An efficient test for product states, with applications to quantum Merlin-Arthur games

Aram W. Harrow
Department of Mathematics, University of Bristol; and Department of Computer Science & Engineering, University of Washington
a.harrow@bris.ac.uk

Ashley Montanaro
Department of Computer Science, University of Bristol; and Department of Applied Mathematics and Theoretical Physics, University of Cambridge
am994@cam.ac.uk.

***Abstract*—We give a test that can distinguish efficiently between product states of n quantum systems and states which are far from product. If applied to a state $|\psi\rangle$ whose maximum overlap with a product state is $1-\epsilon$, the test passes with probability $1-\Theta(\epsilon)$, regardless of n or the local dimensions of the individual systems. The test uses two copies of $|\psi\rangle$. We prove correctness of this test as a special case of a more general result regarding stability of maximum output purity of the depolarising channel.**

A key application of the test is to quantum Merlin-Arthur games with multiple Merlins, where we obtain several structural results that had been previously conjectured, including the fact that soundness amplification is possible and that two Merlins can simulate many Merlins: $\mathsf{QMA}(k) = \mathsf{QMA}(2)$ for $k \geq 2$. Building on a previous result of Aaronson et al, this implies that there is an efficient quantum algorithm to verify 3-SAT with constant soundness, given two unentangled proofs of $\widetilde{O}(\sqrt{n})$ qubits. Among other consequences, this result implies complexity-theoretic obstructions to finding a polynomial-time algorithm to determine separability of mixed quantum states, even up to constant error, and also to proving "weak" variants of the additivity conjecture for quantum channels.

Finally, our test can also be used to construct an efficient test for determining whether a unitary operator is a tensor product, which is a generalisation of classical linearity testing.

I. Introduction

Entanglement of quantum states presents both an opportunity and a difficulty for quantum computing. To describe a pure state of n qudits (d-dimensional quantum systems) requires a comparable number of parameters to a classical probability distribution on d^n items. Effective methods are known for testing properties of probability distributions. However, for quantum states many of these tools no longer work. For example, due to interference, the probability of a test passing cannot be simply written as an average over components of the state. Moreover, measuring one part of a state may induce entanglement between other parts of the state that were not previously entangled with each other.

These counter-intuitive properties of entanglement account for many of the outstanding puzzles in quantum information. In quantum channel coding, the famous additivity violations of [1], [2] reflect how entangled inputs can sometimes have advantages against even uncorrelated noise. For quantum interactive proofs, the primary difficulty is in bounding the ability of provers to cheat using entangled strategies [3]. Even for $\mathsf{QMA}(k)$ (the variant of QMA with k unentangled Merlins [4], [5]), most important open questions could be resolved by finding a way to control entanglement within each proof. Here, the recently discovered failure of parallel repetition for entangled provers [6] is a sort of complexity-theoretic analogue of additivity violations.

The situation is different when we consider quantum states that are *product* across the n systems. In this case, while individual systems of course behave quantumly, the lack of correlation between the systems means that classical tools such as Chernoff bounds can be used. For example, in channel coding with product-state inputs, not only does the single-letter Holevo formula give the capacity, so that there is no additivity problem, but so-called strong converse theorems are known, which prove that attempting to communicate at a rate above the capacity results in an exponentially decreasing probability of successfully transmitting a message [7], [8]. Naturally, many of the difficulties in dealing with entangled proofs and quantum parallel repetition would also go away if quantum states were constrained to be in product form.

A. Our results

In this paper, we present a quantum test to determine whether an n-partite state $|\psi\rangle$ is a product state or far from any product state. We make no assumptions about the local dimensions of $|\psi\rangle$; in fact, the local dimension can even be different for different systems. The test passes with certainty if $|\psi\rangle$ is product, and fails with probability $\Theta(\epsilon)$ if the overlap between $|\psi\rangle$ and the closest product state is $1-\epsilon$. An essential feature of our test (or, as we show, any possible such test) is that it requires two copies of $|\psi\rangle$.

The parameters of our test resemble classical property-testing algorithms [9]. In general, these algorithms make

a small number of queries to some object and accept with high probability if the object has some property P (*completeness*), and with low probability if the object is "far" from having property P (*soundness*). Crucially, the number of queries used and the success probability should not depend on the size of the object. The main result of this paper is a test for a property of a quantum state, in contrast to previous work on quantum generalisations of property testing, which has considered quantum algorithms for testing properties of classical (e.g. [10], [11]) and quantum [12] oracles (a.k.a. unitary operators, although see Section VI for an application to this setting). In this sense, our work is closer to a body of research on determining properties of quantum states directly, without performing full tomography (e.g. the "pretty good tomography" of Aaronson [13]). The direct detection of quantities relating to entanglement has received particular attention; see [14] for an extensive review. However, previous work has generally focused on Bell inequalities and entanglement witnesses, which are typically designed to distinguish a *particular* entangled state from any separable state. By contrast, our product test is generic and will detect entanglement in any entangled state $|\psi\rangle$.

The product test is defined in Protocol 1 below, and illustrated schematically in Figure 1. It uses as a subroutine the *swap test* for comparing quantum states [15]. This test, which can be implemented efficiently, takes two (possibly mixed) states ρ, σ of equal dimension as input, and returns "same" with probability $\frac{1}{2} + \frac{1}{2} \operatorname{tr} \rho\sigma$, otherwise returning "different".

Protocol 1 (Product test).

The product test proceeds as follows.

1) *Prepare two copies of* $|\psi\rangle \in \mathbb{C}^{d_1} \otimes \cdots \otimes \mathbb{C}^{d_n}$*; call these* $|\psi_1\rangle$, $|\psi_2\rangle$.
2) *Perform the swap test on each of the* n *pairs of corresponding subsystems of* $|\psi_1\rangle$, $|\psi_2\rangle$.
3) *If all of the tests returned "same", accept. Otherwise, reject.*

The product test has appeared before in the literature. It was originally introduced in [16] as one of a family of tests for generalisations of the concurrence entanglement measure, and has been implemented experimentally as a means of detecting bipartite entanglement directly [17]. Further, the test was proposed in [12] as a means of determining whether a unitary operator is product. Our contribution here is to prove the correctness of this test for all n, as formalised in the following theorem.

Theorem 1. *Given* $|\psi\rangle \in \mathbb{C}^{d_1} \otimes \cdots \otimes \mathbb{C}^{d_n}$*, let*

$$1 - \epsilon = \max\{|\langle\psi|\phi_1, \ldots, \phi_n\rangle|^2 : |\phi_i\rangle \in \mathbb{C}^{d_i}, 1 \leq i \leq n\}.$$

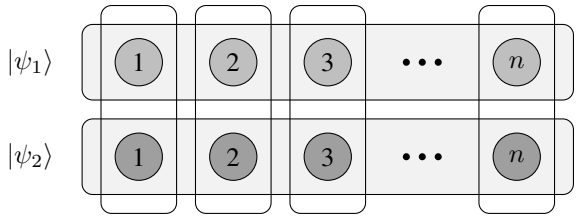

Figure 1. Schematic of the product test applied to an n-partite state $|\psi\rangle$. The swap test (vertical boxes) is applied to the n pairs of corresponding subsystems of two copies of $|\psi\rangle$ (horizontal boxes).

