Testing non-isometry is QMA-complete

Bill Rosgen
Centre for Quantum Technologies
National University of Singapore

1 June, 2010

Abstract

Determining the worst-case uncertainty added by a quantum circuit is shown to be computationally intractable. This is the problem of detecting when a quantum channel implemented as a circuit is close to a linear isometry, and it is shown to be complete for the complexity class QMA of verifiable quantum computation. The main idea is to relate the problem of detecting when a channel is close to an isometry to the problem of determining how mixed the output of the channel can be when the input is a pure state.

1 Introduction

A linear isometry $U\colon \mathcal{H} \to \mathcal{K}$ is a linear map that preserves the inner product of any two elements, or equivalently satisfies $U^*U = \mathbb{1}_{\mathcal{H}}$. These transformations are fundamental in quantum computation: they are exactly the maps that may be realized using unitary quantum circuits with access to ancillary qubits in a known pure state—the standard model of quantum computation. It is an important problem to determine when a computation in a non-unitary model, such as measurement based quantum computing or computation in the presence of noise, approximately implements some operation in the unitary circuit model. In this paper it is shown that this problem is QMA-complete when the input computation is modelled as a quantum circuit consisting of the usual unitary gates, plus the ability to discard qubits as well as introduce ancillary qubits. The circuit model is not essential: the hardness result also applies to any model that can efficiently simulate and be simulated by the mixed-state circuit model.

The complexity class QMA is the quantum analogue of NP: the class corresponding to classically verifiable computation. This concept was first considered in [11], first formally defined in [9], and first studied in [18]. QMA is the class of all problems that can be verified with bounded error by a polynomial-time quantum verifier with access to a quantum proof. This proof is given by a quantum state on a polynomial number of qubits and may depend on the input.

The class QMA has complete (promise) problems: problems in QMA that are computationally at least as hard as any other problem in the class. This implies that an efficient algorithm for any of these complete problems can be used to find an efficient algorithm for any problem in QMA. The simplest of these complete problems is the 2-local Hamiltonian problem, which is informally the quantum version of the circuit satisfiability problem for unitary circuits with gates of constant size. A formal description of this problem, as well as a proof that the 5-local Hamiltonian problem is QMA-complete can be found in [10]. The improvement of this result to the 2-local case is due to Kempe, Kitaev, and Regev [8]. Several other complete problems for QMA are known, such as local consistency [12] (see also [13, 19]), some problems related to the minimum output entropy [2], testing whether unitary circuits are close to the identity [6] (see also [7]), and finding the ground states of some physical systems [16, 17]. In the present paper we add a new complete problem to this list: the problem of determining if a quantum circuit implements an operation that is close to an isometry. As discussed in Section 3, this is equivalent to determining if the channel always maps pure states to states that are approximately pure.

The remainder of the paper is organized as follows. Section 2 introduces notation and background. Section 3 introduces the notion of approximate isometries and makes formal the problem of detecting when a channel is an approximate isometry. The QMA-hardness of this problem is proved in Section 4 and proof of the containment in QMA, the most technical portion of the result, appears in Section 5.

2 Preliminaries

In this section the notation and background that is used throughout the paper are presented. Much of the notation used here is standard and this is in no way a complete introduction to quantum information. See [15] for a more detailed treatment of these topics.

All Hilbert spaces considered in this paper are assumed to be finite-dimensional and are denoted by scripted capital letters $\mathcal{H}, \mathcal{K}, \ldots$. The pure states are the unit vectors in these spaces. The set of density matrices or mixed states on $\mathcal{H}$ is given by $\mathbf{D}(\mathcal{H})$, and the set of all quantum channels mapping $\mathbf{D}(\mathcal{H})$ to $\mathbf{D}(\mathcal{K})$ is $\mathbf{T}(\mathcal{H}, \mathcal{K})$. The quantum channels are exactly the completely positive and trace preserving linear maps. The identity channel in $\mathbf{T}(\mathcal{H}, \mathcal{H})$ is denoted $I_{\mathcal{H}}$, while $\mathbb{1}_{\mathcal{H}}$ is the identity on $\mathcal{H}$.

Given a quantum channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ we make use of two representations. The first of these is the Choi representation [4], which provides a unique representation of a channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ as a linear operator on $\mathcal{K} \otimes \mathcal{H}$. This representation is given by $C(\Phi) = (\Phi \otimes I_{\mathcal{H}})(|\phi^+\rangle\langle\phi^+|)$, where $|\phi^+\rangle = \sum_i |ii\rangle / \sqrt{d}$ is a maximally entangled state in $\mathcal{H} \otimes \mathcal{H}$.

The second representation that we use is the representation of a completely positive map Φ by a set of Kraus operators: matrices A_i such that $\Phi(X) = \sum_i A_i X A_i^*$. This representation is also due to Choi [4]. If in addition the map Φ is trace preserving, then the operators A_i satisfy the property $\sum_i A_i^* A_i = \mathbb{1}$. The number of Kraus operators in a minimal Kraus decomposition is given by the rank of the Choi matrix $C(\Phi)$.

In order to measure how close a state is to being pure we use the operator norm $\|X\|_\infty$, which for a linear operator X is the largest singular value of X. When X is normal, this is simply the largest eigenvalue (in absolute value) of X. Dual to the operator norm is the trace norm, which for a linear operator X is given by $\|X\|_{\text{tr}} = \operatorname{tr} \sqrt{X^* X}$. This is exactly the sum of the singular values of X. When X is a quantum state, this simplifies to the sum of absolute values of the eigenvalues of X, so that $\|\rho\|_{\text{tr}} = 1$ for all density matrices ρ.

