

Computational Distinguishability of Quantum Channels

by

William Rosgen

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Computer Science

Waterloo, Ontario, Canada, 2009

© William Rosgen 2009

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

The computational problem of distinguishing two quantum channels is central to quantum computing. It is a generalization of the well-known satisfiability problem from classical to quantum computation. This problem is shown to be surprisingly hard: it is complete for the class **QIP** of problems that have quantum interactive proof systems, which implies that it is hard for the class **PSPACE** of problems solvable by a classical computation in polynomial space.

Several restrictions of distinguishability are also shown to be hard. It is no easier when restricted to quantum computations of logarithmic depth, to mixed-unitary channels, to degradable channels, or to antidegradable channels. These hardness results are demonstrated by finding reductions between these classes of quantum channels. These techniques have applications outside the distinguishability problem, as the construction for mixed-unitary channels is used to prove that the additivity problem for the classical capacity of quantum channels can be equivalently restricted to the mixed unitary channels.

Acknowledgements

I would like to thank my supervisor John Watrous for years of guidance, support, and insight. Without his help this would not have been possible. I would also like to thank the rest of my committee, Richard Cleve, Stephen Fenner, Achim Kempf, and Ben Reichardt, for providing helpful comments on an earlier draft of this thesis. I would also like to thank Lana for putting up with me during the writing of this thesis and supporting me throughout the process.

Contents

List of Figures	xii
List of Symbols	xiii
1 Introduction	1
1.1 Overview	2
1.2 Quantum information	6
1.3 Classes of quantum channels	18
2 Quantum Computational Complexity	25
2.1 Quantum circuits	26
2.2 Quantum complexity classes	35
3 Measures for Quantum Information	41
3.1 Entropy	42
3.2 Schatten p-norms	45
3.3 The classical capacity of a quantum channel	47
3.4 The trace norm	53
3.5 The diamond norm	57
3.6 Fidelity	64
3.7 Polarization of the diamond norm	69
3.8 Conclusion	76

4	The Close Images Problem	77
4.1	Log-depth mixed-state quantum circuits	78
4.2	QIP completeness of close images	79
4.3	The swap test	83
4.4	Reduction to logarithmic depth	87
4.5	Correctness of the reduction	93
4.6	Conclusion	98
5	Distinguishability of Quantum Computations	99
5.1	Overview of distinguishability problems	100
5.2	Quantum circuit distinguishability	102
5.3	QIP protocol	104
5.4	Reduction from Close Images	107
5.5	Correctness of the reduction	110
5.6	Distinguishing log-depth computations	115
5.7	Conclusion	116
6	Degradable and Antidegradable Channels	117
6.1	Degradable and antidegradable channels	118
6.2	Simulation by a degradable channel	119
6.3	Distinguishing degradable channels	121
6.4	Simulation by an antidegradable channel	123
6.5	Distinguishing antidegradable channels	126
6.6	Conclusion	127
7	Mixed-Unitary Channels	129
7.1	Mixed-unitary channels	130
7.2	Unital channels	132
7.3	Mixed-unitary approximation	134
7.4	Properties of the constructed channel	142

7.5	Multiplicativity of mixed-unitary transformations	145
7.6	Mixed-unitaries and minimum output entropy	147
7.7	Circuit constructions	149
7.8	QIP -completeness of distinguishing mixed-unitary circuits	156
7.9	Conclusion	160
8	Conclusion	161
	Bibliography	163
	Index	177

List of Figures

1.1	Optimal strategy for channel distinguishability	3
1.2	Reductions presented in the thesis	4
2.1	An example quantum circuit	27
2.2	Gates in the unitary circuit model	29
2.3	Simulation of the swap gate with three controlled-not gates	29
2.4	Controlled-U gate	30
2.5	Non-unitary gates in the mixed state circuit model	31
2.6	Simulations of three of the gates in the mixed-state circuit model	31
2.7	Simulating a channel with a unitary circuit	32
2.8	Log-depth implementation of controlled operation on n qubits	34
2.9	Constant depth implementation of controlled operation on n qubits	35
2.10	Known relationships between complexity classes	36
2.11	A three message quantum interactive proof system	38
3.1	Circuits output by the construction in Lemma 3.26	72
3.2	Circuits output by the construction in Lemma 3.28	74
4.1	Transformations in a quantum interactive proof system	80
4.2	Reduction from QIP to CLOSE IMAGES	81
4.3	Constant-depth implementation of a swap gate	84
4.4	Circuit implementation of the swap test	84
4.5	Decomposition of a circuit into constant depth pieces	88

4.6	Testing procedure used in reduction to log-depth circuits	89
4.7	Overview of the output spaces of the constructed log-depth circuits . . .	90
4.8	The output of the reduction to log-depth circuits	91
5.1	Complexity classes and distinguishability problems	101
5.2	Input circuits for the reduction	108
5.3	Circuit to apply either input circuit based on a control qubit	108
5.4	Circuits output by the reduction	109
6.1	Reduction to a degradable channel	120
6.2	Degrading channel for the channel in Figure 6.1	121
6.3	Reduction to an antidegradable channel	124
6.4	Anti-degrading channel for the channel in Figure 6.3	125
7.1	Channel to be approximated by a mixed-unitary	135
7.2	One stage of the circuit for the ancilla simulation procedure	153
7.3	Circuit performing the ancilla simulation procedure	154
7.4	The mixed-unitary circuit that simulates the original circuit	154

List of Symbols

$\tilde{\mathbb{1}}_{\mathcal{H}}$	completely mixed state on \mathcal{H} , $\tilde{\mathbb{1}}_{\mathcal{H}} = \mathbb{1}_{\mathcal{H}} / \dim \mathcal{H}$
δ_{ij}	Kronecker delta function: $\delta_{ij} = 1$ if $i = j$ and 0 otherwise
$F(\rho, \sigma)$	fidelity of states ρ and σ
$F_{\max}(\Phi, \Psi)$	Maximum output fidelity of channels Φ and Ψ
$\mathcal{H}, \mathcal{K}, \dots$	finite dimensional Hilbert spaces
\cong	isomorphism between Hilbert spaces
\log	base-two logarithm
$\ \cdot\ _{\diamond}$	diamond norm
$\ \cdot\ _p$	Schatten p -norm
$\ \cdot\ _{\text{tr}}$	trace norm
$\nu(\cdot)$	Maximum output p -norm
$S(\rho)$	von Neumann entropy of the state ρ
$S_{\min}(\Phi)$	minimum output entropy of the channel Φ
Classes of Operators	
$\mathbf{D}(\mathcal{H})$	density operators on \mathcal{H}
$\mathbf{L}(\mathcal{H}, \mathcal{K})$	set of all linear operators from \mathcal{H} to \mathcal{K}
$\mathbf{T}(\mathcal{H}, \mathcal{K})$	set of all channels from $\mathbf{L}(\mathcal{H})$ to $\mathbf{L}(\mathcal{K})$
$\mathbf{U}(\mathcal{H})$	unitary operators on \mathcal{H}
$\mathbf{U}(\mathcal{H}, \mathcal{K})$	isometries mapping \mathcal{H} to \mathcal{K}

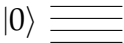
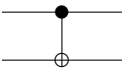
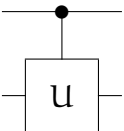
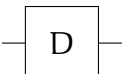
Complexity Classes


BQP	quantum efficiently solvable problems
EXP	classically solvable problems in exponential time
NP	classically efficiently verifiable problems
P	classically efficiently solvable problems
PSPACE	classically solvable problems in polynomial space
QIP	quantum efficiently interactively verifiable problems
QMA	quantum efficiently verifiable problems

Operators


$\mathbb{1}_{\mathcal{H}}$	identity operator on $\mathbf{L}(\mathcal{H}, \mathcal{H})$
H	Hadamard matrix
$I_{\mathcal{H}}$	identity transformation on $\mathbf{L}(\mathcal{H})$
$\text{tr}_{\mathcal{K}}$	Partial trace over the system in \mathcal{K}
W	Swap operation: $W a\rangle b\rangle = b\rangle a\rangle$
X	Pauli X matrix
Y	Pauli Y matrix
Z	Pauli Z matrix

Quantum Circuits


	Introduction of ancillary qubits in $ 0\rangle$ state
	Controlled-not gate
	Controlled U gate
	Completely dephasing channel $D(i\rangle\langle j) = \delta_{ij} i\rangle\langle j $



Hadamard gate




Measurement in the computational basis $\{|0\rangle, |1\rangle\}$




Completely depolarizing channel $N(\rho) = \mathbb{1}/d$




Partial trace




Two-qubit swap gate



Multi-qubit swap gate $W(|a\rangle|b\rangle) = |b\rangle|a\rangle$



Pauli X gate



Pauli Z gate

Computational Problems

CI Close Images Problem

QCD Quantum Circuit Distinguishability problem

Chapter 1

Introduction

Distinguishing two quantum channels is one of the most important tasks in quantum information. This is the problem of determining if there is an input state on which the two channels to produce output states that are distinguishable. When this is phrased as a computational problem it is complete for the complexity class **QIP** of problems that have quantum interactive proof systems. This problem seems to be computationally much more difficult than other variants of the problem, such as distinguishing classical circuits or distinguishing unitary quantum circuits.

In light of this hardness, it is natural to consider restricted versions of the problem. Many of these special cases are also hard: reductions can be found to some of the more interesting classes of quantum channels. These results suggest that this problem is not likely to be tractable even on many of the restricted channels that can be realized by experiment. This is, however, not a surprise: distinguishing two channels is a restricted version of quantum process tomography, which is computationally intractable for large systems.

These reductions provide simulations of general quantum channels by channels in restricted classes. While these simulations do not accurately model all aspects of the original channel, the constructed simulations do share many properties with the original channel. Many of these results can be applied outside the narrow focus of distinguishing quantum channels: it is hoped that these techniques will prove useful for a number of problems in quantum information theory.

Contents

1.1	Overview	2
1.2	Quantum information	6
1.2.1	Hilbert spaces	7

1.2.2	Pure states	9
1.2.3	Linear operators	10
1.2.4	Mixed states	14
1.2.5	State evolution and measurement	15
1.2.6	Channels	16
1.3	Classes of quantum channels	18
1.3.1	Circuit restrictions	18
1.3.2	Degradable and antidegradable channels	20
1.3.3	Entanglement-breaking channels	21
1.3.4	Unital channels	22
1.3.5	Mixed-unitary channels	23

1.1 Overview

This thesis studies the computational problem of distinguishing quantum channels. This problem asks, given two implementations of quantum channels, is there an input on which the implemented channels behave distinctly? One of the main results of the thesis is that this problem is in general extremely difficult: it is complete for the complexity class **PSPACE** of problems that can be solved with a polynomially bounded amount of memory. Since this problem is intractable in general it remains to understand those classes of channels for which the problem has an efficient solution and those classes on which it remains hard. This problem is becoming more significant for quantum computing: as larger practical systems are being studied it becomes more difficult to verify that the implemented transformation is close to an ideal transformation.

One of the other problems considered in the thesis is the question of the additivity of the Holevo capacity for quantum channels. If this quantity were additive, it would significantly simplify the tasks of encoding and decoding for the transmission of classical information through a quantum channel. Specifically, this question asks whether two uses of a channel can send more than twice the classical information that can be sent with only one use of the channel. That this might be possible is because entanglement may be present between the two inputs to the channel. This question stood open for several years until it was recently shown that there exist channels for which the Holevo capacity is super-additive [Has09]. The classes of channels that are




Figure 1.1: The optimal strategy for determining which channel Φ_1, Φ_2 is the unknown channel Φ_i . The strategy is to prepare some state ρ on which the two channels output maximally distinguishable states, send this state through the channel Φ_i , and then make an optimal measurement of the result.

known to have the additivity property are generally quite restricted. It is important to better understand which classes of channels are additive as the use of quantum channels for sending classical information is an important application of quantum information.

The problem of distinguishing channels can be equivalently rephrased as: given a single use of an unknown channel that is one of two known channels Φ_1 and Φ_2 , what is the probability that the optimal strategy can detect which of the two channels it is? In general the best strategy in this case is to prepare an input state ρ on which the output states of Φ_1 and Φ_2 are maximally distinguishable, send ρ through the unknown channel, and then attempt to solve the distinguishability problem for the output states of the channels. In general this strategy requires the preparation of a state on some larger system, only part of which is sent through the unknown channel. This strategy is illustrated in Figure 1.1. One of the main results of the thesis is that this problem, properly formalized, is complete for the complexity class **QIP** of problems that have quantum interactive proof systems. This is a surprising result: the same problem restricted to deterministic classical circuits is a restatement of the canonical **NP**-complete problem satisfiability. These complexity classes are discussed in more detail in Chapter 2. For the reader unfamiliar with computational complexity theory, it is important only to know that the class **QIP** contains the class **NP** and it is thought that **QIP** is much larger than **NP**.

This hardness result is found in Chapter 5, where it is shown using a Karp reduction from the problem of determining if two quantum channels can be made to output states that are close together, where a Karp reduction is simply an efficient procedure that transforms instances of one problem into equivalent instances of another problem. A problem that is the target of a Karp reduction is thus shown to be at least as hard as the starting problem.

The Close Images problem is easily derived from the definition of the class **QIP**,




Figure 1.2: Reductions presented in the thesis. Problems are reduced to those problems above them. CI and QCD are shorthand for CLOSE IMAGES and QUANTUM CIRCUIT DISTINGUISHABILITY. Edges are marked with the chapter the reduction appears in.

which implies that it is also **QIP**-complete. This derivation can be found in Chapter 4, and is originally due to Kitaev and Watrous [KW00].

Given that the distinguishability problem is intractable, much of the remainder of the thesis is a study of several restricted classes of quantum channels, with a focus on the hardness of this distinguishability problem on them. For many of these classes Karp reductions are found from the general problem to the problem on the restricted class. These reductions prove that these restricted versions of the distinguishability problem are also **QIP**-complete. The reductions found in the thesis are illustrated in Figure 1.2.

The problems shown to be **QIP**-complete using these reductions are also hard for the more familiar class **PSPACE**. In fact, it has recently been shown that $\mathbf{QIP} = \mathbf{PSPACE}$, which implies that these two classes are the same [JJUW09]. These problems on quantum channels then provide an interesting characterization of a fundamental classical complexity class. Despite this equivalence, the class is referred to as **QIP** throughout the thesis, as the hardness results presented here all follow from the definition of the class in terms of quantum interactive proof systems.

The first of these reductions, in Chapter 4 concerns not the distinguishability problem, but the close images problem. It is shown that this problem can be equivalently restricted to the channels implemented by circuits of logarithmic depth. These channels are an important class: they can be implemented in parallel in a logarithmic amount of time. This makes these channels interesting for a practical perspective, as a quantum system implementing one of these channels needs to be protected from

decoherence for only a short time.

The second reduction presented in the thesis is the focus of Chapter 5. This is the reduction from the close images problem to the problem of distinguishing two channels, where the channels are given as input to the problem in the form of circuits. This reduction proves the hardness of the distinguishability problem for general quantum channels. One other important property of this reduction is that it adds only logarithmic overhead to the depth of the circuits. This implies that even the log-depth circuit distinguishability problem is **QIP**-complete, which provides powerful evidence that this problem, a restricted case of quantum process tomography, is likely to be difficult in practice even for computations that can be performed in a very limited amount of time.

Chapter 6 extends the hardness of the distinguishability problem in a different direction: to the channels known as the degradable channels. These channels can be thought of as the channels that preserve most of the information about the input, since there exists a second channel that maps the output of a degradable channel to the state of the environment. These channels are described in more detail Section 1.3. This result implies that distinguishing these channels that do not lose very much information to the environment remains hard.

In the other direction, Chapter 6 also contains a reduction of the distinguishability problem to the antidegradable channels. These are the channels for which there exists a second map that takes the environment state to the output state. In particular, this means that an eavesdropper with sufficient resources can reconstruct the output of the channel. At an intuitive level, this implies that an antidegradable channel loses more information to the environment than it preserves in the output state. The fact that distinguishing these channels is **QIP**-hard is evidence that even channels that do not preserve very much information are hard to distinguish.

The final reduction presented in the thesis, in Chapter 7, is a transformation that approximates a general channel by one that is mixed-unitary. The mixed-unitary channels are those channels that can be expressed as the convex mixture of unitary (i.e. noise-free and reversible) channels. The mixed unitary channels have several nice properties that make them interesting in quantum information theory. The approximate simulation of a channel by a mixed-unitary channel performs well only for measures of quality based on the maximized purity of the output of the channel. This, however, suffices to reduce the distinguishability problem to the mixed-unitary channels. This technique is also used to show that the additivity of the Holevo capacity of a general channel can be approximately restated in terms of a mixed-unitary channel.

The Holevo capacity, which can be used to measure the amount of classical information that can be sent through a quantum channel, is introduced in detail in Chapter 3. This technique allows, for instance, the observation that this quantity is additive for all channels if and only if it is additive for mixed-unitary channels (which has recently been shown by Hastings [Has09] through the construction of a mixed-unitary channel that is not additive).

Taken together, these reductions demonstrate the hardness of the distinguishability problem on several distinct classes of channels. As this problem is one of the most interesting problems in quantum computation, these hardness results point to cases of the problem that are not able to be efficiently solved, under the usual complexity theoretic assumptions. It is also hoped that these complete problems for **QIP** will provide a way to further understand this class. The problems shown in Figure 1.2 are among only a few problems in quantum information that are known to be complete for **QIP**.

The technique of reducing a problem to a restricted class by simulating general channels by those of a restricted class can also have applications outside of quantum computational complexity. For instance, the reduction to mixed-unitary channels in Chapter 7 was initially constructed for the distinguishability problem, but the same construction has implications for the additivity of certain capacities. These techniques are powerful and general: any problem defined on quantum channels is a candidate for reduction to these restricted classes. This does not work in general, as these reductions produce channels that do not simulate the general channel in every sense, but for any problem defined using similar notions of distance on quantum channels, these reductions apply. These techniques provide not only a method for the study of the distinguishability problem, the primary application studied in this thesis, but a tool for the more general study of quantum channels and their properties.

1.2 Quantum information

In this section the necessary mathematical framework for the problems outlined in the previous section is introduced. The concepts and notation used here are relatively standard. This is not a complete introduction to quantum information.

More background on quantum information can be found in the books [BŽ06, NC00]. Background on much of the linear algebra introduced here, including a thorough discussion of the tensor product, can be found in [Rom05]. A good general reference

for results from functional analysis that are occasionally useful in quantum information is [Con90], while the books [Bha97, HJ85, HJ91] provide a more focused treatment of the types of operators often found in quantum information

1.2.1 Hilbert spaces

The fundamental backdrop for quantum information is the complex Hilbert space. These spaces are the complete vector spaces over \mathbb{C} with inner product, where the completeness of the space is with respect to the topology induced by the inner product. All such spaces considered in this thesis are finite dimensional and denoted by calligraphic letters $\mathcal{H}, \mathcal{K}, \dots$. Elements of a Hilbert space \mathcal{H} of dimension d space can be represented as vectors in \mathbb{C}^d . These vectors are denoted $|\phi\rangle$. Elements of the dual space \mathcal{H}^* , which are (complex) linear functionals on the space \mathcal{H} , are denoted $\langle\phi| = (|\phi\rangle)^*$. The inner product on these Hilbert spaces is defined, for two vectors $|\phi\rangle$ with elements u_i and $|\psi\rangle$ with elements v_i , by

$$\langle\phi|, |\psi\rangle\rangle = \langle\phi|\psi\rangle = \sum_i \bar{u}_i v_i,$$

where \bar{u} denotes the complex conjugate of u . This inner product is linear in the second argument and conjugate linear in the first: this is the usual convention in physics, but it is opposite to the way things are typically defined in mathematics. The *dimension* of a space has been mentioned several times: this is simply the maximum number of elements in a pairwise orthogonal set. When the elements of a d -dimensional space are viewed as vectors with complex entries, the dimension of the space coincides with the length of these vectors.

Norms are a fundamental tool in quantum information. They provide a means to define a notion of size on quantum states and quantum channels. Throughout the thesis, we will be most interested in using norms to bound the distance between two objects.

Definition 1.1. A norm $\|\cdot\|$ on some linear space V (over the field \mathbb{C}) is a function from V to \mathbb{R} satisfying three basic properties, for all $x, y \in V$:

$$\text{Nonnegativity: } \|x\| \geq 0 \text{ with equality if and only if } x = 0 \quad (1.1)$$

$$\text{Homogeneity: } \|cx\| = |c| \|x\| \text{ for all } c \in \mathbb{C} \quad (1.2)$$

$$\text{Triangle Inequality: } \|x + y\| \leq \|x\| + \|y\| \quad (1.3)$$

The standard Euclidean norm on a Hilbert space can be defined in terms of the inner product given above. This is the norm of a vector $|\phi\rangle$ with elements v_i given by

$$\| |\phi\rangle \| = \sqrt{\langle \phi | \phi \rangle} = \sqrt{\sum_i |v_i|^2}.$$

A vector $|\phi\rangle$ is called normalized if $\| |\phi\rangle \| = 1$.

The standard basis of the Hilbert space \mathcal{H} of dimension d is given by the set of orthonormal (i.e. normalized and pairwise orthogonal) vectors $\{|0\rangle, |1\rangle, \dots, |d-1\rangle\}$. The vector $|i\rangle$ viewed as a vector in \mathbb{C}^d is simply the vector with a one in position $i+1$ and zeroes in all other positions. This basis is also known as the *computational basis*. When no confusion will arise, this basis will also be labelled $\{|1\rangle, |2\rangle, \dots, |d\rangle\}$.

Two finite dimensional Hilbert spaces \mathcal{H}, \mathcal{K} are isomorphic if they are both of the same dimension. This is written $\mathcal{H} \cong \mathcal{K}$. In such a case, the canonical isomorphism between the two spaces simply maps the computational basis of \mathcal{H} to the computational basis of \mathcal{K} . When two spaces are isomorphic, the isomorphism between them will often be used implicitly to consider vectors in one Hilbert space as being vectors in the other space.

