Testing Quantum Circuits and Detecting Insecure Encryption

Bill Rosgen

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

Abstract. We show that the computational problem of testing the behaviour of quantum circuits is hard for the class QMA of problems that can be verified efficiently with a quantum computer. This result generalizes techniques previously used to prove the hardness of other problems on quantum circuits. We use this result to show the QMA-completeness 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.

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. 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 behaviours are known in advance (i.e. are a part of the problem definition).

The class QMA is the set of all (promise) 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 the model. The first of these complete problems is the problem of determining the ground state energy of a local Hamiltonian [11]. The problem of determining if local descriptions of a quantum system are consistent is also known to be QMA-complete [12], though only under Turing reductions. Other problems related to finding ground states of physical systems are also complete for QMA [16,17].

A different set of QMA-complete problems involve quantum circuits. The first of these is the Non-identity check problem [10]: given a unitary quantum circuit

K. Iwama, Y. Kawano, and M. Murao (Eds.): TQC 2012, LNCS 7582, pp. 74–86, 2013.

as input, 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 [15].

In this paper we generalize the hardness proofs of [10,15] to show the QMA-hardness of testing the properties of the outputs of quantum circuits. Specifically, we define the circuit testing problem, which has as parameters two uniformly generated families of quantum circuits $\mathcal{C}_0$ and $\mathcal{C}_1$. The problem is to decide, given an input circuit C, whether C acts like circuits from the family $\mathcal{C}_0$ on a large input subspace, or whether C acts like circuits from $\mathcal{C}_1$ for all input states. It is important to note that the circuit families $\mathcal{C}_0, \mathcal{C}_1$ are part of the problem definition: each choice of circuit families gives a different problem and an algorithm for a specific one of these problems may depend on these families in a non-uniform way. The main result of the paper is a proof that this circuit testing problem is QMA-hard for any circuit families $\mathcal{C}_0, \mathcal{C}_1$ for which the problem is well-defined. Using this result we reprove the QMA-hardness of non-identity check and non-isometry testing as well as proving the hardness of a few other circuit problems. This is done by choosing specific families $\mathcal{C}_0$ and $\mathcal{C}_1$ for which these problems reduce to the associated circuit testing problems.

We then apply the 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 close to a perfect encryption scheme (i.e. a private quantum channel [2,4]), or whether there is a large subspace of input states that the circuit does not encrypt. To prove hardness, 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.

2 Preliminaries

Throughout the paper the set of density matrices on a Hilbert space $\mathcal{H}$ is denoted $\mathbf{D}(\mathcal{H})$ while $\mathbf{T}(\mathcal{H}, \mathcal{K})$ is is the set of channels that map $\mathbf{D}(\mathcal{H})$ to $\mathbf{D}(\mathcal{K})$. To measure the distance between states we will make extensive use of the trace norm, $\|X\|_{\mathrm{tr}}$, which for a linear operator X is given by the sum of the absolute values of the singular values of X. One important property of the trace distance $\|\rho - \sigma\|_{\mathrm{tr}}$ is that it does not increase 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 follows from the fact that an expression involving the trace norm gives the maximum probability that states can be distinguished [9].

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 quantum channels. Such a measure is given by the *diamond norm*, which for a linear map

$\Phi : \mathbf{L}(\mathcal{H}) \to \mathbf{L}(\mathcal{K})$ is defined as $\|\Phi\|_\diamond = \sup_{X \in \mathbf{L}(\mathcal{H}\otimes\mathcal{H})} \|(\Phi \otimes \mathbb{1}_\mathcal{H})(X)\|_{\mathrm{tr}} / \|X\|_{\mathrm{tr}}$. In the case that Φ is the difference of two completely positive maps, we may replace the supremum in the definition with a maximization over pure states in the space $\mathcal{H} \otimes \mathcal{H}$ [14]. As in the case of the trace norm, the diamond norm can be used to characterize the distinguishability of two quantum channels: here the reference system captures the fact that the optimal strategy to distinguish two channels may need entanglement.

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 [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 assume that all Hilbert spaces we encounter are composed of qubits, though this is not essential to our results. We use this circuit model because it can (approximately) represent any quantum channel and in the case of efficient quantum circuits the size of this representation polynomial in the number of input qubits. Using circuits does not (significantly) restrict the applicability of our hardness results: they apply also in any model that can efficiently simulate the circuit model.

