

Testing quantum circuits and detecting insecure encryption

Bill Rosgen
Centre for Quantum Technologies
National University of Singapore
bill.rosgen@nus.edu.sg

August 4, 2011

Abstract

We show that computational problem of testing the behaviour of quantum circuits is hard for the class of problems known as QMA that can be verified efficiently with a quantum computer. This result is a generalization of the techniques previously used to prove the hardness of other problem on quantum circuits. We use this result to show the QMA-hardness of a weak version of the problem of detecting the insecurity of a symmetric-key quantum encryption system, or alternately the problem of determining when a quantum channel is *not* private. We also give a QMA protocol for the problem of detecting insecure encryption to show that it is QMA-complete.

1 Introduction

Testing the behaviour of a computational system is a problem central to the study of quantum computing. This is the problem faced by an experimentalist who has implemented a quantum computation and wants to check that the implementation behaves (approximately) correctly on all input states. An efficient solution to this problem would allow for the verification that a circuit provided by an untrusted party correctly implements some desired operation. Unfortunately we show in a general model that even a weak version of this problem is likely to be computationally intractable and so any solution to this problem will need to make essential use of the structure of the operation that the circuit is supposed to implement. The problem we consider is, given a quantum circuit, to decide between two cases: either the circuit acts in the desired way on all input states, or the circuit misbehaves, acting in some malicious way on a large subspace of input states. This problem is QMA-hard even when both the desired and malicious behaviour are known (i.e. specified by uniform families of quantum circuits).

The class QMA is the set of all problems that can be verified up to bounded error on a quantum computer. Several problems are known to be complete for QMA: these problems can be thought of as alternate characterizations of the class, as they capture exactly the power of this computational model. The first of these complete problems is the problem of determining the ground state energy of a local Hamiltonian. This was first shown to be complete on k -local Hamiltonians [15] for $k \geq 5$, before the problem was shown to remain hard in the 2-local case [14]. The problem of determining if local descriptions of a quantum system are consistent is also known to be QMA-complete [16], though only under Turing reductions. Other problems related to finding ground states of physical systems are also known to be complete for QMA [20, 21].

There are also problems on quantum circuits that are known to be QMA-complete. The first of these is the Non-identity check problem [13], which given as input a unitary quantum circuit, the problem is

to decide if there is an input on which the circuit acts non-trivially or if the circuit is close to the identity for all input states. The problem of determining if a circuit is close to an isometry (i.e. a reversible transformation that maps pure states to pure states) is also known to be QMA-complete [18].

In this paper we generalize the hardness proofs of [13, 18] to show that the QMA-hardness of the problem of testing the properties of the outputs of quantum circuits. More specifically, we define the circuit testing problem, which has as parameters two uniformly generated families of quantum circuits \mathcal{C}_1 and \mathcal{C}_2 . The problem is to decide, given an input circuit C , whether C acts like circuits from the family \mathcal{C}_1 on a large input subspace, or whether C acts like circuits from \mathcal{C}_2 for all input states. Using this result we reprove the QMA-hardness of non-identity check and non-isometry testing by making choices for the families \mathcal{C}_1 and \mathcal{C}_2 . We also show that some other circuit problems are hard, such as a version of finding the minimum output entropy (this is similar in spirit to the results in [5], though our model is incompatible), or determining when a channel has a pure (approximate) fixed point.

It is important to note that, despite the name, this problem is not related to *property testing*. In this problem we have a significantly weaker promise—in one case the circuit only behaves in a certain way on a subspace of the input. For an input space of dimension d , this subspace can be as large as $d^{1-\delta}$ for an arbitrary constant $\delta > 0$ but this subspace is still *far* from the whole input space. Essentially the problem is to detect if the circuit behaves in a certain way only when a specific input state is provided on some subset of the input qubits. Note also that while we can use this problem to show the QMA-hardness of several circuit problems, this technique does *not* show that these problems are in QMA.

We then apply this hardness result to the problem of detecting insecure quantum encryption. This is the problem of deciding, given a quantum circuit that takes as input a quantum state as well as a classical key, whether this circuit is ϵ -close to a perfectly secure encryption scheme (i.e. a private quantum channel [2, 6]), or whether there is a large subspace of input states that the circuit does not encrypt at all (up to error ϵ). To show that this problem is hard, we argue that this problem contains as a special case an instance of the circuit testing problem. Finally, we give a QMA verifier for this problem to prove that it is QMA-complete.

The remainder of the paper is organized as follows: Section 2 contains some mathematical background, a definition of the class QMA, and a discussion of private quantum channels. The hardness of the circuit testing problem is shown in Section 3. Finally, Section 4 contains the proof that the problem of detecting insecure encryption is QMA-complete.

2 Preliminaries

2.1 Background

Throughout the paper we let $\mathcal{H}, \mathcal{K}, \mathcal{X}, \mathcal{Y}, \dots$ represent (finite-dimensional) Hilbert spaces. The pure quantum states are simply the unit vectors in these spaces. The set of density matrices on a space \mathcal{H} is denoted $\mathbf{D}(\mathcal{H})$: these are the positive semidefinite operators with unit trace. We will use the notation $\mathbf{T}(\mathcal{H}, \mathcal{K})$ to represent the set of channels that map states in $\mathbf{D}(\mathcal{H})$ to states in $\mathbf{D}(\mathcal{K})$. More formally, these transformations are exactly the completely positive trace preserving linear maps from $\mathbf{L}(\mathcal{H})$ to $\mathbf{L}(\mathcal{K})$, where we use $\mathbf{L}(\mathcal{H})$ to denote the set of all linear operators on \mathcal{H} .

To measure the distance between quantum states we will make extensive use of the trace norm, which for a linear operator X can be defined as $\|X\|_{\text{tr}} = \text{tr} \sqrt{X^*X}$. A useful alternate characterization is that $\|X\|_{\text{tr}}$ is the sum of the singular values of X , or, in the case of a normal operator, the sum of the absolute values of the eigenvalues. One important property of the trace distance $\|\rho - \sigma\|_{\text{tr}}$ between two

states is that it is monotone nonincreasing under the application of quantum channels.

We will also need the intuitive property that two states that are close together in the trace norm produce similar measurement outcomes. This can be derived from the fact that an expression involving the trace norm gives the maximum probability that two states can be distinguished [12],

Lemma 1. *Let $X \in \mathbf{L}(\mathcal{H})$ satisfy $0 \leq X \leq \mathbb{1}$. Then*

$$\mathrm{tr}(X\rho) \leq \mathrm{tr}(X\sigma) + \|\rho - \sigma\|_{\mathrm{tr}}$$

In addition to the trace norm, we will also need a distance measure on the quantum channels. Such a measure is given by the *diamond norm*, which for a linear map $\Phi : \mathbf{L}(\mathcal{H}) \rightarrow \mathbf{L}(\mathcal{H})$ is given by $\|\Phi\|_{\diamond} = \sup_{X \in \mathbf{L}(\mathcal{H} \otimes \mathcal{H})} \|(\Phi \otimes \mathbb{1}_{\mathcal{H}})(X)\|_{\mathrm{tr}} / \|X\|_{\mathrm{tr}}$. See [15] for an alternate definition and some further properties of this norm. In the case that Φ is the difference of two completely positive maps, we may replace the supremum in the definition of the diamond norm with a maximization over pure states in the space $\mathcal{H} \otimes \mathcal{H}$ [19]. Similarly to the trace norm, the diamond norm can be used to characterize the distinguishability of two quantum channels: here the fact that the definition involves a reference system captures the fact that the optimal strategy to distinguish two channels may involve the use of entangled input states.