One final quantity that we use is the fidelity, which for two density matrices ρ, σ is given by $F(\rho, \sigma) = \operatorname{tr} \sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}$. While it is not obvious from this definition, the fidelity is symmetric in the two arguments. When one of the arguments is a pure state, the fidelity simplifies to $F(\rho, |\psi\rangle\langle\psi|) = \sqrt{\langle\psi|\rho|\psi\rangle}$. An important relationship between the trace norm and the fidelity is

$$2 - 2\,F(\rho, |\psi\rangle\langle\psi|)^2 \leqslant \|\rho - |\psi\rangle\langle\psi|\|_{\text{tr}} \tag{1}$$

that we use to relate different notions of the purity of a quantum state. This inequality can be found in [15, Chapter 9].

We require one final piece of background. In order for a quantum channel to be given as input to a computational problem we need a representation of the channel. Using either the Choi matrix or Kraus operators produces a representation that, in the case of channels implementing efficient quantum algorithms, is exponentially larger than the size of a circuit representation. These channels have circuit representations that are logarithmic in the Hilbert space dimension. For this reason, we use a circuit representation of quantum channels. Such a representation is provided by the mixed-state circuit model of Aharonov et al. [1], which is simply the usual model of unitary quantum circuits with two additional gates. These gates are the gate that introduces ancillary qubits in the $|0\rangle$ state and the gate the traces out (i.e. discards) a qubit. This circuit model can be used to represent any quantum channel, which makes it ideal for the problem that we consider.

3 Isometries and rank non-increasing channels

One important property of the linear isometries is that they do not increase rank. This is essential to the QMA protocol in Section 5, which is able to detect exactly those channels that are rank-increasing. More formally, a channel Φ is *rank non-increasing* if for all states ρ the output of Φ satisfies $\operatorname{rank}(\rho) \geqslant \operatorname{rank}(\Phi(\rho))$. Unfortunately, this property does not characterize the isometries. Consider the channel $\Phi(\rho) = |0\rangle\langle 0|$ that discards the input state and returns a fixed pure state. This channel is not an isometry but it is also rank non-increasing.

This property can be used to characterize the isometries if we make a small adjustment. The channels that are rank non-increasing when adjoined to an auxiliary space of arbitrary dimension are exactly the isometries. We call a channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ *completely* rank non-increasing if for any $\mathcal{F}$ the channel $\Phi \otimes I_{\mathcal{F}}$ is rank non-increasing, i.e. if $\operatorname{rank}[(\Phi \otimes I_{\mathcal{F}})(\rho)] \leqslant \operatorname{rank}(\rho)$ for all ρ. The channel $\Phi(\rho) = |0\rangle\langle 0|$ is not completely rank non-increasing: consider applying it to half of a maximally entangled state $(\Phi \otimes I_{\mathcal{H}})(|\phi^+\rangle\langle\phi^+|) = |0\rangle\langle 0| \otimes \mathbb{1}_{\mathcal{H}}/\dim\mathcal{H}$. As in the case of complete positivity, we need only to verify this property on an auxiliary space of the same dimension as the input space. It is also easy to see that this property characterizes the linear isometries.

Proposition 1. *The following are equivalent for a channel* $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$*:*

1. $\Phi(\rho) = U\rho U^*$ *for some linear isometry* U *from* $\mathcal{H}$ *to* $\mathcal{K}$*,*
2. Φ *is completely rank non-increasing,*
3. $\Phi \otimes I_{\mathcal{H}}$ *is rank non-increasing.*

Proof. The first two implications are immediate. To prove that (3) $\Rightarrow$ (1), let $\Phi \otimes I_{\mathcal{H}}$ be rank non-increasing. This implies that $\operatorname{rank}(C(\Phi)) = 1$. Recalling that the number of Kraus operators in a minimal decomposition is $\operatorname{rank}(C(\Phi))$, it follows that Φ can be expressed as $\Phi(\rho) = A\rho A^*$. The condition that Φ is trace preserving implies that the operator A satisfies $A^*A = \mathbb{1}_{\mathcal{H}}$. □ □

This characterization guides the remainder of the paper. Detecting when the channel $\Phi \otimes I_{\mathcal{H}}$ increases rank provides an operational method to determine when a channel is an isometry.

3.1 Approximately pure states

In order to show that non-isometry detection is QMA-complete we need to consider an approximate version of the problem. This is because a protocol for a QMA language is permitted to fail with small probability. The definition of approximate isometries used here is closely related to the notion of approximately pure states. Several equivalent notions of the purity of a density matrix are considered in this section.

Perhaps the most well-known notion of how close a mixed state ρ is to being pure is the *purity* of ρ, given by $\mathrm{tr}(\rho^2)$. A similar measure is given by $\|\rho\|_\infty$, the largest eigenvalue of ρ. It is not hard to see that these quantities are related. If $\rho = \sum_i \lambda_i |\lambda_i\rangle\langle\lambda_i|$ is the spectral decomposition of ρ, with the eigenvalues λ_i in decreasing order, then $\mathrm{tr}\,\rho^2 = \sum_i \lambda_i^2 \geqslant \lambda_1^2 = \|\rho\|_\infty^2$. In the other direction, since the purity is convex, it is maximized for $1/\lambda_1$ eigenvalues each of value λ_1, i.e. $\mathrm{tr}\,\rho^2 = \sum_i \lambda_i^2 \leqslant \lambda_1^2/\lambda_1 = \|\rho\|_\infty$. Taken together, these two inequalities show that

$$\|\rho\|_\infty^2 \leqslant \mathrm{tr}(\rho^2) \leqslant \|\rho\|_\infty . \tag{2}$$

These quantities are also related to the more familiar trace distance on quantum states.