2.1 QMA

A promise problem $P = (P_{\mathrm{yes}}, P_{\mathrm{no}}) \in \mathsf{QMA}$ if there is a quantum poly-time verifier V such that

1. if $x \in P_{\mathrm{yes}}$, then there exists a witness ρ such that $\Pr[V \text{ accepts } \rho] \geq 1 - \varepsilon$,
2. if $x \in P_{\mathrm{no}}$, then for any state ρ, $\Pr[V \text{ accepts } \rho] \leq \varepsilon$.

The exact value of ε is not significant: any $\varepsilon < 1/2$ that is at least an inverse polynomial suffices [13].

Let P be an arbitrary promise problem in QMA, and let x be an arbitrary input string. Our goal will be to encode the QMA-hard problem of deciding P 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 as follows: the verifier receives a witness state $|\psi\rangle$, applies the unitary V on the witness state and any ancillary qubits 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.

2.2 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 the environment, which, when combined with the output state, allows the input to be recovered. We restrict attention to the private channels that 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. These channels, called *private* channels, were introduced and studied in [2,4].

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. This channel can be efficiently implemented by applying a random Pauli operator to each qubit. In order to use the completely depolarizing channel as a private channel we must add a key. This can be done by applying a key-selected Pauli operator 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 key rate is optimal [2,4,5]. When k is unknown and uniformly distributed, the channel Ω_k is identical to Ω, i.e. if the key k is uniformly distributed in $\{1, \ldots, K\}$ we have $\sum_k \Omega_k / K = \Omega$.

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

Definition 2. *Let E be a channel that takes inputs $k \in \{1, \ldots, K\}$ and a state in $\mathcal{H}$ and produces an output in $\mathcal{K}$, where $\dim \mathcal{H} \leq \dim \mathcal{K}$. For fixed k we write $E_k(\cdot) = E(k, \cdot)$. E is* ε-private *if*

1. *There exists a polynomial-size circuit $D\colon \{1, \ldots, K\} \otimes \mathbf{D}(\mathcal{K}) \to \mathbf{D}(\mathcal{H})$ such that for all k $\| D_k \circ E_k - \mathbb{1}_{\mathcal{H}} \|_\diamond \leq \varepsilon$.*
2. *Without k, the output of E is random, i.e. $\| \sum_k E_k / K - \Omega \|_\diamond \leq \varepsilon$.*

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 encrypted state and remains secure even against an entangled eavesdropper. We use this strong definition because one of the main results of the paper is a hardness result: distinguishing secure and insecure encryption is hard even when the secure encryption is promised to be secure in this model. Our results are also true in the weaker model using 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. The model considered by Hayden et al. [8] uses a stronger bound involving the operator norm under which our hardness result does not apply, as it is ultimately derived from the definition of QMA, and the probability that the Verifier in a QMA protocol can be made to accept is more naturally modelled by the trace norm.

3 Testing Circuits

The problem of testing the behaviour of a circuit can be informally stated as: given a circuit C decide if the circuit acts like some known circuit C_0 on a large

subspace of the input or if the circuit acts like some other known circuit C_1 on the whole input space. We use uniform circuit families $\mathcal{C}_0$ and $\mathcal{C}_1$ as it is important that the circuits C, C_1, and C_2 agree on input and output spaces.

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 is a circuit $C \in \mathbf{T}(\mathcal{X}, \mathcal{Y})$. Let C_0, C_1 be the circuits 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 $\rho \in \mathbf{D}(S \otimes \mathcal{R})$,

$$\|(C \otimes \mathbb{1}_{\mathcal{R}})(\rho) - (C_0 \otimes \mathbb{1}_{\mathcal{R}})(\rho)\|_{\mathrm{tr}} \leq \varepsilon.$$

No: $\|C - C_1\|_{\diamond} \leq \varepsilon$, i.e. for any $\mathcal{R}$, $\rho \in \mathbf{D}(\mathcal{H} \otimes \mathcal{R})$,

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

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

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 $\mathcal{X}$ of size $\dim T \geq \dim \mathcal{X}^{\delta}$ such that for any input states $\rho \in \mathbf{D}(T \otimes \mathcal{R})$ we have $\|(C_0 \otimes \mathbb{1}_{\mathcal{R}})(\rho) - (C_1 \otimes \mathbb{1}_{\mathcal{R}})(\rho)\|_{\mathrm{tr}} \leq 2\varepsilon$, 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 for many families of circuits it is easy to see that they are not too close together. As an example, the application of this hardness result to detecting insecure encryption takes C_0 as the identity and C_1 as the completely depolarizing channel, and these two circuits never agree on pure states. We show that this problem is QMA-hard for any circuit families that satisfy this condition.