Since we consider computational problems on quantum channels, we must specify how they are to be given as input. For this we use the mixed-state circuit model, first defined in [1], where circuits are composed of some (universal) collection of the usual unitary gates, plus a gate that introduces ancillary qubits in the $|0\rangle$ state and a gate that traces out (i.e. discards) qubits. For simplicity we will assume that all Hilbert spaces we encounter are composed of qubits, i.e. that the dimension is always a power of two, though this is not strictly needed.

We use this circuit model because it can (approximately) represent any quantum channel, and in the case of efficient quantum circuits this representation is of size polynomial in the number of input qubits. Using circuits does not (significantly) restrict the applicability of our hardness results: they also apply in any model that can efficiently simulate the circuit model, such as the model of measurement based quantum computation.

2.2 QMA

In order to prove results about the class QMA, we give a formal definition. A language L is in QMA if there is a quantum polynomial-time verifier V such that

1. if $x \in L$, then there exists a witness ρ such that $\Pr[V \text{ accepts } \rho] \geq 1 - \varepsilon$,
2. if $x \notin L$, then for any state ρ , $\Pr[V \text{ accepts } \rho] \leq \varepsilon$,

The exact value of the error parameter ε is not significant: any $\varepsilon < 1/2$ that is at least an inverse polynomial in the input size suffices [15, 17].

Let L be an arbitrary language in QMA, and let x be an arbitrary input string. Our goal will be to encode the QMA-hard problem of deciding if $x \in L$ into the problem of detecting an insecure encryption circuit. To do this it will be convenient to represent the verifier as a unitary circuit V , which represents the algorithm of the verifier in a QMA protocol on some input x . We may “hard-code” the input string x into the circuit for V , since the circuit V needs only to be efficiently generated given x .

The algorithm implemented by the verifier in an arbitrary QMA protocol is given in Figure 1. The verifier receives a witness state $|\psi\rangle$, applies the unitary V on the witness state and any ancillary qubits




Figure 1: Verifier's circuit in a QMA protocol. The verifier accepts the witness state $|\psi\rangle$ if and only if the measurement in the computational basis results in the $|1\rangle$ state.




Figure 2: Example implementation of the completely depolarizing channel Ω on three qubits. In order to obtain a private channel the state the qubits in the $|+\rangle$ state are replaced by a classical key k to obtain the channel Ω_k .

needed, and finally measures the first output qubit to decide whether or not to accept. Any qubits not measured are traced out. One of the main results of this paper is a reduction from an arbitrary QMA verifier to the problem of testing the behaviour of quantum circuits.

2.3 Private Quantum Channels

Quantum channels that are secure against eavesdroppers are those channels for which the input state cannot be determined by the output. These channels can also be viewed as encryption systems: the *key* is simply the environment space of the channel, which, when combined with the output state, allows the input to be recovered. We restrict attention to private channels of a special form: those which allow the input to be recovered not with the quantum state of the environment but instead with a classical key that can be pre-shared between two parties that wish to establish a secure quantum channel. These channels, called, *private* channels, were introduced and studied in [2, 6].

An important example of a private quantum channel is the completely depolarizing channel. This is the channel Ω that maps any input to the completely mixed state. One circuit implementation of this channel is given in Figure 2.

In order to use the completely depolarizing channel as a private channel we must add a key. This can be done to the implementation in Figure 2 by replacing the qubits in the $|+\rangle$ state with a classical string. The result is a channel that applies a key-specified Pauli to each of the input qubits. We will refer to this channel as Ω_k when a specific key is used. Notice that if $\Omega_k \in \mathbf{T}(\mathcal{H})$, then $|k| = 2 \log \dim \mathcal{H}$, i.e. we use

two key bits for each encrypted qubit. In the case of a perfect encryption channel this rate of two key bits per qubit is optimal [2, 6, 7]. When the key k is unknown and uniformly distributed, the channel Ω_k is identical to Ω , i.e. if the key k is uniformly distributed in $\{0, \dots, 2^m - 1\}$ we have

$$\frac{1}{2^m} \sum_k \Omega_k = \Omega. \quad (1)$$

We use the following definition of an approximately private channel (i.e. *secure encryption*).

Definition 2. Let E be a channel that takes two inputs: an integer $k \in \{1, \dots, K\}$ and a quantum state in \mathcal{H} and produces an output in \mathcal{K} , where $\dim \mathcal{H} \leq \dim \mathcal{K}$. For a fixed value of k we write $E_k(\cdot) = E(k, \cdot)$. We call E a ε -private channel if

1. There is a decryption channel, i.e. there exists a channel $D: \{1, \dots, K\} \otimes \mathbf{D}(\mathcal{K}) \rightarrow \mathbf{D}(\mathcal{H})$ such that for all k

$$\|D_k \circ E_k - \mathbb{1}_{\mathcal{H}}\|_{\diamond} \leq \varepsilon,$$

where the size of the circuit for D is bounded by a polynomial in the size of the circuit for E .

2. Without the key k , the output of E has almost no information about the input state, i.e.

$$\left\| \frac{1}{K} \sum_k E_k - \Omega \right\|_{\diamond} \leq \varepsilon$$

where $\Omega \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ is the depolarizing channel that maps all inputs to $\mathbb{1}_{\mathcal{K}} / \dim \mathcal{K}$.

The use of the diamond norm in this definition is significant: we require that both conditions hold even for part of an entangled state. Specifically, a channel satisfying this definition both preserves any entanglement with the transmitted state and remains secure even in the case that an eavesdropper is entangled with the input. We use this strong definition because one of the main results of the paper is a hardness result: distinguishing secure and insecure encryption remains hard even when the secure encryption is promised to be secure in this strong model. Our hardness result remains true for the weaker model of private channels using only the trace norm.

This definition is a strengthened version of the model used by Ambainis and Smith [3], who define security in a similar way, but only against adversaries that are not entangled with the input state. Another similar model is considered by Hayden et al. [11], which also does not consider entangled adversaries, but uses a stronger bound involving the operator norm. The hardness result in this paper does *not* apply with respect to this stronger bound.

Like the perfect encryption schemes found in [2, 6], the encryption scheme constructed by our reduction uses $2 \log d$ key bits to encrypt a state of dimension d . As argued (implicitly) in [4, 11] this is essentially optimal: any scheme using fewer than $2 \log d(1 - \text{poly}(\varepsilon))$ key bits cannot be secure against entangled adversaries.

3 Testing Circuits

The problem of testing the behaviour of a quantum circuit can be informally stated as: given a circuit C , decide between two cases, either the circuit acts like some known circuit C_0 on a large subspace of the input, or the circuit acts like some other known circuit C_1 on the whole input space. We use uniform circuit families \mathcal{C}_0 and \mathcal{C}_1 since it is important that the circuit C , which is provided as input, takes the same number of input and output qubits as the circuits C_0 and C_1 .