Quantum systems of large dimension are often built up of many smaller dimensional system. If \mathcal{H} and \mathcal{K} are Hilbert spaces of dimension $\dim \mathcal{H} = d_{\mathcal{H}}$ and $\dim \mathcal{K} = d_{\mathcal{K}}$, then the Hilbert space of dimension $d_{\mathcal{H}}d_{\mathcal{K}}$ formed by combining them is denoted $\mathcal{H} \otimes \mathcal{K}$. Similarly, the element $|\phi\rangle \otimes |\psi\rangle \in \mathcal{H} \otimes \mathcal{K}$ is formed by combining the two elements $|\phi\rangle \in \mathcal{H}$ and $|\psi\rangle \in \mathcal{K}$. When viewed as complex vectors, these elements are given by the Kronecker product

$$|\phi\rangle = \begin{pmatrix} u_1 \\ u_2 \\ \vdots \\ u_{d_{\mathcal{H}}} \end{pmatrix}, \quad |\psi\rangle = \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_{d_{\mathcal{K}}} \end{pmatrix}, \quad |\phi\rangle \otimes |\psi\rangle = \begin{pmatrix} u_1|\psi\rangle \\ u_2|\psi\rangle \\ \vdots \\ u_{d_{\mathcal{H}}}|\psi\rangle \end{pmatrix}.$$

The notation $|\phi\rangle \otimes |\psi\rangle$ will often be abbreviated $|\phi\rangle|\psi\rangle$ or even $|\phi\psi\rangle$ where no confusion is likely to arise.

The space $\mathcal{H} \otimes \mathcal{K}$ does not consist solely of elements of the form $|\phi\rangle|\psi\rangle$: it also contains linear combinations of these elements. As an example, an element of the tensor product of two systems of dimension two is $|00\rangle + |11\rangle$. This element cannot be written in tensor product form. A basis for $\mathcal{H} \otimes \mathcal{K}$ can be formed by taking the pairwise tensor product of basis elements for the two subsystems, i.e. the set $\{|i\rangle|j\rangle : 0 \leq i < d_{\mathcal{H}}, 0 \leq j < d_{\mathcal{K}}\}$ is a basis for $\mathcal{H} \otimes \mathcal{K}$. When convenient we will also use the standard basis $\{|i\rangle : 0 \leq i < d_{\mathcal{H}}d_{\mathcal{K}}\}$ for this space.

1.2.2 Pure states

The state of a quantum system is described, up to a phase $e^{i\theta}$, by a normalized vector $|\phi\rangle \in \mathcal{H}$, known as a *pure state*. Provided that there is no uncertainty about the system, these pure states suffice to completely describe the state of a quantum system. For this reason they are of fundamental importance in quantum information. Any such state on a d -dimensional Hilbert space can be expressed as

$$|\phi\rangle = \sum_{i=0}^{d-1} a_i |i\rangle,$$

where the *amplitudes* a_i are complex numbers satisfying $\sum_i |a_i|^2 = 1$, which is simply a restatement of the normalization requirement.

The smallest system of interest in quantum computation is the two dimensional Hilbert space. Such a system is often referred to as a *qubit*. On such a system, the two standard basis states are $|0\rangle$ and $|1\rangle$. A second basis that is often extremely useful is given by the two orthogonal states

$$|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle), \quad |-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle).$$

As was previously mentioned, there are elements in a composite Hilbert space $\mathcal{H} \otimes \mathcal{K}$ that cannot be decomposed into a tensor product of an element of \mathcal{H} and an element of \mathcal{K} . When quantum states have this property they are called *entangled*. Up to normalization, we have already met the maximally entangled state. This is the state $|\phi_+\rangle \in \mathcal{H} \otimes \mathcal{H}$, where $d = \dim \mathcal{H}$, given by

$$|\phi_+\rangle = \frac{1}{\sqrt{d}} \sum_{i=0}^{d-1} |i\rangle|i\rangle.$$

Any state that is not entangled is called *separable*.

An important representation of pure states of a composite system $\mathcal{H} \otimes \mathcal{K}$ is the *Schmidt decomposition*. Any state $|\phi\rangle \in \mathcal{H} \otimes \mathcal{K}$ may be expressed as

$$|\phi\rangle = \sum_{i=1}^r \lambda_i |a_i\rangle |b_i\rangle. \quad (1.4)$$

In this decomposition the sets $\{|a_i\rangle\}$ and $\{|b_i\rangle\}$ form orthonormal sets in \mathcal{H} and \mathcal{K} , respectively, and the coefficients λ_i are all positive and real. The number r in Equation (1.4) satisfies $r \leq \min\{\dim \mathcal{H}, \dim \mathcal{K}\}$ and is known as the *Schmidt rank* of $|\phi\rangle$. The numbers λ_i are known as the *Schmidt coefficients*. They satisfy $\sum_i \lambda_i^2 = 1$. Notice that a pure state has Schmidt rank one if and only if it is separable: this is only one example of the utility of this decomposition in quantum information theory.

1.2.3 Linear operators

In order to introduce how states evolve during a quantum computation, we must take a detour through some of the different spaces of linear operators that act on a Hilbert space \mathcal{H} . The most general of these is $\mathbf{L}(\mathcal{H}, \mathcal{K})$, which is the set of all linear operators that map elements of \mathcal{H} to elements of \mathcal{K} . As we assume that all Hilbert spaces appearing in the thesis are finite dimensional, linearity implies boundedness which in turn implies continuity: this space is often referred to as $\mathbf{B}(\mathcal{H}, \mathcal{K})$ for this reason. When the Hilbert space that a linear operator acts on is viewed as the space of vectors $\mathbb{C}^{\dim \mathcal{H}}$, the set $\mathbf{L}(\mathcal{H}, \mathcal{K})$ is exactly the set of $\dim \mathcal{K}$ by $\dim \mathcal{H}$ complex matrices. The notation $\mathbf{L}(\mathcal{H})$ is shorthand for $\mathbf{L}(\mathcal{H}, \mathcal{H})$.

For $A \in \mathbf{L}(\mathcal{H}, \mathcal{K})$, the operator $A^* \in \mathbf{L}(\mathcal{K}, \mathcal{H})$ is the adjoint of A , in the sense that A^* is the unique operator that, for any $|\phi\rangle \in \mathcal{H}$ and any $|\psi\rangle \in \mathcal{K}$, satisfies

$$\langle |\psi\rangle, A|\phi\rangle \rangle = \langle \psi|A|\phi\rangle = \langle A^*|\psi\rangle, |\phi\rangle \rangle.$$

When A is represented by a matrix, A^* is the conjugate transpose of A . Given such a representation, the complex conjugate of A is denoted \bar{A} , and the transpose of A is denoted A^\top .

There are a few more classes of operators that are extremely important in quantum information. One these classes of operators is the class of *Hermitian* operators. These are those operators $A \in \mathbf{L}(\mathcal{H})$ such that $A = A^*$. An important subclass of the Hermitian operators is the set of *positive*, or *positive semidefinite*, operators. These are the Hermitian operators $A \in \mathbf{L}(\mathcal{H})$ such that for any $|\phi\rangle \in \mathcal{H}$

$$\langle \phi|A|\phi\rangle \geq 0.$$

The positive operators can be equivalently characterized as those operators $A \in \mathbf{L}(\mathcal{H})$ that can be expressed as $A = B^*B$ for some $B \in \mathbf{L}(\mathcal{H})$. The notation $A \geq B$ is used to denote that the operator $A - B$ is positive, with the special case $A \geq 0$ used to state that A is positive. The other important operators are the *unitary* operators. These are the invertible operators $U \in \mathbf{L}(\mathcal{H})$ with $U^* = U^{-1}$. It follows from this property that applying a unitary matrix to a pair of element of \mathcal{H} does not change their inner product, which further implies that unitaries do not change the norm. This property implies that the unitary operators are exactly the invertible operators that preserve the pure states, a property which makes them extremely important for quantum computing. One other important characterization is the that unitary operators are exactly those operators that map orthonormal bases to orthonormal bases. The set of all unitary operators in $\mathbf{L}(\mathcal{H})$ is denoted $\mathbf{U}(\mathcal{H})$. The notion of

unitarity can be extended to $V \in \mathbf{L}(\mathcal{H}, \mathcal{K})$ with $\dim \mathcal{K} \geq \dim \mathcal{H}$ by considering those V with the property that $V^*V = \mathbb{1}_{\mathcal{H}}$. Such an operator is called an *isometry*, and the set of all such operators is denoted $\mathbf{U}(\mathcal{H}, \mathcal{K})$. These operators embed the elements of the space \mathcal{H} into the larger space \mathcal{K} .

One of the most important operators in this space is $\mathbb{1}_{\mathcal{H}}$, which is the identity operator on \mathcal{H} . As a matrix, this operator has ones on the main diagonal and zeroes in all other positions. When restricted to qubits, the identity is one of the four *Pauli matrices*. These four matrices belonging to $\mathbf{L}(\mathcal{H})$, where $\dim \mathcal{H} = 2$, are defined by

$$\mathbb{1}_{\mathcal{H}} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

One other matrix that will be consistently useful is the *Hadamard* matrix. This is the unitary operator that converts the basis $\{|0\rangle, |1\rangle\}$ to the basis $\{|+\rangle, |-\rangle\}$ and vice versa. This operator can be expressed in matrix form as

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

An extremely important function on linear operators is the *trace*. This is the operation $\text{tr}: \mathbf{L}(\mathcal{H}) \rightarrow \mathbb{C}$ that, on a matrix representation of an operator A , is simply the sum of the main diagonal. One of the most important properties of the trace is that it is *cyclic*, i.e. for operators A, B, C we have

$$\text{tr}(ABC) = \text{tr}(BCA) = \text{tr}(CAB),$$

whenever the products in the above equation are defined. Note that the trace is not stable under more general commutation of the arguments, i.e. there are operators A, B, C such that $\text{tr}(ABC) \neq \text{tr}(CBA)$.

The space $\mathbf{L}(\mathcal{H}, \mathcal{K})$ equipped with the inner product given by

$$\langle A, B \rangle = \text{tr}(A^*B)$$

is also a Hilbert space. This implies that $\mathbf{L}(\mathcal{H}) \otimes \mathbf{L}(\mathcal{K})$ is well defined. In fact, it is the case that

$$\mathbf{L}(\mathcal{H}, \mathcal{K}) \otimes \mathbf{L}(\mathcal{A}, \mathcal{B}) = \mathbf{L}(\mathcal{H} \otimes \mathcal{A}, \mathcal{K} \otimes \mathcal{B}).$$

When the tensor product is extended to operators it behaves in the same way as it does on vectors. For $A \in \mathbf{L}(\mathcal{H})$ with elements a_{ij} for $1 \leq i, j \leq d$ and $B \in \mathbf{L}(\mathcal{K})$, the operator

$A \otimes B$ has matrix representation given by the block matrix

$$A \otimes B = \begin{pmatrix} a_{11}B & a_{12}B & \cdots & a_{1d}B \\ a_{21}B & a_{22}B & \cdots & a_{2d}B \\ \vdots & \vdots & \ddots & \vdots \\ a_{d1}B & a_{d2}B & \cdots & a_{dd}B \end{pmatrix}.$$

As $L(\mathcal{H})$ is itself a Hilbert space, we can find bases of operators for it. An important orthogonal basis for this space is given by the discrete Weyl operators, also known as the generalized Pauli operators. These operators extend the Pauli operators to the d dimensional space \mathcal{H} , keeping the properties of orthogonality and unitarity, but losing Hermiticity. As these operators will be essential to several of the arguments in the thesis, they are introduced in detail. The discrete Weyl operators are based on generalizations of the X and Z operations, given by

$$X = \sum_{j=1}^d |j+1\rangle\langle j|$$

$$Z = \sum_{j=1}^d \omega_d^j |j\rangle\langle j|,$$

where ω_d is a d -th primitive root of unity (such as $e^{2i\pi/d}$), and in the definition of X the operator $|d+1\rangle\langle d|$ is taken to be $|1\rangle\langle d|$. The operator X simply advances each state of the computational basis to the next, and the operator Z applies a different phase to each basis state. It is clear from the definition that $XX^* = \mathbb{1}_{\mathcal{H}} = ZZ^*$, which implies that these operators are unitary. It is also clear that X and Z fail to commute:

$$ZX = \omega_d XZ. \tag{1.5}$$

Using these operators, the discrete Weyl operator with index $(a, b) \in \mathbb{Z}_d \times \mathbb{Z}_d$ is given by

$$W_{a,b} = X^a Z^b.$$

For two dimensional systems, these operators are, up to phases, exactly the usual Pauli matrices. These operators are unitary, since they are products of the unitary operators X and Z . Equation 1.5 can be directly extended to these operators to obtain

$$W_{a,b} W_{e,f} = X^a Z^b X^e Z^f = \omega_d^{be-af} X^e Z^f X^a Z^b = \omega_d^{be-af} W_{e,f} W_{a,b}. \tag{1.6}$$

To see that these operators form an orthogonal basis for $\mathbf{L}(\mathcal{H})$, notice that, by the cyclic property of the trace

$$\mathrm{tr} W_{a,b}^* W_{e,f} = \mathrm{tr} Z^{-b} X^{-a} X^e Z^f = \mathrm{tr} X^{e-a} Z^{f-b} = \begin{cases} d & \text{if } a = e \text{ and } b = f, \\ 0 & \text{otherwise.} \end{cases} \quad (1.7)$$

These operators can be normalized to obtain an orthonormal basis for $\mathbf{L}(\mathcal{H})$, but this comes at the cost of unitarity, so we will not do this here. The discrete Weyl operators will be used in Chapter 7 to show that several transformations on quantum states can be realized as convex mixtures of unitary transformations.

A linear operator $A \in \mathbf{L}(\mathcal{H})$ is *normal* if $A^*A = AA^*$. By definition the Hermitian and unitary operators are normal. Any normal operator $A \in \mathbf{L}(\mathcal{H})$ has a *spectral decomposition*, which is a representation as

$$A = \sum_i^{\dim \mathcal{H}} \lambda_i |\phi_i\rangle\langle\phi_i|, \quad (1.8)$$

where the vectors $\{|\phi_i\rangle\}$, called the *eigenvectors* of A , are an orthonormal basis for \mathcal{H} . The associated complex numbers λ_i are called the *eigenvalues* of A . The space spanned by the eigenvectors of A corresponding to nonzero eigenvalues is called the *support* of A . This space has dimension equal to the rank of A .

The classes of operators that we have previously encountered can be characterized in terms of the spectral decomposition. A normal matrix U is unitary if and only if all of its eigenvalues have norm one, i.e. if $|\lambda_i| = 1$ for all i . This also implies that U is full rank. A normal operator is Hermitian if and only if all of its eigenvalues are real. This can be seen by considering the adjoint of the representation in Equation (1.8). As a further restriction, an operator is positive if and only if all of it is normal and has only nonnegative real eigenvalues. In addition to this, if A is an operator with eigenvalues λ_i , then the trace of A is given by $\mathrm{tr} A = \sum_i \lambda_i$. This characterization of the trace is extremely useful.

The spectral decomposition also allows functions on the complex numbers to be extended to operators in $\mathbf{L}(\mathcal{H})$. An example of this is square root of a positive matrix. This is defined, for any positive operator A , by taking the square roots of the eigenvalues, i.e.

$$\sqrt{A} = \sum_i \sqrt{\lambda_i} |\phi_i\rangle\langle\phi_i|,$$

where A has spectral decomposition $A = \sum_i \lambda_i |\phi_i\rangle\langle\phi_i|$. It is easy to see from this definition that the operator square root satisfies $\sqrt{A} \sqrt{A} = A$. It is less obvious that

\sqrt{A} defined in this way is the unique square root of the operator A , but this is indeed the case. This technique can be extended to define $\log(A)$ for a positive matrix A as well as e^A and $|A|$ for general normal matrices.

The absolute value of an operator has a different definition when $A \in \mathbf{L}(\mathcal{H}, \mathcal{K})$ is not normal (and potentially not square). In this case, $|A| = \sqrt{A^*A}$, relying on the fact that for any operator A , the operator A^*A is positive. The eigenvalues of $|A|$ play a central role in another important decomposition of an operator. The *singular value decomposition* of $A \in \mathbf{L}(\mathcal{H}, \mathcal{K})$ gives a representation of A that mimics the spectral decomposition, but exists even when A is not normal. This representation is

$$A = \sum_{i=1}^d s_i |\phi_i\rangle\langle\psi_i|$$

where $d = \min\{\dim \mathcal{H}, \dim \mathcal{K}\}$. The values s_i are nonnegative and real, these are called the *singular values* of A . They are equal to the eigenvalues of $|A|$. The vectors $\{|\phi_i\rangle\}$ and $\{|\psi_i\rangle\}$ form orthogonal sets in \mathcal{K} and \mathcal{H} , respectively. The singular values will be very important in Chapter 3 where we consider a collection of operator norms that depend solely on them.

1.2.4 Mixed states

Pure states suffice to model the behaviour of a quantum system in a known state, but they do not completely capture the situation when there is uncertainty about exactly which state a system is in. As an example, if the state a two-dimensional system is $|0\rangle$ or $|1\rangle$ each with probability one-half, then the behaviour of the system is identical to one in which the state is a uniform mixture of $|+\rangle$ or $|-\rangle$, yet these two descriptions differ.

This problem is resolved by resorting to density operators. Given a system that is in the state $|\phi_i\rangle$ with probability p_i (this is called the *ensemble* $\{(p_i, |\phi_i\rangle)\}$), the density operator associated with the system is given by

$$\rho = \sum_i p_i |\phi_i\rangle\langle\phi_i|.$$

A given density matrix may, in general, have an infinite set of ensembles that generate it. This notation resolves the earlier example, since we have, for any two orthonormal bases $|\phi_i\rangle$ and $|\psi_i\rangle$ of \mathcal{H}

$$\frac{1}{\dim \mathcal{H}} \sum_i |\phi_i\rangle\langle\phi_i| = \frac{1}{\dim \mathcal{H}} \sum_i |\psi_i\rangle\langle\psi_i| = \frac{\mathbb{1}_{\mathcal{H}}}{\dim \mathcal{H}} = \tilde{\mathbb{1}}_{\mathcal{H}}$$

where the symbol $\tilde{\mathbb{1}}_{\mathcal{H}}$ denotes $\mathbb{1}_{\mathcal{H}} / \dim \mathcal{H}$, the normalized identity operator on \mathcal{H} . An element $\rho \in \mathbf{L}(\mathcal{H})$ is called a *density operator* (or equivalently a density matrix) if and only if it satisfies the two properties

1. ρ is positive,
2. $\text{tr } \rho = 1$.

The set of all such density operators on \mathcal{H} is denoted $\mathbf{D}(\mathcal{H})$.

The density operators are also referred to as the *mixed states*, as they provide a complete description of a quantum system. Pure states also fit into this framework: the state $|\phi\rangle$ corresponds to the density operator $|\phi\rangle\langle\phi|$, and these two notions of state will be used interchangeably in this case. Notice also that the set of density matrices $\mathbf{D}(\mathcal{H})$ is both compact and convex. The extreme points of this set are simply the rank one projectors $|\phi\rangle\langle\phi|$ corresponding to pure states in \mathcal{H} .

A mixed state in $\mathbf{D}(\mathcal{H} \otimes \mathcal{K})$ is called separable if it is the convex combination of a set of separable states in $\mathcal{H} \otimes \mathcal{K}$. If a mixed state cannot be decomposed in this way, i.e. any ensemble contains an entangled pure state, then it is called entangled.

1.2.5 State evolution and measurement

The evolution of a quantum system in the state $\rho \in \mathbf{D}(\mathcal{H})$ is determined by the action of a unitary operator $U \in \mathbf{U}(\mathcal{H})$. The state of the system after this evolution is $U\rho U^*$. As we shall see in the next section, this does not capture every quantum process, but it is an important special case. As $U^{-1} = U^*$ is also unitary, this implies that any unitary evolution is, in principle, reversible.

Measurements provide a method for retrieving information from a quantum system. The simplest case of measurement is given by a *projective measurement*, which is a set $\{\Pi_i\}$ of orthogonal projectors in $\mathbf{L}(\mathcal{H})$ with the property that $\sum_i \Pi_i = \mathbb{1}_{\mathcal{H}}$. When this measurement is performed on a state $\rho \in \mathbf{D}(\mathcal{H})$ the outcome is i with probability $p_i = \text{tr}(\Pi_i \rho)$ and the state after measurement is $(\Pi_i \rho \Pi_i) / p_i$. In the case that the outcome of the measurement is unknown, i.e. it is discarded or forgotten, the resulting state is given by $\sum_i \Pi_i \rho \Pi_i$, i.e. the weighted mixture of all of the measurement outcomes.

There is a special case of projective measurement that is of particular importance in quantum computing. This is known as measurement in the computational basis, which

is given by the complete set of projectors $\{|i\rangle\langle i| : 0 \leq i < \dim \mathcal{H}\}$. Any measurement using orthogonal rank one projectors can be derived from this measurement by rotating the state ρ to be measured using some unitary operation U .

Projective measurements are not the only case allowed by quantum mechanics. More generally, a POVM measurement is given by a set $\{E_i\}$ of positive operators that sum to $\mathbb{1}_{\mathcal{H}}$. The outcome of such a measurement is i with probability $p_i = \text{tr}(E_i \rho)$. While the state after measurement can be defined as in the case of projective measurements, a simpler model will suffice for the results in this thesis. In this model the outcome after measurement is the state $|i\rangle$ when the result is i . This form of measurement can be quite convenient to work with. POVM measurements can always be realized by projective measurements on the state $\rho \otimes |0\rangle\langle 0|$ in a larger Hilbert space. This result is known as Naimark's theorem [Neu43].

1.2.6 Channels

We have already seen two types of evolution for quantum states: unitary evolution and measurement. Both of these types of evolution are special cases of the most general type of transformation on quantum states. These are the linear transformations that map density matrices to density matrices, known as quantum *channels*. Such a map can capture any process allowed by quantum mechanics. These maps can also be characterized as the linear operators Φ from $\mathbf{L}(\mathcal{H})$ to $\mathbf{L}(\mathcal{K})$ that satisfy two properties

Trace preserving: $\text{tr} \Phi(X) = \text{tr} X$ for all $X \in \mathbf{L}(\mathcal{H})$

Complete positivity: If $X \geq 0$ then $(\Phi \otimes I_{\mathcal{K}})(X) \geq 0$ for all $\mathcal{K}, X \in \mathbf{L}(\mathcal{H} \otimes \mathcal{K})$.

