties of mixed states can be tested significantly more efficiently than this general upper bound. A simple example is the property **Purity**, where ρ satisfies the property if and only if it is a pure state. A natural way to test purity is to apply the swap test (Figure 2) to two copies of ρ . This accepts with probability $(1 + \text{tr}(\rho^2))/2$, which is equal to 1 if and only if ρ is pure. On the other hand, if we let $\rho = \sum_i \lambda_i |\psi_i\rangle\langle\psi_i|$ be the eigendecomposition of ρ , where eigenvalues are listed in non-increasing order, a closest pure state to ρ is $|\psi_1\rangle$. If ρ is ϵ -far away from pure, then $\lambda_1 \leq 1 - \epsilon$. Note that

$$\text{tr}(\rho^2) = \sum_i \lambda_i^2 \leq \max_i \lambda_i \sum_j \lambda_j = \lambda_1.$$

Thus the test accepts with probability at most $1 - \epsilon/2$, implying that **Purity** can be tested with $O(1/\epsilon)$ copies of ρ .

On the other hand, consider the “dual” property of **Mixedness**, where $\rho \in \mathcal{B}(\mathbb{C}^d)$ satisfies the property if and only if it is the maximally mixed state I/d . A strong lower bound has been shown by Childs et al. [54] on the number of copies required to test this property.

Theorem 19 (Childs et al. [54]). *Let d and r be integers such that r strictly divides d . Any algorithm which distinguishes, with probability of success at least $2/3$, between the two cases that $\rho = I/d$ or ρ is maximally mixed on a uniformly random subspace of dimension r must use $\Omega(r)$ copies of ρ . Further, there exists an algorithm which solves this problem using $O(r)$ copies.*

Childs et al. call the problem which they consider the *quantum collision problem*. To see how their result can be applied to **Mixedness**, consider the space of n qubits, whose dimension is $d = 2^n$. As a state ρ which is maximally mixed on a dimension- r subspace of \mathbb{C}^{2^n} satisfies $D(\rho, I/2^n) = 1 - r/2^n$, taking $r = 2^{n-1}$ implies that any algorithm distinguishing between the cases that $\rho = I/2^n$ and ρ is $1/2$ -far from $I/2^n$ must use $\Omega(2^n)$ copies of ρ .²¹ This result also puts strong lower bounds on a number of other property testing problems which one might wish to solve. For example, consider the following three properties:

- **Equality** of pairs of mixed states, where the pair (ρ, σ) satisfies the property if $\rho = \sigma$. This can be seen as the quantum generalization of the classical question of testing whether two probability distributions on d elements are equal or ϵ -far from equal (with respect to the total variation distance), given access to samples from the distributions. A sublinear tester for the classical problem has been given by Batu et al. [29], and recently improved by Chan et al. [52]; for constant ϵ the tester uses $O(d^{2/3})$ samples. By fixing $\sigma = I/d$, the result of [54] implies that the quantum generalization of this problem is more difficult: it requires at least $\Omega(d)$ “samples” (i.e., copies of the states).
- Whether a mixed state ρ has **rank at most** r . Theorem 19 immediately implies that this requires $\Omega(r)$ copies of ρ , which has an interesting implication for testing pure states. Recall that a bipartite state $|\psi\rangle$ on systems AB is said to have Schmidt rank r if it can be written as $|\psi\rangle = \sum_{i=1}^r \sqrt{\lambda_i} |v_i\rangle |w_i\rangle$ for pairwise orthonormal sets of states $\{|v_i\rangle\}$, $\{|w_i\rangle\}$ and non-negative λ_i . If one looks only at the A subsystem, the rank of the reduced state is precisely the Schmidt rank of $|\psi\rangle$. Therefore, Theorem 19 implies that any algorithm which tests whether a pure state $|\psi\rangle$ has Schmidt rank r by producing k copies of $|\psi\rangle$ and acting only on the first subsystems A_1, \dots, A_k of $|\psi\rangle^{\otimes k}$ must satisfy $k = \Omega(r)$. This bound does not apply immediately to general algorithms acting on both the A and B subsystems, leaving the complexity of testing the Schmidt rank open.
- **Separability** of mixed states. A bipartite quantum state $\rho \in \mathcal{B}((\mathbb{C}^d)^{\otimes 2})$ is said to be *separable* if it can be written as a convex combination of product states, and is said to be *entangled* otherwise. Given a classical description of a d -dimensional mixed state as input, determining separability up to accuracy which is inversely polynomial in d is known to be NP-hard [82, 71], and there is some evidence for intractability of the problem even up to constant accuracy [88]. This does not preclude the existence of a tester for separability which is efficient in terms of the number of copies of the input state ρ used; however, Theorem 19 can be used to show that such a tester cannot exist.

The idea is to show that the maximally mixed state on a random subspace of dimension r is far from separable, if r is picked suitably. This can be achieved by combining some previously known results. The *entanglement of formation* of a bipartite state ρ on systems AB , is defined by

$$E_F(\rho) = \min_{\sum_i p_i |\psi_i\rangle\langle\psi_i| = \rho} \sum_i p_i S(\text{tr}_B(|\psi_i\rangle\langle\psi_i|)),$$

²¹Very recently, O’Donnell and Wright [131] strengthened this result, among other things they obtained a tight dependence on ϵ .

where $S(\rho) = -\text{tr}(\rho \log_2 \rho)$ is the von Neumann entropy. Of course, if ρ is separable, $E_F(\rho) = 0$. Let ρ be the maximally mixed state on a random subspace of $\mathbb{C}^d \otimes \mathbb{C}^d$ of dimension $r = \lfloor cd^2 \rfloor$, for some fixed $c \in (0, 1)$. Hayden et al. [89] have shown that, for small enough $c > 0$, there exists a universal constant $C > 0$ such that $E_F(\rho) \geq C \log_2 d$, except with probability exponentially small in d . Also, Nielsen [128] has shown a continuity property for the entanglement of formation:

$$E_F(\rho) - E_F(\sigma) \leq 18(\log_2 d) \sqrt{1 - F(\rho, \sigma)} + 2(\log_2 e)/e.$$

Combining these two properties, and relating the fidelity to the trace distance using (5), we have that ρ is distance $\Omega(1)$ from the set of separable states with high probability. On the other hand, the maximally mixed state I/d^2 is clearly separable. Therefore, any tester which distinguishes separable states from states a constant distance from any separable state can be used to distinguish the maximally mixed state from a random dimension- r subspace; by Theorem 19, this task requires $\Omega(r) = \Omega(d^2)$ copies of the input state.

We remark that the theory of *entanglement witnesses* takes an alternative approach to the direct detection of entanglement (see for example [81, 95] for extensive reviews). An entanglement witness for a state ρ is an observable corresponding to a hyperplane separating ρ from the convex set of separable states; measuring the observable allows one to certify that ρ is entangled. Each such witness will only be useful for certain entangled states, however, so this approach does not provide a means of certifying entanglement of a completely unknown state ρ .

There is a gap between the best known lower and upper bounds for testing the above three properties. We therefore highlight the following open question:

Question 8. *What is the complexity of testing Equality, Separability, and Rank at most r ?*

4.2.1 Testing equality to a fixed pure state

We have seen that testing whether $\rho \in \mathcal{B}(\mathbb{C}^d)$ is the maximally mixed state I/d can require $\Omega(d)$ copies of ρ . By contrast, testing whether ρ is a fixed *pure* state $|\psi\rangle\langle\psi|$ is easy: the obvious test is to perform the measurement $\{|\psi\rangle\langle\psi|, I - |\psi\rangle\langle\psi|\}$, and to accept if the first outcome is returned. The probability of acceptance is $\langle\psi|\rho|\psi\rangle$, which is upper bounded by $1 - D(\rho, |\psi\rangle\langle\psi|)^2$ by (5), so this property can be tested with $O(1/\epsilon^2)$ copies of ρ .

However, there is a more interesting related question of relevance to experimentalists. Imagine we have some experimental apparatus which is claimed to produce a state $|\phi\rangle$ of n qubits, and we would like to certify this fact. In this setting, the above test does not seem to make sense; being able to measure $|\phi\rangle$ is essentially precisely what we wish to certify! We further imagine that n is too large for full state tomography to be efficient. In order to solve this self-certification problem, we would therefore like a procedure which makes a small number of measurements, can easily be implemented experimentally, and certifies that the state produced is approximately equal to $|\phi\rangle$. This question has been considered by da Silva et al. [148], and independently Flammia and Liu [68], who show that certain states $|\phi\rangle$ can be certified using significantly fewer copies of $|\phi\rangle$ than would be required for full tomography, and indeed that any state $|\phi\rangle$ can be certified using quadratically fewer copies ($O(2^n)$ rather than $O(2^{2n})$). The measurements used are also simple: Pauli measurements.

The Pauli matrices $\{\sigma_s\}$ on n qubits form a basis for the space of n -qubit linear operators and

satisfy $\text{tr}(\sigma_s \sigma_t) = 2^n \delta_{st}$. So any state $\rho \in \mathcal{B}(\mathbb{C}^{2^n})$ can be expanded as

$$\rho = \sum_{s \in \{I, X, Y, Z\}^n} \hat{\rho}_s \sigma_s$$

for some real coefficients $\hat{\rho}_s = \text{tr}(\rho \sigma_s) / 2^n$. Writing $\phi := |\phi\rangle\langle\phi|$ for conciseness, the squared fidelity between $|\phi\rangle$ and ρ is

$$\langle\phi|\rho|\phi\rangle = \text{tr}(\rho\phi) = 2^n \sum_{s \in \{I, X, Y, Z\}^n} \hat{\rho}_s \hat{\phi}_s.$$

The works [148, 68] propose the following scheme. First, pick $s \in \{I, X, Y, Z\}^n$ with probability $2^n \hat{\phi}_s^2$; orthonormality of the Pauli matrices implies that this is indeed a valid probability distribution. Then repeatedly measure copies of ρ in the eigenbasis of σ_s , and take the average of the eigenvalues corresponding to the measurement results to produce an estimate $\tilde{\rho}_s$ of $2^n \hat{\rho}_s = \text{tr}(\rho \sigma_s)$. Finally, output $\tilde{\rho}_s / \hat{\phi}_s$ as our guess for the squared fidelity. The expectation of $\tilde{\rho}_s$ is precisely $\text{tr}(\rho \sigma_s)$, and if we assume that this estimate is exact (i.e., $\tilde{\rho}_s = \text{tr}(\rho \sigma_s)$), the expected value of the output is

$$\sum_{s \in \{I, X, Y, Z\}^n} (2^n \hat{\phi}_s^2) \frac{\tilde{\rho}_s}{\hat{\phi}_s} = \text{tr}(\rho\phi).$$

Of course, in general we cannot produce an exact estimate without using an infinite number of copies of ρ . However, to estimate the fidelity up to constant additive error with constant success probability, it suffices to use a finite number of copies. The number of copies required turns out to depend on the quantity $\min_{s, \hat{\phi}_s \neq 0} |\hat{\phi}_s|$; for certain classes of states $|\phi\rangle$ (such as stabilizer states), the number of copies used does not depend on n .

4.2.2 Unitarily invariant properties

Generalizing the properties **Purity** and **Mixedness**, one can consider properties \mathcal{P} of mixed quantum states which are unitarily invariant, in the following sense: If $\rho \in \mathcal{P}$, then $(U\rho U^\dagger) \in \mathcal{P}$ for all $U \in U(d)$, where $U(d)$ denotes the unitary group in d dimensions. Observe that this implies that, if ρ is ϵ -far from \mathcal{P} , then so is $U\rho U^\dagger$, for all ϵ and all $U \in U(d)$. For any ρ , $D(\rho, \mathcal{P})$ must necessarily be a symmetric function of the spectrum of ρ .

We can see unitarily invariant properties as quantum analogs of symmetric properties of classical probability distributions. Quite recently, it has been shown that a particular ‘‘canonical’’ classical tester is close to optimal for all such symmetric properties which satisfy certain continuity constraints [153]. This has allowed strong bounds to be proven on the complexity of testing properties such as equality of probability distributions, and distinguishing high-entropy from low-entropy distributions. We now discuss an analogous ‘‘canonical tester’’ for unitarily invariant properties.

