Computational distinguishability of degradable and antidegradable channels

Bill Rosgen
Centre for Quantum Technologies
National University of Singapore

November 11, 2009

Abstract

A channel is degradable if there exists a second channel that maps the output state of the channel to the environment state. These channels satisfy the property that the output state contains more information about the input than the environment does. A complementary class of channels is the antidegradable channels, which admit channels that map the environment state to the output state of the channel. In this paper we show that the computational problem of distinguishing two channels remains PSPACE-complete when restricted to these classes of channels. This is shown using a construction of Cubitt, Ruskai, and Smith [4] that embeds any channel into a degradable channel, and a related construction for the case of antidegradable channels.

1 Introduction

The task of distinguishing two quantum channels is one of the most fundamental problems in quantum information. This problem, a weak form of process tomography, asks not to completely characterize an unknown quantum channel, but to identify it as one of two known channels. An equivalent formulation asks if there is an input state on which two known channels produce distinct output states. When this problem is phrased computationally, with the two known channels specified as quantum circuits, the resulting problem is complete for the class PSPACE [17], the class of all problems that can be solved in a polynomially bounded amount of space.

In light of this hardness, it is natural to ask if restricted versions of the problem are less difficult. In the case that the channels to be distinguished are unitary, the problem is "only" QMA-complete [11] (see also [12]), where QMA is the class of all those problems that can be efficiently verified with a quantum computer. In the case that the input circuits implement convex mixtures of unitary operation [15] or have short circuit descriptions [16] the problem is known to remain PSPACE-complete. In this paper we add two more restricted classes of channels to this list: the degradable channels and the antidegradable channels.

A quantum channel Φ is degradable if there exists a second channel that maps the output of Φ to the state of the environment after applying Φ. More precisely, a channel $\Phi\colon \mathbf{L}(\mathscr{A}) \to \mathbf{L}(\mathscr{B})$ that is given by

$$\Phi(\rho) = \operatorname{tr}_{\mathscr{E}} U(\rho \otimes |0\rangle\langle 0|)U^*,$$

is called degradable if there exists a channel $\Delta_\Phi\colon \mathbf{L}(\mathscr{B}) \to \mathbf{L}(\mathscr{E})$ such that

$$(\Delta_\Phi \circ \Phi)(\rho) = \operatorname{tr}_{\mathscr{B}} U(\rho \otimes |0\rangle\langle 0|)U^* = \Phi^C(\rho).$$

The channel Φ^C is called the complementary channel to Φ, and it is only defined up to an isometry, since it depends on the Stinespring representation of Φ. This does not affect the notion of degradability, however, as this isometry can be incorporated into the degrading map. These channels were introduced by Shor and Devetak [5] to study the capacity of a channel for transmitting quantum information. Notice that the set of degradable channels is not convex: any unitary channel is degradable, but the completely depolarizing channel is not, and it can be written as a convex combination of unitary channels (see [3] for an example of such a construction).

A channel is called antidegradable if the complementary channel is degradable. Alternately, a channel is Φ antidegradable is there exists a map A_Φ such that $A_\Phi \circ \Phi^C = \Phi$, where once again the channel Φ^C is only defined up to an isometry, but this isometry can also be part of the map A_Φ, so that the antidegradable channels are also well-defined. This class of channels was introduced by Wolf and Pérez-Garcia [19]. These channels can be informally thought of as the very noisy channels that lose more information to the environment than they preserve in the output. A thorough discussion of the degradable and antidegradable channels can be found in [4], where it is shown that, unlike the degradable channels, the set of antidegradable channels is convex.

The degradable and antidegradable channels are interesting from a quantum information perspective. A no-cloning argument implies that the antidegradable channels have zero capacity for the transmission of quantum information [7]. It is also known that the coherent information is additive on degradable channels, which implies that the quantum capacity is given by the coherent information of a single use of the channel, i.e. the formula for the quantum capacity does not require regularization [5].

As the degradable and antidegradable channels have nice properties with respect to the transmission of quantum information, it might be hoped that similar properties extend to the transmission of classical information. In the case of the Holevo (or χ-)capacity, it is shown in [4] that the additivity of this quantity on degradable channels is equivalent to the general case, making use of a result from [6]. As it is also known that this additivity problem is equivalent on the complementary class of channels [9, 13], this implies that the additivity of the antidegradable channels is also equivalent to the general case. Finally, using the recent result of Hastings [8], there are degradable and antidegradable channels that do not have additive Holevo capacity.

Interestingly, we can adapt the construction used by Cubitt, Ruskai, and Smith [4] to show that the quantum circuit distinguishability problem restricted to either the degradable or antidegradable channels remains PSPACE-complete. These results are the focus of this paper.