In the above definition, $I_{\mathcal{K}}$ is the identity transformation on $\mathbf{L}(\mathcal{K})$ and the map $\Phi \otimes \Psi$ is simply the map that applies Φ and Ψ on their respective Hilbert spaces. The set of all channels from $\mathbf{L}(\mathcal{H})$ to $\mathbf{L}(\mathcal{K})$ is denoted $\mathbf{T}(\mathcal{H}, \mathcal{K})$. A linear map taking $\mathbf{L}(\mathcal{H})$ to $\mathbf{L}(\mathcal{K})$ that is not necessarily a channel will occasionally be referred to as a super-operator.

One of the most important quantum channels is the operation known as the *partial trace*. This is the channel in $\mathbf{T}(\mathcal{H} \otimes \mathcal{K}, \mathcal{H})$ that traces out the system in the space \mathcal{K} . This map is defined on $X \otimes Y$ with $X \in \mathbf{L}(\mathcal{H})$ and $Y \in \mathbf{L}(\mathcal{K})$ by

$$\text{tr}_{\mathcal{K}} X \otimes Y = (\text{tr} Y)X,$$

and extended to the whole space $\mathbf{L}(\mathcal{H} \otimes \mathcal{K})$ by linearity. This is the operation that discards the system in \mathcal{K} , and as such, is very useful in quantum information. The

partial trace can also be expressed by explicitly writing out the trace over \mathcal{K} . Let $\{|\phi_i\rangle\}$ be any orthonormal basis for \mathcal{K} , then for any $X \in \mathbf{L}(\mathcal{H} \otimes \mathcal{K})$, the partial trace over \mathcal{K} is

$$\mathrm{tr}_{\mathcal{K}} X = \sum_i (\mathbf{1}_{\mathcal{H}} \otimes \langle \phi_i |) X (\mathbf{1}_{\mathcal{H}} \otimes |\phi_i\rangle).$$

One important feature of mixed states is that they can always be viewed as part of a pure state on a larger Hilbert space. Any $\rho \in \mathbf{D}(\mathcal{H})$ can be expressed as a pure state $|\phi\rangle \in \mathcal{H} \otimes \mathcal{K}$, where $\dim \mathcal{K} \geq \mathrm{rank} \rho$, as

$$\rho = \mathrm{tr}_{\mathcal{K}} |\phi\rangle\langle\phi|.$$

The state $|\phi\rangle$ is referred to as a *purification* of ρ . These purifications will form an integral part of many of the proof techniques used in this thesis. It is also important that any two purifications $|\phi\rangle, |\psi\rangle \in \mathcal{H} \otimes \mathcal{K}$ of a state $\rho \in \mathbf{D}(\mathcal{H})$ are related by a unitary operation on the space \mathcal{K} alone, i.e. there exists a $U \in \mathbf{U}(\mathcal{K})$ such that

$$(\mathbf{1}_{\mathcal{H}} \otimes U)|\phi\rangle = |\psi\rangle.$$

This fact will be used in the definition of the fidelity in Chapter 3.

There are two convenient representations of quantum channels that will be needed. The first of these is the representation of a completely positive map Φ by a set of Kraus operators, which are matrices A_i such that

$$\Phi(X) = \sum_i A_i X A_i^*.$$

This representation is due to Choi [Cho75]. If, in addition, Φ is trace preserving, then the operators A_i satisfy the property

$$\sum_i A_i^* A_i = \mathbf{1}.$$

If $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ then the number of Kraus operators in a minimal Kraus decomposition is at most $(\dim \mathcal{H})(\dim \mathcal{K})$.

The second representation of importance is known as the Stinespring Dilation Theorem, after the 1955 work of Stinespring [Sti55], though the precise statement of the result we use here is given by Hellwig and Kraus [HK70]. This theorem states that any quantum channel can be represented as a unitary operation on a larger space, some of which is traced out. More formally, for a channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ there are spaces \mathcal{A}, \mathcal{B} and a $U \in \mathbf{U}(\mathcal{H} \otimes \mathcal{A}, \mathcal{K} \otimes \mathcal{B}) \cong \mathbf{U}(\mathcal{H} \otimes \mathcal{A})$ such that

$$\Phi(X) = \mathrm{tr}_{\mathcal{B}} U(X \otimes |0\rangle\langle 0|)U^*.$$

Where \mathcal{A} can be chosen so that $\dim \mathcal{A} \leq \dim \mathcal{H} \dim \mathcal{K}$. Such a representation is unique up to an isometry on the space that is traced out. This representation can be used to recover a Kraus representation: see [Sch96] for an overview of this result.

The Stinespring representation implies that in order to model a quantum channel, we need worry about only three parts: introducing ancillary qubits in a known pure state, implementing unitary operations, and implementing the partial trace. This will be extremely helpful for the reductions in this thesis that seek to simulate general quantum channels with channels from restricted classes. More details can be found in Section 2.1 where the model of circuits used in the thesis is formally defined.

1.3 Classes of quantum channels

This section provides an overview of the different classes of quantum channels that will be encountered in this thesis. This overview will be kept somewhat brief, as the classes that will receive detailed treatment are reintroduced more thoroughly in the chapters where they appear.

The classes of channels presented here place different restrictions on the set of channels. Some of these restrictions come from practical notions, like the channels that can be implemented in a small amount of time, and some of the restrictions come from more theoretical concerns, such as the antidegradable channels. The restricted classes studied here are largely incomparable: this is because one of the aims of this thesis is to present simplified versions of the distinguishability problem that are nevertheless just as hard as the general case. In order that these results cover more of the quantum channels that are likely to arise in practice, it is helpful if the distinguishability problem is shown to be hard on several unrelated classes of channels.

The material that appears in this section makes several references to the material that follows in Chapter 3. Because this material is only used in a superficial way in this section, it is not necessary to have read Chapter 3 first, though a familiarity with quantum information will help.

1.3.1 Circuit restrictions

Quantum circuits are a convenient way to provide a quantum channel as input to a computational problem, such as the problem of distinguishing quantum channels. One of the advantages of this representation is that, given a quantum computer, it

allows the channel to be evaluated, but it remains computationally infeasible to find a matrix representation for all but the simplest channels. The circuit model used here is presented in detail in Section 2.1. This level of detail will not be necessary to introduce the classes of channels defined by placing restrictions on this model.

The circuit representation for quantum channels allows restricted classes of channels to be defined by placing restrictions on the types of circuits that are allowed. These channels can be much simpler than general channels. An example of this is the class of channels with *stabilizer circuits*, which are the circuits defined on a restricted set of quantum gates. Given such a circuit, a channel can be efficiently simulated using a deterministic classical computer [AG04], but it is expected that this is not possible for quantum channels given as general quantum circuits, as this would imply the equivalence of classical and quantum computation.

Restricting the input circuits to the distinguishability problem mentioned in Section 1.1 can lead to simpler variants of the problem. One such restriction is to the class of channels that implement unitary operations. These channels can be obtained as a circuit restriction by eliminating the non-unitary gates from the circuit model. Distinguishing these circuits appears to be easier than general mixed-state circuits [JWB05].

One of the more interesting circuit restrictions is the requirement that the input circuits to the distinguishability problem have depth logarithmic in the number of input qubits. These are the circuits that can be performed in logarithmic time with a parallel model of quantum computing. Such a model is not unreasonable in many implementation schemes for quantum computing. These circuits are interesting as they limit the length of time that quantum information needs to be protected from decoherence. For this reason, much of experimental quantum computing is concerned with very short computations, and log depth circuits are an interesting generalization of such computations. Many important quantum algorithms are known to have log depth circuits, such as the approximate quantum Fourier transform [CW00] and the encoding and decoding operations for many quantum error correcting codes [MN02].

One of the results on the thesis is that distinguishing log depth quantum mixed-state circuits is complete for **QIP**, i.e. just as hard as the general case. This is shown by reducing the close images problem, studied in Chapter 4, to a log depth version of itself. The essential idea behind this reduction is to simulate a general quantum circuit by a log-depth one by slicing the circuit into log depth pieces that are performed in parallel. This circuit will perform the same computation as the original circuit only if the input to one piece matches the output of the previous piece. To ensure that this is the case for the circuits constructed in the reduction, tests are applied to force this to be

the case for any outputs of the two circuits that can potentially have close images. This reduction shows that the close images problem remains **QIP**-complete when restricted to log depth circuits. Extending this to the distinguishability problem on log-depth circuits follows from the fact that the reduction in Chapter 5 preserves the log-depth restriction of the circuits.

1.3.2 Degradable and antidegradable channels

The degradable channels are those channels Φ for which there exists a second channel that maps the output of Φ to the environment of Φ , i.e. the space that is traced out in a Stinespring representation. These channels were introduced by Shor and Devetak [DS05] and can be thought of as the channels where the environment contains no information that is not also present in the output of the channel. A more formal definition is: a channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ expressed as $\Phi(\rho) = \text{tr}_{\mathcal{B}} \mathbf{U}(\rho \otimes |0\rangle\langle 0|)\mathbf{U}^*$ is *degradable* if there exists a channel D such that

$$\text{tr}_{\mathcal{K}} \mathbf{U}(\rho \otimes |0\rangle\langle 0|)\mathbf{U}^* = D(\Phi(\rho)).$$

Stinespring representations are not unique, but any two differ by an isometry on the environment space, and this isometry can be absorbed into the channel D , which implies that the notion of degradability does not depend on the choice of representation.

The channel Φ^c given by tracing out the output space \mathcal{K} and not the environment space \mathcal{B} is often referred to as the complementary (or conjugate) channel to Φ , with the caveat that it is only defined up to the Stinespring representation chosen for Φ . A channel is *antidegradable* if the complementary channel is degradable. More plainly, a channel Φ is antidegradable if there exists a second channel that maps the environment of Φ to the output of Φ . Antidegradable channels are also well-defined, since as in the case of degradability, the choice of Stinespring representation can be absorbed into the degrading map. A thorough discussion of the degradable and antidegradable channels can be found in [CRS08].

These channels are discussed in Chapter 6, where it is shown that the problem of computationally distinguishing two channels is made no easier when the channels are promised to be degradable or antidegradable. This is done using a construction similar to one found in [CRS08] that is used to reduce the additivity of the classical capacity to the degradable case.

1.3.3 Entanglement-breaking channels

An entanglement-breaking channel Φ is a channel for which the output $(\Phi \otimes I_{\mathcal{H}})(\rho)$ is separable for any input state ρ . This class of channels contains many of the commonly used channels, such as the completely depolarizing channel and the complete dephasing channel. It is helpful to state a few alternate characterizations of the entanglement-breaking channels.

Proposition 1.2 (Horodecki, Shor, and Ruskai [HSR03]). *Let $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$. The following are equivalent*

1. Φ is entanglement-breaking,
2. $(\Phi \otimes I_{\mathcal{H}})(|\psi\rangle\langle\psi|)$ is separable, for $|\psi\rangle$ a maximally entangled state on $\mathcal{H} \otimes \mathcal{H}$,
3. Φ has a Kraus decomposition using only rank one operators,
4. Φ can be written as

$$\Phi(\rho) = \sum_k \sigma_k \text{tr}(E_k \rho),$$

where the σ_k are density matrices and the set $\{E_k\}$ forms a POVM.

Another property of these channels is that all entanglement-breaking channels are antidegradable [CRS08].

The distinguishability problem on quantum circuits, considered in Chapter 5, is defined in terms of distinguishability with access to a reference system. This method, distinguishing channels by observing their action on part of a larger space, is the most general method for distinguishing channels. There are channels known for which this reference system is required to obtain an optimal distinguishing strategy [Wat08]. This reference system allows for entangled inputs to aid in distinguishing the two channels and it appears to be essential to problem. It might be expected that this entanglement cannot help distinguish entanglement-breaking channels as the output is always separable, but this is not the case. An example on qubit channels has been provided by Sacchi [Sac05a, Sac05b]. When this example is generalized to channels on a d -dimensional space, however, the amount that this entanglement assists the distinguishability goes to zero quadratically with d . This is in contrast to the large difference that this reference system can make in the distinguishability of general channels. Examples of entanglement-breaking channels with this property are not known. Whether or not this is a roadblock to extending the hardness of the distinguishability problem to these channels is an interesting open problem.

It is simple to show that the problem of distinguishing two channels is **QIP**-hard for entanglement-breaking channels that are exponentially close together using a straightforward reduction from the general problem. This can be achieved by simply mixing the channels in a given instance of distinguishability with enough of the completely depolarizing channel that the resulting channels are entanglement-breaking. That this occurs is a consequence of the fact that there exists a ball of separable states around the completely mixed state [GB02]. Unfortunately this ball has radius that is exponentially small in the log of the dimension (i.e. the number of qubits in the original circuits), and so the resulting entanglement-breaking channels are exponentially close together. The polarization technique that can be applied in the general case, which is discussed in Section 3.7, cannot be applied to these circuits as they are too close together, and so this reduction can only be used to show the hardness of distinguishing circuits that are exponentially close together, which is perhaps not a terribly surprising result.

1.3.4 Unital channels

A superoperator $\Phi: \mathbf{L}(\mathcal{H}) \rightarrow \mathbf{L}(\mathcal{K})$ is *unital* if $\Phi(\mathbf{1}_{\mathcal{H}}) = \mathbf{1}_{\mathcal{K}}$. The unital channels are often called *doubly stochastic* as in addition to being unital they are also trace preserving. The trace preserving property of quantum channels requires that any unital channel have input and output spaces of the same dimension. The unital channels have the interesting property that the entropy of the output of the channel is always at least as large as the entropy of the input, as noted by King and Ruskai [KR01].

This property makes the unital channels interesting from the perspective of the additivity of the Holevo capacity, as channels that do not reduce entropy can be used as a natural noise model. Fukuda has shown how to construct a unital channel from a general channel, without changing the minimum output entropy or the maximum output p-norm [Fuk07]. This implies that for a specific class of channels the question of additivity can be restricted to a subclass of the unital channels.

Mendl and Wolf have recently characterized the unital channels as the quantum channels that can be decomposed into the affine combination of a set of unitary channels [MW09]. More explicitly, they have shown that a channel Φ is unital if and only if there exist unitaries U_i and $\lambda_i \in \mathbb{R}$ with $\sum_i \lambda_i = 1$ such that

$$\Phi(X) = \sum_{i=1} \lambda_i U_i X U_i^*.$$

The form of this decomposition is very similar to the next class of channels that we consider.

1.3.5 Mixed-unitary channels

A quantum channel is *mixed-unitary* if it can be decomposed into the probabilistic application of a set of unitary operations. These channels are often referred to as the *random unitary* channels, but this is avoided here because this name often causes confusion with the channels defined by drawing unitary operators from the Haar measure. More formally, Φ is mixed-unitary if there exist unitary operators U_1, \dots, U_n and a probability distribution p_1, \dots, p_n such that

$$\Phi(X) = \sum_{i=1}^n p_i U_i X U_i^*. \quad (1.9)$$

It has been shown by Gregoratti and Werner [GW03] that the mixed-unitary channels describe exactly the noise processes that can be corrected using classical information obtained by measuring the environment. Audenaert and Scheel have recently provided necessary and sufficient conditions for a channel to be mixed-unitary [AS08]. Buscemi has also provided an upper bound on the number of unitaries needed for a mixed-unitary decomposition [Bus06].

The set of mixed-unitary channels is contained in the set of all unital channels; this is a simple consequence of Equation (1.9). For channels on qubits these two sets of channels coincide, but for larger dimensions this is not the case [Tre86, KM87, LS93]. It is known, however, that there exists in the set of unital channels a ball of mixed-unitary channels around the completely depolarizing channel [Wat09a], which is the channel that maps all input states to the completely mixed state. In the case of qubit mixed-unitary channels, both additivity of the Holevo capacity and multiplicativity or the maximum output p-norm are known to hold [Kin02]. For general mixed-unitary channels, both additivity [Has09] and multiplicativity [HW08] are known to fail. These properties are considered in more detail in Chapter 3.

The fact that these properties do not hold in general for mixed-unitary channels does not completely eliminate interest in the additivity properties of specific mixed-unitary channels. One of the contributions of this thesis is a method to approximate a general quantum channel with a mixed-unitary one. This approximation can be made arbitrarily good in the minimum output entropy or the maximum output p-norm by increasing the dimension of the ancillary space used by the approximation. This can be used to show that the additivity and multiplicativity problems for a general channel can be reduced to the same problem on a mixed-unitary approximation, where the approximation error can be made arbitrarily small. Results on this approximation may be easier to prove; mixed-unitary channels have been essential to finding coun-

terexamples to both the additivity and multiplicativity conjectures. These results can then be applied to the original channel by sending the approximation error to zero.

The method for approximating a general channel by a mixed-unitary one is discussed in Chapter 7. Starting with a channel in Stinespring form $\Phi(X) = \text{tr}_{\mathcal{B}} U(X \otimes |0\rangle\langle 0|)U^*$ there are only two operations that are not mixed-unitary: the partial trace on the system \mathcal{B} and the introduction of the auxiliary system in the $|0\rangle$ state. The partial trace is the easy operation to simulate with a mixed-unitary channel as it may be directly replaced by the completely depolarizing channel on the space \mathcal{B} . The auxiliary system in the $|0\rangle$ state is more difficult to replace. The strategy employed is to add this extra system to the input space of the channel and test that the input in this space is close to $|0\rangle$. If this auxiliary input is close to $|0\rangle$ the channel proceeds exactly as does the original channel. If the auxiliary input is far from $|0\rangle$ then the testing procedure sends the input state very close to the maximally mixed state, which results in the output of the channel having very high entropy. As we are concerned with approximating the minimum output entropy this ensures that any input state achieving the minimum is very close to $|0\rangle$ in the auxiliary space. This construction produces a channel with similar minimum output entropy and maximum output p -norm to the original channel, and so it can be used to reduce problems of additivity and multiplicativity to the mixed-unitary case.

This construction can also be performed on circuits in time polynomial in the size of the circuit and so it also has implications for the problem of distinguishing quantum circuits. This can be used to show that the problem of distinguishing two mixed-unitary circuits is as hard as distinguishing two general circuits, which is a **QIP**-complete problem. This result is found in Chapter 7.

Chapter 2

Quantum Computational Complexity

This chapter lays the complexity theoretic groundwork for the remainder of the thesis. This includes a definition of the circuit model that is used throughout the thesis as well as a brief overview of some of the complexity classes that will be encountered later.

The circuit model used here is the mixed-state circuit model of Aharonov et al. [AKN98] that allows measurements and other non-unitary operations to take place during a computation. This model will be essential to the problems considered in the thesis: the distinguishability problem appears to be strictly more difficult on this circuit model than it is on the model of unitary circuits. This is despite the fact that both of these models are computationally equivalent, in the sense that any problem solvable by a circuit in one model can also be solved in the other. This equivalence does not extend to problems that take these circuits as input.

The wide array of complexity classes often encountered in theoretical computer science is not particularly useful or relevant to the results of the thesis. For this reason, the introduction of complexity classes is kept quite brief, with only a few of the most important classes introduced.

Contents

2.1	Quantum circuits	26
2.1.1	Unitary circuits	28
2.1.2	Mixed-state circuits	30
2.1.3	Short quantum circuits	33
2.2	Quantum complexity classes	35

2.1 Quantum circuits

Many questions on quantum channels can be extended to computational problems. This extension leaves one difficulty: what is the correct way to encode a quantum channel as input to a computational problem? One obvious choice is to provide the Kraus operators or the unitary matrix from a Stinespring dilation. Such a representation allows for any quantum channel to be represented approximately, as these matrices can only be specified up to some precision. Viewed computationally, however, this representation is unsatisfying. The reason for this is that the description of the channel is polynomial in the input and output dimensions, which are exponential in the number of qubits needed to represent the input and output. This representation is similar to modelling any classical process as a table of inputs and outputs – this form is convenient, but often exponentially larger than necessary.

Taking a hint from classical complexity theory, we will represent quantum channels using circuits. These circuits will allow for the simulation of a quantum channel from the circuit description, but they will not, in general, allow the efficient solution to most of the computational problems on these channels. This is equivalent to the classical case, where the circuit satisfiability problem is used to represent the problem of determining if a computation can be made to accept. Providing a complete table of outputs as the input to this problem trivializes it, as in the case of a polynomial size circuit, the table encodes the information in the circuit in an exponentially larger description. The problems on quantum computations that we consider in this thesis are similarly trivialized by a representation of quantum channels as matrices of size exponential in the number of input and output qubits.

The most widely used model of quantum computation is the unitary circuit model. In this model a computation is represented by a directed acyclic graph, where the edges represent qubits and the nodes represent gates. In order for a circuit to implement a valid quantum operation, each gate is labelled with a quantum channel that maps the state of the input qubits to the state of the output qubits. The operations that can appear as gates in a circuit depend on exactly which model of quantum computation is being used. As one final restriction, no isolated vertices in the graph are allowed, because these would correspond to gates in the circuit that neither take input nor produce output, and so they cannot affect the computation being performed.

There are two important quantities related to a circuit: *size* and *depth*. If a circuit is represented as a graph, the size of the circuit is the number of vertices, i.e. the number of gates in the circuit. This definition leaves the possibility of very small circuits acting on a large number of input qubits – this undesirable feature is avoided by taking




Figure 2.1: An example quantum circuit.

the size of a circuit to be the maximum of the number of gates and the number of qubits that the circuit acts on. Using this definition, the size of a circuit is essentially equivalent to the number of bits needed to represent the circuit, so long as the number of different types of gates available is constant.

The depth of a circuit is the length of the longest directed path in the graph. As circuits are acyclic, the depth of a circuit can be efficiently computed from a description. Since the transformations implemented by gates acting on different qubits commute, they can be performed in parallel. This implies that the depth of a circuit represents essentially the minimum amount of time used by an implementation of the circuit, provided that gates acting on disjoint sets of qubits can be performed in parallel.