Problem 3 (CIRCUIT TESTING). Let $0 < \varepsilon < 1$, $0 < \delta \leq 1$, and $\mathcal{C}_0, \mathcal{C}_1$ be two uniform families of quantum circuits. The input to the problem is a circuit $C \in \mathbf{T}(\mathcal{X}, \mathcal{Y})$. Let C_0, C_1 be the circuits drawn from \mathcal{C}_0 and \mathcal{C}_1 that take as input states on \mathcal{X} . The promise problem is to decide between:

Yes: There exists a subspace S of \mathcal{X} with $\dim S \geq (\dim \mathcal{X})^{1-\delta}$ such that for any reference space \mathcal{R} and any $\rho \in \mathbf{D}(S \otimes \mathcal{R})$

$$\left\| (C \otimes \mathbb{1}_{\mathcal{R}})(\rho) - (C_0 \otimes \mathbb{1}_{\mathcal{R}})(\rho) \right\|_{\text{tr}} \leq \varepsilon,$$

No: $\|C - C_1\|_{\diamond} \leq \varepsilon$, i.e. for any reference space \mathcal{R} and any $\rho \in \mathcal{H} \otimes \mathcal{R}$

$$\left\| (C \otimes \mathbb{1}_{\mathcal{R}})(\rho) - (C_1 \otimes \mathbb{1}_{\mathcal{R}})(\rho) \right\|_{\text{tr}} \leq \varepsilon.$$

When the values of $\varepsilon, \delta, \mathcal{C}_0$, and \mathcal{C}_1 are significant we will refer to this problem as $\text{CT}(\varepsilon, \delta, \mathcal{C}_0, \mathcal{C}_1)$.

This problem is well-defined only for families \mathcal{C}_0 and \mathcal{C}_1 that do not violate the promise, i.e. any circuits whose output is not too close together. These are the circuits C_0 and C_1 such that there does not exist a subspace T of \mathcal{H} of size $\dim T > \dim \mathcal{H}^\delta$ such that for any input states $\rho \in \mathbf{D}(T \otimes \mathcal{R})$ we have $\left\| (C_0 \otimes \mathbb{1}_{\mathcal{R}})(\rho) - (C_1 \otimes \mathbb{1}_{\mathcal{R}})(\rho) \right\|_{\text{tr}} \leq 2\varepsilon$, i.e. there does not exist a large subspace of pure states on which C_0 and C_1 produce output that is close together. This condition can be difficult to verify, but in many applications it is easy to see that the two circuits do not agree on too many pure states. The application of this hardness result to detecting insecure encryption, for instance, uses C_0 as the identity and C_1 as the completely depolarizing channel, and these two circuits never agree on a pure input state. We are able to prove that this problem is QMA-hard for any circuit families that satisfy this condition.

Notice also the special case $\delta = 1$: here the CT problem asks if there are *any* input states on which the circuit C behaves like C_0 or if it behaves like C_1 for all input states. In this case the problem is well-defined for any families \mathcal{C}_0 and \mathcal{C}_1 that do not agree on the whole space (up to error 2ε).

Concerning the parameters ε and δ , we may take $\varepsilon = 2^{-p}$ for any polynomial p using an amplification result for QMA [15, 17], and we may take δ to be any constant satisfying $0 < \delta \leq 1$.

3.1 Testing Circuits is QMA-hard

To show the hardness of CT we use a reduction from an arbitrary problem in QMA. This involves embedding the verifier in a QMA protocol into an instance of CT with the property that the resulting circuit runs C_0 if the Verifier can be made to accept and runs C_1 if the Verifier cannot be made to accept.

Formalizing this notion, let L be an arbitrary language in QMA and let x be an input string. The QMA-complete problem is to decide whether or not $x \in L$. Since $L \in \text{QMA}$, there exists some unitary circuit $V : \mathcal{H} \otimes \mathcal{A} \rightarrow \mathcal{H}$ which can be constructed efficiently from x such that if $x \in L$, there exists a pure state $|\psi\rangle \in \mathcal{H}$ such that measuring the first qubit of $V(|\psi\rangle \otimes |0\rangle)$ results in $|1\rangle$ with probability at least $1 - \varepsilon$, whereas if $x \notin L$, then for any state $|\psi\rangle$ a measurement of $V(|\psi\rangle \otimes |0\rangle)$ results in $|1\rangle$ with probability at most ε . By using standard error-reduction techniques for QMA, we may take ε to be negligible in the size of the circuit for V [15, 17]. Notice also that the restriction to pure witness states $|\psi\rangle$ can be made without loss of generality using a convexity argument.

Our goal is to show that CT is hard for as many choices of parameters as possible. To this end, let $\delta > 0$ be constant and let \mathcal{C}_0 and \mathcal{C}_1 be uniform circuit families on which the problem $\text{CT}(3\sqrt{\varepsilon}, \delta, \mathcal{C}_0, \mathcal{C}_1)$ is well-defined. These are any families $\mathcal{C}_i = \{C_{i,n} : n \geq 1\}$, where the circuit $C_{i,n}$ takes an n qubit input state, such that for any n the circuits $C_{0,n}$ and $C_{1,n}$ do not produce outputs that are not too close together on some large subspace of pure input states. In particular, we require that for all n , there does not exist a

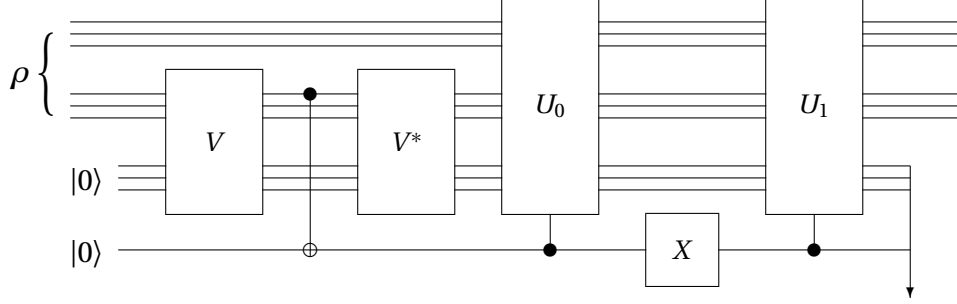


Figure 3: Circuit output by the reduction. The circuit V is the unitary circuit applied by the QMA verifier for the language L . The circuit U_i is the unitary circuit obtained from C_i by removing the gates that introduce ancillary qubits and trace out qubits.

subspace T of the n -qubit input space \mathcal{X} with $\dim T > \dim \mathcal{X}^\delta$ such that for any states $\rho \in \mathbf{D}(T \otimes \mathcal{R})$ we have $\|(C_0 \otimes \mathbb{1}_{\mathcal{R}})(\rho) - (C_1 \otimes \mathbb{1}_{\mathcal{R}})(\rho)\|_{\text{tr}} \leq 6\sqrt{\epsilon}$.