The complexity class PSPACE is the class of all problems that are solvable on a classical computer in a polynomially bounded amount of space. A recent result of Jain et al. shows that this class is equal to QIP, the class of all problems that can be interactively verified [10]. This implies that the problem of distinguishing degradable (or antidegradable) quantum channels exactly captures the difficulty of classical space bounded computation. It is known that PSPACE contains the class QMA, corresponding to quantum one-round verifiable computation. The class QMA in turn contains the class BQP, which captures the notion of efficient quantum computation. It is believed that both of these containments are strict, though proving this for either of them would be a major breakthrough in complexity theory.

The remainder of the paper is organized as follows. Section 2 presents some notation and results that will be used in the rest of the paper. Section 3 presents the construction due to Cubitt et al. [4] that embeds any channel into a degradable channel, as well as a related construction that embeds an arbitrary channel into an antidegradable channel. The main result of the paper is contained in Section 4, where these two constructions are used to show that the problems of distinguishing the degradable channels and the antidegradable channels remain PSPACE-complete.

2 Preliminaries

Here we introduce some technical concepts used throughout the paper. The notation used is standard: the reader familiar with quantum information is invited to skim this section.

Throughout the paper scripted capital letters $\mathscr{A}, \mathscr{B}, \mathscr{C}, \ldots$ will refer to Hilbert spaces, all of which will be finite dimensional. The set of all linear operators mapping $\mathscr{A}$ to $\mathscr{B}$ is denoted $\mathbf{L}(\mathscr{A}, \mathscr{B})$. The set of quantum states, or density matrices, on a space $\mathscr{A}$ is $\mathbf{D}(\mathscr{A})$: these are simply the positive semidefinite operators in $\mathbf{L}(\mathscr{A}, \mathscr{A})$ with unit trace.

A quantum channel is simply a map that takes density operators to density operators, even when applied to part of a larger system. These are given by the completely positive and trace preserving maps from $\mathbf{L}(\mathscr{A})$ to $\mathbf{L}(\mathscr{B})$. These maps are represented using capital Greek letters $\Phi, \Psi, \ldots$, and the set of all such maps is given by $\mathbf{T}(\mathscr{A}, \mathscr{B})$. The notation $I_{\mathscr{A}}$ refers to the identity map on $\mathbf{L}(\mathscr{A})$.

Given one use of an unknown channel $\Phi \in \mathbf{T}(\mathscr{A}, \mathscr{B})$, that is promised to be one of two known channels Φ_1 and Φ_2, what is the maximum probability that the channel can be identified? This quantity is given by the *diamond norm* of $\Phi_1 - \Phi_2$, which is

$$\|\Phi_1 - \Phi_2\|_{\diamond} = \max_{\rho} \left\|(\Phi_1 \otimes I_{\mathscr{F}})(\rho) - (\Phi_2 \otimes I_{\mathscr{F}})(\rho)\right\|_{\mathrm{tr}}, \tag{1}$$

where the space $\mathscr{F}$ is a space of the same dimension as $\mathscr{A}$ and the maximization is taken over all density matrices in $\mathbf{D}(\mathscr{A} \otimes \mathscr{F})$. The maximum probability that an unknown operation in $\{\Phi_1, \Phi_2\}$ can be correctly identified with a single use is given by

$$\frac{1}{2} + \frac{1}{4}\|\Phi_1 - \Phi_2\|_{\diamond},$$

which is one reason that this norm is central to the study of channel distinguishability. The diamond norm can be more generally defined over any linear operator mapping $\mathbf{L}(\mathscr{A})$ to $\mathbf{L}(\mathscr{B})$ (see [14] for such a definition, as well as some properties of this norm). The fact that we may restrict the maximum in Equation (1) to a density matrix in the case of the difference of two completely positive maps can be found in [17]. This norm is closely related to the *completely bounded norm*, in fact, $\|\Phi\|_{\diamond} = \|\Phi^*\|_{\mathrm{cb}}$, where Φ^* is the adjoint of Φ with respect to the Hilbert-Schmidt inner product.

One important property of the diamond norm of two channels is that applying it to several copies increases the norm. Intuitively this is obvious: given k copies of an unknown channel it is expected that it is easier to identify. The diamond norm of $\Phi_1^{\otimes k} - \Phi_2^{\otimes k}$ corresponds to a non-adaptive strategy with access to k copies of the two channels. The following Lemma from [17], which appears there as part of an efficient polarization procedure for the diamond norm, gives simple bounds on the norm as the number of copies increases.