In order to take advantage of the unitary symmetry, one can use a concept known as *Schur-Weyl duality*. We will only briefly summarize this beautiful theory here, and sketch the consequences for property testing; for much more detailed introductions, see the theses [57, 87]. Schur-Weyl duality implies that any linear operator M on $(\mathbb{C}^d)^{\otimes k}$ which commutes with permutations of the k subsystems, and also with local unitaries on each subsystem (i.e., $U^{\otimes k} M (U^{-1})^{\otimes k} = M$ for all $U \in U(d)$) can be written as $M = \sum_{\lambda \vdash k} \alpha_\lambda P_\lambda$ for some coefficients α_λ and projectors P_λ , where the sum is over partitions λ of k (e.g., the partitions of 4 are (4), (3, 1), (2, 2), (2, 1, 1), (1, 1, 1, 1)). Each partition λ corresponds to an irreducible representation (irrep) of S_k , the symmetric group

on k elements; one important irrep is the trivial irrep (k) which maps $\pi \mapsto 1$ for all $\pi \in S_k$. The operators P_λ are defined by

$$P_\lambda := \frac{d_\lambda}{k!} \sum_{\pi \in S_k} \chi_\lambda(\pi) U_\pi.$$

In the above expression, d_λ is the dimension of the corresponding irrep V_λ of S_k , which associates a d_λ -dimensional square matrix with each permutation $\pi \in S_k$. Then χ_λ is the corresponding character $\text{tr}(V_\lambda)$ and U_π is the operator which acts by permuting k d -dimensional systems according to π :

$$U_\pi |i_1\rangle \dots |i_k\rangle = |i_{\pi^{-1}(1)}\rangle \dots |i_{\pi^{-1}(k)}\rangle.$$

One can show that each operator P_λ is indeed a projector, that $P_\lambda P_\mu = \delta_{\lambda\mu} P_\lambda$, and that $\sum_{\lambda \vdash k} P_\lambda = I$. These operators therefore define a measurement (POVM), and performing this measurement is known as *weak Schur sampling* [54]. This can be implemented efficiently via a procedure which is known as generalized phase estimation [87, 54] and generalizes the swap test [46] (cf. Section 4.1.1) and symmetrization [25]. Generalized phase estimation is based on the quantum Fourier transform (QFT) over S_k [30], which is a unitary operation that performs a change of basis from $\{|\pi\rangle : \pi \in S_k\}$ to $\{|\lambda, i, j\rangle : \lambda \vdash k, 1 \leq i, j \leq d_\lambda\}$. It follows from basic representation theory that this makes sense, i.e., that $\sum_{\lambda \vdash k} d_\lambda^2 = k!$.

The generalized phase estimation procedure proceeds as follows:

1. Start with a quantum state $\sigma \in \mathcal{B}((\mathbb{C}^d)^{\otimes k})$.
2. Prepend a $k!$ -dimensional ancilla register whose basis states correspond to triples $|\lambda, i, j\rangle$, initialized in the state $|(k), 1, 1\rangle$ corresponding to the trivial irrep.
3. Apply the inverse quantum Fourier transform over S_k to the ancilla to produce the state $\frac{1}{\sqrt{k!}} \sum_{\pi \in S_k} |\pi\rangle$ (see, e.g., [30] for an explanation of this).
4. Apply the controlled permutation operation $\sum_{\pi \in S_k} |\pi\rangle\langle\pi| \otimes U_\pi$, controlled on the ancilla.
5. Apply the quantum Fourier transform over S_k to the ancilla and measure it, receiving outcome (λ, i, j) .
6. Output λ .

One can show [24, 87] that, on input σ , generalized phase estimation does indeed output λ with probability $\text{tr}(P_\lambda \sigma)$.²²

It turns out that any test for a unitarily invariant property can, essentially, be taken to consist of performing weak Schur sampling and classically post-processing the results.

Lemma 20. *Let $\mathcal{P} \subseteq \mathcal{B}(\mathbb{C}^d)$ be a unitarily invariant property. Assume there exists a tester which uses k copies of the input state ρ , and accepts all states $\rho \in \mathcal{P}$ with probability at least $1 - \delta$, but accepts all states which are ϵ -far from \mathcal{P} with probability at most $1 - f(\epsilon)$ for $\epsilon > 0$. Then there exists a tester with the same parameters which consists of performing weak Schur sampling on $\rho^{\otimes k}$ and classically postprocessing the results.*

²²Some works describe the procedure as instead starting with a QFT and finishing with an inverse QFT [54, 126], but this does not appear correct as the QFT should map from the group algebra of S_k to the space of irreps of S_k [30].

Proof. Let M be the measurement operator corresponding to the tester accepting, and for each ϵ , let ρ_ϵ be a state which is distance ϵ from \mathcal{P} and achieves the worst-case probability of acceptance (so ρ_0 is a state in \mathcal{P} with the lowest probability of acceptance, and for $\epsilon > 0$, ρ_ϵ is a state with the highest probability of acceptance such that $D(\rho_\epsilon, \mathcal{P}) = \epsilon$). Then, by the permutation invariance of $\rho_\epsilon^{\otimes k}$, we have

$$\mathrm{tr}(M\rho_\epsilon^{\otimes k}) = \frac{1}{k!} \sum_{\pi \in S_k} \mathrm{tr}(MU_\pi \rho_\epsilon^{\otimes k} U_\pi^{-1}) =: \mathrm{tr}(\overline{M}\rho_\epsilon^{\otimes k}),$$

where we define $\overline{M} = \frac{1}{k!} \sum_{\pi \in S_k} U_\pi M U_\pi^{-1}$, and by the unitary invariance of \mathcal{P} ,

$$\mathrm{tr}(\overline{M}\rho_0^{\otimes k}) \leq \int \mathrm{tr}(\overline{M}(U\rho_0 U^{-1})^{\otimes k}) dU = \mathrm{tr}\left(\int U^{\otimes k} \overline{M} (U^{-1})^{\otimes k} dU\right) \rho_0 =: \mathrm{tr}(\overline{\overline{M}}\rho_0),$$

where the integral is taken according to Haar measure on $U(d)$, and similarly $\mathrm{tr}(\overline{M}\rho_\epsilon^{\otimes k}) \geq \mathrm{tr}(\overline{\overline{M}}\rho_\epsilon^{\otimes k})$ for $\epsilon > 0$. Therefore, it suffices to implement $\overline{\overline{M}}$ to achieve the same parameters as M . But $\overline{\overline{M}}$ commutes with local unitaries and permutations of the k systems, so by Schur-Weyl duality we can write $\overline{\overline{M}} = \sum_\lambda \alpha_\lambda P_\lambda$ for some coefficients α_λ ; as $\overline{\overline{M}}$ is a measurement operator, for each λ it holds that $0 \leq \alpha_\lambda \leq 1$. So we can implement $\overline{\overline{M}}$ by performing weak Schur sampling, obtaining outcome λ , and then accepting with probability α_λ . \square

Further, one can write down the probability of obtaining each outcome λ as follows: if the input state ρ has eigenvalues (x_1, \dots, x_d) , then

$$\mathrm{tr}(P_\lambda \rho^{\otimes k}) = d_\lambda s_\lambda(x_1, \dots, x_d),$$

where s_λ is a Schur polynomial (see, e.g., [21] for a discussion). In principle, this allows one to calculate the parameters of the optimal test for any unitarily invariant property; in practice, the calculations required are somewhat daunting. Nevertheless, a careful analysis of the output distributions resulting from weak Schur sampling was the approach taken by Childs et al. [54] to prove their bounds on the quantum collision problem. Indeed, their approach is an example of how one can prove lower bounds on quantum property testers more generally: first use symmetry arguments to prove that the optimal test must be of a certain form, then analyse the optimal test directly.

5 Quantum testing of quantum properties: Dynamics

5.1 Unitary operators

In this section, we will consider quantum property testing of quantum dynamics, beginning with unitary dynamics. We will imagine we are given black-box access to a unitary operator U , and we want to test if U either has a certain property or is far from having it, by applying U a small number of times. This setting is more complicated than that of testing properties of quantum states in that, rather than simply performing a measurement on a number of copies of a state, we can consider more involved protocols based on the use of U in a sequential, adaptive fashion.

There are a number of choices one needs to make when defining this model—in particular, what distance measure to use, and whether or not to allow applications of controlled- U and/or U^{-1} as part of the model. In Sections 5.1.1 and 5.1.2 we will discuss the effect of these choices. Next, we will discuss a useful correspondence between quantum states and unitaries—the *Choi-Jamiołkowski*

isomorphism—which allows one to apply many of the algorithms developed for testing quantum states to unitaries. Finally, in sections 5.1.4–5.1.7 we will describe several known results on testing various properties of unitary operators.

We continue to let $U(d)$ denote the unitary group in d dimensions, and let $M(d)$ denote the set of $d \times d$ matrices. A property of unitary operators is simply a (discrete or continuous) subset $\mathcal{P} \subseteq U(d)$.

5.1.1 Distance measures

As compared with the case of pure states, it is less obvious which distance measure between unitary operators is the right one to choose to obtain interesting property testing results. For quantum states, the distinguishability of any two states is controlled by their trace distance. A natural way to generalize this to unitary operations would be to maximize the distinguishability of the output states over all input states²³, to produce

$$D^{\max}(U, V) := \max_{|\psi\rangle} D(U|\psi\rangle, V|\psi\rangle) = \max_{|\psi\rangle} \sqrt{1 - |\langle\psi|U^\dagger V|\psi\rangle|^2}.$$

Unfortunately, there are extremely simple properties which are hard to test with respect to this distance measure. One such example is the **Identity** property: does an input unitary U satisfy $U = e^{i\theta}I$? (Note that, as with the case of pure state properties, we allow an arbitrary phase θ in the definition, as U cannot be distinguished from $e^{i\theta}U$.) Consider the family of n -dimensional unitary operators U_i , $i \in [n]$, where $U_i|j\rangle = (-1)^{\delta_{ij}}|j\rangle$. Each of these has maximal distance from I , according to the distance measure D^{\max} . However, a quantum algorithm which uses the input operator U k times and distinguishes between the case where U is equal to the identity, and the case where $U = U_i$ for some i , would imply a quantum algorithm which computes the OR function of n input bits, promised to have Hamming weight at most 1, using $O(k)$ queries. As this problem is known to require $\Omega(\sqrt{n})$ quantum queries [33], it follows that $k = \Omega(\sqrt{n})$. This is a lower bound on the complexity of identity-testing in an oracular setting; we discuss a lower bound based on computational complexity arguments in Section 6.

It is perhaps not surprising that D^{\max} is not the right measure of distance to choose for property testing problems, as it is a “best-case” rather than “average-case” measure. A suitable such alternative measure can be defined as follows. For any d -dimensional operators $A, B \in M(d)$, let $\langle A, B \rangle$ denote the *normalized Hilbert-Schmidt* inner product

$$\langle A, B \rangle := \frac{1}{d} \text{tr}(A^\dagger B) = \frac{1}{d} \sum_{i,j} A_{ij}^* B_{ij}.$$

Assume that $\langle A, A \rangle = \langle B, B \rangle = 1$ (a property satisfied, for example, if A and B are unitary). Then the distance between A and B is given by

$$D(A, B) := \sqrt{1 - |\langle A, B \rangle|^2}.$$

For $\mathcal{P} \subseteq U(d)$, we analogously define

$$D(U, \mathcal{P}) := \inf_{V \in \mathcal{P}} D(U, V).$$

²³One might wonder whether distinguishability could be improved further by allowing the unknown unitary operator to act on part of an entangled state; it turns out that this is not the case [158].

Note the close analogy to the distance between pure states (3). Indeed, we use the same notation as for the distance $D(|\psi\rangle, |\phi\rangle)$ to highlight the fact that the distance for unitaries is naturally induced by the distance for states. The distance measure $D(A, B)$ seems to have been first explicitly introduced by Low [116]; Wang [156] has defined a closely related alternative measure as $D'(A, B) := \sqrt{1 - |\langle A, B \rangle|}$. As $D(A, B)/\sqrt{2} \leq D'(A, B) \leq D(A, B)$, the two measures are essentially interchangeable. For any operators A and B such that $\langle A, A \rangle = \langle B, B \rangle = 1$, $D(A, B)$ has the following properties.

- $0 \leq D(A, B) \leq 1$, with $D(A, B) = 0$ if and only if $A = e^{i\phi}B$ for some overall phase ϕ . As there exist $A \neq B$ with $D(A, B) = 0$, this implies that $D(\cdot, \cdot)$ is not a metric, but only a “pseudometric.” Further, $D(A, B) = D(WA, WB) = D(AW, BW)$ for any unitary W .
- $D(A, B)$ can alternatively be defined as

$$D(A, B) = \frac{1}{\sqrt{2}} \|A \otimes A^\dagger - B \otimes B^\dagger\|_2,$$

where $\|\cdot\|_2$ is the normalized Schatten 2-norm $\|M\|_2 = \sqrt{\frac{1}{d} \sum_{i,j=1}^d |M_{ij}|^2}$ [116]. Observe that this representation shows that $D(\cdot, \cdot)$ satisfies the triangle inequality.