Note the special case $\delta = 1$: here the CT problem asks if there are *any* inputs on which the circuit C behaves like C_0 or if it behaves like C_1 for all inputs. 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 choose $\varepsilon = 2^{-p}$ for any polynomial p using an amplification result for QMA [13] and we may choose δ any constant satisfying $0 < \delta \leq 1$.

3.1 Testing Circuits Is QMA-Hard

To show the hardness of CT we reduce 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 P be an arbitrary promise problem in QMA and let x be an input string. The QMA-complete problem is to decide whether or not $x \in P_{\text{yes}}$. Since $P \in \mathsf{QMA}$, there exists some unitary circuit $V : \mathcal{H} \otimes \mathcal{A} \to \mathcal{K}$ which can be constructed efficiently from x such that if $x \in P_{\text{yes}}$, 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 $x \in P_{\text{no}}$, 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 [13]. Notice that the restriction to pure witness states $|\psi\rangle$ can be made without loss of generality by 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 $\mathrm{CT}(3\sqrt{\varepsilon}, \delta, \mathcal{C}_0, \mathcal{C}_1)$ 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 too close together on some large subspace of pure input states. In particular, we require that for all n, there does not exist a subspace T of the n-qubit input space $\mathcal{X}$ with $\dim T \geq \dim \mathcal{X}^{\delta}$ such that for any states $\rho \in \mathbf{D}(T \otimes \mathcal{R})$ we have

$$\|(C_0 \otimes \mathbb{1}_{\mathcal{R}})(\rho) - (C_1 \otimes \mathbb{1}_{\mathcal{R}})(\rho)\|_{\text{tr}} \leq 6\sqrt{\varepsilon}.$$

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. The resulting circuit applies C_0 if and only 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 V accepts at least one state then the result is a large subspace of accepted states.

The full construction of the circuit produced by the reduction is shown in Figure 1. Before describing the circuit, we fix notation: 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 P. Further, we may let $\dim \mathcal{F} = \left\lceil \dim \mathcal{H}^{(1-\delta)/\delta} \right\rceil$, 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 $\mathcal{X} \otimes \mathcal{A} \to \mathcal{Y} \otimes \mathcal{G}$, 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

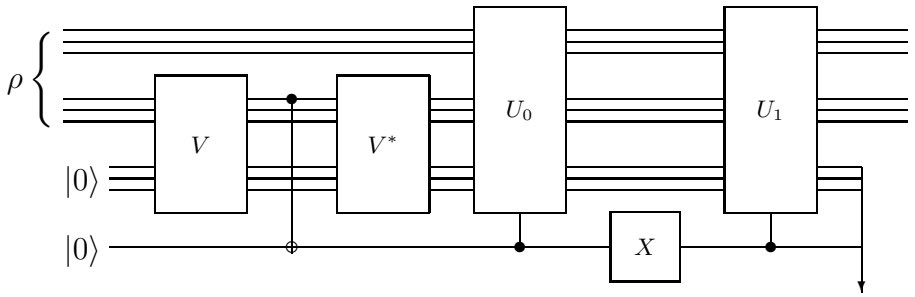

Fig. 1. Circuit output by the reduction. V is the unitary circuit applied by the original QMA verifier and U_i is the unitary circuit obtained from C_i by removing the gates that introduce ancillary qubits and trace out qubits.

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 1. 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 of 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. Applying the unitary V to the portion of $|\psi\rangle$ in $\mathcal{H}$ results in the state

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

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 $\mathcal{K} \otimes \mathcal{F} \otimes \mathcal{R}$ 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 $\mathrm{tr}_{\mathcal{F}} \rho$. 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

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

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

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

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 [18].) 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 $x \in P_{yes}$, 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 $|\psi\rangle \in S \otimes \mathcal{R}$*

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

Proof. If $x \in P_{\mathrm{yes}}$, 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 $\mathrm{tr}_{\mathcal{G}}(U_1 \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle\langle\psi| \otimes |0\rangle\langle 0|)(U_1^* \otimes \mathbb{1}_{\mathcal{R}}) = (C_1 \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle\langle\psi|)$. This is not quite equal to the output of the constructed circuit C, however, as 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, and so by Equation (2) we have

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

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 (3) 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

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

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

The remaining case is when $x \in P_{\mathrm{no}}$, i.e. the Verifier V rejects every state with high probability.