The key idea to the reduction is that we construct a circuit that takes an input state and applies the unitary V to a portion of it, makes a ‘copy’ of the output bit with a controlled-not gate, and then applies V^* . If the result of the QMA protocol would have been the verifier accepting (i.e. the copy of the output qubit is measured in the $|1\rangle$ state), then we apply the circuit C_0 . On the other hand, if the output qubit was in the $|0\rangle$ state, we apply the circuit C_1 . This results in a circuit that applies C_0 if and only if the input is a state the Verifier in the QMA proof system accepts. In order to guarantee that the subspace of accepting states is large enough, we add dummy input qubits that are ignored by the circuit V but are acted on by either C_0 or C_1 . By adding enough of these qubits, we can ensure that if there is at least one state V accepts, then the result is a large subspace of states that are accepted.

The full construction of the circuit produced by the reduction is shown in Figure 3. Before describing the circuit, we fix the notation that we will use. Let C_0 and C_1 be circuits drawn from \mathcal{C}_0 and \mathcal{C}_1 implementing transformations in $\mathbf{T}(\mathcal{X}, \mathcal{Y})$, where $\mathcal{X} = \mathcal{F} \otimes \mathcal{H}$ and $\mathcal{Y} = \mathcal{F} \otimes \mathcal{K}$, using the spaces \mathcal{H}, \mathcal{K} from the QMA Verifier for L . Further, we may let $\dim \mathcal{F} = \lceil \dim \mathcal{H}^{(1-\delta)/\delta} \rceil$, since we are free to take any polynomial number of input qubits to C_0 and C_1 . We also assume without loss of generality that these circuits are implemented by circuits that apply unitary circuits mapping $\mathcal{X} \otimes \mathcal{A} \rightarrow \mathcal{Y} \otimes \mathcal{G}$, where the space \mathcal{A} holds any ancillary qubits needed by the circuit (initially in the $|0\rangle$ state) and the space \mathcal{G} represents the qubits traced out at the end of the computation. Any mixed-state circuit can be efficiently transformed into a circuit of this form by moving the introduction of ancillary qubits to the start of the circuit and delaying any partial traces to the end of the circuit. We may also assume that both the circuit V and the circuits C_0 and C_1 use ancillary spaces \mathcal{A}, \mathcal{G} of the same size, by simply padding the circuits using a smaller space with unused ancillary qubits.

Let C be the circuit in Figure 3. This circuit takes as input a quantum state ρ on the space $\mathcal{X} = \mathcal{F} \otimes \mathcal{H}$. This circuit first applies V to the portion of ρ in \mathcal{H} as well as any needed ancillary qubits in the space \mathcal{A} . Next, the circuit makes a classical copy of the ‘output bit’ of V , which is used as a control for the application of the circuits C_0 and C_1 . The circuit V^* is then applied, so that the result (provided that V accepts or rejects with high probability) is a state that is close to the input state plus a qubit that indicates whether V accepts or rejects the input state. The circuit then applies C_0 if V accepts and C_1 if V rejects. These circuits use the same ancillary space \mathcal{A} as the circuits V and V^* , but as long as the Verifier V either accepts or rejects the input state with high probability, these ancillary qubits will be returned to

the $|0\rangle$ state, up to trace distance $2\sqrt{\varepsilon}$.

Before proving the correctness of the reduction, it will be convenient to write down some of the states produced by running the constructed circuit C . Let ρ be an arbitrary input state in $\mathbf{D}(\mathcal{H} \otimes \mathcal{F})$ and let $|\psi\rangle \in \mathcal{H} \otimes \mathcal{F} \otimes \mathcal{R}$ be a purification of ρ . The order of the spaces \mathcal{H} and \mathcal{F} has been changed for notational convenience. After applying the unitary V to the portion of $|\psi\rangle$ in \mathcal{H} , the state can be written as

$$|\phi\rangle = (V \otimes \mathbb{1}_{\mathcal{F}} \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle \otimes |0\rangle),$$

where the $|0\rangle$ qubits are in the space \mathcal{A} . Then, there exist states $|\phi_0\rangle, |\phi_1\rangle$ on all but the first qubit of $\mathcal{H} \otimes \mathcal{F} \otimes \mathcal{R}$ such that

$$|\phi\rangle = \sqrt{1-p}|0\rangle \otimes |\phi_0\rangle + \sqrt{p}|1\rangle \otimes |\phi_1\rangle$$

where $0 \leq p \leq 1$ is exactly the probability that the Verifier accepts in the original protocol on input $\text{tr}_{\mathcal{F}} \rho$. Applying the controlled-not gate results in

$$|\phi'\rangle = \sqrt{1-p}|00\rangle \otimes |\phi_0\rangle + \sqrt{p}|11\rangle \otimes |\phi_1\rangle.$$

We then bound the trace distance of $|\phi'\rangle$ to $|0\rangle|\psi\rangle$ and $|1\rangle|\psi\rangle$. In the case of $|0\rangle|\psi\rangle$ we have

$$\left\| |\phi'\rangle\langle\phi'| - |0\rangle\langle 0| \otimes |\phi\rangle\langle\phi| \right\|_{\text{tr}} = 2\sqrt{1 - |\langle\phi'|0\rangle|^2} = 2\sqrt{1 - (1-p)^2} < 3\sqrt{p}, \quad (2)$$

and in the similar case of $|1\rangle|\psi\rangle$ we have

$$\left\| |\phi'\rangle\langle\phi'| - |1\rangle\langle 1| \otimes |\phi\rangle\langle\phi| \right\|_{\text{tr}} = 2\sqrt{1 - |\langle\phi'|1\rangle|^2} = 2\sqrt{1 - p^2} < 3\sqrt{1-p}. \quad (3)$$

These two equations show that, when p is close to 0 or 1, the fact that we make a classical copy of the output qubit does not have a large effect on the state of the system. (This fact can also be argued from the Gentle Measurement Lemma [22].) The remainder of the circuit then applies V^* and, depending on the value of the control qubit, one of C_0 and C_1 . We consider two cases, which are argued in two separate propositions.

Proposition 4. *If $x \in L$, then there exists a subspace S of \mathcal{X} with $\dim S \geq \dim \mathcal{X}^{1-\delta}$ such that for any reference system \mathcal{R} and any $\rho \in S \otimes \mathcal{R}$*

$$\left\| (C \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle\langle\psi|) - (C_0 \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle\langle\psi|) \right\|_{\text{tr}} \leq 3\sqrt{\varepsilon}. \quad (4)$$

Proof. If $x \in L$, then there is some input state $|\psi\rangle$ on which the Verifier accepts with probability $p \geq 1 - \varepsilon$. Applying the remainder of the circuit, up to the partial trace, to the state $|1\rangle|\phi\rangle$ results in the state $|1\rangle \otimes (U_1 \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle \otimes |0\rangle)$. Tracing out the space \mathcal{G} as well as the copy of the output qubit, results in exactly the state $\text{tr}_{\mathcal{G}}(U_1 \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle\langle\psi| \otimes |0\rangle\langle 0|)(U_1^* \otimes \mathbb{1}_{\mathcal{R}}) = (C_1 \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle\langle\psi|)$. This is not quite equal to the output of the constructed circuit C , however, as in this evaluation we have replaced the state $|\phi'\rangle$ with the state $|1\rangle|\phi\rangle$. However, using the monotonicity of the trace norm under quantum operations, the remainder of the circuit cannot increase the norm of the two states, and so applying Equation (3), we have

$$\left\| (C \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle\langle\psi|) - (C_0 \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle\langle\psi|) \right\|_{\text{tr}} \leq 3\sqrt{1-p} \leq 3\sqrt{\varepsilon}. \quad (5)$$