- We have $\|M\|_2^2 = \langle M, M \rangle$. Therefore, $\|A - B\|_2^2 = \langle A - B, A - B \rangle = 2 - 2\text{Re}\langle A, B \rangle$. This implies that $D(A, B) \leq \|A - B\|_2$, via the elementary inequality $2\text{Re} z \leq |z|^2 + 1$, valid for any $z \in \mathbb{C}$.

The following justifies the claim that $D(\cdot, \cdot)$ is indeed an “average-case” measure of distance.

Proposition 21. *Fix d -dimensional unitary operators U and V . Then*

$$\int d\psi D(U|\psi\rangle, V|\psi\rangle)^2 = \frac{d}{d+1} D(U, V)^2,$$

where the integral is taken according to Haar measure on pure states $|\psi\rangle \in \mathbb{C}^d$.

Proof. We have

$$\begin{aligned} \int d\psi D(U|\psi\rangle, V|\psi\rangle)^2 &= 1 - \int d\psi |\langle \psi | U^\dagger V | \psi \rangle|^2 \\ &= 1 - \int d\psi \text{tr}[(U^\dagger V \otimes V^\dagger U) |\psi\rangle\langle\psi|^{\otimes 2}] \\ &= 1 - \text{tr} \left[(U^\dagger V \otimes V^\dagger U) \left(\frac{I + F}{d(d+1)} \right) \right] \\ &= \frac{d}{d+1} \left(1 - \left| \frac{\text{tr}(U^\dagger V)}{d} \right|^2 \right) \\ &= \frac{d}{d+1} D(U, V)^2. \end{aligned}$$

In the third equality we use the fact that $\int |\psi\rangle\langle\psi|^{\otimes 2} d\psi = (I + F)/(d(d+1))$, where F is the flip (or swap) operator which interchanges two d -dimensional systems. The fourth equality follows from the facts that, for any d -dimensional operators A, B , $\text{tr}(A \otimes B) = \text{tr}(A)\text{tr}(B)$ and $\text{tr}((A \otimes B)F) = \text{tr}(AB)$. \square

The quantity $\int d\psi |\langle \psi | U^\dagger V | \psi \rangle|^2$ appearing in the proof was previously introduced by Acín [8] as an average-case variant of the fidelity. We will see in Section 5.1.3 below a number of properties, including the **Identity** property, which can be tested efficiently with respect to the distance measure $D(\cdot, \cdot)$.

5.1.2 Controlled and inverse unitaries

As well as being given access to a unitary operator U , we may be given access to the inverse U^{-1} and/or the controlled unitary $c-U$, or in other words the operator $|0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes U$. This may be a reasonable assumption if we would like to apply our property testing algorithm to a unitary operator given in the form of a quantum circuit; on the other hand, it may not be reasonable in an adversarial scenario where we only assume access to U as a black box.

For any U, V we have $\langle c-U, c-V \rangle = (1 + \langle U, V \rangle)/2$, implying

$$\begin{aligned} D(c-U, c-V) &= \sqrt{1 - \left| \frac{1 + \langle U, V \rangle}{2} \right|^2} \\ &= \frac{1}{2} \sqrt{3 - 2\operatorname{Re}\langle U, V \rangle - |\langle U, V \rangle|^2} \\ &= \frac{1}{2} \sqrt{\|U - V\|_2^2 + D(U, V)^2}. \end{aligned}$$

Recalling that $D(U, V) \leq \|U - V\|_2$, we therefore have the inequalities

$$\|U - V\|_2/2 \leq D(c-U, c-V) \leq \|U - V\|_2/\sqrt{2}. \quad (6)$$

Thus, given access to controlled unitaries, one can hope to design tests which are sensitive to the 2-norm distance $\|U - V\|_2$. For example, if we are allowed access to controlled unitaries we can distinguish U from $-U$ (see the next section for how this can be done), whereas this is impossible given access to U alone.

Being given access to U^{-1} can also be powerful. In particular, it allows us to apply the important primitive of amplitude amplification [40] to property testing algorithms, in analogy to Section 2.2.1. Imagine we have a test for a property $\mathcal{P} \subseteq U(d)$ which uses q copies of the input unitary U , and such that for $U \in \mathcal{P}$ the test always accepts (it has perfect completeness), and for U ϵ -far from \mathcal{P} , the test accepts with probability at most $f(\epsilon)$. Then amplitude amplification allows us to test \mathcal{P} with $O(q/\sqrt{f(\epsilon)})$ copies of U , rather than the $O(q/f(\epsilon))$ copies that would be required by simple repetition. For example, we will see below that this gives a square-root speed-up for testing equality of unitary operators. In the complexities we quote below, we assume that amplitude amplification has not been applied.

5.1.3 From properties of states to properties of unitaries

There is a correspondence between pure quantum states and unitary operators, which is known as (a special case of) the Choi-Jamiołkowski isomorphism [56, 100] and will sometimes allow us to translate tests for properties of states to tests for analogous properties of unitaries. Given access to $U \in U(d)$, we first prepare the maximally entangled state of two d -dimensional systems

$$|\Phi\rangle := \frac{1}{\sqrt{d}} \sum_{i=1}^d |i\rangle|i\rangle$$

and then apply U to the first system. We obtain the state $|U\rangle \in (\mathbb{C}^d)^{\otimes 2}$ defined by

$$|U\rangle = \frac{1}{\sqrt{d}} \sum_{i,j=1}^d U_{ji} |j\rangle |i\rangle.$$

The isomorphism is thus simply $U \leftrightarrow |U\rangle$. The state $|U\rangle$ faithfully represents the original operator U ; in particular, it is easy to see that $\langle U|V\rangle = \langle U, V\rangle$ and hence $D(U, V) = D(|U\rangle, |V\rangle)$. So, if we have a tester for some property \mathcal{P} of d^2 -dimensional quantum states, by applying the test to $|U\rangle$ we obtain a tester with the same parameters for an analogous property \mathcal{P}' of d -dimensional unitary operators.

However, one sometimes has to be careful. Imagine we have a tester which accepts states with property \mathcal{P} with certainty, and accepts states which are ϵ -far away from having property \mathcal{P} with probability at most δ . Then, via the Choi-Jamiołkowski isomorphism, this translates into a tester which accepts unitary matrices with property \mathcal{P}' with certainty, and accepts, with probability at most δ , unitaries which are ϵ -far away from any *matrix* M with $\langle M, M\rangle = 1$ such that M has property \mathcal{P}' . Therefore, in principle it could be the case that U is far from any unitary matrix with property \mathcal{P}' , but is close to some non-unitary matrix M which has property \mathcal{P}' . In this situation the tester might incorrectly accept. Nevertheless, in various cases of interest one can show that this situation does not arise. In particular, we have the following lemma (which generalizes similar claims in [88, 156]).

Lemma 22. *Let $\mathcal{P} \subseteq M(d)$, and $U \in U(d)$. For $M \in \mathcal{P}$ such that $\langle M, M\rangle = 1$, let $M = AV$ be a polar decomposition of M , with $A = \sqrt{MM^\dagger}$ and V unitary. Then, if $V \in \mathcal{P}$ and $D(U, M) = \epsilon$, $D(U, \mathcal{P} \cap U(d)) \leq 2\epsilon$.*

Proof. We have

$$\langle M, V\rangle = \frac{1}{d} \text{tr}(\sqrt{MM^\dagger}) = \frac{1}{d} \|M\|_1 = \frac{1}{d} \max_{W \in U(d)} |\text{tr}(WM)| \geq \sqrt{1 - \epsilon^2},$$

using the definition of the trace norm and that $D(U, M) = \epsilon$. Thus

$$D(U, V) \leq D(U, M) + D(M, V) \leq 2\epsilon.$$

□

The following are some examples where one can use the Choi-Jamiołkowski isomorphism to test properties of unitary operators:

- The **Equality to V** property, where U satisfies the property if $U = e^{i\theta}V$, for some θ . The test creates the state $|U\rangle$ and measures in the basis $\{|V\rangle\langle V|, I - |V\rangle\langle V|\}$. Using the analysis of the corresponding property for pure states, this property is testable with $O(1/\epsilon^2)$ uses of U . A simple special case of this is the previously discussed **Identity** property.
- The **Equality** property for pairs of unitary operators, where the pair U, V satisfies the property if $U = e^{i\theta}V$, for some θ . This can be tested by applying the swap test to $|U\rangle$ and $|V\rangle$; again, the analysis of the **Equality** property for states goes through unchanged, implying that this property is testable with $O(1/\epsilon^2)$ uses of U and V .

- The **Inverses** property, where $U, V \in U(d)$ satisfy the property if $U = e^{i\theta}V^{-1}$, for some θ . The test is to create the state $|UV\rangle$ with one use of each of U and V , then to test for equality to $|\Phi\rangle$. The probability of rejection is $D(UV, I)^2 = D(U, V^{-1})^2$, so if $D(U, V^{-1}) = \epsilon$, the test rejects with probability ϵ^2 . Note that there is no need to have access to U^{-1} or V^{-1} .
- The **Product** property for unitary operators, where an operator $U \in U(d^n)$ satisfies the property if $U = U_1 \otimes U_2 \otimes \cdots \otimes U_n$ for some $U_1, \dots, U_n \in U(d)$. This can be tested by applying the product test described in Section 4.1.2 to $|U\rangle$ [88]. One also needs to show that, if U is close to an operator $A \in M(d^n)$ such that $A = A_1 \otimes \cdots \otimes A_n$, U is in fact close to a *unitary* operator of this form; this claim follows from Lemma 22. The final result is that if U is product the test accepts with certainty, whereas if U is ϵ -far from product, the test rejects with probability $\Theta(\epsilon^2)$.

5.1.4 Membership of the Pauli and Clifford groups

Let $B = \{B_1, \dots, B_{d^2}\}$ be a unitary operator basis for the space of linear operators on d dimensions such that B is orthonormal with respect to the normalized Hilbert-Schmidt inner product, i.e., $\langle B_i, B_j \rangle = \delta_{ij}$. Then the set $|B_i\rangle$ forms an orthonormal basis for \mathbb{C}^{d^2} with respect to the standard inner product, implying that one can test membership of a unitary operator U in B with the following procedure, which we call the *operator basis test*.

1. Create two copies of $|U\rangle$.
2. Measure each copy in the basis $\{|B_1\rangle, \dots, |B_{d^2}\rangle\}$.
3. Accept if both measurements give the same result.

The probability of getting outcome i from each measurement is independent and equal to $|\langle U, B_i \rangle|^2$. Thus, if $U = e^{i\theta}B_i$ for some i , then the test will accept with certainty. On the other hand, if $\min_{V \in B} D(U, V) = \epsilon$, the probability of getting the same measurement outcome twice is

$$\sum_{i=1}^{d^2} |\langle U, B_i \rangle|^4 \leq \max_i |\langle U, B_i \rangle|^2 \sum_{i=1}^{d^2} |\langle U, B_i \rangle|^2 = 1 - \epsilon^2.$$

Therefore, by repeating the operator basis test and rejecting if any of the individual tests reject, the property of **Membership in B** can be tested with $O(1/\epsilon^2)$ uses of U .

A natural operator basis to which this test can be applied is the set of Pauli matrices on n qubits [127, 156], which form a basis for the space of linear operators on n qubits. This basis is orthonormal with respect to the normalized Hilbert-Schmidt inner product. We call the corresponding basis for $\mathbb{C}^{2^{2n}}$ obtained via the Choi-Jamiołkowski isomorphism the *Pauli basis*. The operator basis test can be immediately applied to test whether an n -qubit operator is proportional to an n -qubit Pauli matrix, or is far from any such matrix; we call this special case the *Pauli test*. As pointed out in [127], this is a natural quantum generalization of the important classical property of linearity of Boolean functions [37] discussed in Section 2.2.1. Given access to an oracle for $f : \{0, 1\}^n \rightarrow \{0, 1\}$, one can readily construct the diagonal unitary operator U_f where $U_f|z\rangle = (-1)^{f(z)}|z\rangle$, and also the controlled unitary operator $c-U_f$; it is easy to see that f is linear (with respect to addition mod 2) if and only if U_f is a tensor product of identity and Z operators.

Further, if $\ell : \{0, 1\}^n \rightarrow \{0, 1\}$ is a Boolean function, the distance between $c-U_f$ and $c-U_\ell$ is

$$\begin{aligned} D(c-U_f, c-U_\ell) &= \sqrt{1 - \left(\frac{1}{2} + \frac{1}{2^{n+1}} \sum_{z \in \{0,1\}^n} (-1)^{f(z)+\ell(z)} \right)^2} \\ &= \sqrt{1 - (1 - |\{z : f(z) \neq \ell(z)\}|/2^n)^2} \\ &= \sqrt{2d(f, \ell) - d(f, \ell)^2}, \end{aligned}$$

where $d(f, \ell) := |\{x : f(x) \neq \ell(x)\}|/2^n$ is the normalized Hamming distance. This implies that the Pauli test (for the special case of testing diagonal Pauli matrices) can be used to test linearity of Boolean functions, recovering the $O(1/\epsilon)$ complexity of the classical tester discussed in Section 2.2.1, which can be improved to $O(1/\sqrt{\epsilon})$ via amplitude amplification.