Proposition 5. *If $x \in P_{no}$ then $\|C - C_1\|_\diamond \leq 3\sqrt{\varepsilon}$.*

Proof. This proof is similar to the proof of Proposition 4. If $x \in P_{\text{no}}$, 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 the the circuit on this state, but the state $|\phi'\rangle$ which is very close to it. Once again we apply the monotonicity of the trace norm and Equation (1) to show that $\|(C \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle\langle\psi|) - (C_1 \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle\langle\psi|)\|_{\text{tr}} \leq 3\sqrt{p} \leq 3\sqrt{\varepsilon}$. Since this equation applies for all reference systems $\mathcal{R}$ and all states $|\psi\rangle$, this proves that if $x \in P_{\text{no}}$, then we have $\|C - C_1\|_\diamond \leq 3\sqrt{\varepsilon}$. □

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 on any large subspaces of pure inputs. See the discussion following Problem 3 for a technical condition that is equivalent to this requirement. It is straightforward to verify that the reduction is efficient.

Theorem 6. *$\text{CT}(\varepsilon, \delta, \mathcal{C}_0, \mathcal{C}_1)$ is QMA-hard for any $0 < \varepsilon < 1$ with $\varepsilon \geq 2^{-p}$ 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.*

3.2 Applications

In this section we apply Theorem 6 to prove the hardness of some new and old problems.

The first problem we consider is a slightly generalized version of the problem Non-identity Check [10], who show that it is QMA-complete. Our version of the problem differs in that we do not require that the input circuit C is unitary. We do require, however, that if C deviates from the identity, then it does so in a way similar to some efficient unitary circuit U. This restriction is not needed for hardness but it is not clear that the problem is in QMA without it.

Problem 7 (Mixed Non-identity Check [10]). Let $0 < \varepsilon < 1$. On input $C \in \mathbf{T}(\mathcal{X}, \mathcal{X})$:

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 example are the circuits that apply Pauli X to the first 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 [15].

Problem 8 (Non-isometry [15]). Let $0 < \varepsilon < 1/2$. On input a circuit $C \in \mathbf{T}(\mathcal{X}, \mathcal{Y})$:

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 $|\psi\rangle \in \mathcal{X}$, $\|(\Phi \otimes \mathbb{1}_{\mathcal{X}})(|\psi\rangle\langle\psi|)\|_\infty \geq 1 - \varepsilon$.

Theorem 6 shows the QMA-hardness of this problem, as $\mathrm{CT}(\varepsilon, 1, \Omega, \mathbb{1})$ is a special case, where Ω is the completely depolarizing channel. The norm $\|\cdot\|_\infty$ used in this problem is the operator norm.

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 $C \in \mathbf{T}(\mathcal{X}, \mathcal{X})$:

Yes: There exists $|\psi\rangle \in \mathcal{X}$ such that $\|C(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi|\|_{\mathrm{tr}} \leq \varepsilon$
No: For any $|\psi\rangle \in \mathcal{X}$, $\|C(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi|\|_{\mathrm{tr}} \geq 2 - \varepsilon$

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

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.

Problem 10 (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, $\|(E_k \otimes \mathbb{1}_{\mathcal{R}})(\rho) - \rho\|_{\mathrm{tr}} \leq \varepsilon$.
No: E is an ε-private channel, i.e. $\left\|\Omega - \frac{1}{2^m}\sum_{k\in\{0,1\}^m} E_k\right\|_\diamond \leq \varepsilon$, where Ω is the completely depolarizing channel in $\mathbf{T}(\mathcal{H}, \mathcal{K})$, and there exists a polynomial-size quantum circuit D such that for all k we have $\|D_k \circ E_k - \mathbb{1}_{\mathcal{H}}\|_\diamond \leq \varepsilon$.

For specific values of ε and δ, we refer to this problem as $\mathrm{DI}_{\varepsilon,\delta}$.