It remains to show that this occurs on a large subspace of $\mathcal{X} = \mathcal{H} \otimes \mathcal{F}$. Since we have assumed the Verifier V accepts with high probability on the state $|\psi\rangle$, this implies that there is some state $|\gamma\rangle \in \mathcal{H}$ for

which V also accepts with probability at least $1 - \varepsilon$, as V ignores the qubits in \mathcal{F} . Then, since $|\psi\rangle$ was arbitrary, Equation (5) also applies to $|\gamma\rangle \otimes |\xi\rangle \in \mathcal{H} \otimes \mathcal{F}$ for *any* state $|\xi\rangle \in \mathcal{F}$. The subspace S of states whose reduced state on \mathcal{H} is equal to $|\gamma\rangle$ has dimension $\dim \mathcal{F}$. Then, since $\dim \mathcal{F} = \lceil \dim \mathcal{H}^{(1-\delta)/\delta} \rceil$, we have

$$\dim \mathcal{X} = \dim \mathcal{H} \otimes \mathcal{F} = \dim \mathcal{H} \dim \mathcal{F} \leq \dim \mathcal{F}^{\delta/(1-\delta)} \dim \mathcal{F} = \dim \mathcal{F}^{1/(1-\delta)},$$

which implies that $\dim \mathcal{F} \geq \dim \mathcal{X}^{1-\delta}$, as required. Thus, when $x \in L$ the Verifier V can be made to accept, and so the result is a yes instance of CT. \square

The remaining case is when $x \notin L$, i.e. the Verifier V rejects every state with high probability. This proof of this case is extremely similar to the previous one.

Proposition 5. *If $x \notin L$, then for any reference system \mathcal{R} and any $\rho \in \mathcal{X} \otimes \mathcal{R}$, $\|C - C_1\|_{\diamond} \leq 3\sqrt{\varepsilon}$.*

Proof. This proof is similar to the proof of Proposition 4. If $x \notin L$, then V accepts any state $|\psi\rangle$ with probability $p \leq \varepsilon$. If we consider applying V^* and the remainder of the circuit to the state $|0\rangle|\phi\rangle$, the result is $(C_1 \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle\langle\psi|)$, similarly to the previous case. Once again, we do not run this part of the circuit on this state, but the state $|\phi'\rangle$ which is very close to it. Once again we can apply the monotonicity of the trace norm under quantum operations and Equation (2) to show that

$$\left\| (C \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle\langle\psi|) - (C_1 \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle\langle\psi|) \right\|_{\text{tr}} \leq 3\sqrt{p} \leq 3\sqrt{\varepsilon}.$$

Since this equation applies for all reference systems \mathcal{R} and all states $|\psi\rangle$, this proves that if $x \notin L$, then we have $\|C - C_1\|_{\diamond} \leq 3\sqrt{\varepsilon}$. \square

Taken together, these two proposition prove the hardness of the CT problem. Note once again that in order for the CT problem to be well defined (i.e. the set of ‘yes’ instances does not intersect the set of ‘no’ instances) we require that circuits from the two families are not too close together for any large subspaces of pure input states. See the discussion following Problem 3 for a technical condition that is equivalent to this requirement.

Theorem 6. *CT($\varepsilon, \delta, \mathcal{C}_0, \mathcal{C}_1$) is QMA-hard for any $0 < \varepsilon < 1$ such that $\varepsilon \geq 2^{-p}$ for some polynomial p , any constant $0 < \delta \leq 1$, and any uniform circuit families $\mathcal{C}_0, \mathcal{C}_1$ for which the problem is well-defined.*

Proof. The correctness of the reduction is argued in In Propositions 4 and 5. It remains only to verify that the reduction can be performed efficiently. To see that the reduction can be performed in time polynomial in the size of the input x (which is at most polynomially smaller than the size of the circuit V : the only part of the reduction that can cause a problem the size of the space \mathcal{F} , since we have taken $\dim \mathcal{F} = \lceil \dim \mathcal{H}^{(1-\delta)/\delta} \rceil$). This implies that the space \mathcal{F} requires a factor of $(1-\delta)/\delta$ more qubits than the space \mathcal{H} , which is linear in the input dimension so long as δ is a constant. This implies that the reduction can be performed in (classical deterministic) polynomial time. \square

3.2 Applications

In this section we apply Theorem 6 to reprove the hardness of some of the circuit problems that are known to be hard for QMA as well as to show the QMA-hardness of some new circuit problems.

The first problem we consider is a slightly generalized version of the problem NON-IDENTITY CHECK studied by Janzing, Wocjan, and Beth [13], who show that it is QMA-complete. Our version of the problem differs in that we allow the input circuit to be a mixed-state circuit. We do still require, however, that

if the circuit does not act like the identity everywhere, then it acts like some efficient unitary circuit U on some input state for which U is far from the identity. This requirement is not needed to prove that this problem is hard, but it is hard to see how to put the problem into QMA without it.

Problem 7 (MIXED NON-IDENTITY CHECK [13]). Let $0 < \varepsilon < 1$. On input C , a circuit in $\mathbf{T}(\mathcal{X}, \mathcal{X})$, the promise problem is to decide between:

Yes: $\|C - \mathbb{1}\|_{\diamond} \geq 2 - \varepsilon$ and there exists an efficient unitary U such that on some pure state $|\psi\rangle \in \mathcal{X}$ we have $\|C(|\psi\rangle\langle\psi|) - U|\psi\rangle\langle\psi|U^*\|_{\text{tr}} \leq \varepsilon$ and $\|U|\psi\rangle\langle\psi|U^* - |\psi\rangle\langle\psi|\|_{\text{tr}} \geq 2 - \varepsilon$.

No: $\|C - \mathbb{1}\|_{\diamond} \leq \varepsilon$.

The QMA-hardness of this problem follows from Theorem 6 and the fact that $\text{CT}(\varepsilon, 1, \mathcal{U}, \mathbb{1})$ is a special case of the problem, where \mathcal{U} is any uniform family of quantum circuits that are not close to the identity (one such example is the family of circuits that apply Pauli X to the first input qubit).

The next problem we consider is the problem of detecting whether a (mixed-state) circuit is close to an isometry, which was shown to be QMA-complete in [18]. This can be formalized as the problem of detecting if there is a pure input state one which the output state is highly mixed.