Let $P_{test}(|\psi\rangle\langle\psi|)$ *be the probability that the product test passes when applied to* $|\psi\rangle$*. Then*

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

Furthermore, if $\epsilon \geq 11/32 > 0.343$*,* $P_{test}(|\psi\rangle\langle\psi|) \leq 501/512 < 0.979$*.*

More concisely, $P_{test}(|\psi\rangle\langle\psi|) = 1 - \Theta(\epsilon)$*.*

This result is essentially best possible, in a number of ways. First, we show in Section III that the product test itself is optimal: among all tests for product states that use two copies and have perfect completeness, the product test has optimal soundness. We also show that there cannot exist any non-trivial test that uses only one copy of the test state. Second, our analysis of the test cannot be improved too much, without introducing dependence on n and the local dimensions. When ϵ is low, there are examples of states $|\psi\rangle$ which achieve the upper and lower bounds on $P_{\text{test}}(|\psi\rangle\langle\psi|)$, up to leading order. There is also an example of a bipartite state for which ϵ is close to 1, but $P_{\text{test}}(|\psi\rangle\langle\psi|) \approx 1/2$, implying that the constant in our bound cannot be replaced with a function of ϵ that goes to 0 as ϵ approaches 1. (The bounds on this constant obtained from our proof could easily be improved somewhat, but we have not attempted to do this.) Finally, it is unlikely that a similar test could be developed for separability of *mixed* states, as the separability problem for mixed states has been shown to be NP-hard [18], [19] (and indeed we improve on this result, as discussed below).

The proof of Theorem 1 is based on relating the probability of the test passing to the action of the qudit depolarising channel. In fact, we prove a considerably more general result regarding this channel. It is known that the maximum output purity of this channel is achieved for product state inputs [20]; our result, informally, says that any state that is "close" to achieving maximum output purity must in fact be "close" to a product state. This is a *stability* result for this channel, which strengthens the previously known multiplicativity result.

Somewhat more formally, let $\mathcal{D}_\delta$ be the d-dimensional

qudit depolarising channel with noise rate $1-\delta$, i.e.

$$\mathcal{D}_\delta(\rho) = (1-\delta)(\operatorname{tr}\rho)\frac{I}{d} + \delta\,\rho$$

for ρ a arbitrary mixed state of one d-dimensional system, and define the product state output purity to be $P_{\text{prod}}(\delta) = \operatorname{tr}(\mathcal{D}_\delta^{\otimes n}|\phi\rangle\langle\phi|)^2$, where $|\phi\rangle$ is an arbitrary product state. Then our main result is that for small enough $\delta > 0$, if $\operatorname{tr}(\mathcal{D}_\delta^{\otimes n}|\psi\rangle\langle\psi|)^2 \geq (1-\epsilon)P_{\text{prod}}(\delta)$, then there is a product state $|\phi_1,\ldots,\phi_n\rangle$ such that $|\langle\psi|\phi_1,\ldots,\phi_n\rangle|^2 \geq 1 - O(\epsilon)$.

B. Applications and interpretations of the product test

We describe several applications of the product test. The most important of these is that this test can be used to relate $\mathsf{QMA}(k)$ to $\mathsf{QMA}(2)$, as we will discuss in Section IV. The complexity class $\mathsf{QMA}(k)$ is defined to be the class of languages that can be decided with bounded error by a poly-time quantum verifier that receives poly-size witnesses from k unentangled provers[1] [4], [5]. To put $\mathsf{QMA}(k)$ inside $\mathsf{QMA}(2)$ with constant loss of soundness, we can have two provers simulate k provers by each submitting k unentangled proofs, whose lack of entanglement can be verified with our product test. Indeed, this gives an alternate way to understand our test as a method of using bipartite separability to certify k-partite separability.

Surprisingly, using this result as a building block also allows us to prove amplification for $\mathsf{QMA}(k)$ protocols. It has been conjectured [4], [5] that such protocols can be amplified to exponentially small soundness error. We completely resolve this conjecture, showing that $\mathsf{QMA}(k)$ protocols can be simulated in $\mathsf{QMA}(2)$ with exponentially small soundness error, and hence $\mathsf{QMA}(k) = \mathsf{QMA}(2)$ for $k \geq 2$.

As a further corollary, we can improve upon the results of [5], [21] to obtain a protocol in $\mathsf{QMA}(2)$ that verifies 3-SAT with constant soundness gap and $O(\sqrt{n}\operatorname{poly}\log(n))$ qubits (where n is the number of clauses). This in turn has consequences for the difficulty of approximating $\mathrm{SEP}(d,d)$, the set of separable quantum states on $d \times d$ dimensions. It was shown in Ref. [18] that SEP cannot be approximated to precision $\exp(-d)$ in time $\operatorname{poly}(d)$ unless $\mathsf{P} = \mathsf{NP}$. In Refs. [22], [19], this result was improved to show that approximating SEP to precision $1/\operatorname{poly}(d)$ is similarly NP-hard. We show that there is a universal constant $\delta > 0$ such that, if K is a convex set that approximates SEP to within trace distance δ, then membership in K cannot be decided in polynomial time unless 3-SAT $\in \mathsf{DTIME}(\exp(\sqrt{n}\log^{O(1)}(n)))$.

Our result has two further corollaries, under the same assumption on the complexity of 3-SAT. First, we show that the minimum output entropy of a quantum channel cannot be estimated up to a constant in polynomial time. Second, we show a hardness result for approximating the ground-state energy of quantum systems under a mean-field approximation. Our proof that amplification of $\mathsf{QMA}(2)$ protocols is possible implies that one can derive stronger hardness results for these tasks, if one is willing to make stronger assumptions about the hardness of 3-SAT.

Our final application is that the product test can be used to determine whether a unitary operator is a tensor product. This can be seen [12] as one possible generalisation of the well-studied problem of testing whether a boolean function $\{0,1\}^n \to \{0,1\}$ is linear [23]. This application is described in Section VI.

These different applications of the product test reflect the many different interpretations of $P_{\text{test}}(|\psi\rangle\langle\psi|)$. It is related in a precise sense to

- The purity of $|\psi\rangle$ after it is subjected to independent depolarising noise.
- The maximum overlap of $|\psi\rangle$ with any product state. The logarithm of this maximum overlap is an important entanglement measure known as the geometric measure of entanglement (see [24] and references therein).
- The overlap of $|\psi\rangle^{\otimes 2}$ with the tensor product of the symmetric subspaces of $\mathbb{C}^{d_1}\otimes\mathbb{C}^{d_1},\ldots,\mathbb{C}^{d_n}\otimes\mathbb{C}^{d_n}$.
- The average overlap of $|\psi\rangle$ with a *random* product state.
- The average purity of $|\psi\rangle$ across a random partition of $[n]$ into two subsets.

Due to space limitations, discussion of some of these interpretations is deferred to the full version [25], as are many proofs.

II. Overview of the proof of correctness

In this section, we sketch the proof of Theorem 1. We begin with two lemmas which allow the probability of passing the product test to be understood, and to be related to the output purity of the depolarising channel.

Lemma 2. *We have*

$$\operatorname{tr}(\mathcal{D}_\delta^{\otimes n}\rho)^2 = \left(\frac{1-\delta^2}{d}\right)^n \sum_{S\subseteq[n]} \left(\frac{d\delta^2}{1-\delta^2}\right)^{|S|} \operatorname{tr}(\rho_S^2),$$

and in particular

$$\operatorname{tr}(\mathcal{D}_{1/\sqrt{d+1}}^{\otimes n}\rho)^2 = \frac{1}{(d+1)^n}\sum_{S\subseteq[n]}\operatorname{tr}(\rho_S^2),$$

and for pure product states,

$$\begin{aligned} P_{prod}(\delta) &:= \operatorname{tr}(\mathcal{D}_\delta^{\otimes n}(|\psi_1\rangle\langle\psi_1|\otimes\cdots\otimes|\psi_n\rangle\langle\psi_n|))^2 \\ &= \left(\frac{d-1}{d}\delta^2 + \frac{1}{d}\right)^n. \end{aligned}$$

Lemma 3. *Let $P_{test}(\rho,\sigma)$ denote the probability that the product test passes when applied to two mixed states $\rho,\sigma \in \mathcal{B}(\mathbb{C}^{d_1}\otimes\cdots\otimes\mathbb{C}^{d_n})$. Define $P_{test}(\rho) := P_{test}(\rho,\rho)$. Then*

$$P_{test}(\rho,\sigma) = \frac{1}{2^n}\sum_{S\subseteq[n]}\operatorname{tr}\rho_S\sigma_S,$$