Lemma 1. *Let $\Phi_1, \Phi_2 \in \mathbf{T}(\mathscr{A}, \mathscr{B})$ have $\|\Phi_1 - \Phi_2\|_{\diamond} = \delta > 0$. Then for any positive integer k*

$$2 - 2e^{\frac{-k\delta^2}{8}} < \left\|\Phi_1^{\otimes k} - \Phi_2^{\otimes k}\right\|_{\diamond} \leq k\delta.$$

This lemma will be used to show that parallel repetition reduces the error introduced by the reduction of the distinguishability problem to the cases of degradable and antidegradable channels.

In order to capture the difficulty of distinguishing implementations that represent efficient quantum computation, the input channels to the computational problems studied here are given as mixed-state circuits. This model, proposed by Aharonov et al. [1], consists of circuits in the usual unitary model, with two additional gates. These two gates are the introduction of fresh ancillary qubits in the $|0\rangle$ state and the partial trace of a qubit, both of which are non-unitary operations. The resulting circuit model allows for the (approximate) implementation of any quantum channel.

Notice that any circuit in this model can be efficiently converted into a circuit that first introduces any ancillary qubits, then performs a unitary circuit, and finally traces out any qubits that are not part of the output. This is due to the fact that introducing qubits earlier and tracing out qubits later does not affect the rest of the circuit. Thus, any mixed-state circuit can be assumed to be in the form of a Stinespring dilation. This property will be essential to the construction of degradable and antidegradable simulations in Section 3.

Channels are represented using mixed-state circuits as this provides a succinct representation of efficient quantum algorithms. Representing channels using a (potentially) exponentially larger representation, such as a Kraus decomposition, renders the problem solvable in classical polynomial time [2, 18], but in this model we lose the connection to efficient quantum algorithms as the descriptions of the channels are of size exponential in the number of input and output qubits. To maintain relevance in the case of practical distinguishability, such as between two physical implementations of quantum algorithms, we need an input description that scales logarithmically in the Hilbert space dimension, and so we use the mixed-state circuit model.

3 Simulations of channels

In this section we present two related constructions: the first embeds any channel into a degradable channel and the second embeds any channel into an antidegradable channel. Efficient quantum circuits for these problems (as well as the corresponding degrading and antidegrading maps) are presented.

3.1 Degradable channels

Given a quantum channel Φ we seek to simulate Φ by a degradable channel Ψ that has similar properties with respect to distinguishability. This can be done by adapting the construction used by Cubitt, Ruskai, and Smith [4] for the case of the minimum output entropy.

To describe this construction, we assume that $\Phi \in \mathbf{T}(\mathscr{A}, \mathscr{B})$ with $\dim \mathscr{A} = \dim \mathscr{B}$, i.e. that the original channel has identical input and output dimension. This assumption can be made without loss in generality by padding the smaller space with unused qubits, since these qubits will not affect the diamond norm used to define distinguishability. As the spaces $\mathscr{A}$ and $\mathscr{B}$ have the same dimension, they are isomorphic, and so we may view Φ as a channel in $\mathbf{T}(\mathscr{A}, \mathscr{A})$.

The channel Φ, given as input to a computational problem, is specified as a mixed-state quantum circuit. Such a circuit, as discussed in Section 2, can be efficiently transformed into one that first introduces any ancillary qubits, then performs some unitary circuit, and finally traces out any qubits that are not part of the output state. To this end, let the circuit $\Phi \in \mathbf{T}(\mathscr{A}, \mathscr{A})$ use the space $\mathscr{E}$ for ancillary qubits, and let U be the unitary that is applied to the space $\mathscr{A} \otimes \mathscr{E}$. Stated formally, the channel Φ is specified by

$$\Phi(\rho) = \operatorname{tr}_{\mathscr{E}} U(\rho \otimes |0\rangle\langle 0|)U^*. \tag{2}$$

This notation will be used to construct the degradable simulation of Φ.

The idea is to implement the channel

$$\Psi(\rho) = \frac{1}{2}|0\rangle\langle 0| \otimes \rho + \frac{1}{2}|1\rangle\langle 1| \otimes \Phi(\rho), \tag{3}$$

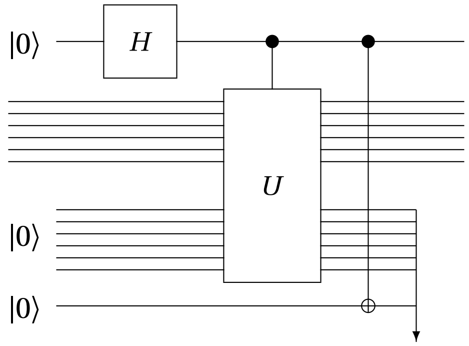

Figure 1: The degradable channel Ψ constructed from Φ, where U is the unitary given in Equation (2).