Problem 8 (NON-ISOMETRY [18]). Let $0 < \varepsilon < 1/2$. On input a circuit $C \in \mathbf{T}(\mathcal{X}, \mathcal{Y})$ the promise problem is to decide between:

Yes: There exists $|\psi\rangle \in \mathcal{X}$ such that $\|(\Phi \otimes \mathbb{1}_{\mathcal{X}})(|\psi\rangle\langle\psi|)\|_{\infty} \leq \varepsilon$,

No: For all $|\psi\rangle \in \mathcal{X}$, $\|(\Phi \otimes \mathbb{1}_{\mathcal{X}})(|\psi\rangle\langle\psi|)\|_{\infty} \geq 1 - \varepsilon$.

The QMA-hardness of this problem follows from Theorem 6, since $\text{CT}(\varepsilon, 1, \Omega, \mathbb{1})$ is a special case.

We can also apply Theorem 6 to show the hardness of the problem of determining if a channel has a pure fixed point. This problem can be stated as follows.

Problem 9 (PURE FIXED POINT). Let $0 < \varepsilon < 1$. On input a circuit $C \in \mathbf{T}(\mathcal{X}, \mathcal{X})$ the promise problem is to decide between:

Yes: There exists $|\psi\rangle \in \mathcal{X}$ such that $\|C(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi|\|_{\text{tr}} \leq \varepsilon$

No: For any $|\psi\rangle \in \mathcal{X}$, $\|C(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi|\|_{\text{tr}} \geq 2 - \varepsilon$

The QMA-hardness of this problem follows from the fact that $\text{CT}(\varepsilon, 1, \mathbb{1}, \Omega)$ is a special case.

A related problem is determining if the minimum output entropy of a quantum channel is small. Related results can be found in [5], though the model used there seems to be incompatible with the model used in the present paper. In order to define this problem, let $S_{\min}(C) = \min_{\rho} S(C(\rho))$ be the minimum output entropy of the channel C (where S is the von Neumann entropy).

Problem 10 (MINIMUM OUTPUT ENTROPY). Let $0 < \varepsilon < 1/2$. On input a circuit $C \in \mathbf{T}(\mathcal{X}, \mathcal{X})$ the promise problem is to decide between:

Yes: $S_{\min}(C) \leq \varepsilon \log \dim \mathcal{X}$

No: $S_{\min}(C) \geq (1 - \varepsilon) \log \dim \mathcal{X}$

As in the previous case, the QMA-hardness of this problem follows from Theorem 6 and the fact that $\text{CT}(\varepsilon/2, 1, \mathbb{1}, \Omega)$ is a special case. The $\log \dim \mathcal{X}$ terms in the statement of the problem are due to the use of Fannes Inequality [10] to transform trace distance bounds to entropy bounds.

4 Detecting Insecure Encryption

In this section we consider the problem of detecting when a two-party symmetric key quantum encryption system is insecure. We first use Theorem 6 to show that this problem is hard, and then give a QMA-verifier to show that it is QMA-complete. The problem can be defined as follows.

Problem 11 (DETECTING INSECURE ENCRYPTION). For $0 < \varepsilon < 1$ and $0 < \delta \leq 1$ an instance of the problem consists of a quantum circuit E that takes as input a quantum state as well as a m classical bits, such that for each $k \in \{0, 1\}^m$ the circuit implements a quantum channel $E_k \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ with $\dim \mathcal{K} \geq \dim \mathcal{H}$. The promise problem is to decide between:

Yes: There exists a subspace S of \mathcal{H} with $\dim S \geq \dim \mathcal{H}^{1-\delta}$ such that for any reference space \mathcal{R} , any $\rho \in \mathbf{D}(S \otimes \mathcal{R})$, and any key k , $\|(E_k \otimes \mathbb{1}_{\mathcal{R}})(\rho) - \rho\|_{\text{tr}} \leq \varepsilon$.

No: E is an ε -private channel, i.e. $\|\Omega - \frac{1}{2^m} \sum_{k \in \{0, 1\}^m} E_k\|_{\diamond} \leq \varepsilon$, where Ω is the completely depolarizing channel in $\mathbf{T}(\mathcal{H}, \mathcal{K})$, and there exists an polynomial-size quantum circuit D such that for all k we have $\|D_k \circ E_k - \mathbb{1}_{\mathcal{H}}\|_{\diamond} \leq \varepsilon$.

When the values of ε and δ are significant, we will refer to this problem as $\text{DI}_{\varepsilon, \delta}$.

Informally, this is the problem of distinguishing two cases: either the channel fails to encrypt a large subspace of the input qubits (for any key), or the channel is very close to a perfect encryption channel.

Theorem 12. $\text{DI}_{\varepsilon, \delta}$ is QMA-hard for all $0 < \varepsilon < 1/2$ and all $0 < \delta \leq 1$.

Proof. Let $\mathcal{E}_k = \{\Omega_{k, n}\}$ where $\Omega_{k, n}$ is the n -qubit channel that applies the k th Pauli operator to the input qubits. As in Equation (1) averaging over all over all keys k results in the completely depolarizing channel on n qubits. Then, Theorem 6 implies that $\text{CT}(\varepsilon, \delta, \mathbb{1}_k, \mathcal{E}_k)$ is hard for QMA, where $\mathbb{1}_k$ is the channel that discards the key k and does nothing to the quantum input.

The problem $\text{CT}(\varepsilon, \delta, \mathbb{1}_k, \mathcal{E}_k)$ involves a slight redefinition of the problem CT to include both a quantum input, as well as a classical input k . This can be done without difficulty by including the classical input as part of the quantum input (to circuits in the families $\mathbb{1}_k$ and \mathcal{E}_k) that is immediately measured in the computational basis (and in the case of $\mathbb{1}_k$, discarded). The problem $\text{CT}(\varepsilon, \delta, \mathbb{1}_k, \mathcal{E}_k)$ remains hard after this modification.

The QMA-hardness of $\text{DI}_{\varepsilon, \delta}$ then follows immediately from the fact that the problem of detecting insecure encryption is simply $\text{CT}(\varepsilon, \delta, \mathbb{1}_k, \mathcal{E}_k)$ with a weakened promise. Since the sets of ‘yes’ instances of the two problems are identical, we need only verify the ‘no’ instances. Let the circuit $C \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ be a ‘no’ instance of $\text{CT}(\varepsilon, \delta, \mathbb{1}_k, \mathcal{E}_k)$ and let $C_k(\cdot) = C(|k\rangle\langle k| \otimes \cdot)$ be the circuit defined by hardcoding the input in the ‘key’ portion of the input space. Then, for any input ρ and any key k , we have $\|C_k - \Omega_k\|_{\diamond} \leq \varepsilon$, since this follows for the versions of these circuits without a hardcoded key (which is just a restriction of the input space). From this equation, the triangle inequality implies that

$$\left\| \Omega - \frac{1}{2^m} \sum_k C_k \right\|_{\diamond} \leq \frac{1}{2^m} \sum_k \|\Omega_k - C_k\|_{\diamond} \leq \varepsilon,$$

which is the property required by ‘no’ instances of DI. To see further that the output of C_k can be decrypted with knowledge of k , observe that $\Omega_k^{-1} \circ \Omega_k = \mathbb{1}$, and so it follows that

$$\|\Omega_k^{-1} \circ C_k - \mathbb{1}\|_{\diamond} \leq \|\Omega_k^{-1} \circ C_k - \Omega_k^{-1} \circ \Omega_k\|_{\diamond} + \|\Omega_k^{-1} \circ \Omega_k - \mathbb{1}\|_{\diamond} \leq \|C_k - \Omega_k\|_{\diamond} \leq \varepsilon,$$

which implies that instances of $\text{CT}(\varepsilon, \delta, \mathbb{1}_k, \mathcal{E}_k)$ are equivalent to instances of $\text{DI}_{\varepsilon, \delta}$, as required. \square

4.1 QMA Protocol

To test the security of an encryption system in QMA the Verifier will need a tool to compare two quantum states. Such a tool is provided by the swap test, introduced in [8], though here we essentially use it to test the purity of quantum states as is done in [9].