Theorem 11. $\mathrm{DI}_{\varepsilon,\delta}$ *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 kth Pauli operator to the input qubits. Averaging over all keys k results in the completely depolarizing channel on n qubits. Then, Theorem 6 implies that $\mathrm{CT}(\varepsilon, \delta, \mathbb{1}_k, \mathcal{E}_k)$ 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 $\mathrm{CT}(\varepsilon, \delta, \mathbb{1}_k, \mathcal{E}_k)$ is a modification of the problem CT to include both a quantum input as well as a classical input k. This is done by including k as part of the quantum input that is immediately measured in the computational basis. $\mathrm{CT}(\varepsilon, \delta, \mathbb{1}_k, \mathcal{E}_k)$ remains hard after this modification.

The QMA-hardness of $\mathrm{DI}_{\varepsilon,\delta}$ then follows from the fact that the problem of detecting insecure encryption is $\mathrm{CT}(\varepsilon, \delta, \mathbb{1}_k, \mathcal{E}_k)$ 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 $\mathrm{CT}(\varepsilon, \delta, \mathbb{1}_k, \mathcal{E}_k)$ and let $C_k(\cdot) = C(|k\rangle\langle k| \otimes \cdot)$ be the circuit defined by hardcoding the '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). The triangle inequality then implies $\|\Omega - \sum_k C_k/2^m\|_\diamond \leq \sum_k \|\Omega_k - C_k\|_\diamond /2^m \leq \varepsilon$, 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

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

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

4.1 QMA Protocol

To test the security of an encryption system in QMA the Verifier needs a tool to compare two quantum states. Such a tool is provided by the swap test, introduced in [6], though here we essentially use it to test the purity of quantum states as is done in [7]. 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, $(\mathbb{1}_{\mathcal{H}\otimes\mathcal{H}} + W)/2$, and the projection onto the antisymmetric subspace, $(\mathbb{1}_{\mathcal{H}\otimes\mathcal{H}} - W)/2$.

Given two pure states $|\psi\rangle, |\phi\rangle$, the swap test returns the symmetric outcome with probability $(1 + |\langle\psi|\phi\rangle|^2)/2$. Applied to mixed states ρ, σ the result is symmetric with probability $(1 + \mathrm{tr}(\rho\sigma))/2$ [7]. This implies that given two copies the swap test can estimate the purity of a state.

The idea behind the protocol is that if the encryption system specified by E is insecure then, regardless of the key, it acts trivially on some subspace of the input. In this case a proof consists of two copies of some pure state in this subspace. The Verifier runs E on both of these states and tests that they have not been changed by with the swap test. In the case that the circuit is insecure,

this proof state will cause the Verifier to obtain the symmetric outcome with probability approaching 1.

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.

Protocol 12. On input a circuit $E\colon \{1,\dots,K\}\otimes\mathbf{D}(\mathcal{H}) \to \mathbf{D}(\mathcal{K})$, an instance of $\mathrm{DI}_{\varepsilon,\delta}$, as well as a quantum proof $|\phi\rangle$ in $\mathbf{D}((\mathcal{H}\otimes\mathcal{R})^{\otimes 2})$ (where $\dim\mathcal{R} = \dim\mathcal{H}$):

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, accepting if the outcome is symmetric.

The space $\mathcal{R}$ appears in this protocol, but Problem 10 places no upper bound on this space, by the properties of the diamond norm, we may take $\dim\mathcal{R} = \dim\mathcal{H}$ without loss of generality.

Proposition 13. *For $0 < \varepsilon < 1/8$, Protocol 12 is a* QMA *protocol for* $\mathrm{DI}_{\varepsilon,\delta}$.

Proof. If E is a 'yes' instance of $\mathrm{DI}_{\varepsilon,\delta}$, 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 $\left\|\hat{E}_k(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi|\right\|_{\mathrm{tr}} \leq \varepsilon$, 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

$$\|\sigma_{k_1}\otimes\sigma_{k_2} - |\psi\rangle\langle\psi|\otimes|\psi\rangle\langle\psi|\|_{\mathrm{tr}} \leq 2\varepsilon,$$

which follows from the triangle inequality. Then, since the state $|\psi\rangle\langle\psi|\otimes|\psi\rangle\langle\psi|$ is symmetric and the swap test performs a projective measurement, Lemma 1 implies 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 $\mathrm{DI}_{\varepsilon,\delta}$ the Verifier does not accept any proof state with high probability. In this case $\left\|\sum_{k=1}^{K} E_k - \Omega\right\|_{\diamond}/K \leq \varepsilon$. Once more, a straightforward argument using the triangle inequality can be used to argue that the tensor product of two copies satisfies the equation $\left\|\sum_{k,j=1}^{K} E_k\otimes E_j - \Omega\otimes\Omega\right\|_{\diamond}/K^2 \leq 2\varepsilon$. 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 $1/2 - \mathrm{tr}[(\mathbb{1}_{\mathcal{K}}/\dim\mathcal{K})^2]/2 + 2\varepsilon = 1/2 - 1/(2\dim\mathcal{K}) + 2\varepsilon$, 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. □