which has been used by Cubitt, Ruskai, and Smith [4] to reduce the (non-)additivity of the minimum output entropy to the degradable case. This is the channel that applies Φ with probability 1/2, does nothing to the input with probability 1/2, and leaves a flag on the output to indicate which case has occurred. The channel Ψ maps states on $\mathscr{A}$ to mixed states on $\mathscr{C} \otimes \mathscr{A}$, where $\mathscr{C}$ is the space of dimension two corresponding to flag state. Using the implementation of Φ in Equation (2), a circuit for the channel Ψ is given in Figure 1. The idea in this implementation is that the top ancillary qubit, corresponding do the space $\mathscr{C}$, is placed in the $|+\rangle$ state, which results in the circuit for Φ being applied with probability one-half. This control qubit is then 'copied' onto one of the environment qubits, so that the resulting output state is the mixture given in Equation (3).

To see that Ψ is a degradable channel, we construct the map that takes the output state of Ψ to the state of the environment. The construction of the channel Ψ, as well as a proof that it is degradable can be found in [4]. This proof is quite simple, and so it is repeated here. Before we construct the degrading map, however, notice that the complementary channel of $\Psi \in \mathbf{T}(\mathscr{A}, \mathscr{C} \otimes \mathscr{A})$, which is given by reversing the output and environment spaces, is

$$\Psi^C(\rho) = \frac{1}{2}|0\rangle\langle 0| \otimes |0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1| \otimes \Phi^C(\rho), \tag{4}$$

where the channel $\Phi^C \in \mathbf{T}(\mathscr{A}, \mathscr{E})$ is the complement of the original channel $\Phi \in \mathbf{T}(\mathscr{A}, \mathscr{A})$, given by

$$\Phi^C(\rho) = \operatorname{tr}_{\mathscr{A}} U(\rho \otimes |0\rangle\langle 0|)U^*.$$

These complementary channels are only defined up to an isometry, since a Stinespring dilation is defined only up to an isometry on the environment space, but for the present purpose, *any* complementary channel suffices.

Given Equation (4), it is not hard to construct the degrading map Δ_Ψ. Starting with the output of Ψ

$$\Psi(\rho) = \frac{1}{2}|0\rangle\langle 0| \otimes \rho + \frac{1}{2}|1\rangle\langle 1| \otimes \Phi(\rho),$$

as given by Equation (3), this channel can, based on a measurement of the flag state in the space $\mathscr{C}$ output one of the two states $|0\rangle\langle 0|$ and $\Phi^C(\rho)$. More formally, when the flag state is $|0\rangle$ the state in $\mathscr{A}$ is the original input ρ, so the channel can apply Φ^C to produce $\Phi^C(\rho)$. On the other hand, when this flag state is $|1\rangle$, the degrading map outputs $|0\rangle\langle 0|$, which can be done by producing the correct number of

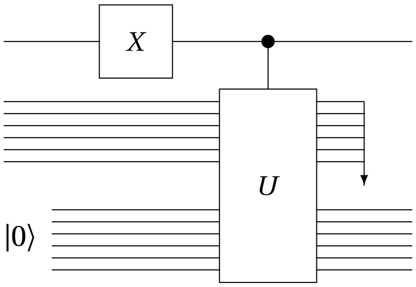

Figure 2: The degrading channel Δ_Ψ corresponding to the channel in Ψ in Figure 1.

untouched ancillary qubits as output. All that remains in to invert the flag qubit to get exactly the output of Ψ^C. A circuit implementation of the channel Δ_Ψ is presented in Figure 2. The fact that this channel has an efficient circuit implementation is not important for the main result: this is merely a simple way to specify the channel. We can verify that this map performs the required operation by observing that

$$\begin{aligned}\Delta_\Psi(\Psi(\rho)) &= \frac{1}{2}\Delta_\Psi\left(|0\rangle\langle 0|\otimes\rho+|1\rangle\langle 1|\otimes\Phi(\rho)\right)\\ &= \frac{1}{2}|1\rangle\langle 1|\otimes\Phi^C(\rho)+\frac{1}{2}|0\rangle\langle 0|\otimes|0\rangle\langle 0|\\ &= \Psi^C(\rho),\end{aligned}$$

where the final equality is Equation (4). This argument, due to Cubitt, Ruskai, and Smith [4] proves that the channel Ψ is degradable. In the next section we adapt this construction to the case of the antidegradable channels.