Proposition 2. *Let $\rho \in \mathbf{D}(\mathcal{H})$ and let $\varepsilon > 0$. There exists a pure state $|\psi\rangle \in \mathcal{H}$ such that $\|\rho - |\psi\rangle\langle\psi|\|_{\mathrm{tr}} \leqslant \varepsilon$ if and only if $\|\rho\|_\infty \geqslant 1 - \varepsilon/2$.*

Proof. Let ρ have spectral decomposition given by $\rho = \sum_i \lambda_i |\lambda_i\rangle\langle\lambda_i|$, with $\lambda_1 \geqslant \lambda_2 \geqslant \ldots \geqslant \lambda_d$. If $\|\rho\|_\infty = \lambda_1 \geqslant 1 - \varepsilon/2$, then

$$\|\rho - |\lambda_1\rangle\langle\lambda_1|\|_{\mathrm{tr}} = (1-\lambda_1) + \sum_{i=2}^{d} \lambda_i = 2(1-\lambda_1) \leqslant 2(\varepsilon/2) = \varepsilon.$$

On the other hand, if $|\psi\rangle \in \mathcal{H}$ is a state such that $\|\rho - |\psi\rangle\langle\psi|\|_{\mathrm{tr}} \leqslant \varepsilon$, then by Equation (1)

$$\varepsilon \geqslant \|\rho - |\psi\rangle\langle\psi|\|_{\mathrm{tr}} \geqslant 2 - 2\,\mathrm{F}(\rho, |\psi\rangle\langle\psi|)^2 = 2 - 2\langle\psi|\rho|\psi\rangle = 2 - 2\sum_i \lambda_i \, |\langle\psi|\lambda_i\rangle|^2 . \tag{3}$$

The final quantity is a convex combination of the λ_i, with weights determined by the state $|\psi\rangle$. This is maximized when $|\psi\rangle = |\lambda_1\rangle$, since λ_1 is the largest eigenvalue of ρ. Combining this with Equation (3) we have $\varepsilon \geqslant 2 - 2\lambda_1 = 2 - 2\|\rho\|_\infty$, which implies that $\|\rho\|_\infty \geqslant 1 - \varepsilon/2$. □ □

Given these notions of purity, we will call a state *ε-pure* if $\|\rho\|_\infty \geqslant 1 - \varepsilon$. By the previous results the purity of such a state satisfies $\mathrm{tr}(\rho^2) \geqslant (1-\varepsilon)^2 \geqslant 1 - 2\varepsilon$, and there is a pure state $|\psi\rangle$ such that $\|\rho - |\psi\rangle\langle\psi|\|_{\mathrm{tr}} \leqslant 2\varepsilon$. For the results of this paper, any of these three measures suffices, as they are equivalent up to polynomial factors in ε.

3.2 Approximate isometries

The focus of this paper is to show that detecting when a channel is far from an isometry is computationally difficult. To do this we need to define the class of channels that are the approximate isometries. Isometries always map pure states to pure states, even in the presence of a reference system. Proposition 1 shows that this condition characterizes the isometries. Weakening this requirement, we call a channel an *ε-isometry* if it maps pure states (over the input space and a reference system) to states that are ε-pure, for some $\varepsilon > 0$.

More formally a channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ is an ε-isometry if for any pure state $|\psi\rangle \in \mathcal{H} \otimes \mathcal{H}$ the output of $\Phi \otimes \mathrm{I}_{\mathcal{H}}$ satisfies $\|(\Phi \otimes \mathrm{I}_{\mathcal{H}})(|\psi\rangle\langle\psi|)\|_\infty \geqslant 1-\varepsilon$, i.e. when applied to part of any pure state the output state is close to pure. This implies that $\Phi \otimes \mathrm{I}_{\mathcal{H}}$ does not reduce the operator norm of any input by more than a factor of $1-\varepsilon$. We use this to define the computational problem that is the main focus of the paper.

Problem 3 (Non-isometry)**.** For $0 \leqslant \varepsilon < 1/2$ and a channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$, given as a mixed-state quantum circuit, the promise problem is to decide between:

Yes: There exists a pure state $|\psi\rangle \in \mathcal{H}$ such that $\|(\Phi \otimes \mathrm{I}_{\mathcal{H}})(|\psi\rangle\langle\psi|)\|_\infty \leqslant \varepsilon$,

No: For all pure states $|\psi\rangle \in \mathcal{H}$, $\|(\Phi \otimes \mathrm{I}_{\mathcal{H}})(|\psi\rangle\langle\psi|)\|_\infty \geqslant 1-\varepsilon$.

When the value of ε is significant, we will refer to this problem as Non-isometry$_\varepsilon$.

Using the equivalence results of Equation (2) and Proposition 2, this problem may be equivalently defined in terms of either the purity or the trace distance to the closest pure state, up to a small increase in ε. The case of the minimum output purity of a channel has been studied in a different context by Zanardi and Lidar [20], though they focus on finding the minimum purity of a channel over a subspace of the inputs. The problem we consider here is equivalent to evaluating the channel purity of $\Phi \otimes \mathrm{I}_{\mathcal{H}}$ over the whole input space.