As an example of size and depth, the circuit in Figure 2.1 takes four qubits as input, produces two qubits as output, and has size four and depth two. In this figure, and in all the circuit diagrams that will appear in the thesis, the gates in the circuit are represented by boxes and the edges by the lines connecting them. The edges in circuits are directed, but by convention, the edges in the diagrams appearing here are always directed from left to right, so that time flows left to right during the evaluation of the circuit. The circuit in the example maps four qubits to two qubits. This can be thought of as a channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$, where $\dim \mathcal{H} = 2^4 = 16$ and $\dim \mathcal{K} = 2^2 = 4$. An alternate view of this circuit is, where \mathcal{A} is a Hilbert space of dimension two, as a transformation in $\mathbf{T}(\mathcal{A}^{\otimes 4}, \mathcal{A}^{\otimes 2})$. We will take whichever view is most convenient, as these two sets of transformations are isomorphic.

As a further notational convenience, throughout the thesis, a circuit C will be identified with the transformation $C \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ that it implements, so that for a state $\rho \in \mathbf{D}(\mathcal{H})$, the state $C(\rho)$ is the output of the circuit when executed on the input state ρ . Each circuit specifies exactly one transformation, but the converse is not true: any quantum channel has (infinitely) many circuit implementations, and so given only a transformation, a circuit implementing it will need to be carefully constructed. In most of the cases we will encounter this is not difficult to do.

Circuits as a model for quantum computation are significantly easier to work with than earlier models of computation, such as the model of quantum Turing machines introduced in [BV97]. When the transformations that can be used as gates are restricted to the right set, it is known that the circuit model of computation is equivalent to the quantum Turing machine model [Yao93]. For this reason, we will use the circuit model of quantum computation, though for most of the results in this thesis the exact model of computation will not be important.

2.1.1 Unitary circuits

The most commonly used model of quantum circuits is the unitary circuit model. In this model every gate implements a unitary transformation on one or two qubits. Unitarity implies that all the gates of this model have the same number of input and output qubits.

It is known that any unitary operation can be approximately represented using only a finite set of one- and two-qubit unitary gates. The set of gates we will use is given in Figure 2.2, and the proof that it is (approximately) universal is due to Boykin et al. [BMP⁺00]. Different universality proofs for slightly different sets of gates can be found in [Sho96, KLZ98, ABO08]. An excellent overview of this and other universal sets of gates, as well a proof that the gate set used here is universal can be found in [NC00].

We have seen a few of these gates before: the Pauli X and Z and Hadamard gates simply apply the corresponding unitary operators to the qubits they act on, where the operators X, Z, and H are as defined in Section 1.2. A few of these operators are new. In matrix form, the swap, controlled-not (CNOT), and $\pi/8$ (T) gates are given by

$$W = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \quad T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}.$$

For the sake of convenience we have added a few gates to the circuit model. The Pauli X and Z and swap gates are not needed for a universal set of gates. They can, however, be constructed exactly using gates from the standard set. Two of these gates are simple to build from the standard set

$$Z = T^4, \quad X = HZH = HT^4H.$$




Figure 2.2: Gates in the unitary circuit model. The Pauli X and Z gates, and the swap gate are not required for universality, but they are included for convenience. The $\pi/8$ gate is needed for universality, but will not be used in any of the circuits constructed outside of this section.




Figure 2.3: Simulation of the swap gate with three controlled-not gates.

The third unnecessary gate, the swap gate, can be implemented in the standard model using no gates at all! This is because the unitary operation that swaps two qubits can be introduced into a circuit by simply redirecting the edges in the underlying graph. In many practical models of computation it is nontrivial to connect gates together in arbitrary directed graphs. One such model is the nearest-neighbour model, where a qubit can only interact with the qubits immediately adjacent to it. This model (with polynomial depth and size overhead) can simulate the more permissive model if the swap gate included in the circuit model, since the required qubits for any operation can always be swapped together. The swap gate can be implemented as a series of three controlled-not gates in this model, as shown in Figure 2.3. We will use W to represent the gate that swaps the input systems, so that $W|a\rangle|b\rangle = |b\rangle|a\rangle$ even when the dimension of the systems to be swapped is larger than two. This gate can be implemented using several two-qubit swap gates. The introduction of these gates does not change the circuit model, as they can be exactly implemented in the model of Boykin et al. [BMP⁺00] using a constant number of gates.




Figure 2.4: Controlled-U gate.

It is often very useful in a quantum circuit to control the application of some unitary operation based on the value of an additional qubit. For a unitary U , this is the operation commonly known as a controlled- U gate. We have already encountered one such gate: the controlled-not gate in the standard model is exactly the controlled application of the Pauli X gate. Given a unitary operation U (as a circuit), the controlled- U operation is the unitary operation that applies U if the control qubit is $|1\rangle$ and does nothing if the control qubit is in the $|0\rangle$ state. The representation of this gate in the circuit model is shown in Figure 2.4. For a unitary $U \in \mathbf{U}(\mathcal{H})$, this gate is represented in block matrix form as

$$\Lambda(U) = \begin{pmatrix} \mathbb{1}_{\mathcal{H}} & 0 \\ 0 & U \end{pmatrix}.$$

Given a circuit for U , it is simple to construct one for the controlled- U operation. Each gate in the circuit can be replaced by a controlled version, so that all of the gates are applied (i.e. U is applied) or none of the gates are applied. The controlled versions of each gate in the basis need to be constructed, but these are guaranteed to (approximately) exist by the fact that we are using a complete basis of gates. Notice, however, that this construction may add significantly to the depth of the circuit: the single control qubit is used many times. A more depth-conscious construction is presented in Section 2.1.3.

2.1.2 Mixed-state circuits

Circuits in the unitary model can (approximately) represent any unitary computation. If these circuits are allowed access to ancillary qubits in a known pure state, then they can perform any efficient quantum computation. There is, however, a drawback to this model: unitary circuits cannot simulate an arbitrary quantum channel. This is because a general completely positive and trace preserving operation may discard information, it may make measurements in the middle of a computation, or it may introduce ancillary qubits in a mixed state. Many of these operations are impossible to implement in the unitary model.

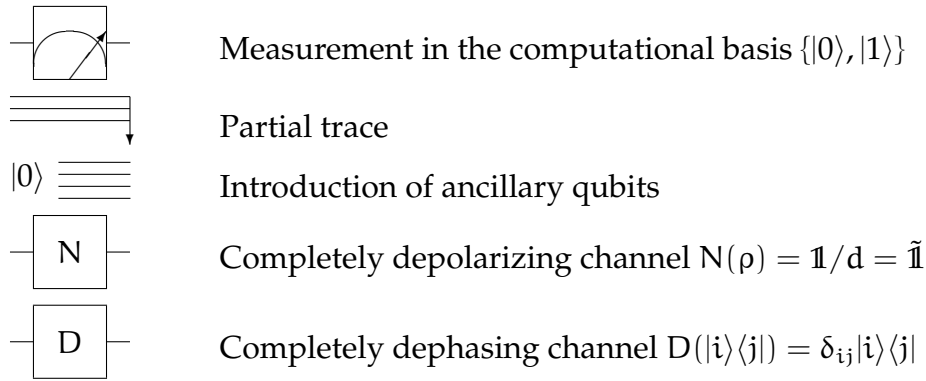


Figure 2.5: Non-unitary gates in the mixed state circuit model.




Figure 2.6: Simulations of the measurement gate, the completely dephasing gate, and the completely depolarizing gate with the other gates in the circuit model.

For this reason we will use the mixed-state circuit model of Aharonov, Kitaev, and Nisan [AKN98]. Circuits in this model can (approximately) represent any quantum channel. This can be thought of as a probabilistic model of quantum computation, as the state of the computation can be a mixed state, whereas the unitary model can be thought of as deterministic computation, since the state during the computation is always pure. The gates available in this model are the standard gates from the unitary model, as well as the additional gates shown in Figure 2.5.

As in the unitary model, we have included a few unnecessary gates for the sake of convenience. The only two gates that are actually required are the gate that introduces ancillary qubits in the $|0\rangle$ state and the gate that traces out a qubit. The gate that makes a measurement in the computational basis can be implemented using an ancillary qubit and a controlled not gate, as shown in Figure 2.6. The output of this measurement gate is viewed as a mixed quantum state in the following way. If the measurement outcome is 0 with probability p and 1 with probability $1 - p$, the outcome of the measurement gate is $p|0\rangle\langle 0| + (1 - p)|1\rangle\langle 1|$. This density matrix is diagonal, and so it may be thought




Figure 2.7: The unitary operation U simulates the admissible operation Φ .

of as a classical probability distribution, but there is no loss in generality in encoding this distribution as a mixed quantum state. Controlled operations will function in exactly the same way regardless of whether they are controlled classically or by the density matrix corresponding to the same probability distribution. This measurement gate, as it turns out, performs exactly the same transformation as the decoherence gate D included in the gate set. The decoherence gate in Figure 2.5 applies to systems of an arbitrary number of qubits, but this operation has the property that when applied individually to a number of qubits, the result is exactly the same as if it had been applied to all of them at once. The completely depolarizing channel also has this property, and so it is sufficient to give an implementation for the channels acting only on a single qubit, as is done in Figure 2.6.

Since the standard model of quantum computation is inherently probabilistic, as we will see in Section 2.2, it is not hard to show that the mixed-state model is equivalent in computational power to the unitary circuit model [AKN98]. The central idea behind the equivalence of the unitary and mixed-state models is the fact that any quantum channel can be implemented in *Stinespring form*, which is the introduction of ancillary qubits first, followed by a unitary operation on the now larger space, and finally tracing out any qubits that are not part of the output. For a channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$, this is exactly the Stinespring representation

$$\Phi(\rho) = \text{tr}_{\mathcal{B}} U(\rho \otimes |0\rangle\langle 0|)U^*$$

where the unitary U is implemented by a unitary circuit. An example of this is illustrated in Figure 2.7. This equivalence is noted in [AKN98], making use of what is known about quantum channels in the physics literature [Sti55, HK70].

Despite this computational equivalence these two models are not identical. The distinguishability problem discussed in Chapter 5 seems to be significantly harder than the distinguishability problem for unitary computations. If this is not the case, there are unexpected consequences in complexity theory [Vya03]. Even stronger evidence is provided by the close images problem studied in Chapter 4. This problem involves

determining the distance between the images of two transformations. If these transformations are unitary, their images always intersect, rendering the problem trivial. For these reasons, the standard model of quantum computations used throughout the thesis is the mixed-state circuit model. Any quantum channel that is given as the input of a computational problem will be in the form of a classical description of a circuit in this model.

2.1.3 Short quantum circuits

A significant challenge in the experimental realization of a quantum computation is the need to keep a quantum system from interacting with the environment. The decoherence caused by these interactions in practice provides a time limit for the computation. One way to ameliorate this difficulty is to find low-depth circuits that solve the problems we are interested in.

Short quantum circuits have been found for several important problems, such as the approximate quantum Fourier transform [CW00] and encoding and decoding operations for many error correcting codes [MN02]. These examples show the significant power of circuits that have depth logarithmic in the size of the circuits. More evidence for the power of short circuits is provided by Terhal and DiVincenzo [TD04] and improved by Fenner et al. [FGHZ05] who show that exactly computing the acceptance probabilities for constant-depth quantum circuits is as hard as simulating general quantum computation. Fenner et al. also show that, under certain restrictions, the acceptance probabilities for these circuits can be efficiently approximated.

The purpose of this section is to give a construction for the controlled version of a log-depth circuit on n qubits that results in a depth $O(\log n)$ circuit. It is not immediately clear how a controlled operation on n qubits, such as a controlled-swap operation can be performed in depth logarithmic in n . The straightforward implementation outlined in Section 2.1.1 requires using one control qubit to control each of the gates in the operation, resulting in a linear depth circuit. Moore and Nilsson [MN02] use a construction from reversible computing to reduce the depth of this technique.

Proposition 2.1 (Moore and Nilsson [MN02]). *Any log-depth operation on n qubits controlled by one qubit can be implemented in $O(\log n)$ depth with $O(n)$ ancillary qubits.*

Moore and Nilsson prove this only for the case of constant-depth operations, but the proof technique used also applies to the log-depth case. They prove this proposition using a tree of $\log n$ controlled-not operations to ‘duplicate’ the control qubit onto n




Figure 2.8: Log-depth implementation of controlled operation on n qubits

ancillary qubits. These copies only capture the information in the computational basis, but this is the exactly the same information that is used by the controlled gates. These extra control qubits can then be used to control the remaining operations, with each control qubit used a logarithmic number of times. Finally, the tree of controlled-not operations is reversed to clean up the ancillary qubits so that they can be traced out without decohering the system. This procedure is demonstrated in Figure 2.8. This implies, as an example, that the $2n$ -qubit controlled swap gate can be implemented in depth $O(\log n)$. This will be critical to the construction used in Chapter 4.

If the unbounded fan-out gate is allowed into the standard basis of gates, then the depth overhead added in this construction can be reduced to a constant. This gate performs a controlled-not operation from one control qubit to any number of target qubits in one computational step. This gate is not in the standard basis for mixed-state quantum computing: it requires a linear number of gates and a logarithmic depth circuit to implement in the standard gate model. Fan-out in classical circuits is simply the operation that copies the value from one bit to several other bits, and is often included in the standard circuit model. When such a gate is included in the usual quantum circuit models, many tasks become much simpler. As an example, this gate allows operations such as sorting, phase estimation, and the quantum Fourier transform to be approximated with constant depth circuits [HŠ05]. This gate will not generally be included in the model, but some of the results in the thesis can be strengthened when it is.

To see how the scheme for implementing controlled operations can be implemented in constant depth using this gate, notice that tree structure of Figure 2.8 can be replaced with a single fan-out gate. This allows n ‘copies’ of the control qubit to be created, which can then be used to control each of the n operations. A final application of the




Figure 2.9: Constant depth implementation of controlled operation on n qubits using the unbounded fan-out gate.

fan-out gate to restore the ancillary qubits to the $|0\rangle$ state. This is demonstrated in Figure 2.9, which implies the following proposition.

Proposition 2.2. *If the unbounded fan-out gate is in the basis of gates, any constant depth operation on n qubits controlled by one qubit can be implemented in $O(1)$ depth with $O(n)$ ancillary qubits.*

2.2 Quantum complexity classes

This section provides a brief overview of the quantum complexity classes related to the topic of this thesis. Many of the technical details related to the definitions of these classes are omitted, as a detailed understanding of complexity theory is not essential to the results that follow. For a more complete reference, see the recent survey of Watrous [Wat09b]. The known relationships between the classes discussed here and some of the more well known classical complexity classes are illustrated in Figure 2.10.

BQP, defined in [BV97], is the quantum complexity class of primary importance. This class is informally the set of all decision problems that are efficiently solvable with a quantum computer. As quantum computation can involve measurements that have inherently probabilistic outcomes, the quantum computation that solves a problem in **BQP** is permitted to fail with some bounded probability. This probability can be made arbitrarily small by using the standard trick of repeating the computation several times in parallel and taking the majority. Error reduction for this class is exactly as for probabilistic classical computations: this is because the inputs and outputs to the decision problems are classical strings that may be copied any number of times.




Figure 2.10: Known relationships between the major quantum and classical complexity classes. Classes are contained in the classes written above them. Only the containment $P \subsetneq EXP$ is known to be proper.

More formally, **BQP** is the set of all languages L for which there exists a uniform family Q of polynomial size quantum circuits, one for each input length, such that

1. if $x \in L$, then $\text{tr}(\Pi Q(x)) \geq \frac{2}{3}$,
2. if $x \notin L$, then $\text{tr}(\Pi Q(x)) \leq \frac{1}{3}$,

where Π is the projector onto the subspace where the first output qubit of Q is $|1\rangle$, i.e. the projector onto the accepting subspace for the circuit. The error bounds $2/3$ and $1/3$ here are not significant: they can be replaced with any $a > b$ that have at least an inverse polynomial gap between them (in the size of the input string x) as noted above.

Several **BQP**-complete promise problems are known. The most famous of these is probably the approximation of the Jones polynomial, which is in **BQP** by an algorithm of Aharonov et al. [AJL06] (or by earlier works of Freedman et al. [FKW02]) and is complete for **BQP** by a result of Freedman et al. [FLW02]. These problems give an important method for the study of quantum computation that is not necessarily connected to quantum information.

Extending the definition of **BQP** to include a single message from a computationally unbounded prover results in the class **QMA**, which is the quantum analogue of **NP**. This concept was first considered in [Kni96], first defined in [Kit99], and first studied in [Wat00]. **QMA** is the class of all problems that can be verified by a polynomial-time quantum verifier with access to a quantum proof. This proof is a quantum state on a polynomial number of qubits and may depend on the input. More formally, a language L is in **QMA** if there is a family of circuits Q such that

1. if $x \in L$, then there exists ρ such that $\text{tr}(\Pi Q(x, \rho)) \geq \frac{2}{3}$,
2. if $x \notin L$, then for any ρ , $\text{tr}(\Pi Q(x, \rho)) \leq \frac{1}{3}$,

where once again Π is the projector onto the accepting subspace of the output of Q . As in the case of **BQP** the error parameters $2/3$ and $1/3$ are not significant. Replacing these with a, b such that $|a - b|$ is at least inverse polynomial in n is also possible in this case, though the argument is not simple [KSV02, MW05].

Similar to **BQP**, the class **QMA** has complete promise problems. The simplest of these is the 2-local Hamiltonian problem, which is informally the quantum version of the 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 [KSV02]. The improvement of this result to the 2-local case is due to Kempe, Kitaev, and Regev [KKR06].

Extending **BQP** further by allowing multiple rounds of interaction with a prover results in the complexity class **QIP**, first defined in [Wat03]. This class is the quantum analogue of the classical class **IP**, which is equal to the more familiar class **PSPACE** [LFKN92, Sha92] of problems solvable with a polynomial amount of space. A recent result has also shown that **QIP = PSPACE** [JJUW09], resolving a major open problem in quantum computational complexity.

As a more formal definition, for any language $L \in \mathbf{QIP}$, there is a polynomial time quantum algorithm V , known as the verifier, that exchanges quantum messages with a prover P . Both the prover and the verifier receive the input string x before the start of the computation. The verifier's algorithm V must be generated from the input string x in polynomial time, but the prover's algorithm is not constrained in this way. Given a pair (V, P) the verifier V will accept that $x \in L$ with some probability after interacting with the prover P . An example of this interaction is shown in Figure 2.11, with the Hilbert spaces available to each party illustrated. For any input x , the verifier V in a **QIP** protocol satisfies

1. if $x \in L$, then there exists a prover P such that, (V, P) accepts with probability at least $\frac{2}{3}$.
2. if $x \notin L$, then for any prover P , (V, P) accepts with probability at most $\frac{1}{3}$.

Once again, the exact parameters used in this definition are not significant. It is known how to use parallel repetition in this model of computation for error reduction, so that as long as the probabilities are an inverse polynomial apart and the probability

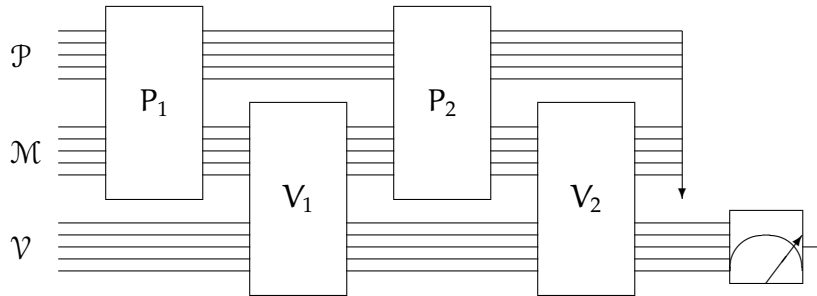


Figure 2.11: A three message quantum interactive proof system. The verifier's polynomial time transformations are V_1 and V_2 and the prover's transformations are given by P_1 and P_2 . All messages are sent through the message space \mathcal{M} and the verifier does not have access to the prover's private space \mathcal{P} . At the end of the interaction, a measurement of the verifier's private space \mathcal{V} determines the acceptance of the computation. All three of these spaces start in the $|0\rangle$ state. No restrictions are made on the size of the space \mathcal{P} , but the spaces \mathcal{M} and \mathcal{V} do not contain more than a polynomial number of qubits. The circuits V_1, V_2, P_1, P_2 may depend on the input x . The circuits V_1, V_2 must be generated from the input x in polynomial time, but the circuits P_1 , and P_2 are not so restricted.

of acceptance in condition 2 is nonzero, the resulting class of problems does not change [KW00].

An interesting property of quantum interactive proof systems is that any quantum interactive proof system can be simulated by one using only three messages [KW00]. This is in contrast to the classical case, where constant round proof systems seem to be much weaker than polynomial round proof systems. For this reason, we may assume that any problem in **QIP** has a proof system as shown in Figure 2.11, in which each of the prover and verifier each perform exactly two transformations, with the verifier acting last.

It is easy to see from the definitions that **QMA** \subseteq **QIP**, as interactive proofs with three messages can only be stronger than those using only one message. It is expected that this containment is strict: if not, unexpected things happen to classical complexity classes (it would imply that **PH** \subseteq **PP**) [Vya03]. There is, however, no proof that this cannot happen, as it would resolve a long-standing open problem complexity (showing that **NP** is properly contained in **EXP**). The case of two message quantum interactive proofs is even more interesting, as quite little is known about this class. It is known that the class of problems two message proofs is contained in **PSPACE** [JUV09], but this result has been subsumed by the result that **QIP** = **PSPACE** using similar

techniques [JJUW09].

The class **QIP** also has complete (promise) problems. The **CLOSE IMAGES** problem is the first such problem known. This problem is implicitly defined and shown to be complete for **QIP** in [KW00] where it is used to show that $\mathbf{QIP} \subseteq \mathbf{EXP}$. This problem is the focus of Chapter 4, where a formal definition can be found.

This thesis adds several new problems to the list of problems that are complete for **QIP**. **CLOSE IMAGES** is effectively a restatement of the acceptance conditions for a quantum interactive proof system, as we will see in Section 4.2, and so these new complete problems provide a method for studying **QIP** that involve quantum information that are not strongly tied to the model of quantum interactive proof systems.

The three quantum complexity classes **BQP**, **QMA**, and **QIP** are the only classes that will be encountered in this thesis. With the exception of Section 4.2 where **CLOSE IMAGES** is shown to be complete for **QIP** from the definition, it is not essential to have a deep understanding of these definitions. More important is to maintain the intuitive picture that problems in **BQP** are easy, problems in **QMA** are difficult, and problems in **QIP** are even more difficult.