[1] We assume throughout this paper that k is at most polynomial in the input size.

and in particular

$$P_{test}(\rho) = \frac{1}{2^n} \sum_{S \subseteq [n]} \operatorname{tr} \rho_S^2.$$

If $d_1 = d_2 = \cdots = d_n = d$, for some d, then

$$P_{test}(\rho) = \left(\frac{d+1}{2}\right)^n \operatorname{tr}(\mathcal{D}_{1/\sqrt{d+1}}^{\otimes n} \rho)^2.$$

The proof itself is split into two parts, beginning with the case where ϵ is low. The difficult part is the upper bound on $P_{\text{test}}(|\psi\rangle\langle\psi|)$. We write $|\psi\rangle = \sqrt{1-\epsilon}|0^n\rangle + \sqrt{\epsilon}|\phi\rangle$ without loss of generality, for some product state $|0^n\rangle$ and arbitrary state $|\phi\rangle$. This allows an explicit expression for $\operatorname{tr} \psi_S^2$ in terms of ϵ and $|\phi\rangle$ to be obtained. Each term of this expression is then carefully upper bounded to give an upper bound in terms of a sum of the amplitudes of $|\phi\rangle$, with weights that decrease exponentially with the Hamming weight of basis states. In order to obtain a non-trivial bound from this expression, the final stage of this 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 of Hamming weight 1. In its most general form, the result we obtain is as follows.

Theorem 4. *Given $|\psi\rangle \in (\mathbb{C}^d)^{\otimes n}$, let*

$$1 - \epsilon = \max\{|\langle\psi|\phi_1, \ldots, \phi_n\rangle|^2 : |\phi_1\rangle, \ldots, |\phi_n\rangle \in \mathbb{C}^d\}.$$

Then

$$\operatorname{tr}(\mathcal{D}_\delta^{\otimes n} |\psi\rangle\langle\psi|)^2 \leq P_{prod}(\delta) \Bigg(1 - 4\epsilon(1-\epsilon) \frac{d\delta^2(1-\delta^2)}{(1+(d-1)\delta^2)^2} + 4\epsilon^{3/2} \left(\frac{(1-\delta^2)^2 + d^2\delta^4}{(1+(d-1)\delta^2)^2} \right)^2 \Bigg).$$

In particular,

$$\operatorname{tr}(\mathcal{D}_{1/\sqrt{d+1}}^{\otimes n} |\psi\rangle\langle\psi|)^2 \leq P_{prod}(1/\sqrt{d+1}) \left(1 - \epsilon + \epsilon^2 + \epsilon^{3/2}\right).$$

Specialising to the particular case of the product test, we deduce the following result.

Theorem 5. *Given $|\psi\rangle \in \mathbb{C}^{d_1} \otimes \cdots \otimes \mathbb{C}^{d_n}$, let*

$$1 - \epsilon = \max\{|\langle\psi|\phi_1, \ldots, \phi_n\rangle|^2 : |\phi_i\rangle \in \mathbb{C}^{d_i}, 1 \leq i \leq n\}.$$

Then

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

Proof: The lower bound holds by general arguments. It is immediate that, if applied to $|\phi_1, \ldots, \phi_n\rangle$, the product test succeeds with probability 1. As the test acts on two copies of $|\psi\rangle$, which has overlap $1 - \epsilon$ with $|\phi_1, \ldots, \phi_n\rangle$, it must succeed when applied to $|\psi\rangle$ with probability at least $(1-\epsilon)^2$. The upper bound follows from Lemma 3 and Theorem 4. The statement of Theorem 4 only explicitly covers the case where the dimensions of all the subsystems are the same; however, we can assume this without loss of generality. ■

In the case where ϵ is high, this result does not yet give a useful upper bound. In the second part of the proof, we derive a constant bound on $P_{\text{test}}(|\psi\rangle\langle\psi|)$ based on considering $|\psi\rangle$ as a k-partite state, for some $k < n$. $P_{\text{test}}(|\psi\rangle\langle\psi|)$ can be shown to be upper bounded by the probability that the test for being product across any partition into k parties passes. Informally speaking, if $|\psi\rangle$ is far from product across the n subsystems, we show that one can find a partition 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.

Theorem 6. *Given $|\psi\rangle \in \mathbb{C}^{d_1} \otimes \cdots \otimes \mathbb{C}^{d_n}$, let*

$$1 - \epsilon = \max\{|\langle\psi|\phi_1, \ldots, \phi_n\rangle|^2 : |\phi_i\rangle \in \mathbb{C}^{d_i}, 1 \leq i \leq n\}.$$

Then, if $\epsilon \geq 11/32 > 0.343$, $P_{test}(|\psi\rangle\langle\psi|) \leq 501/512 < 0.979$.

Between them, Theorems 5 and 6 imply Theorem 1. This completes the overview of the proof; we now demonstrate that the product test itself is essentially optimal.

III. Optimality of the product test

Our test has perfect completeness in the sense that if $|\psi\rangle$ is exactly a product state then it will always pass the product test. It is hard to precisely define soundness, since no state is orthogonal to all product states: however, we can say that our test has constant soundness in that if $|\psi\rangle$ has overlap at most $1 - \epsilon$ with any product state then it will pass the product test with probability at most $1 - \Theta(\epsilon)$.

In fact, if we consider only product-state tests with perfect completeness, then we can show that our test has optimal soundness: that is, it rejects as often as possible given the constraint of always accepting product states. More generally, suppose that a product-state test T is given $|\psi\rangle^{\otimes k}$ as input. Since the outcome of the test is binary, we can say that T is an operator on the nk-qudit Hilbert space with $0 \leq T \leq I$ and that the test accepts with probability $\operatorname{tr} T\psi^{\otimes k}$.

Let S be the set of product states in $\mathbb{C}^{d_1} \otimes \cdots \otimes \mathbb{C}^{d_n}$, and define S^k to be the span of $\{|\phi\rangle^{\otimes k} : |\phi\rangle \in S\}$. For a single system $\mathbb{C}^d$, the span of $\{|\phi\rangle^{\otimes k} : |\phi\rangle \in \mathbb{C}^d\}$ is denoted $\operatorname{Sym}^k \mathbb{C}^d$. This is the symmetric subspace of $(\mathbb{C}^d)^{\otimes k}$, meaning that it can be equivalently defined to be the set of vectors in $(\mathbb{C}^d)^{\otimes k}$ that is invariant under permutation by the symmetric group S_k. This fact allows the projector onto $\operatorname{Sym}^k \mathbb{C}^d$, which we denote $\Pi_{d,k}^{\text{sym}}$, to be implemented efficiently [26]. Also, it implies that $S^k = \operatorname{Sym}^k \mathbb{C}^{d_1} \otimes \cdots \otimes \operatorname{Sym}^k \mathbb{C}^{d_n}$ and that the projector onto S^k, denoted Π_{S^k}, is $\Pi_{d_1,k}^{\text{sym}} \otimes \cdots \otimes \Pi_{d_n,k}^{\text{sym}}$.

Now we return to our discussion of product-state tests. If $\operatorname{tr} T\phi^{\otimes k} = 1$ for all $\phi \in S$, then $T \geq \Pi_{S^k}$. Thus, T will always accept at least as often as Π_{S^k} will on any input, or

equivalently, taking $T = \Pi_{S^k}$ yields the test which rejects as often as possible given the constraint of accepting every state in S^k.

To understand Π_{S^k}, note that the projector onto $\mathrm{Sym}^k \mathbb{C}^d$ is given by $\frac{1}{k!}\sum_{\pi\in\mathcal{S}_k} P(\pi)$, where

$$P(\pi) = \sum_{i_1,\ldots,i_k\in[d]} |i_1,\ldots,i_k\rangle\langle i_{\pi(1)},\ldots,i_{\pi(k)}|.$$

For $k = 1$, $\mathrm{Sym}^1 \mathbb{C}^d$ simply equals $\mathbb{C}^d$, and Π_{S^1} is the identity operator on $(\mathbb{C}^d)^{\otimes n}$. Thus, no non-trivial product-state test is possible when given one copy of $|\psi\rangle$.