3.2 Antidegradable channels

In this section a construction very similar to that used in Section 3.1 is presented that takes any circuit Φ to a circuit Ψ implementing an antidegradable channel. The idea is to (with probability one-half) send the input state to the environment, so that the channel that maps the environment state to the output state will have a copy of the input state. This construction (and the proof that it produces an antidegradable channel) is very similar to the construction used for degradable channels.

Once again we may assume that Φ implements a channel in $\mathbf{T}(\mathscr{A},\mathscr{A})$, i.e. that Φ has the same input and output dimension, by embedding the smaller space into the larger, if necessary. As in Section 3.1, the constructed circuit Ψ will use one additional output qubit, implementing an antidegradable transformation in $\mathbf{T}(\mathscr{A},\mathscr{C}\otimes\mathscr{A})$.

Let Φ implement the transformation given by

$$\Phi(\rho)=\operatorname{tr}_{\mathscr{E}} U(\rho\otimes|0\rangle\langle 0|)U^*,$$

where, as before, the circuit for Φ can be assumed to be in this form by introducing any ancillary qubits first and delaying the tracing out of any qubits to the end of the circuit. The antidegradable channel Ψ will be constructed as

$$\Psi(\rho)=\frac{1}{2}|0\rangle\langle 0|\otimes|0\rangle\langle 0|+\frac{1}{2}|1\rangle\langle 1|\otimes\Phi(\rho). \tag{5}$$

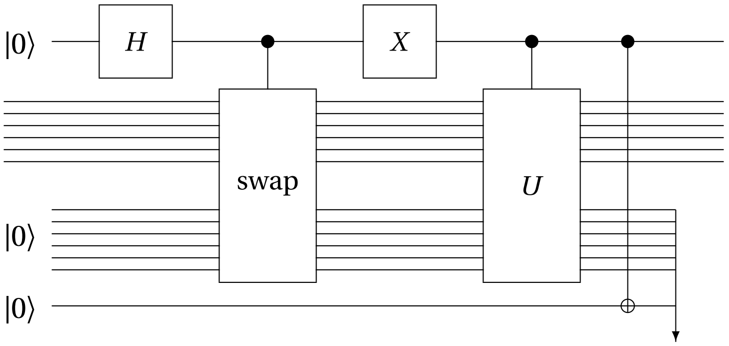

Figure 3: The antidegradable channel Ψ constructed from Φ.

This is just the channel that applies Φ with probability one-half, outputs $|0\rangle$ with probability one-half, and outputs a flag qubit in the space $\mathscr{C}$ to indicate which case has occurred. In a way very similar to the construction in Section 3.1, this channel can be implemented using a controlled-U operation. In this case, however, we will also need the operation that swaps the states in two spaces (i.e. $\mathrm{swap}(|a\rangle|b\rangle) = |b\rangle|a\rangle$). An implementation of the channel Ψ is shown in Figure 3. This circuit will, depending on the value of the control qubit in the space $\mathscr{C}$ either apply Φ or output the pure state $|0\rangle$.

To show that the circuit Ψ implements an antidegradable channel, we explicitly construct the map A_Ψ that maps the environment state of Ψ to the output state. The environment state of Ψ is once again simply the state produced by Ψ^C, the complementary channel to Ψ. This channel is given by

$$\Psi^C(\rho) = \frac{1}{2}|0\rangle\langle 0| \otimes \rho + \frac{1}{2}|1\rangle\langle 1| \otimes \Phi^C(\rho), \tag{6}$$

where as before the channel Φ^C is given by

$$\Phi^C(\rho) = \mathrm{tr}_{\mathscr{A}}\, U(\rho \otimes |0\rangle\langle 0|)U^*.$$

Given the state in Equation (6) it is not hard to see how to map it to the state in Equation (5). This can be done by implementing one of two operations, depending on the value of the flag qubit in the space $\mathscr{C}'$, which is the 'copy' of the control qubit traced out in Figure 3. If this qubit is in the state $|0\rangle$, then the remainder of the input state is ρ, the original input to Ψ, so that applying the circuit for Φ produces the state $\Phi(\rho)$. If the control qubit is in the $|1\rangle$ state, however, the remainder of the input state is $\Phi^C(\rho)$. This state can be discarded (i.e. traced out) and ancillary qubits in the state $|0\rangle$ can be swapped into the output space. As before, the value of the qubit in $\mathscr{C}'$ needs to be flipped with a Pauli X gate so that the state is exactly correct. A circuit implementing this is found in Figure 4.