Chapter 3

Measures for Quantum Information

Distances and other measures give a quantitative method to evaluate how close two states are together, how mixed a state is, or how well a quantum channel can be used to transmit information. These are all tasks that are central to the problems discussed in the thesis, and so we introduce several different techniques for measuring these quantities. It is important that there are several such measures, as most of these measures have an operational meaning that can help to ground the otherwise abstract problems that we consider.

The primary quantities discussed in this chapter include the entropy, the Schatten p -norms, and the trace norm on quantum states. Also included is an overview of some of the extensions of these quantities to the case of channels. A brief overview of the problems related to the additivity of the classical capacity is also provided.

This chapter is largely a collection of these measures, with proofs of the important properties that they have. The results in Sections 3.5.1 and 3.7, are the product of joint work with John Watrous [RW05], and the remainder of the results discussed here are not new and can be found in several of the standard sources.

Contents

3.1 Entropy	42
3.1.1 Minimum output entropy	45
3.2 Schatten p-norms	45
3.2.1 Maximum output p -norm	47
3.3 The classical capacity of a quantum channel	47
3.3.1 Relation to the minimum output entropy	49
3.3.2 Relation to the maximum output p -norm	51

3.3.3	Additivity and multiplicativity on classes of channels	52
3.4	The trace norm	53
3.5	The diamond norm	57
3.5.1	Maximization on a pure state for the difference of channels . . .	61
3.6	Fidelity	64
3.6.1	Relation to the trace norm	66
3.6.2	Maximum output fidelity for channels	68
3.7	Polarization of the diamond norm	69
3.8	Conclusion	76

3.1 Entropy

The entropy of a quantum state can be viewed as a measure of the amount of uncertainty about the value of the state. In support of this intuitive picture, the entropy of a pure state is zero as this represents the case where (in principle) complete knowledge of the state is present. The other extreme is the completely mixed state $\tilde{\mathbb{I}}$, where nothing at all is known about the system, which corresponds to the maximum entropy.

In the case of a classical probability distribution p , the entropy is defined to be

$$S(p) = - \sum_x p(x) \log p(x),$$

where the logarithm is taken base two. The entropy was first used in an information theoretic context by Shannon in 1948 [Sha48], who derived it from axioms that he felt that any such measure of uncertainty should satisfy. By convention $0 \log 0$ is defined to be 0 in this equation.

This quantity has a generalization to quantum systems that was first developed by von Neumann in 1927 [vN27]. This version of the entropy, applied to a density operator ρ with eigenvalues λ_i , is given by

$$S(\rho) = - \text{tr } \rho \log \rho = - \sum_i \lambda_i \log \lambda_i. \tag{3.1}$$

This quantity is often called the von Neumann entropy. In the case of a probability distribution p encoded as a diagonal matrix with diagonal entries $p(x)$, these two quantities agree, which is why the Shannon entropy is usually thought of as a special case of Equation (3.1).

Returning to the previous examples, notice that since a pure state expressed as a density operator has exactly one nonzero eigenvalue, the entropy is given by

$$S(|\psi\rangle\langle\psi|) = -1 \log 1 = 0.$$

In the case of the completely mixed state $\tilde{\mathbb{1}}_{\mathcal{H}}$ on a space \mathcal{H} of dimension d , there are d eigenvalues, each with value $1/d$, computing the entropy, we obtain

$$S(\tilde{\mathbb{1}}_{\mathcal{H}}) = - \sum_{i=1}^d \frac{1}{d} \log \frac{1}{d} = \log d.$$

A good reference on the properties of the entropy and many of the quantities derived from it can be found in [BŻ06] and [NC00]. The exposition of the entropy presented here follows these sources.

One property of the entropy that is useful is that it is additive with respect to the tensor product. To see this, let $\rho = \sum_i \lambda_i |\psi_i\rangle\langle\psi_i|$ and $\sigma = \sum_i \gamma_i |\phi_i\rangle\langle\phi_i|$ be two density operators. Expanding the definition of the entropy, we have

$$S(\rho \otimes \sigma) = - \sum_{ij} \lambda_i \gamma_j \log \lambda_i \gamma_j = - \sum_{ij} \lambda_i \gamma_j \log \lambda_i - \sum_{ij} \lambda_i \gamma_j \log \gamma_j = S(\rho) + S(\sigma), \quad (3.2)$$

where we have made use of the fact that since ρ, σ are density operators $\sum_i \lambda_i = \sum_j \gamma_j = 1$. This implies that for a multipartite quantum system that is not entangled, the entropy of the complete system can be determined locally. This is not true for entangled systems: if two parties each share half of a maximally entangled state, the local entropies are maximized, but the global state is pure, so it has zero entropy.

It is easy to see that the entropy is always nonnegative: the eigenvalues of a density matrix are always in the range $[0, 1]$ as all density matrices are by definition positive operators with unit trace. It is more difficult to see that on a Hilbert space of dimension d , no state can have entropy greater than $\log d$. One way to see this intuitively is to notice that since the logarithm is concave, Equation (3.1) is maximized when there are as many eigenvalues as possible, each of them being as small as possible. Formalizing this argument will require the use of Klein's inequality. The inequality stated here is a special case of Klein's 1931 result, but it is all that will be needed.

Theorem 3.1 (Klein's Inequality [Kle31]). *Let $\rho, \sigma \in \mathbf{D}(\mathcal{H})$, then*

$$\mathrm{tr}(\rho \log \rho) - \mathrm{tr}(\rho \log \sigma) \geq 0,$$

with equality only when $\rho = \sigma$.

This inequality immediately implies that the completely mixed state is the unique state with maximum entropy.

Proposition 3.2. *Let \mathcal{H} be a Hilbert space of dimension d , then for any $\rho \in \mathbf{D}(\mathcal{H})$*

$$0 \leq S(\rho) \leq \log d.$$

Furthermore, $S(\rho) = 0$ implies that ρ is a pure state, and $S(\rho) = \log d$ implies that $\rho = \tilde{\mathbb{1}}_{\mathcal{H}}$.

Proof. The first inequality is simple: if $\text{rank}(\rho) > 1$ then, by the strict concavity of the logarithm on positive values, $S(\rho) > 0$. On the other hand, if ρ is pure a simple calculation reveals that $S(\rho) = 0$.

The second inequality is a direct consequence of Klein's inequality (Theorem 3.1) applied to ρ and $\tilde{\mathbb{1}}_{\mathcal{H}}$:

$$0 \leq \text{tr}(\rho \log \rho - \rho \log \tilde{\mathbb{1}}_{\mathcal{H}}) = \log d - S(\rho).$$

This implies that $S(\rho) \leq \log d$, with equality only when $\rho = \tilde{\mathbb{1}}_{\mathcal{H}}$, by the equality condition of Klein's inequality. \square

Klein's inequality can be used to prove another important property of the entropy: *concavity*. This property is similar to the triangle inequality, except that the inequality goes in the opposite direction.

Proposition 3.3. *Let $\rho, \sigma, \xi \in \mathbf{D}(\mathcal{H})$, with $\rho = q\sigma + (1 - q)\xi$, where $0 \leq q \leq 1$. Then*

$$S(\rho) \geq qS(\sigma) + (1 - q)S(\xi).$$

Proof. Expanding the definition of S , we have

$$S(\rho) = -\text{tr} \rho \log \rho = -q \text{tr} \sigma \log \rho - (1 - q) \text{tr} \xi \log \rho. \quad (3.3)$$

Klein's inequality (Theorem 3.1) implies that $\text{tr} \sigma \log \sigma \geq \text{tr} \sigma \log \rho$, and similarly for ξ in place of σ . Together with Equation (3.3), this implies that

$$S(\rho) \geq -q \text{tr} \sigma \log \sigma - (1 - q) \text{tr} \xi \log \xi = qS(\sigma) + (1 - q)S(\xi),$$

which is the desired inequality. \square

3.1.1 Minimum output entropy

The entropy can be extended to quantum channels in a straightforward way: by minimizing the entropy over the output states of the channel. The resulting quantity is known as the *minimum output entropy*, which is defined for a transformation $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ by

$$S_{\min}(\Phi) = \min_{\rho \in \mathbf{D}(\mathcal{H})} S(\Phi(\rho)). \quad (3.4)$$

This extension of the entropy to quantum channels, as well as the properties of the entropy that have been demonstrated here will be essential to the results of Section 7.6, concerning the additivity of the minimum output entropy on the tensor product of two channels. This is closely related to the capacity of a quantum channel for the transmission of classical information. This is discussed in Section 3.3

3.2 Schatten p -norms

One of the more useful distance measures that can be defined on quantum states comes from the Schatten p -norms [Sch60]. This is the extension to operators in $\mathbf{L}(\mathcal{H}, \mathcal{K})$ of the usual ℓ_p norm of a sequence $\{x_i\}$, which for $1 \leq p < \infty$ is defined by

$$\|x\|_p = \left(\sum_i |x_i|^p \right)^{\frac{1}{p}}. \quad (3.5)$$

For $p = \infty$, this norm is given by $\|x\|_\infty = \sup_i |x_i|$, which can be obtained by taking the limit of Equation (3.5) as $p \rightarrow \infty$. The extension of this norm to an operator $A \in \mathbf{L}(\mathcal{H}, \mathcal{K})$ is done by taking this norm on the singular values of A , so that for $1 \leq p < \infty$

$$\|A\|_p = (\operatorname{tr} |A|^p)^{\frac{1}{p}} = \|s(A)\|_p, \quad (3.6)$$

where $s(A)$ is the (finite) sequence of singular values of A . The extension to the case $p = \infty$ is, as in the vector case, given by $\|A\|_\infty = \max_i s_i(A)$. Two of these norms are widely used in quantum information: the case $p = 1$ corresponds to the trace norm, considered in more detail in Section 3.4, and the case of $p = \infty$ corresponds to the usual operator norm on $\mathbf{L}(\mathcal{H}, \mathcal{K})$. This section discusses some of the most important properties of these norms. A more complete overview of the properties of these norms, as well as the properties of more general classes of norms, can be found in [Bha97].

Recall from Definition 1.1 that any norm satisfies the three properties of nonnegativity, homogeneity, and the triangle inequality. The first two of these properties

are easily to verified for these norms. Equation (3.6) is only zero when all the singular values are zero, which establishes nonnegativity (Equation (1.1)). Homogeneity (Equation (1.2)) follows directly from the definition of the absolute value of a matrix and the linearity of the trace.

Verifying the triangle inequality (Equation (1.3)) for this norm is nontrivial, and so only a brief overview of this result is presented here. The most important part of this proof is a 1951 result of Ky Fan [Fan51], which is a majorization relation on the singular values of $A + B$ in terms of the singular values of A and B . As is common in this thesis, the finite-dimensional result that is presented is a considerable simplification of the known result, which holds in the infinite dimensional case.

Theorem 3.4 (Ky Fan [Fan51]). *Let A, B be $n \times n$ matrices, and let $s(A)$ denote the sequence of singular values of A in decreasing order, then for all $k \in \{1, \dots, n\}$*

$$\sum_{i=1}^k s_i(A + B) \leq \sum_{i=1}^k s_i(A) + \sum_{i=1}^k s_i(B), \quad (3.7)$$

and more generally, if τ is any symmetric gauge function,

$$\tau(s(A + B)) \leq \tau(s(A)) + \tau(s(B)).$$

The triangle inequality for $\|\cdot\|_p$ in the cases that $p = 1, \infty$ is a direct consequence of Equation (3.7) for $k = n, 1$. The triangle inequality for the remaining values of p follow from Fan's theorem and the fact that for a vector $x \in \mathbb{R}^n$, the function $\tau(x) = (\sum_{i=1}^n |x_i|^p)^{1/p}$ is a *symmetric gauge function*. This is a strengthened version of the property that the function $\tau(\cdot)$ is a norm (in the case, the ℓ_p norm). More details, as well as detailed arguments that $\|\cdot\|_p$ is a norm using the theory of symmetric gauge functions can be found in [Bha97, HJ91].

The p -norm satisfies two more properties that will be essential to the results in Chapter 7. The first of these is *unitary invariance*. This is the property that for any unitary operators U, V

$$\|UAV\|_p = \|A\|_p. \quad (3.8)$$

It is easy to prove that this property holds. Consider a singular value decomposition of A given by $\sum_i s_i |\phi_i\rangle\langle\psi_i|$, then a singular value decomposition for UAV is given by

$$\sum_i s_i (U|\phi_i\rangle)(\langle\psi_i|V),$$

from which it can be seen that both A and UAV have the same singular values. Then, as the p -norm is defined in Equation (3.6) solely in terms of the singular values of A , the p -norms of A and UAV must be identical.

The second important property of the p -norm is that it is *multiplicative* with respect to the tensor product of two operators. That is, for operators A, B

$$\|A \otimes B\|_p = \|A\|_p \|B\|_p. \quad (3.9)$$

This property follows directly from the properties of the tensor product. Let singular value decompositions of A and B be given by $A = \sum_i s_i |\phi_i\rangle\langle\psi_i|$ and $B = \sum_i t_i |\gamma_i\rangle\langle\nu_i|$, then the a singular value decomposition of $A \otimes B$ is

$$\left(\sum_i s_i |\phi_i\rangle\langle\psi_i| \right) \otimes \left(\sum_j t_j |\gamma_j\rangle\langle\nu_j| \right) = \sum_{ij} s_i t_j |\phi_i\rangle\langle\psi_i| \otimes |\gamma_j\rangle\langle\nu_j|, \quad (3.10)$$

so that the singular values of the tensor product $A \otimes B$ are simply the products of the singular values of A and B . From this relationship Equations (3.5) and (3.6) imply the desired property.

3.2.1 Maximum output p -norm

These norms have an important extension to channels. This is the *maximum output p -norm*, which, for a channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ is denoted $\nu_p(\Phi)$, or sometimes simply $\|\Phi\|_p$. For $1 \leq p \leq \infty$, this norm is given by

$$\nu_p(\Phi) = \max_{\rho \in \mathbf{D}(\mathcal{H})} \|\Phi(\rho)\|_p. \quad (3.11)$$

This quantity is normally defined by taking the maximum over all inputs $X \in \mathbf{L}(\mathcal{H})$ with $\|X\|_1 = 1$, but a result of Amosov and Holevo [AH03] implies that in the case of Φ completely positive, this maximization can be restricted to the density operators. Note that this simplification cannot be made in the case of the *difference* of two channels, as the resulting operation is not completely positive, by a counterexample found in [Wat05]. In this thesis the maximum output p -norm will only be applied to channels, never the difference of two channels, so the simplification of the definition in Equation (3.11) is justified.

In the next section, this quantity is related to the capacity of a quantum channel for the transmission of classical information, the exact specification of which is currently an important problem in quantum information.

3.3 The classical capacity of a quantum channel

The additivity of the capacity for a quantum channel to communicate classical information is one of the most important unresolved problems in quantum information

theory. Informally, the additivity problem is: given two uses of a quantum channel, is it possible to send more than twice the classical information that could be sent with a single use? This is a common oversimplification; the classical capacity is defined in terms of the average amount of information sent per channel use, asymptotically with the number of uses of the channel [Hol98, SW97]. A more correct statement of the problem is: when encoding for the transmission of classical information, is entanglement across multiple uses of the channel necessary to achieve the best communication rate? This refined problem stood open for over 10 years before a counterexample was recently found by Hastings [Has09].

In this section this problem is given a formal definition and the closely related problems of the additivity of the minimum output entropy and the multiplicativity of the maximum output p -norm are discussed. The minimum output entropy and maximum output p -norm both involve maximizing the purity of the output of a channel, a problem that is intuitively related to the classical capacity by the notion that a channel that is less noisy should be able to send more information. A recent survey of these problems can be found in [Hol06], though it does not include the recent counterexamples to both additivity [Has09] and the related problem of the multiplicativity of the maximum output p -norm [HW08].

The classical capacity of channel Φ , when the input to multiple uses of the channel is restricted to product states, is given by the χ -capacity [Hol98, SW97]

$$C_\chi(\Phi) = \max \left[S\left(\sum_i p_i \Phi(\rho_i)\right) - \sum_i p_i S(\Phi(\rho_i)) \right], \quad (3.12)$$

where the maximum is taken over all convex mixtures $\sum_i p_i \rho_i$ of quantum states. This quantity is also referred to as the Holevo capacity or the “one-shot” or “one-step” capacity of Φ . The question of the additivity of this quantity, i.e. can entangling inputs across multiple uses of the channel be required to increase the capacity, was first raised in [BFS97], and the until recently standing conjecture was that

$$C_\chi(\Phi \otimes \Psi) \stackrel{?}{=} C_\chi(\Phi) + C_\chi(\Psi). \quad (3.13)$$

This is the statement that entangled inputs do not increase the classical information carrying capacity of quantum channels. This conjecture has recently been shown false in general by a result on the additivity of the minimum output entropy [Has09]. This result implies that the maximum rate that classical information can be transmitted using a channel Φ is given by

$$C(\Phi) = \lim_{n \rightarrow \infty} \frac{1}{n} C_\chi(\Phi^{\otimes n}).$$

If it were that case that C_χ were additive, this formula would simplify to $C(\Phi) = C_\chi(\Phi)$, but it is now known that this is not the case. This leaves open the question on many restricted classes of channels: a survey of some of these special cases can be found in [Hol06].

The χ -capacity captures exactly the amount of classical information that can be transmitted per use of the channel when encoding with product states, but it is somewhat awkward to work with. In the effort to resolve the additivity question it has been related to both the minimum output entropy and the maximum output p -norm.

3.3.1 Relation to the minimum output entropy

The minimum output entropy, defined by Equation (3.4), is simpler and often easier to work with than the χ -capacity. The additivity of this quantity, given by

$$S_{\min}(\Phi \otimes \Psi) \stackrel{?}{=} S_{\min}(\Phi) + S_{\min}(\Psi), \quad (3.14)$$

was first conjectured by King and Ruskai [KR01], who attribute the conjecture to Shor. The additivity of this quantity is connected to the additivity of the χ -capacity by a result of Shor [Sho04] that shows that both of these conjectures are globally equivalent to a third conjecture: the strong superadditivity of the entanglement of formation. Hastings has recently given a probabilistic construction that shows that this conjecture is false in general [Has09], which also implies the non-additivity of the χ -capacity.

One direction of Shor's construction in [Sho04] to show that the additivity of C_χ is equivalent to the additivity of S_{\min} is very complicated, but the other direction is quite simple.

Theorem 3.5 (Shor [Sho04]). *The additivity of C_χ implies the additivity of S_{\min} .*

Proof. Let Φ_1, Φ_2 be arbitrary channels in $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$. We will construct channels Φ'_1, Φ'_2 in the larger space $\mathbf{T}(\mathcal{C} \otimes \mathcal{H}, \mathcal{K})$. The channel Φ'_i uses the input space $\mathcal{C} \cong \mathcal{K} \times \mathcal{K}$ to determine which of the discrete Weyl operators to apply to the output of Φ_i , i.e.

$$\Phi'_i(|a, b\rangle\langle a, b| \otimes \rho) = W_{a,b} \Phi_i(\rho) W_{a,b}^*,$$

where the unitaries $W_{a,b}$ are the discrete Weyl operators introduced in Section 1.2. This process is applied decoherently, which is to say that Φ'_i measures the space \mathcal{C} in the computational basis to decide which operator to apply. As we will later show that the uniform mixture of the discrete Weyl operators forms the completely depolarizing

channel (Proposition 7.2), the result of placing a completely mixed state in the input space \mathcal{C} results in the completely mixed state as output, which is

$$\Phi'_i(\tilde{\mathbb{1}}_{\mathcal{C}} \otimes \rho) = \tilde{\mathbb{1}}_{\mathcal{K}}.$$

Let ρ_1 and ρ_2 be states achieving the minimum output entropy for Φ_1 and Φ_2 , respectively. We will assume that C_{χ} is additive, and show that if S_{\min} is not additive on $\Phi_1 \otimes \Phi_2$ then we can find a contradiction. To do this we compute $C_{\chi}(\Phi'_i)$ with the maximization in Equation (3.12) restricted to inputs of the form $\tilde{\mathbb{1}}_{\mathcal{C}} \otimes \rho_i$, which is given by

$$\begin{aligned} S(\tilde{\mathbb{1}}_{\mathcal{K}}) - \frac{1}{\dim \mathcal{C}} \sum_{a,b} S(W_{a,b} \Phi_i(\rho_i) W_{a,b}^*) &= \log \dim \mathcal{K} - \frac{1}{\dim \mathcal{C}} \sum_{a,b} S(\Phi_i(\rho_i)) \\ &= \log \dim \mathcal{K} - S(\Phi_i(\rho_i)) \\ &= \log \dim \mathcal{K} - S_{\min}(\Phi_i). \end{aligned} \quad (3.15)$$

Notice that this is the optimal value of $C_{\chi}(\Phi'_i)$, since the first term is maximized and the second term is minimized. This implies that restricting to inputs of the form $\tilde{\mathbb{1}}_{\mathcal{C}} \otimes \rho_i$ does not reduce the value of $C_{\chi}(\Phi'_i)$. As we have assumed that C_{χ} is additive, this expression also gives the optimal value of $C_{\chi}(\Phi'_1 \otimes \Phi'_2)$. However, if $S_{\min}(\Phi_1 \otimes \Phi_2)$ is not additive, then any state σ on which

$$S((\Phi_1 \otimes \Phi_2)(\sigma)) < S_{\min}(\Phi_1) + S_{\min}(\Phi_2)$$

can be used to increase $C_{\chi}(\Phi'_1 \otimes \Phi'_2)$, exactly as in the derivation of Equation (3.15), contradicting the (assumed) additivity of C_{χ} . \square

This result of Shor [Sho04], coupled with Hastings' results [Has09] shows that C_{χ} is non-additive in general. As the proof is constructive, it also shows that if S_{\min} is not additive on a class of channels, then C_{χ} is also not additive for the related class of channels given by applying the construction in the proof to all the channels in the class. This result is included here as it is of a similar flavour to many of the results in the thesis: reducing one problem to another, by embedding arbitrary channels into channels with specific properties, is a powerful method for obtaining results. In fact, Fukuda has used the same construction to show that the additivity of the Holevo capacity and the minimum output entropy can be restricted to the unital channels without loss of generality [Fuk07]. As in the case of the χ -capacity, these non-additivity results leave open the question of on which classes of channels does the additivity of the minimum output entropy hold? This is currently an important open problem in quantum information, as a deeper understanding of the channels for which the minimum output entropy is not additive may lead to a deeper understanding of those channels for which C_{χ} is not additive.