When $k = 2$, $\mathrm{Sym}^2 \mathbb{C}^d$ is the $+1$ eigenspace of $(I+\mathcal{F})/2$, which is the space that passes the swap test. Thus, the product test (in Protocol 1) performs the projection onto S^2 and therefore rejects non-product states as often as possible for a test on $|\psi\rangle^{\otimes 2}$ that always accepts when $|\psi\rangle$ is a product state. These arguments also imply that given $|\psi\rangle^{\otimes k}$, projecting onto S^k yields an optimal k-copy product-state test of $|\psi\rangle$. The strength of these tests is strictly increasing with k, but we leave the problem of analysing them carefully to future work.

Finally, this interpretation of the product test allows us to consider generalisations to testing membership in other sets S. The general prescription for a test that is given k copies of a state is simply to project onto the span of $\{|\psi\rangle^{\otimes k} : |\psi\rangle \in S\}$. However, we will not explore these possibilities further in this paper.

IV. $\mathsf{QMA}(2)$ vs. $\mathsf{QMA}(k)$

In this section, we apply the product test to a problem in quantum complexity theory: whether k unentangled provers are better than 2 unentangled provers. This question can be formalised as whether the complexity classes $\mathsf{QMA}(k)$ and $\mathsf{QMA}(2)$ are equal [4], [5]. These classes are defined as follows.

Definition 1. *A language L is in $\mathsf{QMA}(k)_{s,c}$ if there exists a polynomial-time quantum algorithm $\mathcal{A}$ such that, for all inputs $x \in \{0,1\}^n$:*

1) **Completeness:** *If $x \in L$, there exist k witnesses $|\psi_1\rangle,\ldots,|\psi_k\rangle$, each a state of $\mathrm{poly}(n)$ qubits, such that $\mathcal{A}$ outputs "accept" with probability at least c on input $|x\rangle|\psi_1\rangle\ldots|\psi_k\rangle$.*
2) **Soundness:** *If $x \notin L$, then $\mathcal{A}$ outputs "accept" with probability at most s on input $|x\rangle|\psi_1\rangle\ldots|\psi_k\rangle$, for all states $|\psi_1\rangle,\ldots,|\psi_k\rangle$.*

We use $\mathsf{QMA}(k)$ as shorthand for $\mathsf{QMA}(k)_{1/3,2/3}$, and QMA as shorthand for $\mathsf{QMA}(1)$. We always assume $1 \leq k \leq \mathrm{poly}(n)$.

We also define $\mathsf{QMA}_m(k)_{s,c}$ to indicate that $|\psi_1\rangle,\ldots,|\psi_k\rangle$ each involve m qubits, where m may be a function of n other than $\mathrm{poly}(n)$.

Two of the major open problems related to $\mathsf{QMA}(k)_{s,c}$ are to determine how the size of the complexity class depends on k and on s, c. It has been conjectured for some years [4], [5] that in fact $\mathsf{QMA}(k) = \mathsf{QMA}(2)$ for $2 \leq k \leq \mathrm{poly}(n)$, and that the soundness and completeness can be amplified by parallel repetition in a way similar to BPP, BQP, MA, QMA and other complexity classes with bounded error. In fact, these conjectures are related: $2k$ independent provers can simulate k independent realisations of a $\mathsf{QMA}(2)$ protocol in order to amplify the soundness-completeness gap, and conversely, [4], [5] proved that $\mathsf{QMA}(2)$ amplification implies that $\mathsf{QMA}(2) = \mathsf{QMA}(\mathrm{poly})$. In this section, we will fully resolve these conjectures, proving that $\mathsf{QMA}(2) = \mathsf{QMA}(\mathrm{poly})$ and that $\mathsf{QMA}(k)$ can have its soundness and completeness amplified by a suitable protocol.

The most direct way of putting $\mathsf{QMA}(k)$ inside $\mathsf{QMA}(2)$ is to ask two provers to each send the k unentangled proofs that correspond to a $\mathsf{QMA}(k)$ protocol. If $k = \mathrm{poly}(n)$, then each prover is still sending only polynomially many qubits. Then the product test can be used to verify that the states sent were indeed product states and can be used as valid inputs to a $\mathsf{QMA}(k)$ protocol. The specific protocol is described in Protocol 2.

Protocol 2 ($\mathsf{QMA}(k)$ **to** $\mathsf{QMA}(2)$)**.**

The $\mathsf{QMA}(2)$ protocol proceeds as follows.

1) *Each of the two Merlins sends $|\psi\rangle := |\psi_1\rangle \otimes \ldots \otimes |\psi_k\rangle$ to Arthur.*
2) *Arthur runs the product test with the two states as input.*
3) *If the test fails, Arthur rejects. Otherwise, Arthur runs the algorithm $\mathcal{A}$ on the state sent by the first Merlin, and outputs the result.*

First observe that for YES instances (instances in the language), the two Merlins can achieve at least the same success probability that k Merlins can in the original protocol, so the completeness can only increase. Now consider NO instances. Assume for now that the two Merlins always send the same state. Then according to Theorem 1, if the Merlins send states that are far from product, they are likely to fail the product test, whereas basic continuity arguments can show that if they send states that are nearly product then the success probability will not be much larger than the soundness of the original protocol. Thus, the soundness does not become too much worse. These ideas establish

Lemma 7. *For any m, k, $0 \leq s < c \leq 1$,*

$$\mathsf{QMA}_m(k)_{s,c} \subseteq \mathsf{QMA}_{km}(2)_{s',c}$$

where $s' = 1 - \Omega((1-s)^2)$.

This is already strong enough to achieve amplification up to constant soundness. However, Protocol 2 has a salutary

side effect that will allow us to achieve stronger amplification. To see this, we will first introduce a further generalisation of the $\mathsf{QMA}(k)$ family. Let $\mathbb{M}$ be a set of Hermitian operators M satisfying $0 \leq M \leq I$. Each $M \in \mathbb{M}$ defines a binary measurement with M corresponding to the "accept" outcome and $I - M$ corresponding to the "reject" outcome. Then define $\mathsf{QMA}_m^{\mathbb{M}}(k)_{s,c}$ to be the class $\mathsf{QMA}_m(k)_{s,c}$ with Arthur restricted to performing measurements from $\mathbb{M}$. For example, if $\mathbb{M} = \text{PROD}$ is the set of product measurements followed by post-processing, then $\mathsf{QMA}^{\text{PROD}}(k)$ is the class $\text{BellQMA}(k)$ that was proved equal to QMA (for constant k) by Brandão [27], [5]. We will be interested in taking $\mathbb{M} = \text{SEP}$. Note that for such measurements, M is a separable operator, but $I - M$ is not necessarily separable.

Armed with the definition of $\mathsf{QMA}^{\text{SEP}}$, we can now see that Protocol 2 produces a protocol that is not only in $\mathsf{QMA}(2)$, but also $\mathsf{QMA}^{\text{SEP}}(2)$. More formally, we can strengthen Lemma 7 to:

Lemma 8. *For any m, k, $0 \leq s < c \leq 1$,*

$$\mathsf{QMA}_m(k)_{s,c} \subseteq \mathsf{QMA}_{km}^{\text{SEP}}(2)_{s',c}$$

where $s' = 1 - \Omega((1-s)^2)$.