To see that A_Ψ correctly implements the anti-degrading map for Ψ, we compute

$$\begin{aligned} A_\Psi(\Psi^C(\rho)) &= \frac{1}{2}A_\Psi\left(|0\rangle\langle 0| \otimes \rho + |1\rangle\langle 1| \otimes \Phi^C(\rho)\right) \\ &= \frac{1}{2}|1\rangle\langle 1| \otimes \Phi(\rho) + \frac{1}{2}|0\rangle\langle 0| \otimes |0\rangle\langle 0| \\ &= \Psi(\rho), \end{aligned}$$

where the final equality is Equation 5. This demonstrates that the channel Ψ constructed from Φ is antidegradable. In the following section the implications of this construction for the hardness of distinguishing antidegradable channels are considered.

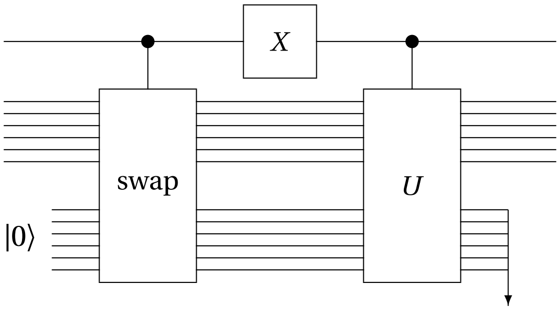

Figure 4: The anti-degrading channel corresponding to the channel in Ψ in Figure 3.

4 Distinguishing degradable and antidegradable channels

In this section we consider the implications of the constructions in Sections 3.1 and 3.2 for the computational problem of distinguishing quantum channels. These constructions essentially embed any channel into either a degradable or antidegradable channel. This can be used to show that distinguishing these channels is no easier than distinguishing general channels.

To observe this in a more formal setting, we introduce the problem of distinguishing two quantum channels when they are provided as quantum circuits. The diamond norm $\|\Phi_1 - \Phi_2\|_\diamond$ determines the maximum probability that an unknown channel $\Phi \in \{\Phi_1, \Phi_2\}$ can be identified with a single use. For this reason, we may formalize the computational distinguishability problem in terms of evaluating the diamond norm of the difference of two known channels.

Problem 2 (Quantum Circuit Distinguishability)**.** For constants $0 \le b < a \le 2$, the input consists of quantum mixed-state circuits Φ_1 and Φ_2 that implement transformations in $\mathbf{T}(\mathscr{A}, \mathscr{A})$. The promise problem is to distinguish the two cases:

Yes: $\|\Phi_1 - \Phi_2\|_\diamond \ge a$,

No: $\|\Phi_1 - \Phi_2\|_\diamond \le b$.

This problem is introduced and shown to be PSPACE-complete for all $0 < b < a < 2$ in [17]. The distinguishability problem as originally defined allows the input and output dimensions of the channels to differ, but as discussed in Section 3 this can be avoided by padding the smaller of the two spaces. For conciseness, this problem will be abbreviated $\mathrm{QCD}_{a,b}$. Restricting the channels Φ_1, Φ_2 in Problem 2 to degradable channels results in the problem DEGRADABLE $\mathrm{QCD}_{a,b}$, whereas the restriction to antidegradable channels gives the problem ANTIDEGRADABLE $\mathrm{QCD}_{a,b}$. The main result of this paper is that these problems remain PSPACE-complete.

To show the hardness of these restricted distinguishability problems, we show that the constructions of Sections 3.1 and 3.2 reduce the general QCD problem to these two restricted problems. These two constructions can be efficiently implemented when the input channels are given as quantum circuits: this can be seen from the circuit representations in Figures 1 and 3. It will be shown that these reductions prove the hardness of the restricted distinguishability problems, which suffices to prove that they are PSPACE-complete: this is because they are contained in PSPACE by the same algorithm used for the general QCD problem, which can be found in [17].

The primary ingredient in the proof that these constructions reduce the distinguishability problem to degradable and antidegradable channels is a proof that when applied to each of a pair of channels, the constructions preserve the diamond norm of the difference of the two channels. This is not difficult to see from the output of the constructions, given by Equations (3) and (5), but for completeness this is argued formally in the following lemma.