The Pauli test can also be used as a subroutine in an algorithm for testing membership in the *Clifford group*. The Clifford group \mathcal{C}_n on n qubits is the normalizer of the Pauli group \mathcal{P}_n , or in other words the set $\mathcal{C}_n := \{C \in U(2^n) : \forall P \in \mathcal{P}_n, CPC^{-1} \in \mathcal{P}_n\}$. The Clifford group plays an important role in many areas of quantum information theory, including quantum error-correction and simulation of quantum circuits [77, 129]. Wang [156] has shown that, given access to a unitary U and its inverse U^{-1} , whether U is a member of the Clifford group can be tested with $O(1/\epsilon^2)$ uses of U and U^{-1} ; this result improves a previous test of Low [116] by removing any dependence on n , and can in turn be improved to $O(1/\epsilon)$ using amplitude amplification [40].

Wang's test is very natural: pick a Pauli matrix $P \in \mathcal{P}_n$ uniformly at random, and apply the Pauli test to the operator UPU^{-1} . If $U \in \mathcal{C}_n$, this test will always accept. Intuitively, if U is far from any Clifford operator, then we expect that for most Pauli operators P , UPU^{-1} will be far from being a Pauli operator, so repeating this test a constant number of times would suffice to detect this. Making this intuition precise requires some work; see [156] for the details.

Question 9. *Is there an efficient test for the property of membership in the Clifford group which does not require access to U^{-1} ?*

5.1.5 Testing commutativity

Say that $U, V \in U(d)$ satisfy the **Commuting** property if $UV = VU$. Assuming that we are given access to the controlled operators $c-U$ and $c-V$, consider the following tester for this property:

1. Create the states $|c-Uc-V\rangle, |c-Vc-U\rangle$ by applying controlled- U and controlled- V operations to the first half of each of two maximally entangled states.
2. Apply the swap test to these states and accept if the test accepts.

If U and V commute, then $c-U$ and $c-V$ also commute, so $|c-Uc-V\rangle = |c-Vc-U\rangle$ and hence the swap test accepts with certainty. On the other hand, if $\|UV - VU\|_2 = \epsilon$, then by (6) the test rejects with probability at least $\epsilon^2/8$. In order for this to be a good test for commutativity, we therefore need to show that, if $\|UV - VU\|_2 = \epsilon$, U and V are close to a pair of unitary operators \tilde{U}, \tilde{V} such that \tilde{U} and \tilde{V} commute. Precisely this result has recently been shown by Glebsky [72] in the form of the following theorem, whose proof we omit.

Theorem 23 (Glebsky [72]). *Let $U, V \in U(d)$ satisfy $\|UV - VU\|_2 = \epsilon$. Then there exist $\tilde{U}, \tilde{V} \in U(d)$ such that \tilde{U} and \tilde{V} commute and $\|U - \tilde{U}\|_2 \leq 30\epsilon^{1/9}$, $\|V - \tilde{V}\|_2 \leq 30\epsilon^{1/9}$.*

The consequence is that the above tester rejects pairs (U, V) such that U and V are ϵ -far from a pair of commuting matrices with probability $\Omega(\epsilon^{18})$. By repeating the test $\text{poly}(1/\epsilon)$ times, we obtain a tester which rejects such pairs with constant probability.

Question 10. *Is there an efficient test for commutativity which does not require access to the controlled unitaries $c-U$, $c-V$, but just uses U and V ?*

5.1.6 Testing quantum juntas

Analogously to the classical case of Boolean functions $f : \{0, 1\}^n \rightarrow \{0, 1\}$, a unitary operation on n qubits is said to be a k -junta if it acts non-trivially on at most k of the qubits, or in other words is of the form $U_S \otimes I_{S^c}$, where $U \in U(2^k)$ and S is a k -subset of $[n]$. Wang [156] has given a tester for whether a unitary operator U is a k -junta, which turns out to be a direct generalization of the tester of Atııcı and Servedio [20] for the classical property of a Boolean function being a k -junta (Section 2.2.3). The work [127] had previously studied a different tester for being a 1-junta (“dictator”), but did not prove correctness. Wang’s tester proceeds as follows:

1. Set $W = \emptyset$.
2. Repeat the following procedure T times, for some T to be determined:
 - (a) Create the state $|U\rangle$ and measure in the Pauli basis, obtaining outcome $s \in \{I, X, Y, Z\}^n$.
 - (b) Update $W \leftarrow W \cup \{i : s_i \neq I\}$.
 - (c) If $|W| > k$, reject.
3. Accept.

To show correctness of this test, it suffices to prove the following claim:

Theorem 24 (Wang [156]). *If U is ϵ -far from any k -junta, and $T = \Theta(k/\epsilon^2)$, the above procedure accepts with probability at most $1/3$.*

The result originally shown by Wang [156] was a somewhat worse bound of $T = \Theta(k \log(k/\epsilon)/\epsilon^2)$, but the bound can be improved to $\Theta(k/\epsilon^2)$ via a straightforward generalization of the analysis of Atııcı and Servedio [20], as we now show (cf. Section 2.2.3). If we are given access to U^{-1} as well, the bound can be improved further to $T = \Theta(k/\epsilon)$ via amplitude amplification.

Proof. As the Pauli matrices form a basis for the space of n -qubit operators, we can expand

$$U = \sum_{s \in \{I, X, Y, Z\}^n} \widehat{U}_s \sigma_s,$$

where σ_s is the n -qubit Pauli operator corresponding to the string s , and $\widehat{U}_s \in \mathbb{C}$. Pauli matrices are orthonormal with respect to the normalized Hilbert-Schmidt inner product, implying that $\sum_{s \in \{I, X, Y, Z\}^n} |\widehat{U}_s|^2 = 1$. Assume that U is ϵ -far from any unitary operator V that is a k -junta, and for $s \in \{I, X, Y, Z\}^n$, let $\text{supp}(s) := \{i : s_i \neq I\}$. Then, for any subset $W \subseteq [n]$ of size at most k ,

$$w_W := \sum_{s: \text{supp}(s) \subseteq W} |\widehat{U}_s|^2 \leq 1 - \epsilon^2/4.$$

To see this, assume the opposite and consider the operator $M_W = w_W^{-1/2} \sum_{s, \text{supp}(s) \subseteq W} \widehat{U}_s \sigma_s$. Then M_W is a k -junta, $\langle M_W, M_W \rangle = 1$, and $D(U, M_W) = (1 - w_W)^{1/2} < \epsilon/2$. Further, the unitary matrix V occurring in a polar decomposition of M is also a k -junta. So, by Lemma 22, $D(U, V) \leq \epsilon$, contradicting that U is ϵ -far from any unitary k -junta.

For each measurement, the probability that a string s is returned such that $\text{supp } s \not\subseteq W$ is therefore at least $\epsilon^2/4$. Thus the expected number of measurements required to find $k + 1$ such indices is at most $4(k + 1)/\epsilon^2$. The theorem then follows from Markov's inequality. \square

5.1.7 Other properties of unitary matrices

We finish this section by mentioning a few other properties of unitary matrices which have fairly straightforward testers. Say that a unitary matrix U satisfies the **Diagonality** property if $U_{ij} = 0$ for $i \neq j$. Consider the following easy tester for this property: Apply U to a uniformly random computational basis state $|i\rangle$, measure in the computational basis, and accept if and only if the outcome is i . Writing $U_{kk} = r_k e^{i\theta_k}$ for $r_k \geq 0$ and $0 \leq \theta_k < 2\pi$, we have

$$\max_D \max_{\text{diagonal}} |\langle U, D \rangle| = \frac{1}{d} \max_D \max_{\text{diagonal}} \left| \sum_{k=1}^d U_{kk}^* D_{kk} \right| = \frac{1}{d} \max_{\theta_k} \left| \sum_{k=1}^d r_k e^{i(\theta_k - \gamma_k)} \right| = \frac{1}{d} \left| \sum_{k=1}^d r_k \right| = \frac{1}{d} \sum_{k=1}^d |U_{kk}|.$$

On the other hand, the probability of accepting is precisely

$$\frac{1}{d} \sum_{k=1}^d |U_{kk}|^2 \leq \frac{1}{d} \max_k |U_{kk}| \sum_{k=1}^d |U_{kk}| \leq \frac{1}{d} \sum_{k=1}^d |U_{kk}|.$$

Thus, if the test accepts with probability $1 - \delta$, U is distance at most $\sqrt{2\delta}$ from a diagonal unitary matrix D , implying that **Diagonality** can be ϵ -tested with $O(1/\epsilon^2)$ uses of U .

This tester is simple, but can be applied to the following more general problem: Given a basis B for \mathbb{C}^d , is every vector in B an eigenvector of U ? This is equivalent to asking whether VUV^{-1} is diagonal, where V is the change of basis matrix for B . This problem can be solved by applying the test for diagonality to VUV^{-1} , noting that the distance of VUV^{-1} from the nearest diagonal matrix is the same as the distance of U from the nearest matrix \tilde{U} such that every vector in B is an eigenvector of \tilde{U} . For example, this allows us to test U for being a **Circulant** matrix (i.e., a matrix of the form $U_{xy} = f(x - y)$ for some $f : \{0, \dots, d - 1\} \rightarrow \mathbb{C}$, where subtraction is understood modulo d) as such matrices are characterized by being diagonalized by the quantum Fourier transform over \mathbb{Z}_d .

Finally, Wang [156] has proven that membership of a unitary operator $U \in U(d)$ in the orthogonal group $O(d) := \{M \in M(d) : MM^T = I\}$ can be ϵ -tested with $O(1/\epsilon^2)$ uses of U . The tester is based on applying $U \otimes U$ to $|\Phi\rangle$, which produces the state $|UU^T\rangle$, then performing the measurement $\{|\Phi\rangle\langle\Phi|, I - |\Phi\rangle\langle\Phi|\}$. (Recall that $|\Phi\rangle = \frac{1}{\sqrt{d}} \sum_{i=1}^d |i\rangle|i\rangle$.) If $U \in O(d)$, the test always accepts; Wang shows that if the test accepts with high probability, then U is close to an orthogonal matrix.

5.2 Properties of quantum channels

Not all physical processes which occur in quantum mechanics are reversible. The mathematical framework in which the most general physically realizable operations are studied is the formalism

of *quantum channels*. A quantum channel (or “superoperator”) is a completely positive, trace-preserving linear map $\mathcal{E} : \mathcal{B}(\mathbb{C}^{d_{\text{in}}}) \rightarrow \mathcal{B}(\mathbb{C}^{d_{\text{out}}})$. Here “completely positive” means that the operator $\mathcal{E} \otimes \text{id}$ preserves positivity, where id is the identity map on some ancilla system of arbitrary dimension. A comprehensive introduction to the world of quantum channels is provided by lecture notes of Watrous [158].

There has been less work on testing properties of quantum channels than the other types of properties considered above, although the problem of discriminating between quantum channels has been considered by a number of authors (e.g. [145, 64, 135]).

5.2.1 A distance measure on channels

In the context of property testing, the first task when considering quantum channels is to define a suitable measure of distance. One approach is to use the same idea as for unitary operators, and take the distance induced by the Choi-Jamiołkowski isomorphism [56, 100]. In the case of channels, this isomorphism states that there is a bijection between the set of quantum channels $\mathcal{E} : \mathcal{B}(\mathbb{C}^{d_{\text{in}}}) \rightarrow \mathcal{B}(\mathbb{C}^{d_{\text{out}}})$ and the set of bipartite density matrices ρ on a $(d_{\text{out}} \times d_{\text{in}})$ -dimensional system such that applying the partial trace to the first subsystem of ρ leaves the maximally mixed state I/d_{in} . The bijection can be explicitly given as

$$\mathcal{E} \leftrightarrow \frac{1}{d_{\text{in}}} \sum_{i,j=1}^{d_{\text{in}}} \mathcal{E}(|i\rangle\langle j|) \otimes |i\rangle\langle j| =: \chi_{\mathcal{E}}.$$

Then one distance measure that can be put on quantum channels \mathcal{E}, \mathcal{F} is

$$D(\mathcal{E}, \mathcal{F}) := D(\chi_{\mathcal{E}}, \chi_{\mathcal{F}}).$$

As with the correspondence between unitary operators and pure states, this distance measure allows one to translate tests for properties of mixed states to properties of channels. For example, consider the property **Unitarity**, where $\mathcal{E} : \mathcal{B}(\mathbb{C}^d) \rightarrow \mathcal{B}(\mathbb{C}^d)$ satisfies the property if and only if it is a unitary operator. \mathcal{E} is unitary if and only if $\chi_{\mathcal{E}}$ is a pure state (and hence maximally entangled). In order to test this property, we can apply the test for **Purity** of mixed states to $\chi_{\mathcal{E}}$. From the analysis of Section 4.2, we see that if the test accepts with probability $1 - \delta$, there exists a pure state $|\psi\rangle$ such that $D(\chi_{\mathcal{E}}, |\psi\rangle\langle\psi|) = O(\delta)$. We still need to show that $\chi_{\mathcal{E}}$ is in fact close to a pure state which is maximally entangled. To do so, first write $|\psi\rangle = \sum_{i=1}^d \sqrt{\lambda_i} |v_i\rangle |w_i\rangle$ for the Schmidt decomposition of $|\psi\rangle$, and define the maximally entangled state $|\eta\rangle = \frac{1}{\sqrt{d}} \sum_{i=1}^d |v_i\rangle |w_i\rangle$.