Proof: Suppose the first Merlin sends systems $A_1, \ldots, A_k$ and the second Merlin sends systems $B_1, \ldots, B_k$. The "accept" outcome of the product test corresponds to the tensor product of projectors onto the symmetric subspaces of $A_1B_1, A_2B_2, \ldots, A_kB_k$. These are all separable across the A:B cut, and so their tensor product is as well. The second step is to simply apply a measurement entirely on $A_1, \ldots, A_k$, which is automatically separable. Finally, performing two separable measurements in a row creates a composite measurement which is also separable. ∎

The advantage of $\mathsf{QMA}^{\text{SEP}}(k)$ is that it removes the chief difficulty with $\mathsf{QMA}(k)$ amplification, which is that conditioning on measurement outcomes can induce entanglement between systems we have not yet measured. This phenomenon is known as entanglement swapping. However, if we condition on the outcome of a measurement being M, for some $M \in \text{SEP}$, then no entanglement will be produced in the unmeasured states. As a result, cheating provers cannot gain any advantage by sending entangled proofs, and we obtain the following lemma.

Lemma 9. *For any $\ell \geq 1$,*

$$\mathsf{QMA}_m^{\text{SEP}}(k)_{s,c} \subseteq \mathsf{QMA}_{\ell m}^{\text{SEP}}(k)_{s^\ell, c^\ell}.$$

The idea is to simply repeat the original protocol ℓ times in parallel and to accept iff each subprotocol accepts. Since we are considering $\mathsf{QMA}^{\text{SEP}}$ protocols, obtaining a YES outcome on one proof will not induce any entanglement on the remaining proofs.

From Lemma 8 and Lemma 9, we can almost conclude that strong amplification is possible. Indeed, when we start with protocols with perfect completeness, we can apply Protocol 2, repeat $p(n)$ times, and reduce the soundness from s to $s^{O(p(n))}$. For the case of $c < 1$, we need one additional argument to keep the completeness from being reduced too much at the same time. Here we will use a method for completeness amplification proved in both [4, Lemma 5] and [5, Lemma 6].

Lemma 10 ([4], [5])**.** *For any $\ell \geq 1$,*

$$\mathsf{QMA}_m(k)_{s,c} \subseteq \mathsf{QMA}_{\ell m}(k)_{1-\frac{c-s}{3}, 1-\exp(-\ell(c-s)^2/2)}.$$

Our amplification procedure for general $c < 1$ is then to

1) Use Lemma 10 to bring the completeness exponentially close to 1.
2) Use Lemma 8 to convert a general $\mathsf{QMA}(k)$ protocol to a $\mathsf{QMA}^{\text{SEP}}(2)$ protocol.
3) Repeat the protocol polynomially many times to make the soundness exponentially small.

This procedure then achieves

Theorem 11. 1) *If $s \leq 1 - 1/\operatorname{poly}(n)$, $k = \operatorname{poly}(n)$ and $p(n)$ is an arbitrary polynomial, then $\mathsf{QMA}(k)_{s,1} = \mathsf{QMA}^{\text{SEP}}(2)_{\exp(-p(n)),1}$.*
2) *If $c - s \geq 1/\operatorname{poly}(n)$, $k = \operatorname{poly}(n)$ and $p(n)$ is an arbitrary polynomial, then $\mathsf{QMA}(k)_{s,c} = \mathsf{QMA}^{\text{SEP}}(2)_{\exp(-p(n)),1-\exp(-p(n))}$.*

There are obvious variants of Theorem 11 to cover the case of limited message size, whose statements we leave implicit.

V. Complexity-theoretic implications

A key application of Theorem 11 is to the protocol of Ref. [5] that puts 3-SAT on n clauses inside the complexity class $\mathsf{QMA}_{\log}(\sqrt{n}\operatorname{poly}\log(n))_{1-\Omega(1),1}$. Applying Theorem 11 lets us simulate this using two provers with perfect completeness and arbitrary soundness, so that we obtain

Corollary 12. *Let $\ell : \mathbb{N} \to \mathbb{N}$ be polynomially bounded. Then there is a universal constant $0 < s < 1$ such that*

$$\textit{3-SAT} \in \mathsf{QMA}_{\ell(n)\sqrt{n}\operatorname{poly}\log(n)}(2)_{s^{\ell(n)},1}.$$

In other words, there is a 3-SAT protocol with two provers, $\ell(n)\sqrt{n}\operatorname{poly}\log(n)$-qubit proofs, perfect completeness and soundness $s^{\ell(n)}$.

Therefore, making assumptions about the hardness of 3-SAT allows us to prove hardness results for the complexity class $\mathsf{QMA}_{\log}(2)$, and stronger assumptions naturally imply stronger hardness results. We formalise this correspondence as the following corollary.

Corollary 13. *Assume that, for some function $\ell : \mathbb{N} \to \mathbb{N}$ such that $\ell(n) = o(n)$, 3-SAT $\notin \mathsf{DTIME}(\exp(O(\ell(n))))$. Then, defining $d = 2^{\ell(n)\sqrt{n}\,\mathrm{polylog}(n)}$,*

$$\mathsf{QMA}_{\log(d)}(2)_{s^{\ell(n)},1} \nsubseteq \mathsf{DTIME}(\mathrm{poly}(d))$$

for some universal constant $0 < s < 1$. In particular, taking $\ell(n) = 1$, we have

$$\mathsf{QMA}_{\log(d)}(2)_{s,1} \nsubseteq \mathsf{DTIME}(\mathrm{poly}(d))$$

assuming that 3-SAT $\notin \mathsf{DTIME}(\exp(\sqrt{n}\log^{O(1)}(n)))$.

Note that the (not implausible) *Exponential Time Hypothesis* of Impagliazzo and Paturi [28] states that 3-SAT $\notin \mathsf{DTIME}(\exp(\ell(n)))$ for *any* $\ell(n) = o(n)$. We conclude this section by discussing three natural $\mathsf{QMA}_{\log}(2)$-complete problems whose hardness is implied by this hypothesis. For simplicity, we focus on the weakest possible hardness assumption in Corollary 13, but it should be easy to see how to strengthen this assumption to obtain stronger results.

First observe that the acceptance probability in a $\mathsf{QMA}_m(2)$ protocol can be expressed as $\max_{\rho \in \mathrm{SEP}(2^m,2^m)} \mathrm{tr}\, M\rho$, where $0 \leq M \leq I$ is the measurement resulting from the verifier's quantum circuit and $\mathrm{SEP}(d_A, d_B)$ denotes the set of separable density matrices on $d_A \times d_B$ dimensions. In other words, $\mathsf{QMA}_m(2)$ is equivalent to optimising a linear objective function over the convex set $\mathrm{SEP}(2^m, 2^m)$.

In some cases, it may be useful to obtain an explicit description of M. This can be achieved up to error ϵ by running the verifier's circuit $\mathrm{poly}(2^m, 1/\epsilon)$ times and performing tomography. As a result, we trivially obtain that $\mathsf{QMA}_m(2)_{s,c} \subseteq \mathsf{NTIME}(\mathrm{poly}(2^m, n, 1/(c-s)))$. In particular, $\mathsf{QMA}(2) \subseteq \mathsf{NEXP}$. Unfortunately this cannot be scaled down to place $\mathsf{QMA}_{\log}(2)$ in NP. This is because the verifier in a $\mathsf{QMA}_{\log}(2)$ protocol still can perform a poly-time quantum computation. Thus, we only have that $\mathsf{QMA}_{\log}(2) \subseteq \mathsf{NP}^{\mathsf{BQP}}$.

Application 1: Separability-testing. A folk theorem of convex optimisation [29] states that the problem of optimising a linear function over a convex set, such as SEP, is equivalent to determining membership in that set. Thus, we should be able to relate $\mathsf{QMA}_{\log}(2)$ to the problem of determining membership in SEP. To make this precise, for any convex $K \subseteq \mathbb{R}^d$ we define $B(K, \epsilon)$ to be $\{x : \exists y \in K, \|x - y\| \leq \epsilon\}$ when $\epsilon > 0$ and $\{x : \nexists y \notin K, \|x - y\| \leq -\epsilon\}$ when $\epsilon < 0$. The weak membership problem for K, $\mathrm{WMEM}_\epsilon(K)$, is to determine whether a point x belongs to $B(K, \epsilon)^c$ or $B(K, -\epsilon)$ given the promise that one of these is the case. The weak optimisation problem for K, $\mathrm{WOPT}_\epsilon(K)$ is to maximize a linear objective function over any set L satisfying $B(K, -\epsilon) \subseteq L \subseteq B(K, \epsilon)$. Given some mild conditions on K, we can reduce $\mathrm{WOPT}_\epsilon(K)$ to $\mathrm{WMEM}_{\epsilon/\,\mathrm{poly}(d)}(K)$ in polynomial time [29]. This fact has been used to show the NP-hardness of $\mathrm{WMEM}_{1/\,\mathrm{poly}}(\mathrm{SEP})$ in Refs. [22], [19], [30] and, previously, of $\mathrm{WMEM}_{1/\exp}(\mathrm{SEP})$ by Gurvits [18] (although the connection to $\mathsf{QMA}_{\log}(2)$ was only observed by [30]).