The difficulty of the Non-isometry problem does not change if the dimension of the ancillary system is permitted to be larger than the size of the input system, so long as the number of qubits needed to represent the ancillary system is polynomial in the number of input qubits.

The notion of approximate isometry that we consider here is *not* equivalent to the channel being completely rank non-increasing on average. This property is modelled by the distance between the Choi matrix of a channel and a pure state. While it is true that the Choi matrix is pure if and only if the channel is an isometry, it is close to pure in the trace distance when the channel is close to an isometry *on average*. In this paper we consider the worst-case, i.e. we consider a channel to be close to an isometry if and only if the output of $\Phi \otimes \mathrm{I}_{\mathcal{H}}$ is close to pure for *any* pure state input. A simplification of the protocol presented in Section 5 yields a polynomial-time quantum algorithm for the problem of determining how close the Choi matrix of a channel is to a pure state. This is because $C(\Phi)$ can be generated efficiently, and given two copies the swap test can be used to test the purity of a quantum state as shown in [5].

4 QMA hardness

In order to prove the hardness of Non-isometry we modify an arbitrary QMA protocol to obtain a circuit that can output a mixed state exactly when the verifier would have accepted in the original protocol. This yields a circuit that is far from an isometry if and only if there is a witness that causes the verifier in the original protocol to accept. Deciding whether or not there is such a witness is QMA-hard, by the definition of the complexity class. More formally, 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 $\Pr[V \text{ accepts } \rho] \geqslant 1-\varepsilon$,
2. if $x \notin L$, then for any state ρ, $\Pr[V \text{ accepts } \rho] \leqslant \varepsilon$,

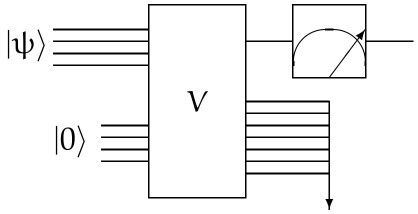

Figure 1: Verifier's circuit in a QMA protocol. The verifier accepts the witness state $|\psi\rangle$ if and only if the measurement in the computational basis results in the $|1\rangle$ state.

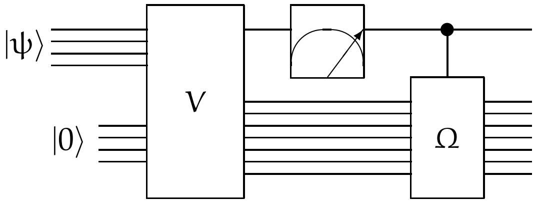

Figure 2: Constructed instance of Non-isometry. The output state is mixed by the completely depolarizing channel Ω only if the state $|\psi\rangle$ is a valid witness to the original QMA protocol.

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 [10, 14].

Let L be an arbitrary language in QMA, and let x be an arbitrary input string. The goal is to embed the QMA-hard problem of deciding if $x \in L$ into the problem of testing whether a mixed-state quantum circuit is close to an isometry. Let V be the isometry representing the algorithm of the verifier in a QMA protocol for L on input x. We may "hard-code" the input string x into V because the circuit needs only to be efficiently generated from x. The algorithm implemented by the verifier is shown in Figure 1. The verifier first receives a witness state $|\psi\rangle$, applies the isometry V, and then makes a measurement on one of the qubits, the result of which determines whether or not the verifier accepts. Any qubits not measured are traced out.

For concreteness, let V act on the input spaces $\mathcal{W}$ and $\mathcal{A}$, which hold the witness state and the $|0\rangle$ state of the ancilla respectively. Let $\mathcal{M}$ be the space corresponding to the measured output qubit in the protocol and let $\mathcal{G}$ represent the 'garbage' qubits that are traced out at the end of the protocol. The probability that verifier accepts the witness state $|\psi\rangle \in \mathcal{W}$ is

$$\Pr[V \text{ accepts } |\psi\rangle] = \langle 1|\, \mathrm{tr}_{\mathcal{G}} \left[V(|\psi\rangle\langle\psi| \otimes |0\rangle\langle 0|)V^* \right] |1\rangle. \tag{4}$$

Deciding if there is some $|\psi\rangle$ such that this expectation is close to one is complete for QMA.

From Figure 1 it is simple to construct a circuit that produces highly mixed output exactly when there exists such a $|\psi\rangle$. The idea is add a controlled application of the completely depolarizing channel Ω on the space $\mathcal{G}$, instead of tracing it out. The resulting circuit is shown in Figure 2. In the case that the verifier accepts with negligible probability for every input state $|\psi\rangle$, then both the measurement and the controlled depolarizing channel have little effect, leaving the state of the system close to a pure state. If, on the other hand, there is a state on which the verifier accepts with high probability, then on this input the circuit in Figure 2 produces a highly mixed state.

Formalizing this notion proves that Non-isometry is QMA-hard.