Then we have the sequence of inequalities and equalities

$$\begin{aligned} D(\chi_{\mathcal{E}}, |\psi\rangle\langle\psi|) &\geq D(I/d, \text{tr}_B(|\psi\rangle\langle\psi|)) \geq 1 - F(I/d, \text{tr}_B(|\psi\rangle\langle\psi|)) = 1 - \frac{1}{\sqrt{d}} \sum_{i=1}^d \sqrt{\lambda_i} \\ &= 1 - |\langle\psi|\eta\rangle| \geq D(|\psi\rangle\langle\psi|, |\eta\rangle\langle\eta|)^2/2. \end{aligned}$$

The first inequality holds because the trace norm does not increase under partial trace [129, Theorem 9.2], and the second is (5). Therefore, if the test accepts with probability $1 - \delta$,

$$D(\chi_{\mathcal{E}}, |\eta\rangle\langle\eta|) \leq D(\chi_{\mathcal{E}}, |\psi\rangle\langle\psi|) + D(|\psi\rangle\langle\psi|, |\eta\rangle\langle\eta|) = O(\delta + \sqrt{2\delta}) = O(\sqrt{\delta}),$$

implying that **Unitarity** of a quantum channel can be ϵ -tested with $O(1/\epsilon^2)$ uses of the channel.

5.2.2 Testing quantum measurements

An important special case of quantum channels is the case of *quantum measurements*. In full generality, a measurement on a d -dimensional quantum mechanical system is defined by a sequence of linear operators $M = (M_1, \dots, M_k)$ such that $\sum_{i=1}^k M_i^\dagger M_i = I$. If M is performed on the state ρ , the probability of receiving outcome i is $\text{tr}(M_i \rho M_i^\dagger)$, and the resulting state of the system, given that outcome i occurred, is

$$\rho_i = \frac{M_i \rho M_i^\dagger}{\text{tr}(M_i \rho M_i^\dagger)}.$$

The quantum channel corresponding to performing the measurement M and storing the outcome in a new register is the map \mathcal{M} where

$$\mathcal{M}(\rho) = \sum_{i=1}^k M_i \rho M_i^\dagger \otimes |i\rangle\langle i|,$$

so the Choi-Jamiołkowski state is

$$\chi_{\mathcal{M}} = \frac{1}{d} \sum_{i,j=1}^d \left(\sum_{\ell=1}^k M_\ell |i\rangle\langle j| M_\ell^\dagger \otimes |\ell\rangle\langle \ell| \right) \otimes |i\rangle\langle j|$$

which, by reordering subsystems, is equivalent to

$$\sum_{\ell=1}^k \left(\frac{1}{\sqrt{d}} \sum_{i=1}^d M_\ell |i\rangle\langle i| \right) \left(\frac{1}{\sqrt{d}} \sum_{j=1}^d M_\ell^\dagger |j\rangle\langle j| \right) \otimes |\ell\rangle\langle \ell| =: \sum_{\ell=1}^k |\psi_M^{(\ell)}\rangle\langle \psi_M^{(\ell)}| \otimes |\ell\rangle\langle \ell|.$$

For any two measurements M and N with at most k outcomes, we can therefore compute the distance between the corresponding channels as

$$D(\mathcal{M}, \mathcal{N}) = \sum_{\ell=1}^k D \left(|\psi_M^{(\ell)}\rangle, |\psi_N^{(\ell)}\rangle \right),$$

where if M (resp. N) has $\ell < k$ outcomes, we set $M_i = 0$ (resp. $N_i = 0$) for $\ell < i \leq k$. Observe that, using this measure of distance, we take into account the distance of the post-measurement states as well as the distance between the probability distributions corresponding to the measurement outcomes. One can explicitly calculate that, for any (potentially unnormalized) vectors $|\psi\rangle, |\phi\rangle$,

$$D(|\psi\rangle, |\phi\rangle) = \sqrt{\frac{1}{4} (\langle \psi | \psi \rangle + \langle \phi | \phi \rangle)^2 - |\langle \psi | \phi \rangle|^2},$$

which implies that

$$D(\mathcal{M}, \mathcal{N}) = \frac{1}{2} \sum_{i=1}^k \sqrt{(\langle M_i, M_i \rangle + \langle N_i, N_i \rangle)^2 - 4|\langle M_i, N_i \rangle|^2}.$$

Recent work by Wang [157] has given efficient tests for a number of properties of quantum measurements, but with respect to a measure of distance which appears somewhat different to the measure $D(\cdot, \cdot)$. Given two measurements M and N with at most k outcomes, Wang's distance measure is

$$\Delta(M, N) := \sqrt{\frac{1}{2} \sum_{i=1}^k \langle M_i, M_i \rangle + \langle N_i, N_i \rangle - 2|\langle M_i, N_i \rangle|}.$$

Wang demonstrates that $\Delta(\cdot, \cdot)$ has a number of desirable properties, including satisfying the triangle inequality and being an “average-case” measure of distance [157]. It turns out that $\Delta(\cdot, \cdot)$ is in fact closely related to $D(\cdot, \cdot)$, which we encapsulate as the following lemma.

Lemma 25. *Given two measurements M and N , let \mathcal{M} and \mathcal{N} be the corresponding channels. Then*

$$D(\mathcal{M}, \mathcal{N})/\sqrt{2} \leq \Delta(M, N) \leq D(\mathcal{M}, \mathcal{N})^{1/2}.$$

Proof. To prove the upper bound part of the lemma, it suffices to show that, for each i ,

$$(\langle M_i, M_i \rangle + \langle N_i, N_i \rangle - 2|\langle M_i, N_i \rangle|)^2 \stackrel{?}{\leq} (\langle M_i, M_i \rangle + \langle N_i, N_i \rangle)^2 - 4|\langle M_i, N_i \rangle|^2.$$

Setting $x_i = \langle M_i, M_i \rangle + \langle N_i, N_i \rangle$, $y_i = 2|\langle M_i, N_i \rangle|$ and rearranging terms, we get the claimed inequality

$$(x_i - y_i)^2 \stackrel{?}{\leq} (x_i - y_i)(x_i + y_i),$$

which holds because $y_i \leq x_i$ by Cauchy-Schwarz or the inequality of arithmetic and geometric means. For the lower bound, we need to show

$$\frac{1}{2\sqrt{2}} \sum_{i=1}^k (x_i - y_i)^{1/2} (x_i + y_i)^{1/2} \stackrel{?}{\leq} \sqrt{\frac{1}{2} \sum_{i=1}^k (x_i - y_i)}.$$

Indeed, by Cauchy-Schwarz,

$$\begin{aligned} \frac{1}{2\sqrt{2}} \sum_{i=1}^k (x_i - y_i)^{1/2} (x_i + y_i)^{1/2} &\leq \frac{1}{2\sqrt{2}} \sqrt{\sum_{i=1}^k (x_i - y_i)} \sqrt{\sum_{i=1}^k (x_i + y_i)} \\ &\leq \sqrt{\frac{1}{2} \sum_{i=1}^k (x_i - y_i)} \sqrt{\frac{1}{2} \sum_{i=1}^k x_i} \\ &= \sqrt{\frac{1}{2} \sum_{i=1}^k (x_i - y_i)} \end{aligned}$$

as required, using $\sum_{i=1}^k \langle M_i, M_i \rangle = \sum_{i=1}^k \langle N_i, N_i \rangle = 1$. \square

Lemma 25 implies that Wang’s results with respect to the distance measure $\Delta(\cdot, \cdot)$ can be translated into results with respect to $D(\cdot, \cdot)$. In particular, Wang [157] gives efficient testers for the following properties of quantum measurements:

- The property of being a **Pauli** measurement (called “stabilizer measurement” in [157]): M is a Pauli measurement if it is a two-outcome projective measurement onto the ± 1 eigenspaces of an n -qubit Pauli operator σ_s , for some $s \in \{I, X, Y, Z\}^n$. Wang shows that this property can be ϵ -tested with $O(1/\epsilon^4)$ measurements.
- The property of being an ℓ -**local** measurement of n qubits, i.e., acting non-trivially on at most ℓ qubits. Wang gives an ϵ -tester for this property which uses $O(\ell \log(\ell/\epsilon)/\epsilon^2)$ measurements.

- The property of being a **Permutation invariant** measurement of n d -dimensional systems, i.e., a measurement which is unchanged when the n systems are permuted arbitrarily. This property can be ϵ -tested with $O(1/\epsilon^2)$ measurements.
- Being contained within any finite set of measurements $\mathcal{S} = \{M_i\}$ with k outcomes on a d -dimensional system. If $\Delta(M_i, M_j) \geq \gamma$ for all $i \neq j$, and we set $\delta = \min\{\gamma, \epsilon\}$, membership in \mathcal{S} can be ϵ -tested with $O(k^2(\log k)/\delta^8 + (\log |S|)/\delta^2)$ measurements.
- **Equality** of measurements, which can be ϵ -tested with $O(k^5(\log k)/\epsilon^{12})$ measurements. This is based on a more general algorithm for estimating the distance between measurements.

All of the above testers are based on constructing multiple copies of the Choi-Jamiołkowski state corresponding to the measurement to be tested, and performing some measurements on the states. As remarked in [157], it is an interesting question whether efficient testers can be designed in a setting where one is not allowed access to entanglement.

Question 11. *Can efficient testers for the properties of unitary operators and quantum channels discussed above be designed which do not require entanglement with an ancilla system?*

It is possible to use quantum process tomography to completely characterize any quantum channel without the use of entanglement [129, §8.4.2], so the question is only whether the above properties can still be tested *efficiently* in this setting.

6 Quantum properties and computational complexity

Classically, the field of property testing has had close connections to computational complexity. In this section, we briefly discuss three ways in which quantum property testing can be related to quantum computational complexity. First, we discuss how, if we change the setting in which we work, testing certain natural properties can be proven computationally hard. Second, we mention how quantum property testers can be used to prove complexity class inclusions. Finally, we consider potential connections between quantum property testing and a proposed quantum PCP conjecture.

6.1 Computational hardness of testing quantum properties

A different perspective from which to study the question of testing properties of quantum systems is to consider problems where, instead of being given access to a quantum object, we are given a concise *classical* description of that object (for example, a quantum circuit on n qubits with $\text{poly}(n)$ gates). Our aim is to efficiently determine whether the corresponding quantum object has some property, or is far from having that property, in terms of some distance measure. The distance measure used may be quite different to the distances we discuss in the rest of the survey, leading to qualitatively different results. This type of problem turns out to be naturally addressed via the framework of computational complexity.

In particular, many problems related to testing properties of quantum circuits turn out to be QMA-complete.²⁴ These hardness results provide an interesting counterpoint to the largely positive results obtained in the “average-case” scenarios considered by property testing. A prototypical example of this phenomenon is “non-identity-check,” which was proven to be QMA-complete by

²⁴QMA is the quantum analog of NP [107]; see [38] for a recent survey.

Janzing et al. [101]. Here the input is a quantum circuit implementing a unitary U , and two numbers a, b such that $b - a \geq 1/\text{poly}(n)$, and the problem is to distinguish between the two cases that $\min_{\theta \in \mathbb{R}} \|U - e^{i\theta} I\| \leq a$ and $\min_{\theta \in \mathbb{R}} \|U - e^{i\theta} I\| \geq b$. Observe that, if we replace the operator norm with the normalized 2-norm in this definition, this problem is in BQP by the efficient tester for the **Equality to V** property discussed in Section 5.1.3.