Unfortunately, many of these techniques break down in the setting of constant error. We believe that it should not be possible to approximate $\mathrm{SEP}(d, d)$ to within a (sufficiently small) constant accuracy in time $\mathrm{poly}(d)$. However, we are able to rule out only algorithms that have the further restriction of recognizing a nearly convex set that in turn approximates SEP to constant accuracy.

Corollary 14. *Let K be a convex subset of the space of $d^2 \times d^2$ Hermitian matrices such that $K \subseteq B(\mathrm{SEP}(d, d), \delta)$ and $\mathrm{SEP}(d, d) \subseteq B(K, \delta)$, where $B(\cdot, \cdot)$ is defined relative to the trace norm and $\delta > 0$ is a universal constant. Then, assuming 3-SAT $\notin \mathsf{DTIME}(\exp(\sqrt{n}\log^{O(1)}(n)))$, $\mathrm{WMEM}_{1/\,\mathrm{poly}}(K)$ cannot be decided in time $\mathrm{poly}(d)$.*

(As with the other hardness results in this section, the precise value of δ is determined by the protocol in Ref. [5].)

Proof: Solving $\mathrm{WMEM}_{1/\,\mathrm{poly}(d)}(K)$ in polynomial time would allow us to solve $\mathrm{WOPT}_{1/\,\mathrm{poly}(d)}(K)$ in poly time, which in turn would give a poly-time algorithm for $\mathrm{WOPT}_{\delta+1/\,\mathrm{poly}(d)}(\mathrm{SEP})$. This last claim, together with Corollary 13, would contradict the hypothesis on the complexity of 3-SAT. ∎

Application 2: Minimum output entropy of quantum channels. Our results also have implications for quantum information theory. Let $\mathcal{N}$ denote a quantum channel with d-dimensional input and output. Define the minimum output Rényi α-entropy of $\mathcal{N}$ to be $S_\alpha^{\min}(\mathcal{N}) = \min_\rho S_\alpha(\mathcal{N}(\rho))$, where $S_\alpha(\sigma) = \frac{1}{1-\alpha}\log \mathrm{tr}\,\sigma^\alpha$ and the minimum is taken over all quantum states ρ. Note that $S_\alpha^{\min}(\mathcal{N})$ is also equivalent to $\frac{\alpha}{1-\alpha}\log\|\mathcal{N}\|_{1\to\alpha}$, where $\|\cdot\|_{1\to\alpha}$ (also called ν_α in e.g. [20]) is the $\ell_1 \to \ell_p$ norm. When $\alpha = 0, 1, \infty$, we define $S_\alpha(\sigma)$ by letting α approach these values, obtaining $S_0(\sigma) = \log \mathrm{rank}\,\sigma$, $S_1(\sigma) = -\mathrm{tr}\,\sigma\log\sigma$ and $S_\infty(\sigma) = -\log\|\sigma\|_\infty$. Additivity of $S_1^{\min}(\mathcal{N})$, the minimum output entropy, is intimately connected to additivity of the Holevo capacity [31], [32].

It was observed by Matsumoto [33] (citing a personal communication from Watrous) that the maximum acceptance probability of a $\mathsf{QMA}_m(2)$ protocol is precisely $\|\mathcal{N}\|_{1\to\infty}$ for some quantum channel $\mathcal{N}$ acting on $d = 2^m$ dimensions. This implies that determining whether $S_\infty^{\min}(\mathcal{N})$ is $\geq \log(1/s)$ or $\leq \log(1/c)$ is a complete problem for $\mathsf{QMA}_{\log(d)}(2)_{s,c}$ under BQP reductions.

The $\mathsf{QMA}(2)$-completeness of estimating $S_\infty^{\min}$ implies that other information-theoretic quantities that are close to $S_\infty^{\min}$ are also hard to approximate. For example, for any $\alpha \geq 0$, we have $S_\alpha^{\min}(\mathcal{N}) \geq S_\infty^{\min}(\mathcal{N})$ but also $S_\alpha^{\min}(\mathcal{N}) = 0$ iff $S_\infty^{\min}(\mathcal{N}) = 0$. Thus, our hardness result for approximating $S_\infty^{\min}$ immediately translates to a hardness result for approximating $S_\alpha^{\min}$.

Corollary 15. *There exists a universal constant $\delta > 0$ such that for any $\alpha \geq 0$, if 3-SAT $\notin$ $\mathsf{DTIME}(\exp(\sqrt{n}\log^{O(1)}(n)))$ then it is impossible to determine whether $S_\alpha^{\min}(\mathcal{N}) = 0$ or $S_\alpha^{\min}(\mathcal{N}) \geq \delta$ in worst-case time $\mathrm{poly}(d)$.*

Beigi and Shor previously showed that it is NP-hard to compute the minimum output entropy up to $1/\mathrm{poly}(d)$ accuracy [32]. Our result improves theirs, but under a more restrictive complexity assumption. Another major goal in information theory is to estimate the regularised minimum output entropies of quantum channels, which are defined to be

$$S_\alpha^{R,\min}(\mathcal{N}) := \liminf_{n\to\infty} \frac{1}{n} S_\alpha^{\min}(\mathcal{N}^{\otimes n}).$$

The $S_\alpha^{R,\min}(\mathcal{N})$ are relevant to determining the ultimate channel capacity, to proving strong converse theorems [34] and to cryptographic protocols [35]. Our hardness result for $S_\alpha^{\min}$ immediately gives us the equivalent hardness result for $S_\alpha^{R,\min}$. The reason is that our proof of amplification for $\mathsf{QMA}(2)$ protocols (see Lemma 9) essentially works by constructing a channel $\mathcal{N}$ for which $S_\infty^{R,\min}(\mathcal{N}) = S_\infty^{\min}(\mathcal{N})$ by design.

For general channels, we automatically have $S_\alpha^{R,\min}(\mathcal{N}) \leq S_\alpha^{\min}(\mathcal{N})$; however, the famous failures of the additivity conjecture imply that sometimes this inequality can be strict, with examples known for $\alpha \geq 1$ [2], [36] and for α near 0 [37]. Still, these examples only demonstrate that $S_\alpha^{R,\min}$ can deviate very slightly from $S_\alpha^{\min}$. On the other hand, various lower bounds for $S_\alpha^{R,\min}$ are known [38], [39], [40], [41], and it may be that one of these bounds could be related to $S_\alpha^{\min}$, thereby proving that $S_\alpha^{R,\min}$ cannot be far from $S_\alpha^{\min}$. Our results do not rule out the possibility that $S_\alpha^{\min}$ may be fruitfully related to $S_\alpha^{R,\min}$. However, they do imply that these lower bounds on $S_\alpha^{R,\min}$ (and thereby on $S_\alpha^{\min}$) are unlikely to be efficiently computable, or if they are, they are likely to be extremely loose bounds in general.