3.3.2 Relation to the maximum output p-norm

In an effort to better understand the question of the additivity of the minimum output entropy, yet another question has been raised. This problem is the multiplicativity of the maximum output p-norm, which was first conjectured by Amosov, Holevo, and Werner [AHW00]. The conjecture corresponding to this quantity was that it is multiplicative with respect to the tensor product of two channels, i.e. that

$$\nu_p(\Phi \otimes \Psi) \stackrel{?}{=} \nu_p(\Phi)\nu_p(\Psi). \quad (3.16)$$

This conjecture is not true in general: channels can be constructed for any fixed $p > 1$ that falsify this conjecture [HW08].

Amosov, Holevo, and Werner have related the multiplicativity of ν_p to the additivity of S_{\min} [AHW00]. This relationship can be used to show that the minimum output entropy is additive for a pair of channels if and only if the maximum output p-norm is additive on the same channels, for values of p close to 1.

Theorem 3.6 (Amosov, Holevo, Werner [AHW00]). *Let $\Phi_1, \Phi_2 \in \mathbf{T}(\mathcal{H}, \mathcal{K})$. If for some sequence $p_i \rightarrow 1$, with $p_i \geq 1$ for all i , it holds that*

$$\nu_{p_i}(\Phi_1 \otimes \Phi_2) = \nu_{p_i}(\Phi_1)\nu_{p_i}(\Phi_2),$$

then the minimum output entropy satisfies

$$S_{\min}(\Phi_1 \otimes \Phi_2) = S_{\min}(\Phi_1) + S_{\min}(\Phi_2)$$

Proof. Let the sequence $\{p_i\}$ be as in the statement of the theorem. The result is easy to verify by introducing the quantum Rényi entropy of order p of a density matrix σ

$$R_p(\sigma) = \frac{1}{1-p} \log \text{tr } \sigma^p. \quad (3.17)$$

The important property is that taking the limit as $i \rightarrow \infty$ (so that $p \rightarrow 1$ from above) results in the usual entropy (up to a factor due to the base of the logarithm)

$$\lim_{i \rightarrow \infty} R_{p_i}(\sigma) = (\log e)S(\sigma).$$

This can be verified by invoking l'Hopital's rule on Equation (3.17). By observing that the logarithm of the p-norm can be used to recover R_p , we may conclude from the previous equation that

$$\lim_{i \rightarrow \infty} \frac{p_i}{1-p_i} \log \|\sigma\|_{p_i} = \lim_{i \rightarrow \infty} R_{p_i}(\sigma) = (\log e)S(\sigma).$$

From which the multiplicativity of the maximum output p -norm immediately implies the additivity of the minimum output entropy. This follows from the fact that $\log \nu_{p_i}(\Psi) \leq 0$ as $\nu_{p_i}(\Psi) \leq 1$ for any channel Ψ . More concretely, we have

$$\begin{aligned} S_{\min}(\Phi_1 \otimes \Phi_2) &= \frac{1}{\log e} \lim_{i \rightarrow \infty} \frac{p_i}{1 - p_i} \log \nu_{p_i}(\Phi_1 \otimes \Phi_2) \\ &= \frac{1}{\log e} \lim_{i \rightarrow \infty} \frac{p_i}{1 - p_i} \log \nu_{p_i}(\Phi_1) + \frac{1}{\log e} \lim_{i \rightarrow \infty} \frac{p_i}{1 - p_i} \log \nu_{p_i}(\Phi_2) \\ &= S_{\min}(\Phi_1) + S_{\min}(\Phi_2), \end{aligned}$$

as desired. □

The minimum output entropy and the maximum output p -norm will be encountered again in Chapter 7, where it is shown that the additivity (respectively multiplicativity) of a channel can be rephrased in terms of the additivity (multiplicativity) of a related set of mixed-unitary channels.

3.3.3 Additivity and multiplicativity on classes of channels

The questions of the additivity of the Holevo capacity (also called the χ -capacity) and the multiplicativity of the maximum output p -norm have been resolved for many of the restricted classes of channels that are studied in this thesis.

It is shown in [CRS08] that the additivity of degradable channels is equivalent to the additivity of general channels, using a result from [FW07]. Combining this result with the result that the additivity problems are equivalent on a class of channels and the class of complementary channels [Hol07, KMNR07] shows that the additivity of the antidegradable channels is also equivalent to the general case. The recent result of Hastings [Has09] can then be used to show that there exist degradable and antidegradable channels that are not additive. It is perhaps not a surprise that these channels are well-behaved: the degradable and antidegradable channels cannot be used to transmit quantum information to the environment and receiver, respectively, because to do so would violate the principle of no-cloning.

It is also perhaps not a surprise that on the entanglement breaking channels the minimum output entropy is additive [Sho02] and the maximum output p -norm is multiplicative [Kin03]. The problem of distinguishing channels of this class, however, is quite interesting and remains open.

The unital channels cannot decrease the entropy [KR01]. This property makes them interesting from the perspective of additivity, as channels that do not reduce entropy

would seem to be a natural noise model. Fukuda has shown how to construct a unital channel from a general channel, without changing the minimum output entropy or the maximum output p-norm [Fuk07], using a the same construction used by Shor to show that the additivity of S_{\min} implies the additivity of C_X [Sho04]. This implies that for a set of channels the question of additivity can be rephrased in terms of the additivity of a related set of unital channels.

The mixed-unitary channels are a subclass of the unital channels. In the case of qubit mixed-unitary channels, both additivity and multiplicativity are known to hold [Kin02]. For general mixed-unitary channels, both additivity [Has09] and multiplicativity [HW08] are known to fail. In fact, all of the recent counterexamples to additivity and multiplicativity are obtained by choosing a random mixed-unitary channel from some distribution. This makes these channels very interesting. It is shown in Chapter 7 that the additivity or multiplicativity for a general channel can be reduced to the approximate additivity or multiplicativity of a mixed-unitary channel.

3.4 The trace norm

The trace norm is perhaps the most important measure of size and distance in quantum information. The trace norm of the difference of two states measures how distinguishable the two states are, which makes this an essential quantity for the problems considered in this thesis. The remainder of this section surveys some properties of this norm, further background on the trace norm can be obtained in the books [NC00] and [Bha97].

We have already encountered the trace norm: it is simply the $p = 1$ case of the Schatten p-norm discussed in Section 3.2. This implies that $\|X\|_{\text{tr}}$ is given by the sum of the singular values of X , though it is often useful to define this norm by an explicit formula. Such a formula can be obtained by noticing that the singular values of X are exactly the square roots of the eigenvalues of the positive operator X^*X . Using this observation, the trace norm can equivalently be defined as

$$\|X\|_{\text{tr}} = \|X\|_1 = \text{tr} \sqrt{X^*X}. \quad (3.18)$$

The trace norm inherits many properties from the p-norm. The triangle inequality is the $k = n$ case of Fan's Theorem (3.4), and unitary invariance is given by Equation (3.8). One other convenient property is that $\|\rho\|_{\text{tr}} = 1$ for any density matrix ρ , which is implied by the fact that density operators are positive semidefinite operators with unit

trace, so that the eigenvalues are all positive and sum to one. Several other properties of the trace norm can be easily derived from the following characterization.

Lemma 3.7. *For any $X \in \mathbf{L}(\mathcal{H})$,*

$$\|X\|_{\text{tr}} = \max_{U \in \mathbf{U}(\mathcal{H})} |\text{tr} XU|$$

Proof. Let X have a singular value decomposition given by $X = \sum_i s_i |\phi_i\rangle\langle\psi_i|$, where $\{|\phi_i\rangle\}$ and $\{|\psi_i\rangle\}$ are orthonormal bases for \mathcal{H} . Then for any unitary $U \in \mathbf{U}(\mathcal{H})$

$$|\text{tr} XU| = \left| \sum_i s_i \langle\psi_i|U|\phi_i\rangle \right| \leq \sum_i s_i |\langle\psi_i|U|\phi_i\rangle| \leq \sum_i s_i = \|X\|_{\text{tr}}. \quad (3.19)$$

If the unitary U is chosen such that $U|\phi_i\rangle = |\psi_i\rangle$, then

$$\langle\psi_i|U|\phi_i\rangle = \langle\psi_i|\psi_i\rangle = 1,$$

for all i . In this case equality is achieved in Equation (3.19). \square

From this characterization it is easy to see that the trace norm does not increase under the partial trace, and in fact, does not increase under the application of any channel. This is intuitively obvious: applying any potentially noisy operation to two states cannot help to distinguish them.

Theorem 3.8. *Let $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$, then for any $\rho, \sigma \in \mathbf{D}(\mathcal{H})$*

$$\|\Phi(\rho) - \Phi(\sigma)\|_{\text{tr}} \leq \|\rho - \sigma\|_{\text{tr}}.$$

Proof. We first prove this for the case that $\Phi \in \mathbf{T}(\mathcal{K} \otimes \mathcal{B}, \mathcal{K})$ is the partial trace over the space \mathcal{B} . For any $\rho, \sigma \in \mathbf{D}(\mathcal{K} \otimes \mathcal{B})$, this follows directly from Lemma 3.7, since

$$\|\text{tr}_{\mathcal{B}} \rho - \text{tr}_{\mathcal{B}} \sigma\|_{\text{tr}} = \max_{U \in \mathbf{U}(\mathcal{K})} |\text{tr}[(\sigma - \rho)(U \otimes \mathbf{1}_{\mathcal{B}})]| \leq \max_{U \in \mathbf{U}(\mathcal{K} \otimes \mathcal{B})} |\text{tr}(\sigma - \rho)U| = \|\rho - \sigma\|_{\text{tr}}.$$

To see the general case, let $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ have Strinespring representation given by $\Phi(\rho) = \text{tr}_{\mathcal{B}} U(\rho \otimes |0\rangle\langle 0|)U^*$. Then the previous equation and the unitary invariance of the trace norm imply that

$$\begin{aligned} \|\Phi(\rho) - \Phi(\sigma)\|_{\text{tr}} &= \|\text{tr}_{\mathcal{B}} U[(\rho - \sigma) \otimes |0\rangle\langle 0|] U^*\|_{\text{tr}} \\ &\leq \|U[(\rho - \sigma) \otimes |0\rangle\langle 0|] U^*\|_{\text{tr}} \\ &= \|(\rho - \sigma) \otimes |0\rangle\langle 0|\|_{\text{tr}} \\ &= \|\rho - \sigma\|_{\text{tr}}, \end{aligned}$$

as required. \square

The following theorem due to Helstrom [Hel67] formalizes the notion that the trace norm of the difference of two density matrices represents how well they can be distinguished by a measurement. This result underlies the definition of the quantum circuit distinguishability problem in Chapter 5, as the problem of distinguishing channels is exactly the problem of distinguishing the outputs of the channels. This result is easy to generalize to the case that the two density matrices are not chosen with equal probabilities, but doing so unnecessarily complicates the argument.

Theorem 3.9 (Helstrom [Hel67]). *The optimal probability that an unknown state $\xi \in \mathbf{D}(\mathcal{H})$ that is chosen uniformly at random from the set $\{\rho, \sigma\}$ can be correctly identified is given by*

$$\frac{1}{2} + \frac{\|\rho - \sigma\|_{\text{tr}}}{4}.$$

Proof. The optimal strategy consists of some two-outcome POVM measurement. By Naimark's theorem it may be assumed that the optimal measurement is a projective measurement Π_ρ, Π_σ with $\Pi_\rho + \Pi_\sigma = \mathbb{1}_{\mathcal{H}}$, since the operation that embeds ρ and σ into a larger space is an isometry, which will not affect the trace norm, by unitary invariance. The probability that this measurement succeeds is

$$p_{\text{succ}} = \frac{1}{2} \text{tr}(\Pi_\rho \rho) + \frac{1}{2} \text{tr}(\Pi_\sigma \sigma).$$

Similarly, the probability of failure is

$$p_{\text{fail}} = \frac{1}{2} \text{tr}(\Pi_\sigma \rho) + \frac{1}{2} \text{tr}(\Pi_\rho \sigma).$$

Subtracting the probability of failure from the probability of success gives the bound

$$p_{\text{succ}} - p_{\text{fail}} = \frac{1}{2} \text{tr}((\Pi_\rho - \Pi_\sigma)(\rho - \sigma)) \leq \frac{1}{2} \|\rho - \sigma\|_{\text{tr}}, \quad (3.20)$$

where the inequality follows from Lemma 3.7 by the fact that $\Pi_\rho - \Pi_\sigma$ is a unitary operator. Adding this equation to the equation $p_{\text{succ}} + p_{\text{fail}} = 1$ results in the bound

$$2p_{\text{succ}} \leq 1 + \frac{1}{2} \|\rho - \sigma\|_{\text{tr}}. \quad (3.21)$$

This is the probability given in the statement of the theorem, and so it remains only to show that it can be achieved.

To see this, consider the projectors Π_+ and Π_- that project onto the positive and non-positive eigenspaces of $\rho - \sigma$. Re-examining Equation (3.20) with this measurement results in

$$p_{\text{succ}} - p_{\text{fail}} = \frac{1}{2} \text{tr}((\Pi_+ - \Pi_-)(\rho - \sigma)) = \frac{1}{2} \text{tr}|\rho - \sigma| = \frac{1}{2} \|\rho - \sigma\|_{\text{tr}},$$

since the eigenvalues of the Hermitian operator $(\Pi_+ - \Pi_-)(\rho - \sigma)$ are the absolute values of the eigenvalues of $\rho - \sigma$. This demonstrates that this measurement achieves the bound in Equation (3.21). \square

There is one final property of the trace norm that is needed in Section 3.5.1 and Chapter 5. This property relates the trace norm of an operator to a Hermitian block matrix that is closely related to it. This construction is often useful: it is one way to take a general linear operator and construct a Hermitian operator on a space including one additional qubit. The proof of this relationship is not difficult or particularly illuminating, but it is included as this result is critical to some of the proofs that follow.

Lemma 3.10. *Let $A \in \mathbf{L}(\mathcal{H})$ be a linear operator, then*

$$\| |0\rangle\langle 1| \otimes A + |1\rangle\langle 0| \otimes A^* \|_{\text{tr}} = 2 \|A\|_{\text{tr}}.$$

Proof. Let r be the rank of A and let $n = \dim \mathcal{H}$. If $r = 0$ the result is trivial, so we may assume that $r > 0$. Let $\hat{A} = |0\rangle\langle 1| \otimes A + |1\rangle\langle 0| \otimes A^*$. Written as a block matrix, this operator is

$$\hat{A} = \begin{pmatrix} 0 & A \\ A^* & 0 \end{pmatrix}.$$

For the evaluation of the trace norm of \hat{A} it suffices to consider the eigenvalues, as this operator is Hermitian by construction. To compute these eigenvalues, let A have singular value decomposition $A = \sum_{i=1}^r s_i |\psi_i\rangle\langle \phi_i|$, where $\{|\psi_i\rangle\}$ and $\{|\phi_i\rangle\}$ are orthonormal sets of vectors. Let the notation $[\phi, \psi]^\top$ denote the vector of length $2n$ whose first n entries are the entries of ϕ and whose second n entries are the entries of ψ . As observed in [Bha97, Section II.1], the $2r$ nonzero eigenvalues of \hat{A} are given by

$$\begin{pmatrix} 0 & A \\ A^* & 0 \end{pmatrix} \begin{pmatrix} \phi_i \\ \psi_i \end{pmatrix} = s_i \begin{pmatrix} \phi_i \\ \psi_i \end{pmatrix} \quad \begin{pmatrix} 0 & A \\ A^* & 0 \end{pmatrix} \begin{pmatrix} -\phi_i \\ \psi_i \end{pmatrix} = -s_i \begin{pmatrix} -\phi_i \\ \psi_i \end{pmatrix},$$

for $i \in \{1, \dots, r\}$. This implies that

$$\|\hat{A}\|_{\text{tr}} = 2 \sum_{i=1}^r |s_i| = 2 \|A\|_{\text{tr}},$$

as desired. \square

The trace norm can be extended from states to channels, though some care must be taken in doing so to ensure that the resulting norm retains the desirable properties of the trace norm, such as the relationship to the optimal distinguishing probability. This extension is the focus of the next section.

3.5 The diamond norm

In this section the diamond norm is introduced and studied. This norm defines the distance measure that is central to the results that follow on the distinguishability of channels. The norm is introduced and some of the basic properties of the norm that can be found in the literature are discussed. Further background on the diamond norm can be found in [Kit97, AKN98] and [KSV02].

The diamond norm is a norm on channels with similar properties to the trace norm on quantum states. It will give a numerical value to the distance between quantum channels, and, more importantly, as in the case of the trace norm, it will be closely related to how well two channels can be distinguished.

The straightforward way to extend the trace norm to quantum channels is to do as we have done for the entropy: optimize the output over all input states. Doing this for the trace norm results in the norm given defined by

$$\|\Phi\|_{\text{tr}} = \max_{X \in \mathcal{L}(\mathcal{H}), X \neq 0} \frac{\|\Phi(X)\|_{\text{tr}}}{\|X\|_{\text{tr}}}. \quad (3.22)$$

This results in a norm, as it inherits many of the properties of the trace norm directly. As an example, it is easy to see that this norm obeys the triangle inequality. One of the other properties inherited by this norm is *submultiplicativity*. This property will be useful, and so a short proof of this fact is given below.

Lemma 3.11. *For any $\Phi \in \mathbf{T}(\mathcal{K}, \mathcal{F})$ and any $\Psi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$,*

$$\|\Phi \circ \Psi\|_{\text{tr}} \leq \|\Phi\|_{\text{tr}} \|\Psi\|_{\text{tr}}.$$

Proof. Let Φ, Ψ be as in the statement of the theorem. Using the definition of the trace norm on channels in Equation (3.22) we have

$$\begin{aligned} \|\Phi\|_{\text{tr}} \|\Psi\|_{\text{tr}} &= \max_{X \in \mathcal{L}(\mathcal{K})} \max_{Y \in \mathcal{L}(\mathcal{H})} \frac{\|\Phi(X)\|_{\text{tr}} \|\Psi(Y)\|_{\text{tr}}}{\|X\|_{\text{tr}} \|Y\|_{\text{tr}}} \\ &\geq \max_{Y \in \mathcal{L}(\mathcal{H})} \frac{\|\Phi(\Psi(Y))\|_{\text{tr}} \|\Psi(Y)\|_{\text{tr}}}{\|\Psi(Y)\|_{\text{tr}} \|Y\|_{\text{tr}}} \\ &= \max_{Y \in \mathcal{L}(\mathcal{H})} \frac{\|\Phi(\Psi(Y))\|_{\text{tr}}}{\|Y\|_{\text{tr}}} \\ &= \|\Phi \circ \Psi\|_{\text{tr}}, \end{aligned}$$

as in the statement of the lemma. □

One other useful fact related to this norm is that it is always achieved on an input operator of the form $|\phi\rangle\langle\psi|$. This follows from a direct convexity argument.

Proposition 3.12. *For any $\Phi: \mathbf{L}(\mathcal{H}) \rightarrow \mathbf{L}(\mathcal{K})$, there are pure states $|\phi\rangle$ and $|\psi\rangle$ such that*

$$\|\Phi\|_{\text{tr}} = \|\Phi(|\phi\rangle\langle\psi|)\|_{\text{tr}}.$$

Proof. Let $X \in \mathbf{L}(\mathcal{H})$ achieve the maximum in the definition of the norm, and let $X = \sum_i s_i |\phi_i\rangle\langle\psi_i|$ be a singular value decomposition of X . Applying the triangle inequality to the definition of the operator trace norm (Equation (3.22)) results in

$$\|\Phi\|_{\text{tr}} = \frac{\|\Phi(X)\|_{\text{tr}}}{\|X\|_{\text{tr}}} = \frac{\|\sum_i s_i \Phi(|\phi_i\rangle\langle\psi_i|)\|_{\text{tr}}}{\|X\|_{\text{tr}}} \leq \frac{\sum_i s_i \|\Phi(|\phi_i\rangle\langle\psi_i|)\|_{\text{tr}}}{\|X\|_{\text{tr}}}$$

Then, since the definition of the trace norm implies that $\|X\|_{\text{tr}} = \sum_i s_i$, we may view this as a weighted average over terms $\|\Phi(|\phi_i\rangle\langle\psi_i|)\|_{\text{tr}}$. Since at least one of these terms must be at least as large as the average, there exists an i for which

$$\|\Phi\|_{\text{tr}} \leq \|\Phi(|\phi_i\rangle\langle\psi_i|)\|_{\text{tr}},$$

which implies that the trace norm is achieved on the state $|\phi_i\rangle\langle\psi_i|$. \square

Unfortunately, this extension of the trace norm to channels does not have all of the properties that we might like it to have. The most important of these is *stability*, which is, the norm of an operator should not depend on the existence of a reference system, or more concretely, the norm of the map Φ should not be smaller than the norm of the map $\Phi \otimes \mathbf{I}$. An example due to Watrous [Wat08] provides two channels Φ, Ψ on d -dimensional states such that $\|\Phi \otimes \mathbf{I} - \Psi \otimes \mathbf{I}\|_{\text{tr}} = 2$ for an appropriate reference system, but $\|\Phi - \Psi\|_{\text{tr}} \in O(1/d)$. Phrased in terms of distinguishability, these are two channels that are perfectly distinguishable with a reference system but almost identical without it.

For this reason, we make use of the *diamond norm*, introduced by Kitaev [Kit97]. This norm defines in the reference system, which stabilizes the super-operator trace norm.

Definition 3.13. For a linear map Φ taking $\mathbf{L}(\mathcal{H})$ to $\mathbf{L}(\mathcal{K})$, the diamond norm of Φ is

$$\|\Phi\|_{\diamond} = \|\Phi \otimes \mathbf{I}_{\mathcal{H}}\|_{\text{tr}} = \max_{X \in \mathbf{L}(\mathcal{H} \otimes \mathcal{H}), X \neq 0} \frac{\|(\Phi \otimes \mathbf{I}_{\mathcal{H}})(X)\|_{\text{tr}}}{\|X\|_{\text{tr}}}.$$

This norm is closely related to the completely bounded norm studied in operator algebra. If Φ maps $\mathbf{L}(\mathcal{H})$ to $\mathbf{L}(\mathcal{K})$ and Φ^* is the adjoint map defined by $\text{tr}(A^*\Phi(B)) = \text{tr}((\Phi^*(A))^*B)$, then $\|\Phi\|_\diamond = \|\Phi^*\|_{\text{cb}}$. More information on the completely bounded norm can be found in [Pau02].