Lemma 3. *Let Φ_1, Φ_2 be quantum circuits implementing transformations in $\mathbf{T}(\mathscr{A}, \mathscr{A})$. If the channels $\Psi_1, \Psi_2 \in \mathbf{T}(\mathscr{A}, \mathscr{C} \otimes \mathscr{A})$ are given by*

$$\Psi_i(\rho) = \frac{1}{2}|0\rangle\langle 0| \otimes \rho + \frac{1}{2}|1\rangle\langle 1| \otimes \Phi_i(\rho),$$

and the channels $\Lambda_1, \Lambda_2 \in \mathbf{T}(\mathscr{A}, \mathscr{C} \otimes \mathscr{A})$ are given by

$$\Lambda_i(\rho) = \frac{1}{2}|0\rangle\langle 0| \otimes |0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1| \otimes \Phi_i(\rho),$$

for $i \in \{1,2\}$, then

$$||\Psi_1 - \Psi_2||_\diamond = ||\Lambda_1 - \Lambda_2||_\diamond = \frac{1}{2}\,||\Phi_1 - \Phi_2||_\diamond.$$

Proof. Let $\rho \in \mathbf{D}(\mathscr{A} \otimes \mathscr{F})$ be an arbitrary state. Then

$$\begin{aligned}\left\|(\Psi_1 \otimes I_{\mathscr{F}} - \Psi_2 \otimes I_{\mathscr{F}})(\rho)\right\|_{\mathrm{tr}} &= \frac{1}{2}\left\||0\rangle\langle 0| \otimes (\rho - \rho) + |1\rangle\langle 1| \otimes ([\Phi_1 \otimes I_{\mathscr{F}} - \Phi_2 \otimes I_{\mathscr{F}}](\rho))\right\|_{\mathrm{tr}} \\ &= \frac{1}{2}\left\|(\Phi_1 \otimes I_{\mathscr{F}} - \Phi_2 \otimes I_{\mathscr{F}})(\rho)\right\|_{\mathrm{tr}}.\end{aligned}$$

Since the diamond norm may be defined as the maximization over all such states ρ, this implies the equality $||\Psi_1 - \Psi_2||_\diamond = ||\Phi_1 - \Phi_2||_\diamond/2$. The same argument implies that $||\Lambda_1 - \Lambda_2||_\diamond = ||\Phi_1 - \Phi_2||_\diamond/2$. □

This lemma implies that distinguishing degradable or antidegradable channels is PSPACE-complete for all $0 < b < a < 1$, using the hardness result known for the general problem, $\mathrm{QCD}_{2a,2b}$ [17]. Notice that we lose a factor of two in the parameters a and b: this is because the diamond norm of the constructed channels is half the norm of the original channels.

This result can be strengthened with parallel repetition. The strategy is to take an instance (Ψ_1, Ψ_2) of Degradable $\mathrm{QCD}_{1-\varepsilon,\varepsilon}$ and construct the instance $(\Psi_1^{\otimes k}, \Psi_2^{\otimes k})$. This second instance will have outputs that are more distinguishable, for the simple reason that there are more copies of the states to be distinguished available. This will send the norm for 'yes' instances of the problem from $1-\varepsilon$ to a value close to 2, but it also has the property that the norm of 'no' instances is not made too large. This is a straightforward consequence of the bounds in Lemma 1, which appears in [17] as part of an efficient procedure for polarizing the diamond norm. This technique can be applied to both Degradable QCD and Antidegradable QCD as the classes of degradable and antidegradable channels are closed under parallel repetition.

Theorem 4. *For any choice of constants $0 < b < a < 2$, both of the problems* Degradable $\mathrm{QCD}_{a,b}$ *and* Antidegradable $\mathrm{QCD}_{a,b}$ *are* PSPACE*-complete.*

Proof. These problems are in PSPACE as they are restrictions of the general QCD problem [17]. To see that it they are also PSPACE-hard, take an instance (Φ_1, Φ_2) of the $\mathrm{QCD}_{2-2\varepsilon,2\varepsilon}$, for $\varepsilon > 0$ a small constant. We will reduce this problem to Degradable QCD.

Applying the construction of Section 3.1 to (Φ_1, Φ_2) results in a pair of circuits (Ψ_1, Ψ_2) that form an instance of Degradable $\mathrm{QCD}_{1-\varepsilon,\varepsilon}$, by Lemma 3. As the degradable channels are closed under tensor products, $(\Psi_1^{\otimes k}, \Psi_2^{\otimes k})$ gives a pair of circuits implementing degradable channels. By Lemma 1, we have

$$\|\Psi_1 - \Psi_2\|_\diamond \geq 1 - \varepsilon \implies \left\|\Psi_1^{\otimes k} - \Psi_2^{\otimes k}\right\|_\diamond \geq 2 - 2e^{-k(1-2\varepsilon)/8},$$
$$\|\Psi_1 - \Psi_2\|_\diamond \leq \varepsilon \implies \left\|\Psi_1^{\otimes k} - \Psi_2^{\otimes k}\right\|_\diamond \leq k\varepsilon.$$