Application 3: mean-field approximation. Finally, we discuss an application from condensed-matter physics. Consider a system of n d-dimensional quantum systems arranged in a lattice with identical nearest-neighbour pairwise interactions. The mean-field approximation replaces the true nearest-neighbour graph with the complete graph. When the number of spatial dimensions is 3 (or more), this is often a reasonable approximation. If K is a fixed two-qudit Hamiltonian and $K_{i,j}$ denotes the action of K on systems i, j and the identity on the other systems, then the total Hamiltonian is $H = \sum_{i\neq j} K_{i,j}$. To set the overall scale of the problem, assume that $0 \leq K \leq I$. One of the more important physical questions about H is to determine its ground-state energy; that is, its smallest eigenvalue.

In Ref. [42], the quantum de Finetti theorem was used to show that when $n \gg d^2$, then the ground state of H is very close to a product state. In this case, finding the ground-state energy of H is equivalent to minimising $\mathrm{tr}\,\rho K$ over all $\rho \in \mathrm{SEP}(d, d)$. Again applying Corollary 13, we obtain:

Corollary 16. *Assuming that 3-SAT $\notin$ $\mathsf{DTIME}(\exp(\sqrt{n}\log^{O(1)}(n)))$ and with H defined as above, it is impossible to estimate $\min\{\mathrm{tr}\,\rho K : \rho \in \mathrm{SEP}(d, d)\}$ in time $\mathrm{poly}(d)$ to within $o(1)$ error. Equivalently, it is impossible to estimate the ground-state energy of H to within additive error $o(n^2)$.*

Previous work on the hardness of approximating ground-state energy of quantum systems generally had d constant and only ruled out the possibility of $1/\mathrm{poly}(n)$ approximation error. In terms of approximation errors, our result achieves one of the goals of the conjectured quantum PCP theorem [43]. However, we require d to grow asymptotically, and we achieve a hardness result much weaker than QMA-hardness. Indeed, due to the *classical* PCP theorem combined with the Exponential Time Hypothesis, finding the ground state of a system of $d^2\log(d)$ bits (without any symmetry constraint) is likely to require time $\exp(d^2\log(d))$, while our results merely imply an $\Omega(\exp(\log^2(d)))$ lower bound. Still, our result provides a superpolynomial bound on an important class of Hamiltonians that had been previously considered to be computationally easy to work with.

VI. Testing for product unitaries

As well as being useful for testing quantum states, the product test has applications to testing properties of unitary operators. The results we obtain will be in terms of the normalised Hilbert-Schmidt inner product, which is defined as $\langle M, N\rangle := \frac{1}{d}\,\mathrm{tr}\,M^\dagger N$ for $M, N \in M(d)$, where $M(d)$ denotes the set of $d \times d$ matrices. Note that, with this normalisation, $|\langle U, V\rangle| \leq 1$ for unitary operators U, V. The following correspondence (also known as the Choi-Jamiołkowski isomorphism), underlies our ability to apply the product test to unitaries.

Let $|\Phi\rangle$ be a maximally entangled state of two d-dimensional qudits, written as $\frac{1}{\sqrt{d}}\sum_{i=1}^d |i, i\rangle$ in terms of some basis $\mathcal{B} = (|1\rangle, \ldots, |d\rangle)$. For any matrix $M \in M(d^n)$, define $|v(M)\rangle := (M \otimes I)|\Phi\rangle^{\otimes n}$. Then $\langle j|\langle k|v(M)\rangle = \frac{\langle j|M|k\rangle}{\sqrt{d^n}}$. In particular, for any matrices $M, N \in M(d^n)$, $\langle M, N\rangle = \langle v(M)|v(N)\rangle = \mathrm{tr}\,M^\dagger N/d^n$.

We consider the problem of testing whether a unitary operator is a tensor product. That is, we are given access to a unitary U on the space of n qudits (for simplicity, restricting to the case where each of the qudits has the same dimension d), and we would like to decide whether $U = U_1 \otimes \cdots \otimes U_n$. This is one possible generalisation of the classical problem of testing linearity of functions $f : \{0, 1\}^n \to \{0, 1\}$ [23]; the classical special case is obtained by restricting U to be diagonal in the computational basis and to have diagonal entries all equal to ± 1.

In Protocol 3 we give a test that solves this problem using the product test. The test always accepts product unitaries, and rejects unitaries that are far from product, as measured by the normalised Hilbert-Schmidt inner product.

Protocol 3 (Product unitary test).

The product unitary test proceeds as follows.

1) *Prepare two copies of the state* $|\Phi\rangle^{\otimes n}$*, then in both cases apply* U *to the* n *first halves of each pair of qudits to create two copies of the state* $|v(U)\rangle \in (\mathbb{C}^{d^2})^{\otimes n}$.
2) *Return the result of applying the product test to the two copies of* $|v(U)\rangle$*, with respect to the partition into* n d^2*-dimensional subsystems.*

Let the probability that this test passes when applied to some unitary U be $P_{\text{test}}(U)$. Then we have the following theorem, which proves a conjecture from [12].

Theorem 17. *Given* $U \in U(d^n)$*, let*

$$1 - \epsilon = \max\{|\langle U, V_1 \otimes \cdots \otimes V_n\rangle|^2 : V_1, \ldots, V_n \in U(d)\}.$$

Then, if $\epsilon = 0$*,* $P_{test}(U) = 1$*. If* $\epsilon \lesssim 0.106$*, then* $P_{test}(U) \leq 1 - \frac{1}{4}\epsilon + \frac{1}{16}\epsilon^2 + \frac{1}{8}\epsilon^{3/2}$*. If* $0.106 \lesssim \epsilon \leq 1$*,* $P_{test}(U) \leq 501/512$*. More concisely,* $P_{test}(U) = 1 - \Theta(\epsilon)$.

The proof is not quite immediate from the previous results; the key problem is that the closest product state to $|v(U)\rangle$ may not correspond to the closest unitary operator to U.

Our test is sensitive to the Hilbert-Schmidt distance of a unitary from the set of product unitaries. One might hope to design a similar test that instead uses a notion of distance based on the operator norm. However, this is not possible. For example, if we could detect a constant difference in the operator norm between an arbitrary unitary U and the set of product unitaries then we could find a single marked item in a set of size d^n. By the optimality of Grover's algorithm, this requires $\Omega(d^{n/2})$ queries to U. More generally, any test that uses only a constant number of black-box queries to U can only detect an $\Omega(1)$ difference in an $\Omega(1)$ fraction of the d^n dimensions that U acts upon.

VII. Conclusion

Our main result can be seen as a "stability" theorem for the output purity of the depolarising channel. It is an interesting problem to determine whether a similar result holds for all output Rényi entropies for the depolarising channel, or even for all channels where additivity holds. As a more modest open question, can Theorem 1 be tightened further, perhaps by improving the constant in the $\epsilon^{3/2}$ term? It would also be interesting to improve the constants in Theorem 1 in the regime of large ϵ, as at present they are extremely pessimistic. The regime of large ϵ is generally somewhat mysterious: for example, we do not know the minimum value of P_{test}, or the largest distance from any product state that can be achieved by a state of n qudits.

Acknowledgements

AM was supported by the EC-FP6-STREP network QICS and an EPSRC Postdoctoral Research Fellowship. AWH was supported by the EPSRC grant "QIP-IRC". We would like to thank Salman Beigi, Toby Cubitt, Julia Kempe, Thomas Vidick and Andreas Winter for inspiring discussions.

References

[1] D. P. DiVincenzo, P. W. Shor, and J. A. Smolin, "Quantum channel capacity of very noisy channels," *Phys. Rev. A.*, vol. 57, p. 830, 1998, quant-ph/9706061.

[2] M. B. Hastings, "A counterexample to additivity of minimum output entropy," *Nature Physics*, vol. 5, 2009, arXiv:0809.3972.

[3] T. Ito, H. Kobayashi, and K. Matsumoto, "Oracularization and two-prover one-round interactive proofs against nonlocal strategies," 2008, arXiv:0810.0693.

[4] H. Kobayashi, K. Matsumoto, and T. Yamakami, "Quantum Merlin-Arthur proof systems: are multiple Merlins more helpful to Arthur?" in *Proc. ISAAC '03*, 2003, pp. 189–198, quant-ph/0306051.