The swap test is an efficient procedure that makes the projective measurement onto the symmetric and antisymmetric subspaces of a bipartite space. Let W be the swap operation on $\mathcal{H} \otimes \mathcal{H}$, i.e. $W(|\psi\rangle \otimes |\phi\rangle) = |\phi\rangle \otimes |\psi\rangle$ for all $|\psi\rangle, |\phi\rangle \in \mathcal{H}$. The swap test performs the two-outcome projective measurement given by the projection onto the symmetric subspace, given by $(\mathbb{1}_{\mathcal{H} \otimes \mathcal{H}} + W)/2$, and the projection onto the antisymmetric subspace, given by $(\mathbb{1}_{\mathcal{H} \otimes \mathcal{H}} - W)/2$.

Given two pure states $|\psi\rangle, |\phi\rangle$, the swap test returns the symmetric outcome with probability $(1 + |\langle\psi|\phi\rangle|^2)/2$. When applied to mixed states ρ, σ , the swap test can also be used to estimate the overlap, as the result is symmetric with probability $(1 + \text{tr}(\rho\sigma))/2$, as observed in [9]. Notice that this implies that the swap test can be used to estimate the purity of a state, given two copies.

The idea behind the protocol is that if the encryption system specified by E is insecure then, regardless of the key chosen, it acts trivially on some subspace of the input states. In this case a proof can consist simply of two copies of some pure state in this subspace. The Verifier runs E on both of these states in parallel and tests that they have not been changed by performing the swap test. In the case that the circuit is insecure, this proof state will cause the Verifier to obtain the symmetric outcome of the swap test with probability approaching 1. Note that this protocol does *not* check that the input state is unchanged, only that the output states of the two applications of E are (close to) the same pure state.

If E represents a secure encryption system, then without knowledge of the key, the output of E is close to the completely mixed state, regardless of the input state. In this case the Verifier performs the swap test on two highly mixed states and the result is antisymmetric with probability close to $1/2$.

This protocol can be formalized as follows. A circuit implementation can be found in Figure 4.

Protocol 13. On input a circuit $E: \{1, \dots, K\} \otimes \mathbf{D}(\mathcal{H}) \rightarrow \mathbf{D}(\mathcal{H})$, an instance of $\text{DI}_{\varepsilon, \delta}$, as well as a quantum proof $|\phi\rangle$ in $\mathbf{D}((\mathcal{H} \otimes \mathcal{R})^{\otimes 2})$ (where $\dim \mathcal{R} = \dim \mathcal{H}$), the Verifier performs the following protocol.

1. The Verifier generates random keys $k_1, k_2 \in \{1, \dots, K\}$.
2. The Verifier applies $(E_{k_1} \otimes \mathbb{1}_{\mathcal{R}}) \otimes (E_{k_2} \otimes \mathbb{1}_{\mathcal{R}})$ to the state $|\phi\rangle$.
3. The Verifier applies the swap test to the resulting state, accepting if the outcome is symmetric.

The reference space \mathcal{R} appears in this protocol, but Problem 11 places no upper bound on the size of this space, and the value of the norm being verified may increase with the size of the space \mathcal{R} . Fortunately, this process stabilizes when $\dim \mathcal{R} = \dim \mathcal{H}$, and so we may assume that this space is of this size, which at most doubles the number of input qubits to the protocol.

A straightforward argument based on the continuity of measurement probabilities (here given as Lemma 1) can be used to show that this protocol is correct.

Proposition 14. For $0 < \varepsilon < 1/8$, Protocol 13 is a QMA protocol for $\text{DI}_{\varepsilon, \delta}$.

Proof. If E is a ‘yes’ instance of $\text{DI}_{\varepsilon, \delta}$, then there exists a state $|\psi\rangle \in \mathcal{H} \otimes \mathcal{R}$ such that for any key $k \in \{1, \dots, K\}$ we have $\|\hat{E}_k(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi|\|_{\text{tr}} \leq \varepsilon$, where throughout this proof we use the shorthand notation $\hat{E}_k = E_k \otimes \mathbb{1}_{\mathcal{R}}$. Let the input state be $|\phi\rangle = |\psi\rangle \otimes |\psi\rangle$. Fixing notation further, let $\hat{E}_k(|\psi\rangle\langle\psi|) = \sigma_k$. Applying $\hat{E}_{k_1} \otimes \hat{E}_{k_2}$ to $|\psi\rangle \otimes |\psi\rangle$ results in a state $\sigma_{k_1} \otimes \sigma_{k_2}$ that satisfies

$$\|\sigma_{k_1} \otimes \sigma_{k_2} - |\psi\rangle\langle\psi| \otimes |\psi\rangle\langle\psi|\|_{\text{tr}} \leq 2\varepsilon, \quad (6)$$




Figure 4: The Verifier's circuit in the QMA protocol.

which follows from the triangle inequality. Then, since the state $|\psi\rangle\langle\psi| \otimes |\psi\rangle\langle\psi|$ is symmetric and we can view the swap test can be viewed as a projective measurement, Lemma 1 shows that the swap test returns the symmetric outcome on $\sigma_{k_1} \otimes \sigma_{k_2}$ with probability at least $1 - 2\varepsilon$. This implies that when the circuit E is not secure the Verifier accepts with high probability.

It remains to show that when the circuit E is a 'no' instance of $\text{DI}_{\varepsilon, \delta}$ the Verifier does not accept any proof state with high probability. In this case we know that $\|\sum_{k=1}^K E_k - \Omega\|_{\diamond} / K \leq \varepsilon$. Once more, a straightforward argument using the triangle inequality can be used to argue that the tensor product of two copies satisfies the equation $\|\sum_{k,j=1}^K E_k \otimes E_j - \Omega \otimes \Omega\|_{\diamond} / K^2 \leq 2\varepsilon$. This implies that regardless of the proof state $|\psi\rangle$ the input to the swap test is within trace distance 2ε of the completely mixed state. On such a state, Lemma 1 implies that the swap test returns the symmetric outcome with probability at most

$$\frac{1}{2} - \frac{1}{2} \text{tr} \left[\left(\frac{\mathbb{1}_{\mathcal{K}}}{\dim \mathcal{K}} \right)^2 \right] + 2\varepsilon = \frac{1}{2} - \frac{1}{2 \dim \mathcal{K}} + 2\varepsilon,$$

and so the probability the Verifier accepts is bounded above by $1/2 + 2\varepsilon$. Thus, when $\varepsilon < 1/8$, there is a constant gap between the acceptance probabilities in the two cases, and so $\text{DI}_{\varepsilon, \delta} \in \text{QMA}$. \square

Combining the previous Proposition with Theorem 12 we obtain the main result.