As in the case of the super-operator trace norm, this norm inherits many properties from the trace norm, such as the triangle inequality and invariance under unitary operations. From a computational perspective, the optimization in the definition can be cast as a semidefinite program [Wat09c] or as a more general convex optimization problem [BATS09]. The paper [JKP09] also gives a heuristic for evaluating this norm.

It is not too difficult to see that this norm is stable. This was first noted by Kitaev [Kit97] for the diamond norm, and by Smith [Smi83] for the equivalent case of the completely bounded norm. The simpler proof used here can be found in [Wat05], though a similar argument appears in [GLN05], for the case that the maximization in the definition of the diamond norm is restricted to the density operators.

Theorem 3.14 (Kitaev [Kit97], Smith [Smi83]). *Let Φ be a linear map from $\mathbf{L}(\mathcal{H})$ to $\mathbf{L}(\mathcal{K})$. For any space \mathcal{F}*

$$\|\Phi\|_\diamond = \|\Phi \otimes I_{\mathcal{H}}\|_{\text{tr}} \geq \|\Phi \otimes I_{\mathcal{F}}\|_{\text{tr}}.$$

Proof. In the case that $\dim \mathcal{F} < \dim \mathcal{H}$ the statement of the theorem is clear: the maximization in the definition of the super-operator trace norm is being taken over a smaller space.

In the case that $\dim \mathcal{F} \geq \dim \mathcal{H}$, Proposition 3.12 implies that there exist vectors $|\phi\rangle$ and $|\psi\rangle$ such that

$$\|\Phi \otimes I_{\mathcal{F}}\|_{\text{tr}} = \|(\Phi \otimes I_{\mathcal{F}})(|\phi\rangle\langle\psi|)\|_{\text{tr}}.$$

If we take Schmidt decompositions of these vectors, they can have at most $d = \min\{\dim \mathcal{F}, \dim \mathcal{H}\} = \dim \mathcal{H}$ terms. Doing so, we have

$$|\phi\rangle = \sum_{i=1}^d \lambda_i |\mathbf{a}_i\rangle |\mathbf{x}_i\rangle, \quad |\psi\rangle = \sum_{i=1}^d \gamma_i |\mathbf{b}_i\rangle |\mathbf{y}_i\rangle, \quad (3.23)$$

where $\{|\mathbf{a}_i\rangle\}, \{|\mathbf{b}_i\rangle\}$ are orthonormal bases for \mathcal{H} , and $\{|\mathbf{x}_i\rangle\}, \{|\mathbf{y}_i\rangle\}$ are bases for d -dimensional subspaces of \mathcal{F} . The remainder of the proof involves the straightforward but technical argument based on the Schmidt decomposition that we can embed these subspaces into a space of dimension d with no loss in the value of the norm. This is simply a formalization of the observation that since the states $|\phi\rangle$ and $|\psi\rangle$ live in a d -dimensional subspace of \mathcal{F} , we do not need the auxiliary space in the diamond norm to have more than this dimension.

To formalize this, let $U, V \in \mathbf{U}(\mathcal{H}, \mathcal{F})$ be the isometries that take the standard basis $\{|i\rangle\}$ of \mathcal{H} to the bases $\{x_i\}$ and $\{y_i\}$, respectively. The maps UU^* and VV^* are then the projections onto the spaces spanned by $\{x_i\}$ and $\{y_i\}$, which do not affect $|\phi\rangle$ and $|\psi\rangle$ by Equation (3.23). Using this notation

$$\begin{aligned}
\|\Phi \otimes I_{\mathcal{H}}\|_{\text{tr}} &\geq \|(\Phi \otimes I_{\mathcal{H}})[(\mathbf{1}_{\mathcal{H}} \otimes U^*)|\phi\rangle\langle\psi|(\mathbf{1}_{\mathcal{H}} \otimes V)]\|_{\text{tr}} \\
&= \|(\mathbf{1}_{\mathcal{H}} \otimes U)(\Phi \otimes I_{\mathcal{H}})[(\mathbf{1}_{\mathcal{H}} \otimes U^*)|\phi\rangle\langle\psi|(\mathbf{1}_{\mathcal{H}} \otimes V)](\mathbf{1}_{\mathcal{H}} \otimes V^*)\|_{\text{tr}} \\
&= \|(\Phi \otimes I_{\mathcal{F}})[(\mathbf{1}_{\mathcal{H}} \otimes UU^*)|\phi\rangle\langle\psi|(\mathbf{1}_{\mathcal{H}} \otimes VV^*)]\|_{\text{tr}} \\
&= \|(\Phi \otimes I_{\mathcal{F}})(|\phi\rangle\langle\psi|)\|_{\text{tr}} \\
&= \|\Phi \otimes I_{\mathcal{F}}\|_{\text{tr}},
\end{aligned}$$

as desired. □

In addition to stability, this norm has several other convenient properties. One of these is multiplicativity with respect to the tensor product. This is not true of the maximum output p -norm for any $p > 1$ [HW08], but in the case of the diamond norm, it is a direct consequence of the submultiplicativity of the trace norm.

Theorem 3.15. *Let $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ and $\Psi \in \mathbf{T}(\mathcal{F}, \mathcal{G})$. Then*

$$\|\Phi \otimes \Psi\|_{\diamond} = \|\Phi\|_{\diamond} \|\Psi\|_{\diamond}.$$

Proof. One direction follows immediately from the multiplicativity of the trace norm with respect to the tensor product (Equation (3.10)), since

$$\begin{aligned}
\|\Phi \otimes \Psi\|_{\diamond} &= \max_{X \in \mathbf{L}(\mathcal{H} \otimes \mathcal{F})} \frac{\|(\Phi \otimes \Psi)(X)\|_{\text{tr}}}{\|X\|_{\text{tr}}} \\
&\geq \max_{X \in \mathbf{L}(\mathcal{H}), Y \in \mathbf{L}(\mathcal{F})} \frac{\|\Phi(X) \otimes \Psi(Y)\|_{\text{tr}}}{\|X \otimes Y\|_{\text{tr}}} = \|\Phi\|_{\diamond} \|\Psi\|_{\diamond}.
\end{aligned}$$

The other direction is a consequence of Theorem 3.14 and Lemma 3.11:

$$\begin{aligned}
\|\Phi \otimes \Psi\|_{\diamond} &= \|\Phi \otimes \Psi \otimes I_{\mathcal{H} \otimes \mathcal{F}}\|_{\text{tr}} \\
&= \|\Phi \otimes I_{\mathcal{F} \otimes \mathcal{H} \otimes \mathcal{F}} \circ \Psi \otimes I_{\mathcal{H} \otimes \mathcal{H} \otimes \mathcal{F}}\|_{\text{tr}} \\
&\leq \|\Phi \otimes I_{\mathcal{F} \otimes \mathcal{H} \otimes \mathcal{F}}\|_{\text{tr}} \|\Psi \otimes I_{\mathcal{H} \otimes \mathcal{H} \otimes \mathcal{F}}\|_{\text{tr}} \\
&= \|\Phi \otimes I_{\mathcal{H}}\|_{\text{tr}} \|\Psi \otimes I_{\mathcal{F}}\|_{\text{tr}} \\
&= \|\Phi\|_{\diamond} \|\Psi\|_{\diamond},
\end{aligned}$$

which completes the proof. □

As promised, the diamond norm of $\Phi - \Psi$ gives the probability that an unknown channel in $\{\Phi, \Psi\}$ can be correctly identified with only a single use of the channel. This gives an important operational characterization of the diamond norm that has many useful applications to quantum error correction and other fields. The proof follows directly from the definition of the diamond norm and Helstrom's result on the minimum error distinguishability for two states (Theorem 3.9).

Corollary 3.16. *The optimal probability that an unknown channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ chosen uniformly at random from $\{\Phi_1, \Phi_2\}$ can be correctly identified given a single use is given by*

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

Proof. By Theorem 3.17 that follows (and is proven independently of this result) the maximization in the definition of the diamond norm may be taken over a pure state, so that for some space \mathcal{F} , there exists a pure state $|\psi\rangle \in \mathcal{H} \otimes \mathcal{F}$ such that, the value in the statement of the theorem is equal to

$$\frac{1}{4} (2 + \|(\Phi_1 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|) - (\Phi_2 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|)\|_{\text{tr}}). \quad (3.24)$$

This expression is simply the optimal probability of identifying the state

$$(\Phi \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|)$$

from the set

$$\{(\Phi_1 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|), (\Phi_2 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|)\},$$

by Theorem 3.9.

Given only one use of Φ , there is no other strategy than applying it to (a portion of) some optimal distinguishing state and then measuring the result. Theorem 3.17 implies that there is such a state, and so the optimal probability is given by Equation (3.24), as required. \square

With this corollary in mind, we will define the computational problem of distinguishing two channels in terms of the diamond norm of their difference. This is the main problem studied in this thesis, and so the properties of the diamond norm that we have so far defined will be useful throughout.

3.5.1 Maximization on a pure state for the difference of channels

In this section it is shown that when applied to the difference of two channels the diamond norm is achieved on a pure state. The result is technical, but it has many

applications in the remainder of the thesis. This is the product of joint work with John Watrous [RW05].

The theorem applies to the maps Φ that are the difference of two completely positive maps. This property implies another simple property: there exist completely positive Ψ, Γ such that $\Phi = \Psi - \Gamma$ if and only if $\Phi(X^*) = \Phi(X)^*$ for all X . We will only need one direction of this equivalence. If $\Phi = \Psi - \Gamma$, taking Kraus operator decompositions of the completely positive maps Ψ and Γ implies that

$$\Phi(X^*) = \Psi(X^*) - \Gamma(X^*) = \sum_i A_i X^* A_i^* - B_i X^* B_i^* = \left(\sum_i A_i X A_i^* - B_i X B_i^* \right)^* = \Phi(X)^*.$$

This implication will be used in the proof of the main theorem of the section. In the construction used for the proof of this theorem, the space \mathcal{F} has dimension $2 \dim \mathcal{H}$. As explained following the theorem, this can be achieved with a space of dimension $\dim \mathcal{H}$, though the argument is not included in the proof of the theorem for simplicity.

Theorem 3.17. *Let $\Phi: \mathbf{L}(\mathcal{H}) \rightarrow \mathbf{L}(\mathcal{K})$ be the difference of two completely positive maps. There exists a Hilbert space \mathcal{F} and a unit vector $|\psi\rangle \in \mathcal{H} \otimes \mathcal{F}$ such that*

$$\|\Phi\|_{\diamond} = \|(\Phi \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|)\|_{\text{tr}}.$$

Proof. By the definition of the diamond norm

$$\|\Phi\|_{\diamond} = \|\Phi \otimes I_{\mathcal{H}}\|_{\text{tr}} = \max\{\|(\Phi \otimes I_{\mathcal{H}})(X)\|_{\text{tr}} : \|X\|_{\text{tr}} = 1\}.$$

Let $X \in \mathbf{L}(\mathcal{H} \otimes \mathcal{H})$ be a state that achieves this maximum and let \mathcal{C} be a Hilbert space of dimension two (i.e. a single qubit). Consider the Hermitian operator $Y \in \mathbf{L}(\mathcal{H} \otimes \mathcal{H} \otimes \mathcal{C})$ given by

$$Y = \frac{1}{2}X \otimes |0\rangle\langle 1| + \frac{1}{2}X^* \otimes |1\rangle\langle 0|.$$

Notice also that $\|Y\|_{\text{tr}} = \|X\|_{\text{tr}} = 1$ by Lemma 3.10.

As observed above, the condition that Φ is the difference of two complete positive transformations implies that $\Phi(X^*) = \Phi(X)^*$ for all X . Using this, as well as Lemma 3.10

$$\begin{aligned} \|(\Phi \otimes I_{\mathcal{H} \otimes \mathcal{C}})(Y)\|_{\text{tr}} &= \frac{1}{2} \|(\Phi \otimes I_{\mathcal{H}})(X) \otimes |0\rangle\langle 1| + (\Phi \otimes I_{\mathcal{H}})(X^*) \otimes |1\rangle\langle 0|\|_{\text{tr}} \\ &= \frac{1}{2} \|(\Phi \otimes I_{\mathcal{H}})(X) \otimes |0\rangle\langle 1| + (\Phi \otimes I_{\mathcal{H}})(X)^* \otimes |1\rangle\langle 0|\|_{\text{tr}} \\ &= \|(\Phi \otimes I_{\mathcal{H}})(X)\|_{\text{tr}} \\ &= \|\Phi\|_{\diamond}. \end{aligned}$$

This implies that the maximum is achieved on a Hermitian matrix Y .

It is not hard to see that this implies that the maximum is achieved on a pure state. To do so, note that since Y is Hermitian, it has a spectral decomposition. Let such a decomposition be given by

$$Y = \sum_i \lambda_i |\psi_i\rangle\langle\psi_i|,$$

where $\{|\psi_i\rangle\}$ is an orthonormal basis of eigenvectors with real eigenvalues $\{\lambda_i\}$. In addition, because $\|Y\|_{\text{tr}} = 1$, it is the case that $\sum_i |\lambda_i| = 1$. By the linearity of Φ , as well as the triangle inequality and the homogeneity of the trace norm,

$$\|(\Phi \otimes I_{\mathcal{H} \otimes \mathcal{C}})(Y)\|_{\text{tr}} \leq \sum_i |\lambda_i| \|(\Phi \otimes I_{\mathcal{H} \otimes \mathcal{C}})(|\psi_i\rangle\langle\psi_i|)\|_{\text{tr}}.$$

Because $\sum_i |\lambda_i| = 1$, it follows that at least one term in the average achieves the bound, i.e. that

$$\|(\Phi \otimes I_{\mathcal{H} \otimes \mathcal{C}})(|\psi_i\rangle\langle\psi_i|)\|_{\text{tr}} \geq \|\Phi\|_{\diamond}$$

for some value of i . For this value of i we have, by Theorem 3.14,

$$\|(\Phi \otimes I_{\mathcal{H} \otimes \mathcal{C}})(|\psi_i\rangle\langle\psi_i|)\|_{\text{tr}} \leq \|\Phi\|_{\diamond},$$

which implies that $\|(\Phi \otimes I_{\mathcal{H} \otimes \mathcal{C}})(|\psi_i\rangle\langle\psi_i|)\|_{\text{tr}} = \|\Phi\|_{\diamond}$ as required. \square

This theorem does not hold for the trace norm on super-operators, by an example due to Watrous [Wat05]. It may seem odd that in the proof of this theorem the space $\mathcal{F} = \mathcal{H} \otimes \mathcal{C}$ has larger dimension than is required to achieve the maximum, since $\dim \mathcal{H} \otimes \mathcal{C} = 2 \dim \mathcal{H}$. By examining the proof of Theorem 3.14, however, it can be seen that this need not be the case. Applying this theorem in the case that the maximum is Hermitian produces a Hermitian state in $\mathbf{D}(\mathcal{H} \otimes \mathcal{H})$ that achieves the maximum. Applying the convexity argument made at the end of the proof of Theorem 3.17 yields a pure state in $\mathcal{H} \otimes \mathcal{H}$ on which the maximum is achieved.

One convenient consequence of this theorem is that the diamond norm of any quantum channel is equal to one. This is implied by the definition, since for any $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ the theorem implies that there is a state $|\psi\rangle$ such that

$$\|\Phi\|_{\diamond} = \|(\Phi \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|)\|_{\text{tr}} = \|\rho\|_{\text{tr}} = 1, \quad (3.25)$$

where ρ is the density matrix that results from applying $\Phi \otimes I_{\mathcal{F}}$, which has trace norm one because it is normalized. An alternate proof of this fact, not using the above result, can be found in [AKN98].

There is one further result on the diamond norm demonstrated in Section 3.7. This result is a procedure for polarizing this norm, in the sense that if the diamond norm of the difference of the input channels is small, the result is two channels with extremely small norm. Similarly, if the input channels have large norm, then the resulting channels are almost perfectly distinguishable. Results of this type can have powerful applications for error reduction. The discussion of this result is postponed until Section 3.7 so that the fidelity can be introduced, as it is used in the proof of the polarization result. The fidelity is the topic of the next section.

3.6 Fidelity

One of the most important tools in quantum information is the fidelity, which provides a way to determine how close two states are together. For pure states $|\psi\rangle$ and $|\phi\rangle$ the fidelity has a simple expression

$$F(|\psi\rangle, |\phi\rangle) = |\langle\psi|\phi\rangle|.$$

This can be generalized to the case of mixed quantum states $\rho, \sigma \in \mathbf{D}(\mathcal{H})$ by

$$F(\rho, \sigma) = \text{tr} \sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}. \quad (3.26)$$

This quantity and its generalization to mixed states are due to Uhlmann [Uhl76]. The fidelity ranges between $F(\rho, \sigma) = 1$ when $\rho = \sigma$ and $F(\rho, \sigma) = 0$ when ρ and σ have orthogonal support. The remainder of this section is a survey of some of the most important properties of the fidelity. A more complete introduction to this quantity can be found in [NC00].

One property that is convenient to show from Equation (3.26) is multiplicativity with respect to the tensor product. Following [Joz94], this is an easy consequence of the fact that $\sqrt{\rho \otimes \sigma} = \sqrt{\rho} \otimes \sqrt{\sigma}$. This implies that

$$\begin{aligned} F(\rho_1 \otimes \rho_2, \sigma_1 \otimes \sigma_2) &= \text{tr} \sqrt{\sqrt{\rho_1 \otimes \rho_2}(\sigma_1 \otimes \sigma_2)\sqrt{\rho_1 \otimes \rho_2}} \\ &= \left(\text{tr} \sqrt{\sqrt{\rho_1}\sigma_1\sqrt{\rho_1}} \right) \left(\text{tr} \sqrt{\sqrt{\rho_2}\sigma_2\sqrt{\rho_2}} \right) \\ &= F(\rho_1, \sigma_1) F(\rho_2, \sigma_2). \end{aligned} \quad (3.27)$$

This is one of the few properties that is easy to prove from Equation (3.26). For example, it is not clear from this equation that $F(\rho, \sigma) = F(\sigma, \rho)$. This property follows directly from a characterization of the fidelity known as Uhlmann's theorem. This

characterization is extremely useful, often being used as the definition of the fidelity, and is presented as the following theorem.

Theorem 3.18 (Uhlmann's Theorem [Uhl76]). *Let $\rho, \sigma \in \mathbf{D}(\mathcal{H})$, and let \mathcal{K} be any space large enough to admit purifications of ρ and σ . Then*

$$F(\rho, \sigma) = \max\{|\langle \phi | \psi \rangle| : |\phi\rangle, |\psi\rangle \in \mathcal{H} \otimes \mathcal{K}, \text{tr}_{\mathcal{K}} |\phi\rangle\langle\phi| = \rho, \text{tr}_{\mathcal{K}} |\psi\rangle\langle\psi| = \sigma\}.$$

Uhlmann's theorem restricted to the finite dimensional case (once again, the cited result is much more general than has been applied here) allows the derivation of several nice properties of the fidelity. These properties are summarized in the following proposition, many of which are observed by Jozsa [Joz94].

Proposition 3.19. *For $\rho, \sigma \in \mathbf{D}(\mathcal{H})$ and $U \in \mathbf{U}(\mathcal{H}, \mathcal{K})$, the fidelity satisfies*

- (a) $0 \leq F(\rho, \sigma) \leq 1$
- (b) $F(\rho, \sigma) = F(\sigma, \rho)$
- (c) $F(U\rho U^*, U\sigma U^*) = F(\rho, \sigma)$
- (d) For $\rho, \sigma \in \mathbf{D}(\mathcal{H} \otimes \mathcal{K})$, $F(\text{tr}_{\mathcal{K}} \rho, \text{tr}_{\mathcal{K}} \sigma) \geq F(\rho, \sigma)$

Proof. All four of these properties are simple corollaries of Theorem 3.18. Properties (a) and (b) follow immediately. Property (c) follows from the fact that if $|\psi\rangle, |\phi\rangle \in \mathcal{H} \otimes \mathcal{F}$ are purifications of ρ and σ achieving the maximum in the theorem, then $(U \otimes \mathbf{1}_{\mathcal{F}})|\psi\rangle$ and $(U \otimes \mathbf{1}_{\mathcal{F}})|\phi\rangle$ are purifications of $U\rho U^*$ and $U\sigma U^*$, respectively, and

$$|\langle \psi | (U \otimes \mathbf{1}_{\mathcal{F}})^* (U \otimes \mathbf{1}_{\mathcal{F}}) | \phi \rangle| = |\langle \psi | \phi \rangle| = F(\rho, \sigma).$$

Property (d) is a consequence of the fact that if $|\psi\rangle \in \mathcal{H} \otimes \mathcal{K} \otimes \mathcal{F}$ is a purification of ρ , then it is also a purification of $\text{tr}_{\mathcal{K}} \rho$. \square

One further property of the fidelity will be quite important: the monotonicity under the application of a quantum channel. We have seen two special cases of this, unitary operations and the partial trace, as part of the previous proposition. Extending these cases to the set of all channels is a simple consequence of the Stinespring representation for channels.

This proof is due to Jozsa [Joz94] (see also [NC00]), though it is not difficult to derive from Theorem 3.18.