Theorem 4. *Let $\varepsilon > 0$ be a constant, and let p be the maximum acceptance probability of the protocol V. Let $\Phi \in \mathbf{T}(\mathcal{W}, \mathcal{M} \otimes \mathcal{G})$ be the circuit in Figure 2. Then if $\dim \mathcal{R} = \dim \mathcal{W}$*

$$p \leqslant \varepsilon \implies \min_{|\psi\rangle} \| (\Phi \otimes I_{\mathcal{R}})(|\psi\rangle\langle\psi|) \|_{\infty} \geqslant 1 - \varepsilon,$$
$$p \geqslant 1 - \varepsilon \implies \min_{|\psi\rangle} \| (\Phi \otimes I_{\mathcal{R}})(|\psi\rangle\langle\psi|) \|_{\infty} \leqslant \varepsilon.$$

Proof. Notice that we may assume that the output dimension of Φ is $\dim \mathcal{M} \otimes \mathcal{G} = 2d > 2/\varepsilon$ by padding the circuit for V with $\log 1/\varepsilon$ unused ancillary qubits, if necessary.

As a first step, we evaluate the output state of the channel $\Phi \otimes I_{\mathcal{R}}$. Applied to a pure state $|\psi\rangle \in \mathcal{W} \otimes \mathcal{R}$ this channel first adds the ancillary $|0\rangle$ qubits in the space $\mathcal{A}$ and then applies the isometry V from the QMA protocol. This results in the pure state $|\phi\rangle = (V \otimes \mathbb{1}_{\mathcal{R}})(|\psi\rangle \otimes |0\rangle)$. We may decompose this state in terms of the qubit in the space $\mathcal{M}$, obtaining for some $0 \leqslant p \leqslant 1$

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

The value of p is exactly the probability that the measurement result is $|1\rangle$, i.e. the probability that the verifier will accept the input state $\operatorname{tr}_{\mathcal{R}} |\psi\rangle\langle\psi|$ in the original protocol. Using this, the state after the measurement and the controlled depolarizing channel on $\mathcal{G}$ is

$$(1-p)|0\rangle\langle 0| \otimes |\phi_0\rangle\langle\phi_0| + (p/d)|1\rangle\langle 1| \otimes \mathbb{1}_{\mathcal{G}} \otimes \rho, \tag{5}$$

where ρ is the residual state on $\mathcal{R}$ after this channel has been applied ($\rho = \operatorname{tr}_{\mathcal{G}} |\phi_1\rangle\langle\phi_1|$, but this will not be important). Evaluating the largest eigenvalue of this state we find that

$$\| (\Phi \otimes I_{\mathcal{R}})(|\phi\rangle\langle\phi|) \|_{\infty} = \max\{1-p, \frac{p}{d} \| \rho \|_{\infty}\}. \tag{6}$$

We analyze the maximum in Equation (6) in two cases. The first of these cases is when there is no input the verifier accepts with probability larger than ε. In this case the output of the channel $\Phi \otimes I_{\mathcal{R}}$ is given by Equation (5) where $p \leqslant \varepsilon$. Here Equation (6) shows that the output has an eigenvalue of magnitude at least $\min_{|\mu\rangle} \| (\Phi \otimes I_{\mathcal{R}})(|\mu\rangle\langle\mu|) \|_{\infty} \geqslant 1 - p \geqslant 1 - \varepsilon$.

The second case is when there exists a state $|\psi\rangle$ that verifier to accepts with probability at least $1 - \varepsilon$. In this case we take the input state to $\Phi \otimes I_{\mathcal{R}}$ to be $|\gamma\rangle = |\psi\rangle \otimes |0\rangle$, i.e. we set the reference system to be any pure state that is not entangled with the rest of the input. The output is given by Equation (5) with $p \geqslant 1 - \varepsilon$ and $\rho = |0\rangle\langle 0|$. Equation (6) yields

$$\min_{|\mu\rangle} \| (\Phi \otimes I_{\mathcal{R}})(|\mu\rangle\langle\mu|) \|_{\infty} \leqslant \| (\Phi \otimes I_{\mathcal{R}})(|\gamma\rangle\langle\gamma|) \|_{\infty} = \max\left\{1-p, \frac{p}{d} \| \rho \|_{\infty}\right\} \leqslant \max\left\{\varepsilon, \frac{1}{d}\right\} = \varepsilon,$$

as we have taken $1/d < \varepsilon$ (by adding $O(\log 1/\varepsilon)$ unused ancillary qubits if necessary). □ □

This theorem shows that determining how far the output $\Phi \otimes I_{\mathcal{R}}$ is from a pure state is as computationally difficult as determining whether or not the verifier can be made to accept with high probability in a QMA protocol. Since the construction of the circuit shown in Figure 2 can be performed efficiently, this implies the hardness of this problem.

Corollary 5. *For any constant $0 \leqslant \varepsilon < 1/2$,* Non-isometry *is* QMA-*hard.*

Using the equivalences between notions of purity in of Section 3.1, this also implies that evaluating the purity of a quantum channel, as defined by Zanardi and Lidar [20] is QMA-hard.

5 QMA protocol

In order to show that Non-isometry is QMA-complete, it remains only to construct a QMA protocol for the problem. The key idea behind this protocol is that when two copies of a channel Φ are applied in parallel to the input state $|\psi\rangle \otimes |\psi\rangle$ the output lies in the antisymmetric subspace if and only if $\Phi(|\psi\rangle\langle\psi|)$ is a mixed state. This provides a probabilistic test that can detect when a channel is far from an isometry.

Unfortunately, in a QMA protocol the verifier cannot assume the witness is given by two non-entangled pure states. It suffices, however, for the verifier to require that the input state lies in the symmetric subspace of the input space $(\mathcal{H} \otimes \mathcal{R})^{\otimes 2}$. To show that the channel is not an isometry in QMA, the prover can provide a symmetric state that a parallel application of the channel maps into the antisymmetric space of the output space $(\mathcal{K} \otimes \mathcal{R})^{\otimes 2}$.