Then, for any $0 < b < a < 2$, choosing $k \geq -16\ln(1 - a/2)$ and $\varepsilon \leq \min\{1/4, b/k\}$ implies the desired inequalities $2 - 2e^{-k(1-2\varepsilon)/8} > a$ and $k\varepsilon < b$. This shows the PSPACE hardness of Degradable $\mathrm{QCD}_{a,b}$. The case of Antidegradable $\mathrm{QCD}_{a,b}$ is completely symmetric, with the exception that we use the construction of Section 3.2 to obtain antidegradable channels Ψ_1 and Ψ_2. □

5 Conclusion

This paper has presented a construction for embedding an arbitrary channel into a degradable channel due to Cubitt, Ruskai, and Smith [4], as well as a closely related construction for antidegradable channels. These constructions can be efficiently implemented on quantum circuits, so that instances of the quantum circuit distinguishability problem can be mapped to degradable or antidegradable channels. The main result is that the distinguishability problem on quantum circuits remains hard when restricted to either the class of degradable channels or the class of antidegradable channels.

Acknowledgements

I am grateful for discussions with John Watrous, as well as helpful comments from Richard Cleve, Achim Kempf, and Steve Fenner. This work has been supported by the Centre for Quantum Technologies at the National University of Singapore, as well the Bell Family Fund for Quantum Computing, while the author was a student 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] A. Ben-Aroya and A. Ta-Shma. On the complexity of approximating the diamond norm. *Quantum Information and Computation*, to appear, 2010. EPRINT: arXiv:0902.3397 [quant-ph].

[3] 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.

[4] T. S. Cubitt, M. B. Ruskai, and G. Smith. The structure of degradable quantum channels. *Journal of Mathematical Physics*, 49(10):102104, 2008. DOI: 10.1063/1.2953685. EPRINT: arXiv:0802.1360 [quant-ph].

[5] I. Devetak and P. W. Shor. The capacity of a quantum channel for simultaneous transmission of classical and quantum information. *Communications in Mathematical Physics*, 256(2):287–303, 2005. DOI: 10.1007/s00220-005-1317-6. EPRINT: arXiv:quant-ph/0311131.

[6] M. Fukuda and M. M. Wolf. Simplifying additivity problems using direct sum constructions. *Journal of Mathematical Physics*, 48(7):072101, 2007. DOI: 10.1063/1.2746128. EPRINT: arXiv:0704.1092 [quant-ph].

[7] V. Giovannetti and R. Fazio. Information-capacity description of spin-chain correlations. *Physical Review A*, 71(3):032314, 2005. DOI: 10.1103/PhysRevA.71.032314. EPRINT: arXiv:quant-ph/0405110.

[8] M. B. Hastings. Superadditivity of communication capacity using entangled inputs. *Nature Physics*, 5:255–257, 2009. DOI: 10.1038/nphys1224. EPRINT: arXiv:0809.3972 [quant-ph].

[9] A. S. Holevo. Complementary channels and the additivity problem. *Theory of Probability and its Applications*, 51(1):92–100, 2007. DOI: 10.1137/S0040585X97982244. EPRINT: arXiv:quant-ph/0509101.

[10] R. Jain, Z. Ji, S. Upadhyay, and J. Watrous. QIP = PSPACE, 2009. EPRINT: arXiv:0907.4737 [quant-ph].

[11] 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.

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

[13] C. King, K. Matsumoto, M. Nathanson, and M. B. Ruskai. Properties of conjugate channels with applications to additivity and multiplicativity. *Markov Processes and Related Fields*, 13(2):391–423, 2007. EPRINT: arXiv:quant-ph/0509126.

[14] 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.

[15] B. Rosgen. Additivity and distinguishability of random unitary channels. *Journal of Mathematical Physics*, 49(10):102107, 2008. DOI: 10.1063/1.2992977. EPRINT: arXiv:0804.1936 [quant-ph].

[16] B. Rosgen. Distinguishing short quantum computations. In *Proceedings of the 25th Symposium on Theoretical Aspects of Computer Science*, pp. 597–608. 2008. EPRINT: arXiv:0712.2595 [quant-ph], URL `http://hal.archives-ouvertes.fr/hal-00255825/en/`.

[17] 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.

[18] J. Watrous. Semidefinite programs for completely bounded norms, 2009. EPRINT: arXiv:0901.4709 [quant-ph].

[19] M. M. Wolf and D. Pérez-García. Quantum capacities of channels with small environment. *Physical Review A*, 75(1):012303, 2007. DOI: 10.1103/PhysRevA.75.012303. EPRINT: arXiv:quant-ph/0607070.