If one generalizes to quantum circuits acting on mixed states, where each elementary gate is a quantum channel, some natural problems even become PSPACE-complete. In particular, Rosgen and Watrous [144] showed that PSPACE-completeness holds for the problem of testing whether two mixed-state quantum circuits are distinguishable, and it remains hard when the quantum circuits are restricted to be logarithmic depth [142], degradable or anti-degradable [143]. In this case, distinguishability is measured in the so-called diamond norm for quantum channels [107]; the diamond norm of a linear operator $\Phi : \mathcal{B}(\mathbb{C}^{d_{\text{in}}}) \rightarrow \mathcal{B}(\mathbb{C}^{d_{\text{out}}})$ is defined to be

$$\|\Phi\|_{\diamond} := \max_{X, \|X\|_1=1} \|(\Phi \otimes \text{id})(X)\|_1, \quad (7)$$

where id is the identity map acting on an ancilla system, which may be taken to be of dimension at most d_{in} . Then the Quantum Circuit Distinguishability problem is to determine, given two mixed-state quantum circuits Q_0, Q_1 and constants $a < b$, whether $\|Q_0 - Q_1\|_{\diamond} \leq a$ or $\|Q_0 - Q_1\|_{\diamond} \geq b$. As with the trace distance between quantum states, $\|Q_0 - Q_1\|_{\diamond}$ can be seen as measuring the distinguishability of Q_0 and Q_1 in a “best-case” scenario. This contrasts with the “average-case” distance measure $D(Q_0, Q_1)$ introduced in Section 5.2.1.

These distinguishability problems were originally shown to be hard for the complexity class QIP of languages decided by quantum interactive proof systems, but this class was later proven to equal PSPACE [99]. The proof technique of [144] starts by using a result of Kitaev and Watrous [108], which states that all quantum interactive proofs can be parallelized to three rounds. A mathematical reformulation of this result is that the Close Images problem is QIP-hard. This problem is defined as follows: given two quantum circuits Q_0, Q_1 and constants $a < b$, distinguish between the cases that there is an input ρ such that $F(Q_0(\rho), Q_1(\rho)) \geq b$, or that for all inputs ρ , $F(Q_0(\rho), Q_1(\rho)) \leq a$. Hardness of Quantum Circuit Distinguishability is then shown by a reduction from Close Images [144].

6.2 From quantum property testers to complexity class inclusions

By contrast to the results in the previous section, work by Hayden et al. [90] demonstrates that quantum property testers can be used to prove positive results (i.e., upper bounds) regarding the complexity of testing properties of quantum circuits. The problem considered by these authors is a variant of the separability-testing problem (cf. Sections 4.1.2 and 4.2). In this variant the input is the description of a mixed-state quantum circuit Q on n qubits, and one considers the output of the circuit as a bipartite state by dividing these qubits into two disjoint sets. The problem is to distinguish between the two cases that: (a) the output of Q , when applied to the input $|0^n\rangle$, is close to separable; (b) the output is far from separable. Hayden et al. [90] show that this problem can be solved by a quantum interactive proof system with two messages (i.e., a message from verifier to prover, followed by a reply from prover to verifier), and hence sits in the complexity class QIP(2). The protocol is based on the verifier applying the permutation test discussed in Section 4.1.1. This result is somewhat subtle in that “close” and “far” are defined asymmetrically (the former in terms of the trace distance, the latter in terms of the so-called “1-way LOCC” distance); see [90] for details.

More recently Gutoski et al. [83] generalized this work: for almost every complexity class defined by quantum interactive proofs, they give a version of the separability testing problem which is complete for that class. This shows that property testing problems can be used to characterize many quantum complexity classes. For example, they use the product test of [88] (see Section 4.1.2) to show that testing whether the output of a pure-state quantum circuit is a product state is in BQP.

6.3 The quantum PCP conjecture

A classic and important problem in quantum computational complexity is the *local Hamiltonian problem*. Here we are given as input a Hamiltonian H on n qubits, described by a set of Hermitian operators H_i such that $H = \sum_{i=1}^m H_i$, with each operator H_i acting non-trivially on at most $k = O(1)$ qubits and satisfying $\|H_i\| = O(1)$. We are also given two real numbers a and b such that $b - a \geq 1/\text{poly}(n)$. We are promised that the lowest eigenvalue of H is either smaller than a , or larger than b ; our task is to determine which of these is the case.

This problem was proven QMA-complete for $k = 5$ by Kitaev [107], which was later improved to $k = 2$ by Kempe et al. [105] (the case where $k = 1$ is easily seen to be in P). One way in which this hardness result could potentially be improved is in the scaling of the gap between b and a . Indeed, it could be the case that the local Hamiltonian problem remains QMA-hard if we have the promise $b - a \geq cm$ for some constant $0 < c < 1$. This is (one formulation of) the quantum PCP conjecture; see a recent survey of Aharonov et al. [10] for much more on this conjecture and its implications. Classically, one version of the famous PCP Theorem states that there exist constraint satisfaction problems for which it is hard to distinguish between there existing an assignment to the variables that satisfies all of the constraints, and there being no assignment that satisfies more than a constant fraction of them; the quantum PCP conjecture would be a direct quantization of this result. One way of looking at this is as the conjecture that the local Hamiltonian problem remains hard in a “property-testing-type” scenario where there is a large gap between “yes” and “no” instances.

Question 12. *Is there a quantum PCP theorem?*

Classically, the proof of the PCP Theorem relied on efficient property testers, so it seems plausible that property testing could be useful in proving a quantum generalization. Indeed, the analysis of a *classical* property tester in a quantum setting has recently been central to establishing a quantum complexity-theoretic result. MIP is the class of languages decided by multiple-prover interactive proof systems, which was shown to be equal to NEXP by Babai et al. [23]. Recently Ito and Vidick [98] have shown that the quantum generalization MIP*, where the provers are allowed to share entanglement, is at least as powerful: $\text{MIP} \subseteq \text{MIP}^*$. Their proof is based on proving soundness of the classical multilinearity test of Babai et al. [23] in the presence of entanglement. Another application of quantum property testing to quantum complexity is the use of the analysis of an efficient quantum property tester to prove the complexity class equality $\text{QMA}(k) = \text{QMA}(2)$ [88], as discussed in Section 4.1.2.

Yet another connection is explored in recent work of Aharonov and Eldar [11] on a quantum generalization of *locally testable codes* (LTCs). Classically, LTCs are codes for which the property of being a codeword can be tested efficiently by means of a few local checks; such codes played a crucial role in the original proof of the PCP Theorem [19]. The “qLTCs” studied in [11] are the kernel (zero eigenspace) of k -local Hamiltonians $H = \sum_i H_i$, such that containment of a state in the eigenspace can be tested with good accuracy by performing measurements corresponding to

only a few of the individual k -local terms H_i . Aharonov and Eldar [11] prove some surprising upper bounds on the soundness for qLTCs that are *stabilizer* codes, showing that they do not exist in certain regimes of parameters where classical LTCs do exist.

7 Conclusion

The goal of property testing is to design efficient algorithms (“testers”) to decide whether a given object has a property or is somehow “far” from that property, and to determine in which cases such algorithms can exist. When the objects that need to be tested are very large, exact algorithms that are also required to work for objects that “almost” have the property become infeasible, and property testing is often the best we can hope for. Classical property testing is by now a very well-developed area, but *quantum* property testing is just starting out. In this paper we surveyed what is known about this:

1. Quantum testers for classical properties (Section 2).
2. Classical testers for quantum properties (Section 3).
3. Quantum testers for quantum properties (Sections 4 and 5).

We hope the overview given here, as well as the open questions mentioned along the way, will give rise to much more research in this area. Besides the properties mentioned here, there are many other properties which have been of great interest in the classical property testing literature, and whose quantum complexity is unknown. Examples include monotonicity of Boolean functions, membership of error-correcting codes, and almost all properties of graphs. In the case of quantum properties, natural targets include testing whether a unitary operator is implemented by a small circuit, and whether a Hamiltonian is k -local (which would be yet another variant of junta testing).

Another very broad open question not discussed previously is to what extent one can *characterize* the properties (classical or quantum) that have efficient quantum testers. This may seem a hopelessly ambitious goal; nevertheless, in the case of classical algorithms it has already been achieved in some important cases, such as graph properties [13] and symmetric properties of probability distributions [153]. Such a characterization could have importance far beyond property testing, by shedding light on the structure of problems that have efficient quantum algorithms.

Acknowledgements

We thank Scott Aaronson, Aleksandrs Belovs, Robin Blume-Kohout, Sourav Chakraborty, Wim van Dam, Aram Harrow, Frédéric Magniez, Marcelo Marchioli, Miguel Navascués, Falk Unger, Lev Vaidman, Andreas Winter and Tzyh Haur Yang for helpful comments, answers to questions and pointers to literature. We also thank the anonymous ToC referees for their exceptionally constructive feedback.

References

- [1] S. Aaronson. QMA/qpoly \subseteq PSPACE/poly: De-Merlinizing quantum protocols. In *Proceedings of 21st Annual IEEE Conference on Computational Complexity (CCC)*, pages 261–273, 2006. [quant-ph/0510230](https://arxiv.org/abs/quant-ph/0510230).
- [2] S. Aaronson. BQP and the Polynomial Hierarchy. In *Proceedings of 42nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 141–150, 2010. [arXiv:0910.4698](https://arxiv.org/abs/0910.4698).

- [3] S. Aaronson and A. Ambainis. The need for structure in quantum speedups. *Theory of Computing*, 10:133–166, 2014. Earlier version in ICS’11. [arXiv:0911.0996](#).
- [4] S. Aaronson and A. Ambainis. Forrelation: A problem that optimally separates quantum from classical computing, 21 Nov 2014. [arXiv:1411.5729](#).
- [5] S. Aaronson, S. Beigi, A. Drucker, B. Fefferman, and P. Shor. The power of unentanglement. *Theory of Computing*, 5(1):1–42, 2009. [arXiv:0804.0802](#).
- [6] S. Aaronson and D. Gottesman. Identifying stabilizer states, 2008. <http://pirsa.org/08080052/>.
- [7] S. Aaronson and Y. Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM*, 51(4):595–605, 2004. Earlier versions in STOC’02 and FOCS’02.
- [8] A. Acín. Statistical distinguishability between unitary operations. *Physical Review Letters*, 87(17):177901, 2001. [quant-ph/0102064](#).
- [9] A. Acín, N. Brunner, N. Gisin, S. Massar, S. Pironio, and V. Scarani. Device-independent security of quantum cryptography against collective attacks. *Physical Review Letters*, 98:230501, 2008.
- [10] D. Aharonov, I. Arad, and T. Vidick. The quantum PCP conjecture. *ACM SIGACT News*, 44(2):47–79, 2013. Enhanced version at [arXiv:1309.7495](#).
- [11] D. Aharonov and L. Eldar. Quantum locally testable codes, Oct 21 2013. [arXiv:1310.5664](#).
- [12] D. Aharonov, A. Harrow, Z. Landau, D. Nagaj, M. Szegedy, and U. Vazirani. Local tests of global entanglement and a counterexample to the generalized area law. In *Proceedings of 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 246–255, 2014. [arXiv:1410.0951](#).
- [13] N. Alon, E. Fischer, I. Newman, and A. Shapira. A combinatorial characterization of the testable graph properties: it’s all about regularity. *SIAM Journal on Computing*, 39(1):143–167, 2009.
- [14] N. Alon, T. Kaufman, M. Krivelevich, S. Litsyn, and D. Ron. Testing Reed-Muller codes. *IEEE Transactions on Information Theory*, 51(11):4032–4038, 2005.
- [15] A. Ambainis. Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64(4):750–767, 2002. Earlier version in STOC’00. [quant-ph/0002066](#).
- [16] A. Ambainis. Quantum walk algorithm for element distinctness. *SIAM Journal on Computing*, 37(1):210–239, 2007. Earlier version in FOCS’04. [quant-ph/0311001](#).
- [17] A. Ambainis, A. Belovs, O. Regev, and R. de Wolf. Better quantum junta testers, 2014. Manuscript.
- [18] A. Ambainis, A. Childs, and Y-K. Liu. Quantum property testing for bounded-degree graphs. In *Proceedings of 15th RANDOM*, volume 6845 of *Lecture Notes in Computer Science*, pages 365–376, 2011. [arXiv:1012.3174](#).
- [19] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998. Earlier version in FOCS’92.
- [20] A. Atıcı and R. Servedio. Quantum algorithms for learning and testing juntas. *Quantum Information Processing*, 6(5):323–348, 2009. [arXiv:0707.3479](#).
- [21] K. Audenaert. A digest on representation theory of the symmetric group, 2006. Available at http://personal.rhul.ac.uk/usah/080/QITNotes_files/Irreps_v06.pdf.
- [22] K. Audenaert, M. Nussbaum, A. Szkoła, and F. Verstraete. Asymptotic error rates in quantum hypothesis testing. *Communications in Mathematical Physics*, 279:251–283, 2008. [arXiv:0708.4282](#).
- [23] L. Babai, L. Fortnow, and C. Lund. Nondeterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1:3–40, 1991. Earlier version in FOCS’90.
- [24] D. Bacon, I. Chuang, and A. Harrow. Efficient quantum circuits for Schur and Clebsch-Gordan transforms. *Physical Review Letters*, 97(17):170502, 2006. [quant-ph/0407082](#).
- [25] A. Barenco, A. Berthiaume, D. Deutsch, A. Ekert, R. Jozsa, and C. Macchiavello. Stabilisation of quantum computations by symmetrisation. *SIAM Journal on Computing*, 26(5):1541–1557, 1997. [quant-ph/9604028](#).
- [26] S. Barnett and S. Croke. Quantum state discrimination. *Advances in Optics and Photonics*, 1(2):238–278, 2009. [arXiv:0810.1970](#).