The verifier in such a protocol needs a test to determine when a state is symmetric or antisymmetric. Such a test is provided by the swap test, which was introduced in the context of communication complexity in [3], though we make use of it to test purity using an idea from [5].

The swap test can be characterized as the projection onto the symmetric and antisymmetric subspaces of a bipartite space. If W is the swap operation on a space $\mathcal{H} \otimes \mathcal{H}$, then the symmetric measurement outcome of the swap test corresponds to the the projector $(\mathbb{1}_{\mathcal{H}\otimes\mathcal{H}} + W)/2$, and the projector $(\mathbb{1}_{\mathcal{H}\otimes\mathcal{H}} - W)/2$ corresponds to the antisymmetric outcome.

The main idea behind the protocol for Non-isometry is that the swap test can be used to measure the purity of a state. As observed in [5], when applied two to copies of a state $\rho = \sum_i \lambda_i |\psi_i\rangle\langle\psi_i|$ the swap test returns the antisymmetric outcome with probability

$$\frac{1}{2}\operatorname{tr}((\mathbb{1} - W)(\rho \otimes \rho)) = \frac{1}{2} - \frac{1}{2}\sum\nolimits_i \lambda_i^2 = \frac{1}{2} - \frac{1}{2}\operatorname{tr}(\rho^2). \tag{7}$$

This implies that the swap test on two copies of a state can be used to test purity and, by extension, when a channel is far from an isometry.

A straightforward protocol for Non-isometry on a channel Φ is then to receive a witness state $|\psi\rangle \otimes |\psi\rangle$, apply the channel to obtain $[(\Phi \otimes I)(|\psi\rangle\langle\psi|)]^{\otimes 2}$, and finally apply the swap test. The result is the antisymmetric outcome with high probability only when the state $(\Phi \otimes I)(|\psi\rangle\langle\psi|)$ is highly mixed. Such a protocol detects the channels that are far from isometries.

Unfortunately, the verifier in a QMA protocol cannot assume that the witness state is of the form $|\psi\rangle \otimes |\psi\rangle$. The verifier *can* check that he has received some state in the symmetric subspace and then use the fact that this subspace is closed under the parallel application of a rank non-increasing channel. The verifier in the following protocol uses the swap test to both check the symmetry of the input state and the antisymmetry of the output state.

Protocol 6 (Non-isometry)**.** On an input channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$:

1. Receive a witness state $\rho \in \mathbf{D}((\mathcal{H} \otimes \mathcal{R})^{\otimes 2})$, where $\mathcal{R}$ is a reference space such that $\dim \mathcal{R} = \dim \mathcal{H}$. Apply the swap test to ρ, rejecting if the outcome is antisymmetric.

2. Use the channel Φ to obtain $\sigma = (\Phi \otimes I_{\mathcal{R}})^{\otimes 2}(\rho)$.

3. Apply the swap test to σ, accepting if the outcome is symmetric and rejecting otherwise.

A diagram of this protocol can be found in Figure 3. The correctness of this protocol is argued in the following theorem.

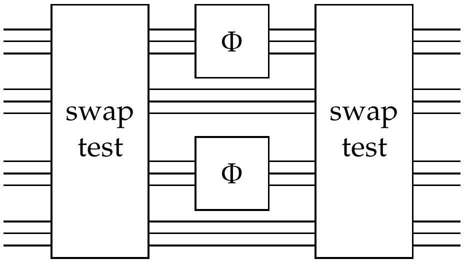

Figure 3: QMA protocol for Non-isometry. The verifier accepts only if the first swap test results the symmetric outcome and the second swap test results in an antisymmetric outcome.

Theorem 7. *Let $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$, and let $p(\rho)$ be the probability that the verifier described in Protocol 6 accepts the input state $\rho \in \mathbf{D}((\mathcal{H} \otimes \mathcal{R})^{\otimes 2})$, then*

1. *If $\min_{|\psi\rangle} \|(\Phi \otimes I_{\mathcal{R}})(|\psi\rangle\langle\psi|)\|_\infty \leqslant \varepsilon$, then there exists a witness ρ such that $p(\rho) \geqslant (1-\varepsilon)/2$.*
2. *If $\min_{|\psi\rangle} \|(\Phi \otimes I_{\mathcal{R}})(|\psi\rangle\langle\psi|)\|_\infty \geqslant 1-\varepsilon$, then for any witness ρ, $p(\rho) \leqslant 9\varepsilon$.*

Proof. For the sake of brevity, let $\hat{\Phi} = \Phi \otimes I_{\mathcal{R}}$ throughout. To prove the first assertion, let $|\psi\rangle$ be a pure state in $\mathcal{H} \otimes \mathcal{R}$ for which $\|\hat{\Phi}(|\psi\rangle\langle\psi|)\|_\infty \leqslant \varepsilon$, and let the witness state $\rho = |\psi\rangle\langle\psi| \otimes |\psi\rangle\langle\psi|$. This state is invariant under the swap operation and so the swap test in Step 1 passes and does not change the state. Step 2 results in the state $[\hat{\Phi}(|\psi\rangle\langle\psi|)]^{\otimes 2}$. Using Equations (2) and (7), the final swap test returns the antisymmetric outcome with probability

$$\frac{1}{2} - \frac{1}{2} \operatorname{tr}\left[\hat{\Phi}(|\psi\rangle\langle\psi|)^2\right] \geqslant \frac{1}{2} - \frac{1}{2}\left\|\hat{\Phi}(|\psi\rangle\langle\psi|)\right\|_\infty \geqslant \frac{1-\varepsilon}{2},$$

and so the verifier accepts ρ with probability approaching one-half for small ε.