[5] S. Aaronson, S. Beigi, A. Drucker, B. Fefferman, and P. Shor, "The power of unentanglement," *Theory of Computing*, vol. 5, no. 1, pp. 1–42, 2009, arXiv:0804.0802.

[6] J. Kempe and O. Regev, "No strong parallel repetition with entangled and non-signaling provers," 2009, arXiv:0911.0201.

[7] T. Ogawa and H. Nagaoka, "Strong converse to the quantum channel coding theorem," *IEEE Trans. Inform. Theory*, vol. 45, no. 7, pp. 2486–2489, 1999, quant-ph/9808063.

[8] A. Winter, "Coding theorem and strong converse for quantum channels," *IEEE Trans. Inform. Theory*, vol. 45, no. 7, pp. 2481–2485, 1999.

[9] E. Fischer, "The art of uninformed decisions: A primer to property testing," *Bulletin of the European Association for Theoretical Computer Science*, vol. 75, pp. 97–126, 2001.

[10] H. Buhrman, L. Fortnow, I. Newman, and H. Röhrig, "Quantum property testing," *SIAM J. Comput.*, vol. 37, no. 5, pp. 1387–1400, 2008, quant-ph/0201117.

[11] A. Atici and R. A. Servedio, "Quantum algorithms for learning and testing juntas," *Quantum Information Processing*, vol. 6, pp. 323–348, 2007, arXiv:0707.3479.

[12] A. Montanaro and T. Osborne, "Quantum boolean functions," *Chicago Journal of Theoretical Computer Science*, vol. to appear, 2010, arXiv:0810.2435.

[13] S. Aaronson, "The learnability of quantum states," *Proceedings of the Royal Society A*, vol. 463, p. 2088, 2007, quant-ph/0608142.

[14] O. Gühne and G. Toth, "Entanglement detection," *Physics Reports*, vol. 471, no. 1, 2009, arXiv:0811.2803.

[15] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf, "Quantum fingerprinting," *Phys. Rev. Lett.*, vol. 87, no. 16, p. 167902, 2001, quant-ph/0102001.

[16] F. Mintert, M. Kuś, and A. Buchleitner, "Concurrence of mixed multipartite quantum states," *Phys. Rev. Lett.*, vol. 95, no. 26, p. 260502, 2005, quant-ph/0411127.

[17] S. Walborn, P. Ribeiro, L. Davidovich, F. Mintert, and A. Buchleitner, "Experimental determination of entanglement with a single measurement," *Nature*, vol. 440, no. 7087, pp. 1022–1024, 2006.

[18] L. Gurvits, "Classical deterministic complexity of Edmonds' problem and quantum entanglement," in *Proc. 35th Annual ACM Symp. Theory of Computing*, 2003, pp. 10–19, quant-ph/0303055.

[19] S. Gharibian, "Strong NP-hardness of the quantum separability problem," *Quantum Inf. Comput.*, vol. 10, no. 3&4, pp. 343–360, 2010, arXiv:0810.4507.

[20] G. G. Amosov, A. S. Holevo, and R. F. Werner, "On some additivity problems in quantum information theory," *Problems Inform. Transmission*, vol. 36, no. 4, pp. 305–313, 2000, arXiv:math-ph/0003002.

[21] H. Blier and A. Tapp, "All languages in NP have very short quantum proofs," in *First International Conference on Quantum, Nano, and Micro Technologies*. Los Alamitos, CA, USA: IEEE Computer Society, 2009, pp. 34–37.

[22] Y.-K. Liu, "The complexity of the consistency and N-representability problems for quantum states," Ph.D. dissertation, Univ. of California, San Diego, 2007.

[23] M. Blum, M. Luby, and R. Rubinfeld, "Self-testing/correcting with applications to numerical problems," *J. Comput. Syst. Sci.*, vol. 47, no. 3, pp. 549–595, 1993.

[24] T. Wei and P. Goldbart, "Geometric measure of entanglement and applications to bipartite and multipartite quantum states," *Phys. Rev. A.*, vol. 68, no. 4, p. 42307, 2003, quant-ph/0307219.

[25] A. W. Harrow and A. Montanaro, "An efficient test for product states, with applications to quantum Merlin-Arthur games," 2010, arXiv:1001.0017.

[26] A. Barenco, A. Berthiaume, D. Deutsch, A. Ekert, R. Jozsa, and C. Macchiavello, "Stabilisation of quantum computations by symmetrisation," *SIAM J. Comput.*, vol. 26, no. 5, pp. 1541–1557, 1997, quant-ph/9604028.

[27] F. Brandão, "Entanglement theory and the quantum simulation of many-body physics," Ph.D. dissertation, Imperial College, London, 2008, arXiv:0810.0026.

[28] R. Impagliazzo and R. Paturi, "On the complexity of k-SAT," *J. Comput. Syst. Sci.*, vol. 62, no. 2, pp. 367–375, 2001.

[29] M. Grötschel, L. Lovász, and A. Schrijver, *Geometric algorithms and combinatorial optimization*. Springer-Verlag, 1993.

[30] S. Beigi, "NP vs QMA_log(2)," *Quantum Inf. Comput.*, vol. 10, no. 1&2, 2010, arXiv:0810.5109.

[31] P. W. Shor, "Equivalence of additivity questions in quantum information theory," *Comm. Math. Phys.*, vol. 246, no. 3, pp. 453–472, 2004, quant-ph/0305035.

[32] S. Beigi and P. Shor, "On the complexity of computing zero-error and Holevo capacity of quantum channels," 2007, arXiv: 0709.2090.

[33] K. Matsumoto, "Some new results and applications of additivity problem of quantum channel," Poster at QIP'05 conference, 2005.

[34] R. König and S. Wehner, "A strong converse for classical channel coding using entangled inputs," *Phys. Rev. Lett.*, vol. 103, p. 070504, 2009, arXiv:0903.2838.

[35] R. König, S. Wehner, and J. Wullschleger, "Unconditional security from noisy quantum storage," 2009, arXiv:0906.1030.

[36] P. Hayden and A. Winter, "Counterexamples to the maximal p-norm multiplicativity conjecture for all p $>$ 1," *Comm. Math. Phys.*, vol. 284, no. 1, pp. 263–280, 2008.

[37] T. Cubitt, A. W. Harrow, D. Leung, A. Montanaro, and A. Winter, "Counterexamples to additivity of minimum output p-Renyi entropy for p close to 0," *Comm. Math. Phys.*, vol. 284, pp. 281–290, 2008, arXiv:0712.3628.

[38] G. Vidal and J. I. Cirac, "Irreversibility in asymptotic manipulations of entanglement," *Phys. Rev. Lett.*, vol. 86, p. 022308, 2001, quant-ph/0102036.

[39] D. Yang, M. Horodecki, R. Horodecki, and B. Synak-Radtke, "Irreversibility for all bound entangled states," *Phys. Rev. Lett.*, vol. 95, p. 190501, 2005, quant-ph/0506138.

[40] I. Devetak, M. Junge, C. King, and M. B. Ruskai, "Multiplicativity of completely bounded p-norms implies a new additivity result," *Comm. Math. Phys.*, vol. 266, pp. 37–63, 2006, quant-ph/0506196.

[41] D. Yang, M. Horodecki, and Z. D. Wang, "An additive and operational entanglement measure: conditional entanglement of mutual information," *Phys. Rev. Lett.*, vol. 101, p. 140501, 2008, arXiv:0804.3683.

[42] M. Fannes and C. Vandenplas, "Finite size mean-field models," *J. Phys. A: Math. Gen.*, vol. 39, no. 45, p. 13843, 2006, quant-ph/0605216.

[43] D. Aharonov, I. Arad, Z. Landau, and U. Vazirani, "The detectability lemma and quantum gap amplification," in *Proc. 41st Annual ACM Symp. Theory of Computing*. New York, NY, USA: ACM, 2009, pp. 417–426, arXiv:0811.3412.