Theorem 15. *For $0 < \varepsilon < 1/8$ and $0 < \delta \leq 1$, the problem $\text{DI}_{\varepsilon, \delta}$ is QMA-complete.*

5 Discussion

We have shown the QMA-hardness of a general version of the problem of testing the behaviour of a quantum circuit. This result generalizes the proofs of hardness for many of the known circuit problems that are QMA-hard [13, 18], as well as allows for simple proofs of hardness for new circuit problems. As an application of this result we have shown that the problem of detecting insecure encryption is complete for QMA by in addition finding an efficient QMA verifier for the problem.

An open problem related to this is to find a QMA verifier for the PURE FIXED POINT problem, or an argument that the problem is likely to lie outside of the class. The direct approach to construct a verifier using the swap test on (ideally) two copies of the fixed-point state, similar to the verifier in [18], does not seem to work: the circuit that measures a qubit in the computational basis and then applies the Pauli X gate, when applied to half of the input space, maps the symmetric state $|01\rangle + |10\rangle$ to a symmetric state. This circuit, however, does not have any pure (approximate) fixed points.

Acknowledgements

I am grateful for discussions with Markus Grassl, Matthew McKague, and Lana Sheridan. This work has been supported by the Centre for Quantum Technologies, which is funded by the Singapore Ministry of Education and the Singapore National Research Foundation.

References

- [1] D. Aharonov, A. Kitaev, and N. Nisan. Quantum circuits with mixed states. In *Proceedings of the 30th ACM Symposium on the Theory of Computing*, pp. 20–30. 1998. DOI: 10.1145/276698.276708. EPRINT: arXiv:quant-ph/9806029.
- [2] A. Ambainis, M. Mosca, A. Tapp, and R. de Wolf. Private quantum channels. In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science*, pp. 547–553. 2000. DOI: 10.1109/SFCS.2000.892142. EPRINT: arXiv:quant-ph/0003101.
- [3] A. Ambainis and A. Smith. Small pseudo-random families of matrices: Derandomizing approximate quantum encryption. In *Proceedings of the 8th International Workshop on Randomization and Computation*, volume LNCS 3122, pp. 249–260. 2004. DOI: 10.1007/978-3-540-27821-4_23. EPRINT: arXiv:quant-ph/0404075.
- [4] H. Barnum, C. Crépeau, D. Gottesman, A. Smith, and A. Tapp. Authentication of quantum messages. In *Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science*, pp. 449 – 458. 2002. DOI: 10.1109/SFCS.2002.1181969. EPRINT: arXiv:quant-ph/0205128.
- [5] S. Beigi and P. W. Shor. On the complexity of computing zero-error and Holevo capacity of quantum channels, 2007. EPRINT: arXiv:0709.2090v3 [quant-ph].
- [6] P. O. Boykin and V. Roychowdhury. Optimal encryption of quantum bits. *Physical Review A*, 67(4):042317, 2003. DOI: 10.1103/PhysRevA.67.042317. EPRINT: arXiv:quant-ph/0003059.
- [7] S. Braunstein, H.-K. Lo, and T. Spiller. Forgetting qubits is hot to do. Unpublished manuscript, 1999.
- [8] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16):167902, 2001. DOI: 10.1103/PhysRevLett.87.167902. EPRINT: arXiv:quant-ph/0102001.
- [9] A. K. Ekert, C. M. Alves, D. K. Oi, M. Horodecki, P. Horodecki, and L. C. Kwek. Direct estimations of linear and nonlinear functionals of a quantum state. *Physical Review Letters*, 88(21):217901, 2002. DOI: 10.1103/PhysRevLett.88.217901. EPRINT: arXiv:quant-ph/0203016.

- [10] M. Fannes. A continuity property of the entropy density for spin lattice systems. *Communications in Mathematical Physics*, 31(4):291–294, 1973. DOI: 10.1007/BF01646490.
- [11] P. Hayden, D. Leung, P. W. Shor, and A. Winter. Randomizing quantum states: constructions and applications. *Communications in Mathematical Physics*, 250:371–391, 2004. DOI: 10.1007/s00220-004-1087-6. EPRINT: arXiv:quant-ph/0307104.
- [12] C. W. Helstrom. Detection theory and quantum mechanics. *Information and Control*, 10(3):254–291, 1967. DOI: 10.1016/S0019-9958(67)90302-6.
- [13] D. Janzing, P. Wocjan, and T. Beth. “Non-identity-check” is QMA-complete. *International Journal of Quantum Information*, 3(3):463–473, 2005. DOI: 10.1142/S0219749905001067. EPRINT: arXiv:quant-ph/0305050.
- [14] J. Kempe, A. Kitaev, and O. Regev. The complexity of the local Hamiltonian problem. *SIAM Journal on Computing*, 35(5):1070–1097, 2006. DOI: 10.1137/S0097539704445226. EPRINT: arXiv:quant-ph/0406180.
- [15] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*, volume 47 of *Graduate Studies in Mathematics*. American Mathematical Society, 2002.
- [16] Y.-K. Liu. Consistency of local density matrices is QMA-complete. In *Proceedings of the 10th International Workshop on Randomization and Computation*, volume 4110 of *Lecture Notes in Computer Science*, pp. 438–449. Springer, 2006. DOI: 10.1007/11830924_40. EPRINT: arXiv:quant-ph/0604166.
- [17] C. Marriott and J. Watrous. Quantum Arthur-Merlin games. *Computational Complexity*, 14(2):122–152, 2005. DOI: 10.1007/s00037-005-0194-x. EPRINT: arXiv:cs/0506068.
- [18] B. Rosgen. Testing non-isometry is QMA-complete. In *Proceedings of the 5th Conference on the Theory of Quantum Computation, Communication, and Cryptography*, pp. 63–76. 2010. DOI: 10.1007/978-3-642-18073-6_6. EPRINT: arXiv:0910.3740 [quant-ph].
- [19] B. Rosgen and J. Watrous. On the hardness of distinguishing mixed-state quantum computations. In *Proceedings of the 20th Conference on Computational Complexity*, pp. 344–354. 2005. DOI: 10.1109/CCC.2005.21. EPRINT: arXiv:cs/0407056.
- [20] N. Schuch, I. Cirac, and F. Verstraete. Computational difficulty of finding matrix product ground states. *Physical Review Letters*, 100(25):250501, 2008. DOI: 10.1103/PhysRevLett.100.250501. EPRINT: arXiv:0802.3351 [quant-ph].
- [21] N. Schuch and F. Verstraete. Computational complexity of interacting electrons and fundamental limitations of density functional theory. *Nature Physics*, 5(10):732 – 735, 2009. DOI: doi:10.1038/nphys1370. EPRINT: arXiv:0712.0483 [quant-ph].
- [22] A. Winter. Coding theorem and strong converse for quantum channels. *IEEE Transactions on Information Theory*, 45(7):2481–2485, 1999. DOI: 10.1109/18.796385.