To show the second assertion, we take $\hat{\Phi}$ is an ε-isometry and analyze the probability that the verifier can be made to accept. We may assume that the witness state lies in the symmetric subspace of $(\mathcal{H} \otimes \mathcal{R})^{\otimes 2}$, as the verifier either rejects in Step 1 or projects the witness onto this subspace. To complete the proof, we show that $(\hat{\Phi})^{\otimes 2}$ leaves ρ approximately symmetric.

To do this, we approximate $\hat{\Phi}$ by an operator that preserves the symmetry of input states. Let $\{|i\rangle : 1 \leqslant i \leqslant \dim \mathcal{H}\}$ be an orthonormal basis for the spaces $\mathcal{H}, \mathcal{R}$ (this is possible because they have the same dimension). The states $\{|ij\rangle : 1 \leqslant i, j \leqslant \dim \mathcal{H}\}$ are an orthonormal basis for $\mathcal{H} \otimes \mathcal{R}$. Since $\hat{\Phi}$ approximately preserves rank, there are states $|\psi_i\rangle \in \mathcal{K}$ such that

$$\|(\Phi \otimes I_{\mathcal{R}})(|ij\rangle\langle ij|) - |\psi_i\rangle\langle\psi_i| \otimes |j\rangle\langle j|\|_{\mathrm{tr}} \leqslant \varepsilon \tag{8}$$

for all i and j. We define a linear operator $A\colon \mathcal{H} \to \mathcal{K}$ by the equation $A|i\rangle = c_i|\psi_i\rangle$, where the $c_i \in \mathbb{C}$ with $|c_i| = 1$. The introduction of the phases c_i is necessary because Equation (8) only defines the states $|\psi_i\rangle$ up to a phase. Note that the operator A is not necessarily unitary as we may not assume that the states $|\psi_i\rangle$ are orthogonal. The next step is to show that, for some choice of the phases c_i, conjugation by A approximates the channel Φ in the trace norm. This is the most technical portion of the proof.

Consider the output of $\hat{\Phi}$ on the entangled state $(|ii\rangle + |jj\rangle)/\sqrt{2}$ in $\mathcal{H} \otimes \mathcal{R}$, given by

$$\rho = \frac{1}{2} \sum_{a,b \in \{i,j\}} \hat{\Phi}(|aa\rangle\langle bb|) = \frac{1}{2} \sum_{a,b \in \{i,j\}} \Phi(|a\rangle\langle b|) \otimes |a\rangle\langle b|. \tag{9}$$

Since $\hat{\Phi}$ maps pure states to states that are nearly pure, we know that the purity of ρ satisfies $\operatorname{tr}(\rho^2) \geqslant (1-\varepsilon)^2 \geqslant 1 - 2\varepsilon$. Evaluating the purity using Equation (9) gives

$$\begin{aligned} 1 - 2\varepsilon \leqslant \operatorname{tr}(\rho^2) &= \frac{1}{4}\left(\operatorname{tr}\Phi(|i\rangle\langle i|)^2 + \operatorname{tr}\Phi(|j\rangle\langle j|)^2 + 2\operatorname{tr}\Phi(|i\rangle\langle j|)\Phi(|j\rangle\langle i|)\right) \\ &\leqslant \frac{1}{2} + \frac{1}{2}\operatorname{tr}\left((\Phi(|i\rangle\langle j|)\Phi(|i\rangle\langle j|)^*\right). \end{aligned} \tag{10}$$

Interpreting the expression $\operatorname{tr} XX^*$ as the sum of the squared singular values of X, Equation (10) implies that the operator $\Phi(|i\rangle\langle j|)$ has largest singular value at least $1 - 4\varepsilon$. Since the sum of the singular values of this operator cannot exceed one (as the trace norm does not increase under the application of a channel), this implies that it can be decomposed as

$$\Phi(|i\rangle\langle j|) = (1 - 4\varepsilon)|\phi_i\rangle\langle\phi_j| + 4\varepsilon Y, \tag{11}$$

where $|\phi_i\rangle, |\phi_j\rangle \in \mathcal{K}$ are pure and Y is a linear operator on $\mathcal{K}$ with $\|Y\|_{\mathrm{tr}} = 1$. It remains to show that the vectors $|\phi_i\rangle$ and $|\phi_j\rangle$ are, up to a phase, approximately equal to the vectors $|\psi_i\rangle$ and $|\psi_j\rangle$ defined in Equation (8). To do this, we consider the action of Φ on $(|i\rangle + |j\rangle)/\sqrt{2}$. Since Φ is an ε-isometry, the output of Φ on this state is within trace distance 2ε of some pure state $|\gamma\rangle$. Combining Equations (8) and (11) and applying the triangle inequality yields

$$\left\| |\gamma\rangle\langle\gamma| - \frac{1}{2}\left(|\psi_i\rangle\langle\psi_i| + |\phi_i\rangle\langle\phi_j| + |\phi_j\rangle\langle\phi_i| + |\psi_j\rangle\langle\psi_j|\right)\right\|_{\mathrm{tr}} \leqslant 5\varepsilon.$$

Since $|\gamma\rangle$ is pure, for some phases c_i and c_j we have $\||\phi_i\rangle\langle\phi_j| - c_i c_j^* |\psi_i\rangle\langle\psi_j|\|_{\mathrm{tr}} \leqslant 5\varepsilon$, which in turn implies that $\|\Phi(|i\rangle\langle j|) - c_i c_j^* |\psi_i\rangle\langle\psi_j|\|_{\mathrm{tr}} \leqslant 9\varepsilon$, using Equation (11). Finally, since this is true for any $i \neq j$, and the case of $i = j$ is Equation (8), the previous equation implies that $\max_\rho \|\Phi(\rho) - A\rho A^*\|_{\mathrm{tr}} \leqslant 9\varepsilon$, where A is the operator defined by $A|i\rangle = c_i|\psi_i\rangle$ for all i.