- [27] J. Barrett, L. Hardy, and A. Kent. No signaling and quantum key distribution. *Physical Review Letters*, 95(1):010503, June 2005.
- [28] T. Batu, L. Fortnow, E. Fischer, R. Kumar, R. Rubinfeld, and P. White. Testing random variables for independence and identity. In *Proceedings of 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 442–451, 2001.
- [29] T. Batu, L. Fortnow, R. Rubinfeld, W. Smith, and P. White. Testing closeness of discrete distributions. *Journal of the ACM*, 60(1), 2013. Previous version in FOCS’00. [arXiv:1009.5397](#).
- [30] R. Beals. Quantum computation of Fourier transforms over symmetric groups. In *Proceedings of 29th Annual ACM Symposium on Theory of Computing (STOC)*, pages 48–53, 1997.
- [31] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf. Quantum lower bounds by polynomials. *Journal of the ACM*, 48(4):778–797, 2001. Earlier version in FOCS’98. [quant-ph/9802049](#).
- [32] M. Bellare, D. Coppersmith, J. Høastad, M. Kiwi, and M. Sudan. Linearity testing in characteristic two. *IEEE Transactions on Information Theory*, 42(6), 1996.
- [33] C. Bennett, E. Bernstein, G. Brassard, and U. Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26(5):1510–1523, 1997. [quant-ph/9701001](#).
- [34] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997. Earlier version in STOC’93.
- [35] E. Blais. Testing juntas nearly optimally. In *Proceedings of 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 151–158, 2009.
- [36] E. Blais, J. Brody, and K. Matulef. Property testing lower bounds via communication complexity. *Computational Complexity*, 21(2):311–358, 2012. Earlier version in Complexity’11.
- [37] M. Blum, M. Luby, and R. Rubinfeld. Self-testing/correcting with applications to numerical problems. *Journal of Computer and System Sciences*, 47(3):549–595, 1993. Earlier version in STOC’90.
- [38] A. Bookatz. QMA-complete problems. *Quantum Information and Computation*, 14(5-6):361–383, 2014. [arXiv:1212.6312](#).
- [39] G. Brassard and P. Høyer. An exact quantum polynomial-time algorithm for Simon’s problem. In *Proceedings of the 5th Israeli Symposium on Theory of Computing and Systems (ISTCS’97)*, pages 12–23, 1997. [quant-ph/9704027](#).
- [40] G. Brassard, P. Høyer, M. Mosca, and A. Tapp. Quantum amplitude amplification and estimation. In *Quantum Computation and Quantum Information: A Millennium Volume*, volume 305 of *AMS Contemporary Mathematics Series*, pages 53–74. 2002. [quant-ph/0005055](#).
- [41] S. Braunstein, A. Mann, and M. Revzen. Maximal violation of Bell inequalities for mixed states. *Physical Review Letters*, 68(22):3259–3261, 1992.
- [42] S. Bravyi, A. Hassidim, and A. Harrow. Quantum algorithms for testing properties of distributions. *IEEE Transactions on Information Theory*, 57(6):3971–3981, 2011. Earlier version in STACS’10. [arXiv:0907.3920](#).
- [43] T. Brun. Measuring polynomial functions of states. *Quantum Information and Computation*, 4:401, 2004. [quant-ph/0401067](#).
- [44] N. Brunner, D. Cavalcanti, S. Pironio, V. Scarani, and S. Wehner. Bell nonlocality. *Reviews of Modern Physics*, 86:419–478, 2014. [arXiv:1303.2849](#).
- [45] D. Bruß and C. Macchiavello. Optimal state estimation for d -dimensional quantum systems. *Physical Review A*, 253(5–6):249–251, 1999. [quant-ph/9812016](#).
- [46] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16):167902, 2001. [quant-ph/0102001](#).
- [47] H. Buhrman, L. Fortnow, I. Newman, and H. Röhrig. Quantum property testing. *SIAM Journal on Computing*, 37(5):1387–1400, 2008. Earlier version in SODA’03. [quant-ph/0201117](#).
- [48] H. Buhrman, D. García-Soriano, A. Matsliah, and R. de Wolf. The non-adaptive query complexity of testing k -parities. *Chicago Journal of Theoretical Computer Science*, 6, 2013. [arXiv:1209.3849](#).

- [49] H. Buhrman and R. de Wolf. Complexity measures and decision tree complexity: A survey. *Theoretical Computer Science*, 288(1):21–43, 2002.
- [50] K. Chakraborty and S. Maitra. Improved quantum test for linearity of a Boolean function, 2013. [arXiv:1306.6195](#).
- [51] 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](#).
- [52] S.-O. Chan, I. Diakonikolas, G. Valiant, and P. Valiant. Optimal algorithms for testing closeness of discrete distributions. In *Proceedings of 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1193–1203, 2014. [arXiv:1308.3946](#).
- [53] A. Chefles. Quantum state discrimination. *Contemporary Physics*, 41(6):401–424, 2001. [quant-ph/0010114](#).
- [54] A. Childs, A. Harrow, and P. Wocjan. Weak Fourier-Schur sampling, the hidden subgroup problem, and the quantum collision problem. In *Proceedings of 24th Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 598–609, 2007. [quant-ph/0609110](#).
- [55] H. Chockler and D. Gutfreund. A lower bound for testing juntas. *Information Processing Letters*, 90(6):301–305, 2004.
- [56] M.-D. Choi. Completely positive linear maps on complex matrices. *Linear Algebra and its Applications*, 10(3):285–290, 1975.
- [57] M. Christandl. *The Structure of Bipartite Quantum States – Insights from Group Theory and Cryptography*. PhD thesis, University of Cambridge, 2006. [quant-ph/0604183](#).
- [58] B. S. Cirel’son. Quantum generalizations of Bell’s inequality. *Letters in Mathematical Physics*, 4(2):93–100, 1980.
- [59] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt. Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23(15):880–884, 1969.
- [60] R. Colbeck. *Quantum and relativistic protocols for secure multi-party computation*. PhD thesis, University of Cambridge, 2006. [arXiv:0911.3814](#).
- [61] M. Cramer, M. Plenio, S. Flammia, R. Somma, D. Gross, S. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu. Efficient quantum state tomography. *Nature Communications*, 1(9):49, 2010. [arXiv:1101.4366](#).
- [62] W. van Dam, F. Magniez, M. Mosca, and M. Santha. Self-testing of universal and fault-tolerant sets of quantum gates. *SIAM Journal on Computing*, 37(2):611–629, 2007. Earlier version in STOC’00. [quant-ph/9904108](#).
- [63] A. Dasgupta, R. Kumar, and D. Sivakumar. Sparse and lopsided set disjointness via information theory. In *Proceedings of 16th RANDOM*, pages 517–528, 2012.
- [64] R. Duan, Y. Feng, and M. Ying. Perfect distinguishability of quantum operations. *Physical Review Letters*, 103:210501, 2009. [arXiv:0908.0119](#).
- [65] E. Fischer. The art of uninformed decisions. *Bulletin of the EATCS*, 75:97, 2001.
- [66] E. Fischer, G. Kindler, D. Ron, S. Safra, and A. Samorodnitsky. Testing juntas. *Journal of Computer and System Sciences*, 68(4):753–787, 2004. Earlier version in FOCS’02.
- [67] S. Flammia, D. Gross, Y.-K. Liu, and J. Eisert. Quantum tomography via compressed sensing: Error bounds, sample complexity, and efficient estimators. *New Journal of Physics*, 14:095022, 2012. [arXiv:1205.2300](#).
- [68] S. Flammia and Y.-K. Liu. Direct fidelity estimation from few Pauli measurements. *Physical Review Letters*, 106:230501, 2011. [arXiv:1104.4695](#).
- [69] K. Friedl, G. Ivanyos, and M. Santha. Efficient testing of groups. In *Proceedings of 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 157–166, 2005.
- [70] K. Friedl, F. Magniez, M. Santha, and P. Sen. Quantum testers for hidden group properties. *Fundamenta Informaticae*, 91(2):325–340, 2009. Earlier version in MFCS’03.
- [71] S. Gharibian. Strong NP-hardness of the quantum separability problem. *Quantum Information and Computation*, 10(3&4):343–360, 2010. [arXiv:0810.4507](#).
- [72] L. Glebsky. Almost commuting matrices with respect to normalized Hilbert-Schmidt norm, 2010. [arXiv:1002.3082](#).

- [73] O. Goldreich, editor. *Property Testing: Current Research and Surveys*, volume 6390 of *Lecture Notes in Computer Science*. Springer, 2010.
- [74] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, 45(4):653–750, 1998.
- [75] O. Goldreich and D. Ron. On testing expansion in bounded-degree graphs. *Electronic Colloquium on Computational Complexity (ECCC)*, 7(20), 2000.
- [76] O. Goldreich and D. Ron. Property testing in bounded degree graphs. *Algorithmica*, 32(2):302–343, 2002.
- [77] D. Gottesman. *Stabilizer Codes and Quantum Error Correction*. PhD thesis, Caltech, 1999. [quant-ph/9705052](#).
- [78] D. Greenberger, M. Horne, and A. Zeilinger. Going beyond Bell’s theorem. In M. Kafatos, editor, *Bell’s Theorem, Quantum Theory, and Conceptions of the Universe*, pages 69–72. Kluwer, 1989. [arXiv:0712.0921](#).
- [79] D. Gross, Y.-K. Liu, S. Flammia, S. Becker, and J. Eisert. Quantum state tomography via compressed sensing. *Physical Review Letters*, 105:150401, 2010. [arXiv:0909.3304](#).
- [80] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of 28th Annual ACM Symposium on Theory of Computing (STOC)*, pages 212–219, 1996. [quant-ph/9605043](#).
- [81] O. Gühne and G. Toth. Entanglement detection. *Physics Reports*, 471(1), 2009. [arXiv:0811.2803](#).
- [82] L. Gurvits. Classical deterministic complexity of Edmonds’ problem and quantum entanglement. In *Proceedings of 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 10–19, 2003. [quant-ph/0303055](#).
- [83] G. Gutoski, P. Hayden, K. Milner, and M. Wilde. Quantum interactive proofs and the complexity of separability testing. *Theory of Computing*, 2015. To appear. [arXiv:1308.5788](#).
- [84] H. Häffner, W. Hänsel, C. Roos, J. Benhelm, D. Chek-al-kar, M. Chwalla, T. Körber, U. Rapol, M. Riebe, P. Schmidt, C. Becher, O. Gühne, W. Dür, and R. Blatt. Scalable multiparticle entanglement of trapped ions. *Nature*, 438:643–646, 2005. [quant-ph/0603217](#).
- [85] L. Hales. *The Quantum Fourier Transform and Extensions of the Abelian Hidden Subgroup Problem*. PhD thesis, University of California, Berkeley, 2002. [quant-ph/0212002](#).
- [86] L. Hales and S. Hallgren. An improved quantum Fourier transform algorithm and applications. In *Proceedings of 41st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 515–525, 2000.
- [87] A. Harrow. *Applications of coherent classical communication and the Schur transform to quantum information theory*. PhD thesis, Massachusetts Institute of Technology, 2005. [quant-ph/0512255](#).
- [88] A. Harrow and A. Montanaro. Testing product states, quantum Merlin-Arthur games and tensor optimization. *Journal of the ACM*, 60(1), 2013. Earlier version in FOCS’10. [arXiv:1001.0017](#).
- [89] P. Hayden, D. Leung, and A. Winter. Aspects of generic entanglement. *Communications in Mathematical Physics*, 265(1):95–117, 2006. [quant-ph/0407049](#).
- [90] P. Hayden, K. Milner, and M. Wilde. Two-message quantum interactive proofs and the quantum separability problem. In *Proceedings of 28th Annual IEEE Conference on Computational Complexity (CCC)*, pages 156–167, 2013. [arXiv:1211.6120](#).
- [91] C. W. Helstrom. *Quantum detection and estimation theory*. Academic Press, New York, 1976.
- [92] M. Hillery and E. Andersson. Quantum tests for the linearity and permutation invariance of Boolean functions. *Physical Review A*, 84:062329, 2011. [arXiv:1106.4831](#).
- [93] A. S. Holevo. Bounds for the quantity of information transmitted by a quantum communication channel. *Problemy Peredachi Informatsii*, 9(3):3–11, 1973. English translation *Problems of Information Transmission*, vol. 9, pp. 177–183, 1973.
- [94] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the AMS*, 43:439–561, 2006.
- [95] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki. Quantum entanglement. *Reviews of Modern Physics*, 81:865–942, 2009. [quant-ph/0702225](#).