Combining the previous Proposition with Theorem 11 we obtain the main result.

Theorem 14. *For $0 < \varepsilon < 1/8$ and $0 < \delta \leq 1$, the problem* $\mathrm{DI}_{\varepsilon,\delta}$ *is* QMA*-complete.*

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

References

1. Aharonov, D., Kitaev, A., Nisan, N.: Quantum circuits with mixed states. In: Proc. 30th STOC, pp. 20–30 (1998)
2. Ambainis, A., Mosca, M., Tapp, A., de Wolf, R.: Private quantum channels. In: Proc. 41st FOCS, pp. 547–553 (2000)
3. Ambainis, A., Smith, A.: Small Pseudo-random Families of Matrices: Derandomizing Approximate Quantum Encryption. In: Jansen, K., Khanna, S., Rolim, J.D.P., Ron, D. (eds.) RANDOM 2004 and APPROX 2004. LNCS, vol. 3122, pp. 249–260. Springer, Heidelberg (2004)
4. Boykin, P.O., Roychowdhury, V.: Optimal encryption of quantum bits. Phys. Rev. A 67(4), 042317 (2003)
5. Braunstein, S., Lo, H.K., Spiller, T.: Forgetting qubits is hot to do (1999) (unpublished manuscript)
6. Buhrman, H., Cleve, R., Watrous, J., de Wolf, R.: Quantum fingerprinting. Phys. Rev. Lett. 87(16), 167902 (2001)
7. Ekert, A.K., Alves, C.M., Oi, D.K., Horodecki, M., Horodecki, P., Kwek, L.C.: Direct estimations of linear and nonlinear functionals of a quantum state. Phys. Rev. Lett. 88(21), 217901 (2002)
8. Hayden, P., Leung, D., Shor, P.W., Winter, A.: Randomizing quantum states: constructions and applications. Commun. Math. Phys. 250, 371–391 (2004)
9. Helstrom, C.W.: Detection theory and quantum mechanics. Inform. Control 10(3), 254–291 (1967)
10. Janzing, D., Wocjan, P., Beth, T.: "Non-identity-check" is QMA-complete. Int. J. Quantum Inf. 3(3), 463–473 (2005)
11. Kempe, J., Kitaev, A., Regev, O.: The complexity of the local Hamiltonian problem. SIAM J. Comput. 35(5), 1070–1097 (2006)
12. Liu, Y.-K.: Consistency of Local Density Matrices Is QMA-Complete. In: Díaz, J., Jansen, K., Rolim, J.D.P., Zwick, U. (eds.) APPROX 2006 and RANDOM 2006. LNCS, vol. 4110, pp. 438–449. Springer, Heidelberg (2006)
13. Marriott, C., Watrous, J.: Quantum Arthur-Merlin games. Comp. Compl. 14(2), 122–152 (2005)
14. Rosgen, B., Watrous, J.: On the hardness of distinguishing mixed-state quantum computations. In: Proc. 20th CCC, pp. 344–354 (2005)
15. Rosgen, B.: Testing Non-isometry Is QMA-Complete. In: van Dam, W., Kendon, V.M., Severini, S. (eds.) TQC 2010. LNCS, vol. 6519, pp. 63–76. Springer, Heidelberg (2011)
16. Schuch, N., Cirac, I., Verstraete, F.: Computational difficulty of finding matrix product ground states. Phys. Rev. Lett. 100(25), 250501 (2008)
17. Schuch, N., Verstraete, F.: Computational complexity of interacting electrons and fundamental limitations of density functional theory. Nat. Phys. 5(10), 732–735 (2009)
18. Winter, A.: Coding theorem and strong converse for quantum channels. IEEE T. Inform. Theory 45(7), 2481–2485 (1999)