It remains only to show that the operator $A \otimes A$ preserves symmetric states. To see this, take $|ij\rangle + |ji\rangle$ an arbitrary basis element of the symmetric subspace of $\mathcal{H}^{\otimes 2}$. By a simple calculation

$$(A \otimes A)(|ij\rangle + |ji\rangle) = c_i c_j |\psi_i\rangle \otimes |\psi_j\rangle + c_i c_j |\psi_j\rangle \otimes |\psi_i\rangle,$$

which remains invariant under swapping the two spaces. By linearity, conjugation by $A \otimes \mathbb{1}_{\mathcal{R}}$ also preserves the symmetry of states on $(\mathcal{H} \otimes \mathcal{R})^{\otimes 2}$. It follows that $\hat{\Phi}$ preserves symmetry up to an error of 9ε in the trace distance. This implies that the swap test on the output of $\hat{\Phi} \otimes \hat{\Phi}$ applied to a symmetric state returns the symmetric outcome with probability at least $1 - 9\varepsilon$. □ □

This theorem shows that Non-isometry$_\varepsilon$ is in QMA for any constant ε satisfying $(1-\varepsilon)/2 > 9\varepsilon$. Together with the QMA-hardness of the problem shown in Theorem 4 this gives the main result.

Corollary 8. *For any constant $\varepsilon < 1/19$,* Non-isometry$_\varepsilon$ *is* QMA*-complete.*

This also implies that problem of computing the channel purity, as defined by Zanardi and Lidar [20], over the whole input space is QMA-complete.

6 Conclusion

We have shown the computational intractability of the problem of detecting when a quantum channel is far from an isometry, or equivalently, when a channel can be made to output a highly mixed state. These results show that it is extremely difficult to characterize the worst-case behaviour of a quantum computation. This is similar to the classical case, where the problem of determining if a circuit can produce a specific output is known to be intractable.

We have also added to the short but growing list of problems that are known to be complete for the complexity class QMA. The Non-isometry problem provides a new way to study this class, as it exactly characterizes the difficulty of the problems in the class. It is hoped that this will lead to new results about the power of this model of computation.

Acknowledgements

I am grateful for discussions with Markus Grassl, Masahito Hayashi, Lana Sheridan, and John Watrous, from which I have learnt a great deal. 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, as well as as well the Bell Family Fund, while the author was at the Institute for Quantum Computing at the University of Waterloo.

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] 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].

[3] 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.

[4] M.-D. Choi. Completely positive linear maps on complex matrices. *Linear Algebra and its Applications*, 10(3):285–290, 1975. doi: 10.1016/0024-3795(75)90075-0.

[5] 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.

[6] 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.

[7] Z. Ji and X. Wu. Non-identity check remains QMA-complete for short circuits, 2009. eprint: arXiv:0906.5416 [quant-ph].

[8] 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.

[9] A. Y. Kitaev. Quantum NP. Talk at the 2nd Workshop on Algorithms in Quantum Information Processing (AQIP), DePaul University, 1999.

[10] 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.

[11] E. Knill. Quantum randomness and nondeterminism. Techical Report LAUR-96-2186, Los Alamos National Laboratory, 1996. EPRINT: arXiv:quant-ph/9610012.

[12] Y.-K. Liu. Consistency of local density matrices is QMA-complete. In *Proceedings of the 10th International Workshop on Randomized Techniques in Computation*, volume 4110 of *Lecture Notes in Computer Science*, pp. 438–449. Springer, 2006. DOI: 10.1007/11830924_40. EPRINT: arXiv:quant-ph/0604166.

[13] Y.-K. Liu, M. Christandl, and F. Verstraete. Quantum computational complexity of the N-representability problem: QMA complete. *Physical Review Letters*, 98(11):110503, 2007. DOI: 10.1103/PhysRevLett.98.110503. EPRINT: arXiv:quant-ph/0609125.

[14] 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.

[15] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[16] 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].

[17] 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].

[18] J. Watrous. Succinct quantum proofs for properties of finite groups. *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science*, pp. 537 – 546, 2000. DOI: 10.1109/SFCS.2000.892141. EPRINT: arXiv:cs/0009002.

[19] T.-C. Wei, M. Mosca, and A. Nayak. Interacting boson problems can be QMA hard. *Physical Review Letters*, 104(4):040501, 2010. DOI: 10.1103/PhysRevLett.104.040501. EPRINT: arXiv:0905.3413 [quant-ph].

[20] P. Zanardi and D. A. Lidar. Purity and state fidelity of quantum channels. *Physical Review A*, 70(1):012315, 2004. DOI: 10.1103/PhysRevA.70.012315. EPRINT: arXiv:quant-ph/0403074.