- [96] P. Høyer, T. Lee, and R. Špalek. Negative weights make adversaries stronger. In *Proceedings of 39th Annual ACM Symposium on Theory of Computing (STOC)*, pages 526–535, 2007. [quant-ph/0611054](#).
- [97] Y. Inui and F. Le Gall. Quantum property testing of group solvability. *Algorithmica*, 59(1):35–47, 2011. Earlier version in LATIN’08. [arXiv:0712.3829](#).
- [98] T. Ito and T. Vidick. A multi-prover interactive proof for NEXP sound against entangled provers. In *Proceedings of 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 243–252, 2012. [arXiv:1207.0550](#).
- [99] R. Jain, Z. Ji, S. Upadhyay, and J. Watrous. QIP=PSPACE. *Journal of the ACM*, 58(6):30, 2011. Earlier version in STOC’10. [arXiv:0907.4737](#).
- [100] A. Jamiolkowski. Linear transformations which preserve trace and positive semidefiniteness of operators. *Reports on Mathematical Physics*, 3(4):275–278, 1972.
- [101] D. Janzing, P. Wocjan, and T. Beth. Non-identity check is QMA-complete. *International Journal of Quantum Information*, 3:463, 2005. [quant-ph/0305050](#).
- [102] M. Kada, H. Nishimura, and T. Yamakami. The efficiency of quantum identity testing of multiple states. *Journal of Physics A: Mathematical and Theoretical*, 41:395309, 2008. [arXiv:0809.2037](#).
- [103] B. Kalyanasundaram and G. Schnitger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5(4):545–557, 1992. Earlier version in Structures’87.
- [104] D. Kane and S. Kutin. Quantum interpolation of polynomials. *Quantum Information and Computation*, 1(1&2):95–103, 2011. [arXiv:0909.5683](#).
- [105] J. Kempe, A. Kitaev, and O. Regev. The complexity of the local Hamiltonian problem. *SIAM Journal on Computing*, 35(5):1070–1097, 2006. [quant-ph/0406180](#).
- [106] M. Keyl and R. Werner. Estimating the spectrum of a density operator. *Physical Review A*, 64(5):052311, 2001. [quant-ph/0102027](#).
- [107] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*, volume 47 of *Graduate Studies in Mathematics*. AMS, 2002.
- [108] A. Yu Kitaev and J. Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof systems. In *Proceedings of 32nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 608–617, 2000.
- [109] H. Kobayashi, K. Matsumoto, and T. Yamakami. Quantum Merlin-Arthur proof systems: Are multiple Merlins more helpful to Arthur? *Chicago Journal of Theoretical Computer Science*, 2009. Earlier version in ISAAC’03. [quant-ph/0306051](#).
- [110] P. Koiran, V. Nesme, and N. Portier. A quantum lower bound for the query complexity of Simon’s problem. In *Proceedings of 32nd International Colloquium on Automata, Languages and Programming (ICALP)*, volume 3580 of *Lecture Notes in Computer Science*, pages 1287–1298, 2005. [quant-ph/0501060](#).
- [111] R. Krauthgamer and O. Sasson. Property testing of data dimensionality. In *Proceedings of 14th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 18–27, 2003.
- [112] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [113] O. Lachish and I. Newman. Testing periodicity. *Algorithmica*, 60(2):401–420, 2011. Earlier version in RANDOM’05.
- [114] F. Le Gall and Y. Yoshida. Property testing for cyclic groups and beyond. *Journal of Combinatorial Optimization*, 26(4):636–654, 2013. Earlier version in COCOON’11. [arXiv:1105.1842](#).
- [115] T. Lee, R. Mittal, B. Reichardt, R. Špalek, and M. Szegedy. Quantum query complexity of state conversion. In *Proceedings of 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 344–353, 2011. [arXiv:1011.3020](#).
- [116] R. Low. Learning and testing algorithms for the Clifford group. *Physical Review A*, 80:052314, 2009. [arXiv:0907.2833](#).
- [117] F. J. MacWilliams and N. J. A. Sloane. *The Theory of Error-Correcting Codes*. North-Holland, Amsterdam, 1983.

- [118] F. Magniez, D. Mayers, M. Mosca, and H. Ollivier. Self-testing of quantum circuits. In *Proceedings of 33rd International Colloquium on Automata, Languages and Programming (ICALP)*, volume 4051 of *Lecture Notes in Computer Science*, pages 72–83, 2006. [quant-ph/0512111](#).
- [119] K. Majewski and N. Pippenger. Attribute estimation and testing quasi-symmetry. *Information Processing Letters*, 109(4):233–237, 2007. [arXiv:0708.2105](#).
- [120] D. Mayers and A. Yao. Quantum cryptography with imperfect apparatus. In *Proceedings of 39th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 503–509, 1998. [quant-ph/9809039](#).
- [121] D. Mayers and A. Yao. Self testing quantum apparatus. *Quantum Information and Computation*, 4(4):273–286, 2004. [quant-ph/0307205](#).
- [122] M. McKague. Self-testing graph states. In *Proceedings of 6th Conference on Theory of Quantum Computation, Communication, and Cryptography (TQC)*, pages 104–120, 2011. [arXiv:1010.1989](#).
- [123] M. McKague, T. H. Yang, and V. Scarani. Robust self-testing of the singlet. *Journal of Physics A: Mathematical and Theoretical*, 45(455304), 2012. [arXiv:1203.2976](#).
- [124] C. Miller and Y. Shi. Optimal robust quantum self-testing by binary nonlocal XOR games. In *Proceedings of 8th Conference on Theory of Quantum Computation, Communication, and Cryptography (TQC)*, 2013. [arXiv:1207.1819](#).
- [125] F. Mintert, M. Kuś, and A. Buchleitner. Concurrence of mixed multipartite quantum states. *Physical Review Letters*, 95(26):260502, 2005. [quant-ph/0411127](#).
- [126] A. Montanaro. Symmetric functions of qubits in an unknown basis. *Physical Review A*, 79(6):062316, 2009. [arXiv:0903.5466](#).
- [127] A. Montanaro and T. Osborne. Quantum boolean functions. *Chicago Journal of Theoretical Computer Science*, 1, 2010. [arXiv:0810.2435](#).
- [128] M. Nielsen. Continuity bounds for entanglement. *Physical Review A*, 61(6):064301, 2000. [quant-ph/9908086](#).
- [129] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [130] R. O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.
- [131] R. O’Donnell and J. Wright. Quantum spectrum testing, 2015. [arXiv:1501.05028](#).
- [132] T. Ogawa and H. Nagaoka. A new proof of the channel coding theorem via hypothesis testing in quantum information theory. In *Proceedings of 2002 IEEE International Symposium on Information Theory*, page 73, 2002. [quant-ph/0208139](#).
- [133] M. Paris and J. Řeháček (eds.). *Quantum State Estimation*, volume 649 of *Lecture Notes in Physics*. Springer-Verlag, 2004.
- [134] D. Pérez-García, F. Verstraete, M. Wolf, and J. Cirac. Matrix product state representations. *Quantum Information and Computation*, 7:401, 2007. [quant-ph/0608197](#).
- [135] M. Piani and J. Watrous. All entangled states are useful for channel discrimination. *Physical Review Letters*, 102:250501, 2009. [arXiv:0901.2118](#).
- [136] S. Popescu and D. Rohrlich. Which states violate Bell’s inequality maximally? *Physics Letters A*, 166:411–414, 1992.
- [137] R. Raussendorf, D. Browne, and H. Briegel. Measurement-based quantum computation with cluster states. *Physical Review A*, 68:022312, 2003. [quant-ph/0301052](#).
- [138] A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385–390, 1992.
- [139] B. Reichardt. Span programs and quantum query complexity: The general adversary bound is nearly tight for every Boolean function. In *Proceedings of 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 544–551, 2009. [arXiv:0904.2759](#).
- [140] B. Reichardt, F. Unger, and U. Vazirani. Classical command of quantum systems. *Nature*, 496:456–460, 2013. [arXiv:1209:0448](#) and [arXiv:1209:0449](#).
- [141] D. Ron. Property testing: A learning theory perspective. *Foundations and Trends in Machine Learning*, 1(3):307–402, 2008.
- [142] W. Rosgen. Distinguishing short quantum computations. In *Proceedings of 25th Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 597–608, 2008. [arXiv:0712.2595](#).

- [143] W. Rosgen. Computational distinguishability of degradable and antidegradable channels. *Quantum Information and Computation*, 10:735–746, 2010. [arXiv:0911.2109](#).
- [144] W. Rosgen and J. Watrous. On the hardness of distinguishing mixed-state quantum computations. In *Proceedings of 20th Annual IEEE Conference on Computational Complexity (CCC)*, pages 344–354, 2005. [cs/0407056](#).
- [145] M. Sacchi. Optimal discrimination of quantum operations. *Physical Review A*, 71:062340, 2005. [quant-ph/0505183](#).
- [146] M. Santha. Quantum walk based search algorithms. In *Proceedings of 5th conference on Theory and Applications of Models of Computation (TAMC)*, pages 31–46, 2008. [arXiv:0808.0059](#).
- [147] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997. Earlier version in FOCS’94. [quant-ph/9508027](#).
- [148] M. da Silva, O. Landon-Cardinal, and D. Poulin. Practical characterization of quantum devices without tomography. *Physical Review Letters*, 107:210404, 2011. [arXiv:1104.3835](#).
- [149] D. Simon. On the power of quantum computation. *SIAM Journal on Computing*, 26(5):1474–1483, 1997. Earlier version in FOCS’94.
- [150] R. Špalek and M. Szegedy. All quantum adversary methods are equivalent. *Theory of Computing*, 2(1):1–18, 2006. Earlier version in ICALP’05, [quant-ph/0409116](#).
- [151] S. Summers and R. Werner. Maximal violation of Bell’s inequalities is generic in quantum field theory. *Communications in Mathematical Physics*, 110(2):247–259, 1987.
- [152] B. S. Tsirelson. Some results and problems on quantum Bell-type inequalities. *Hadronic Journal Supplement*, 8:329–345, 1993. Available at <http://www.tau.ac.il/~tsirel/download/hadron.pdf>.
- [153] P. Valiant. Testing symmetric properties of distributions. *SIAM Journal on Computing*, 40(6):1927–1968, 2011.
- [154] U. Vazirani and T. Vidick. Certifiable quantum dice: or, true random number generation secure against quantum adversaries. In *Proceedings of 44th Annual ACM Symposium on Theory of Computing (STOC)*, pages 61–76, 2012. [arXiv:1111.6054](#).
- [155] U. Vazirani and T. Vidick. Fully device-independent quantum key distribution. *Physical Review Letters*, 113:140501, 2014. [arXiv:1210.1810](#).
- [156] G. Wang. Property testing of unitary operators. *Physical Review A*, 84:052328, 2011. [arXiv:1110.1133](#).
- [157] G. Wang. Property testing of quantum measurements, 2012. [arXiv:1205.0828](#).
- [158] J. Watrous. *Theory of Quantum Information* lecture notes, 2008. <http://www.cs.uwaterloo.ca/~watrous/quant-info/>.
- [159] A. Winter. Coding theorem and strong converse for quantum channels. *IEEE Transactions on Information Theory*, 45(7):2481–2485, 1999.
- [160] R. de Wolf. A brief introduction to Fourier analysis on the Boolean cube. *Theory of Computing*, 2008. ToC Library, Graduate Surveys 1.
- [161] T. H. Yang and M. Navascués. Robust self testing of unknown quantum systems into any entangled two-qubit states. *Physical Review A*, 87(5):050102(R), 2013. [arXiv:1210.4409](#).
- [162] T. H. Yang, T. Vértesi, J-D. Bancal, V. Scarani, and M. Navascués. Opening the black box: how to estimate physical properties from non-local correlations, 26 Jul 2013. [arXiv:1307.7053](#).
- [163] A. C-C. Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of 18th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 222–227, 1977.
- [164] A. C-C. Yao. Some complexity questions related to distributive computing. In *Proceedings of 11th Annual ACM Symposium on Theory of Computing (STOC)*, pages 209–213, 1979.