Theorem 3.20. *Let $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$, then for any $\rho, \sigma \in \mathbf{D}(\mathcal{H})$*

$$F(\rho, \sigma) \leq F(\Phi(\rho), \Phi(\sigma)).$$

Proof. By Theorem 3.18, let $|\phi\rangle, |\psi\rangle \in \mathcal{H} \otimes \mathcal{F}$ be purifications of ρ and σ such that

$$F(\rho, \sigma) = |\langle \phi | \psi \rangle|.$$

Additionally, let Φ have a Stinespring representation given by $\Phi(X) = \text{tr}_{\mathcal{B}} U(X \otimes |0\rangle\langle 0|)U^*$. For brevity, let $\hat{U} = U \otimes \mathbf{1}_{\mathcal{F}}$. Notice that $\hat{U}|\phi\rangle|0\rangle$ purifies $\Phi(\rho)$ and that $\hat{U}|\psi\rangle|0\rangle$ purifies $\Phi(\sigma)$. Using this notation along with Theorem 3.18,

$$\begin{aligned} F(\Phi(\rho), \Phi(\sigma)) &\geq F(\hat{U}(|\phi\rangle\langle\phi| \otimes |0\rangle\langle 0|)\hat{U}^*, \hat{U}(|\psi\rangle\langle\psi| \otimes |0\rangle\langle 0|)\hat{U}^*) \\ &= |\langle 0|\langle\phi|\hat{U}^*\hat{U}|\psi\rangle|0\rangle| = |\langle\phi|\psi\rangle| = F(\rho, \sigma), \end{aligned}$$

as required. □

The special case of the monotonicity of the fidelity for the partial trace (item (d) in Proposition 3.19) will be particularly important. This case implies that the fidelity can only increase when it is taken over only part of a system, where the remainder of the system has been traced out. This property will be essential to the results in Chapter 4.

3.6.1 Relation to the trace norm

The fidelity and the trace norm are two of the most useful quantities for determining how close two quantum states are to each other. Despite the similarities between them, it is often much more convenient to work with one or the other of these quantities, and so relationships between them are very useful. This section presents two such relationships that we will make use of later in the thesis.

Uhlmann's Theorem (3.18) can also be used to characterize the fidelity using the trace norm. This result is not hard to prove, but will be central to a couple of proofs that appear later.

Lemma 3.21. *Let $\rho, \xi \in \mathbf{D}(\mathcal{H})$. Then for arbitrary purifications $|\psi\rangle, |\phi\rangle \in \mathcal{H} \otimes \mathcal{A}$ of ρ and ξ , respectively, we have $\|\text{tr}_{\mathcal{H}} |\psi\rangle\langle\phi|\|_{\text{tr}} = F(\rho, \xi)$.*

Proof. Using Lemma 3.7 together with Theorem 3.18 and the fact that all the purifications of ρ, σ are unitarily equivalent using a unitary on the space \mathcal{A} , we have

$$\begin{aligned}
\|\mathrm{tr}_{\mathcal{H}} |\psi\rangle\langle\phi|\|_{\mathrm{tr}} &= \max_{\mathbf{U} \in \mathbf{U}(\mathcal{A})} |\mathrm{tr}(\mathrm{tr}_{\mathcal{H}} |\psi\rangle\langle\phi|) \mathbf{U}| \\
&= \max_{\mathbf{U} \in \mathbf{U}(\mathcal{A})} |\mathrm{tr} |\psi\rangle\langle\phi|(\mathbf{1}_{\mathcal{H}} \otimes \mathbf{U})| \\
&= \max_{\mathbf{U} \in \mathbf{U}(\mathcal{A})} |\langle\phi|(\mathbf{1}_{\mathcal{H}} \otimes \mathbf{U})|\psi\rangle| \\
&= F(\rho, \xi)
\end{aligned}$$

as claimed. \square

A very useful relationship between the fidelity and the trace norm is given by the Fuchs-van de Graaf Inequalities that relate the trace norm and the fidelity. These inequalities show that, up to polynomial factors, the fidelity and the trace norm are equivalent. This is helpful, since it is often much easier to work with one or the other of these quantities.

Theorem 3.22 (Fuchs and van de Graaf [FvdG99]). *For any $\rho, \sigma \in \mathbf{D}(\mathcal{H})$*

$$1 - F(\rho, \sigma) \leq \frac{1}{2} \|\rho - \sigma\|_{\mathrm{tr}} \leq \sqrt{1 - F(\rho, \sigma)^2}.$$

The second inequality is not hard to prove. Following [NC00], let $|\phi\rangle, |\psi\rangle \in \mathcal{H} \otimes \mathcal{A}$ be purifications of ρ and σ achieving the bound $F(\rho, \sigma) = |\langle\phi|\psi\rangle|$ in Uhlmann's Theorem (3.18). Using the monotonicity of the trace norm under the partial trace (Theorem 3.8), we have

$$\begin{aligned}
\|\rho - \sigma\|_{\mathrm{tr}} &= \|\mathrm{tr}_{\mathcal{A}}(|\phi\rangle\langle\phi| - |\psi\rangle\langle\psi|)\|_{\mathrm{tr}} \\
&\leq \|(|\phi\rangle\langle\phi| - |\psi\rangle\langle\psi|)\|_{\mathrm{tr}} \\
&= 2\sqrt{1 - |\langle\phi|\psi\rangle|^2} \\
&\leq 2\sqrt{1 - F(\rho, \sigma)^2}.
\end{aligned}$$

The first inequality is more difficult: it requires either characterizing the fidelity and the trace norm in terms of classical variants on the outcomes of measurements, as is done in [FvdG99, NC00], or proving a technical result on the trace norm of the difference of two positive operators, as is done in [KSV02]. Neither of these techniques are used in the remainder of the thesis, and so the proof of this inequality is omitted.

The Fuchs-van de Graaf Inequalities may be equivalently rephrased in terms of upper and lower bounds on the fidelity in terms of the trace norm. These bounds are

$$1 - \frac{1}{2} \|\rho - \sigma\|_{\mathrm{tr}} \leq F(\rho, \sigma) \leq \sqrt{1 - \frac{1}{4} \|\rho - \sigma\|_{\mathrm{tr}}^2}, \quad (3.28)$$

and can be derived by simple manipulations of Theorem 3.22.

3.6.2 Maximum output fidelity for channels

The fidelity can be extended to quantum channels in much the same way as the previous quantities that we have considered. The *maximum output fidelity* of two channels $\Phi_1, \Phi_2 \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ can be defined as

$$F_{\max}(\Phi_1, \Phi_2) = \max_{\rho, \sigma \in \mathbf{D}(\mathcal{H})} F(\Phi_1(\rho), \Phi_2(\sigma))$$

This quantity will be essential to the `CLOSE IMAGES` problem that is the focus of Chapter 4. The property of primary importance for this application is the multiplicativity of F_{\max} with respect to the tensor product of two channels. This will be essential for error reduction on instances of `CLOSE IMAGES`. This result is used implicitly by Kitaev and Watrous [KW00], and the main thrust of it can also be found in [KSV02] (see Problem 11.10). Due to its importance, a complete proof is presented here. The method of proof used here is due to John Watrous¹, though it is similar to the techniques used in [KW00, KSV02]. This proof makes use of the diamond norm, and specifically the multiplicativity of the diamond norm with respect to tensor products, which was introduced in Section 3.5.

The first part of the proof is a relationship between the maximum output fidelity of two channels and the diamond norm of a certain completely positive super-operator.

Lemma 3.23 (Kitaev and Watrous [KW00]). *Let $\Phi, \Psi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$, and the linear map $\Gamma: \mathbf{L}(\mathcal{H}) \rightarrow \mathbf{L}(\mathcal{B})$ be given by*

$$\begin{aligned}\Phi(X) &= \text{tr}_{\mathcal{B}} UXU^* \\ \Psi(X) &= \text{tr}_{\mathcal{B}} VXV^* \\ \Gamma(X) &= \text{tr}_{\mathcal{K}} UXV^*,\end{aligned}$$

where $U, V \in \mathbf{U}(\mathcal{H}, \mathcal{B} \otimes \mathcal{K})$. Using this notation,

$$F_{\max}(\Phi, \Psi) = \|\Gamma\|_{\diamond}.$$

Proof. Let \mathcal{A} be a space with $\dim \mathcal{A} = \dim \mathcal{H}$ to allow purifications of states in $\mathbf{D}(\mathcal{H})$, and let $\hat{U} = U \otimes \mathbf{1}_{\mathcal{A}}$ and $\hat{V} = V \otimes \mathbf{1}_{\mathcal{A}}$, then

$$\begin{aligned}F_{\max}(\Phi, \Psi) &= \max_{\rho, \sigma \in \mathbf{D}(\mathcal{H})} F(\Phi(\rho), \Psi(\sigma)) \\ &= \max_{|\phi\rangle, |\psi\rangle \in \mathcal{H} \otimes \mathcal{A}} F(\text{tr}_{\mathcal{A} \otimes \mathcal{B}} \hat{U}|\phi\rangle\langle\phi|\hat{U}^*, \text{tr}_{\mathcal{A} \otimes \mathcal{B}} \hat{V}|\psi\rangle\langle\psi|\hat{V}^*),\end{aligned}$$

¹John Watrous, private communication

where $|\phi\rangle$ and $|\psi\rangle$ are purifications of ρ and σ . Applying Lemma 3.21 to this, since $\hat{U}|\phi\rangle$ purifies $\Phi(\rho)$ and $\hat{V}|\psi\rangle$ purifies $\Psi(\sigma)$, results in

$$\max_{|\phi\rangle, |\psi\rangle \in \mathcal{H} \otimes \mathcal{A}} \left\| \text{tr}_{\mathcal{K}} \hat{U}|\phi\rangle\langle\psi|\hat{V}^* \right\|_{\text{tr}} = \max_{|\phi\rangle, |\psi\rangle \in \mathcal{H} \otimes \mathcal{A}} \left\| (\Gamma \otimes \mathbf{I}_{\mathcal{A}})(|\phi\rangle\langle\psi|) \right\|_{\text{tr}} = \|\Gamma\|_{\diamond},$$

where the last inequality is an application of Proposition 3.12. \square

Using this lemma the desired result on the multiplicativity of F_{\max} follows immediately from the multiplicativity of the diamond norm with respect to the tensor product.

Theorem 3.24 (Kitaev and Watrous [KW00]). *For any $\Phi_1, \Psi_1 \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ and any $\Phi_2, \Psi_2 \in \mathbf{T}(\mathcal{J}, \mathcal{L})$*

$$F_{\max}(\Phi_1 \otimes \Phi_2, \Psi_1 \otimes \Psi_2) = \prod_{i=1,2} F_{\max}(\Phi_i, \Psi_i).$$

Proof. Let $\Phi_i(X) = \text{tr}_{\mathcal{B}} U_i X U_i^*$, $\Psi_i = \text{tr}_{\mathcal{B}} V_i X V_i^*$ be Stinespring representations of the channels Φ_i, Ψ_i for $i = 1, 2$, where for notational convenience the introduction of ancillary qubits has been merged into the isometries U_i and V_i . Then, setting $\Gamma_i(X) = \text{tr}_{\mathcal{K}} U_i X V_i^*$ we are in exactly the situation of Lemma 3.23. Applying this lemma, as well as the multiplicativity of the diamond norm, gives

$$F_{\max}(\Phi_1 \otimes \Phi_2, \Psi_1 \otimes \Psi_2) = \|\Gamma_1 \otimes \Gamma_2\|_{\diamond} = \|\Gamma_1\|_{\diamond} \|\Gamma_2\|_{\diamond} = F_{\max}(\Phi_1, \Psi_1) F_{\max}(\Phi_2, \Psi_2),$$

as claimed. \square

3.7 Polarization of the diamond norm

This section describes a method for “polarizing” the diamond norm of two channels. This is a technique that, starting with two channels $\Phi_1, \Phi_2 \in \mathbf{T}(\mathcal{H}, \mathcal{K})$, and constants $0 < b < a < 2$ such that $2b < a^2$, creates channels Ψ_1 and Ψ_2 satisfying

$$\begin{aligned} \|\Phi_1 - \Phi_2\|_{\diamond} \leq b &\implies \|\Psi_1 - \Psi_2\|_{\diamond} \leq 2^{-k} \\ \|\Phi_1 - \Phi_2\|_{\diamond} \geq a &\implies \|\Psi_1 - \Psi_2\|_{\diamond} \geq 2 - 2^{-k}. \end{aligned}$$

The constructed channels Ψ_1 and Ψ_2 belong to $\mathbf{T}(\mathcal{H}^{\otimes r}, \mathcal{K}^{\otimes r})$, where $r \in O(k)$, i.e. the size of the resulting channels depends only linearly on the error parameter k . This provides a powerful technique for reducing the error in many settings. It provides one way to see that any promise problem defined with a promise on the diamond norm

difference of two channels can be reduced to the same problem with a weaker gap, since the instance with the weaker gap can be polarized using this technique. The method is not perfect, however, as it depends on the technical condition that $2b < a^2$.

This construction generalizes the polarization technique of Sahai and Vadhan for the case of the ℓ_1 norm of efficiently samplable probability distributions [SV03]. This construction was generalized by Watrous to the case of quantum states that can be efficiently prepared [Wat02]. The further generalization given here to quantum channels does not require any conceptual changes: the details work out in almost exactly the same way as in the case of states. This result is the product of joint work with John Watrous, and has been published in [RW05].

In order that the polarization technique is useful in the setting of computational hardness it must satisfy one further significant property. The construction must be efficient. That is, given access to polynomial-time circuits (or black boxes) for the original channels Φ_1, Φ_2 , circuits that implement the output channels Ψ_1 and Ψ_2 can be efficiently constructed. That the polarization technique has this property will be easy to observe from the construction given in the proof.

The proof of the polarization theorem makes use of two constructions. One of these constructions increases the diamond norm and the other reduces it. Applying these constructions in the correct sequence will result in transformations with the desired properties. These two constructions mirror the proof of the classical result due to Sahai and Vadhan [SV03].

The first construction is a technique for increasing the diamond norm of two channels. The idea is simple: it is much easier to distinguish k copies of the channels than it is to distinguish one copy. The channels constructed using this procedure are simply $\Phi_i^{\otimes k}$. The argument must be carefully made, however, to show that entanglement across the multiple uses of the channels does not increase the diamond norm too much. The following direct product lemma gives bounds on the diamond norm for the difference of k copies of the two channels.

Lemma 3.25. *Let $\Phi_1, \Phi_2 \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ have $\|\Phi_1 - \Phi_2\|_\diamond = \delta > 0$. Then for any positive integer k*

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

Proof. To prove the first inequality, let $\rho \in \mathbf{D}(\mathcal{H} \otimes \mathcal{H})$ achieve the maximum in the diamond norm, i.e. let

$$\|(\Phi_1 \otimes I_{\mathcal{H}})(\rho) - (\Phi_2 \otimes I_{\mathcal{H}})(\rho)\|_{\text{tr}} = \|\Phi_1 - \Phi_2\|_\diamond = \delta.$$

Such a state exists by Theorem 3.17. As the trace norm is multiplicative with respect to the tensor product (by Equation 3.10), $\|\rho^{\otimes k}\|_{\text{tr}} = \|\rho\|_{\text{tr}}^k = 1$. Evaluating the maximum in the definition of the diamond norm on this state, we find that

$$\|\Phi_1^{\otimes k} - \Phi_2^{\otimes k}\|_{\diamond} \geq \left\| ((\Phi_1 \otimes I_{\mathcal{H}})(\rho))^{\otimes k} - ((\Phi_2 \otimes I_{\mathcal{H}})(\rho))^{\otimes k} \right\|_{\text{tr}}. \quad (3.29)$$

We can then apply the bound from [Wat02] to the two states $\rho = (\Phi_1 \otimes I_{\mathcal{H}})(\rho)$ and $\sigma = (\Phi_2 \otimes I_{\mathcal{H}})(\rho)$ having trace distance δ to obtain the desired inequality. For completeness, the proof of this bound follows.

Since the fidelity is multiplicative with respect to the tensor product of two states (Equation (3.27)) we can use the Fuchs-van de Graaf inequalities (Theorem 3.22) to obtain

$$\begin{aligned} \|\rho^{\otimes k} - \sigma^{\otimes k}\|_{\text{tr}} &\geq 2(1 - F(\rho^{\otimes k}, \sigma^{\otimes k})) = 2(1 - F(\rho, \sigma)^k) \\ &\geq 2 - 2 \left(\sqrt{1 - \|\rho - \sigma\|_{\text{tr}}^2 / 4} \right)^k = 2 - 2(1 - (\delta/2)^2)^{k/2}. \end{aligned}$$

We can then bound this quantity using the inequality $(1 - x)^k < e^{-kx}$, which holds for all nonzero $-1 < x < 1$. This can be verified by taking logarithms and considering a Taylor series for $\ln(1 - x)$. In our case, $x = \delta/2 < 1$, so we have

$$2 - 2(1 - (\delta/2)^2)^{k/2} > 2 - 2 \exp\left(\frac{-\delta^2}{4} \cdot \frac{k}{2}\right) = 2 - 2e^{-\frac{k\delta^2}{8}}.$$

Combining this with Equation (3.29) proves the first inequality.

The second inequality follows by induction on k . The case of $k = 1$ leaves nothing to prove. For $k > 1$, let $\Psi_i = \Phi_i^{\otimes(k-1)}$ for simplicity. Using this notation, as well as the triangle inequality, we have

$$\begin{aligned} \|\Phi_1^{\otimes k} - \Phi_2^{\otimes k}\|_{\diamond} &= \|\Psi_1 \otimes \Phi_1 - \Psi_2 \otimes \Phi_2\|_{\diamond} \\ &= \|\Psi_1 \otimes \Phi_1 - \Psi_2 \otimes \Phi_1 + \Psi_2 \otimes \Phi_1 - \Psi_2 \otimes \Phi_2\|_{\diamond} \\ &\leq \|(\Psi_1 - \Psi_2) \otimes \Phi_1\|_{\diamond} + \|\Psi_2 \otimes (\Phi_1 - \Phi_2)\|_{\diamond} \\ &= \|\Psi_1 - \Psi_2\|_{\diamond} \|\Phi_1\|_{\diamond} + \|\Psi_2\|_{\diamond} \|\Phi_1 - \Phi_2\|_{\diamond}. \end{aligned}$$

The final equality follows from the multiplicativity of the diamond norm, given by Theorem 3.15. Since the diamond norm of any channel is one (Equation (3.25)), the inductive hypothesis implies that

$$\|\Psi_1 - \Psi_2\|_{\diamond} \|\Phi_1\|_{\diamond} + \|\Psi_2\|_{\diamond} \|\Phi_1 - \Phi_2\|_{\diamond} \leq (k-1)\delta + \delta = k\delta$$

as required. \square




Figure 3.1: Circuits C_1 and C_2 output by the construction in Lemma 3.26. Each circuit C_i contains r independent copies of the circuit Q_i .

The lemma implies the existence of an efficient procedure to increase the diamond norm of two channels, which will form half of the construction used to polarize the diamond norm. The key to this procedure is that while this procedure increases the norm, it does so much faster when the norm of the original circuits is large. The circuits produced by this procedure are demonstrated in Figure 3.1.

Lemma 3.26. *There is a polynomial-time deterministic procedure that, on input $(Q_1, Q_2, 1^r)$, where Q_1, Q_2 are descriptions of mixed-state quantum circuits, produces as output descriptions of two quantum circuits, (C_1, C_2) satisfying*

$$2 - 2 \exp\left(-\frac{r}{8} \|Q_1 - Q_2\|_\diamond^2\right) < \|C_1 - C_2\|_\diamond \leq r \|Q_1 - Q_2\|_\diamond.$$

Proof. For $i = 1, 2$, the circuit C_i is constructed from r parallel copies of the circuit Q_i . This results in $C_i = Q_i^{\otimes r}$, so that the bounds in the statement of the lemma follow from Lemma 3.25. □

This procedure to increase the diamond norm of the difference of two channels is used in Chapter 6, as it preserves the degradability or antidegradability of the input channels. This will not be true of the remainder of the polarization procedure.

The second procedure that is used in the polarization construction is used to reduce the diamond norm of the difference of two channels. Before outlining the procedure, however, we prove the following simple property of the norm.

Proposition 3.27. Let $\Phi_1, \Phi_2 \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ and $\Psi_1, \Psi_2 \in \mathbf{T}(\mathcal{F}, \mathcal{G})$. Let

$$\begin{aligned}\Xi_1 &= \frac{1}{2}\Phi_1 \otimes \Psi_1 + \frac{1}{2}\Phi_2 \otimes \Psi_2, \\ \Xi_2 &= \frac{1}{2}\Phi_1 \otimes \Psi_2 + \frac{1}{2}\Phi_2 \otimes \Psi_1.\end{aligned}$$

Then $\|\Xi_1 - \Xi_2\|_\diamond = \frac{1}{2} \|\Phi_1 - \Phi_2\|_\diamond \|\Psi_1 - \Psi_2\|_\diamond$.

Proof. The diamond norm is multiplicative with respect to tensor products (Theorem 3.15), so that

$$\|\Xi_1 - \Xi_2\|_\diamond = \left\| \frac{1}{2}(\Phi_1 - \Phi_2) \otimes (\Psi_1 - \Psi_2) \right\|_\diamond = \frac{1}{2} \|\Phi_1 - \Phi_2\|_\diamond \|\Psi_1 - \Psi_2\|_\diamond$$

as required. \square

This property is useful in the proof that the technique for reducing the diamond norm works correctly. The idea behind this procedure is that even if Q_1 and Q_2 are easy to distinguish, then the channel C_1 is constructed by taking the tensor product of r channels, each chosen from $\{Q_1, Q_2\}$ uniformly at random, with the restriction that Q_1 appears an odd number of times should be very hard to distinguish from the channel C_2 constructed in the same way, except that Q_1 appears an even number of times in C_2 . In effect, a procedure that distinguishes C_1 and C_2 must succeed for all r embedded channels: this is because the goal is to determine the parity of the number of times that Q_1 appears, and the parity is affected by even a single mistake made by the distinguishing procedure. This construction mirrors that used on states in [Wat02], which itself mirrors that used on probability distributions in [SV03]. The circuits produced by this procedure are illustrated in Figure 3.2.

Lemma 3.28. *There is a deterministic polynomial-time procedure that, on input $(Q_1, Q_2, 1^r)$, where Q_1, Q_2 are descriptions of mixed-state quantum circuits, produces as output descriptions of two quantum circuits (C_1, C_2) satisfying*

$$\|C_1 - C_2\|_\diamond = 2 \left(\frac{\|Q_1 - Q_2\|_\diamond}{2} \right)^r.$$

Proof. We use the circuits C_1 and C_2 outlined above. The circuit C_1 performs the transformation defined as

$$C_1 = \frac{1}{2^{r-1}} \sum_{\substack{x_1, \dots, x_r \in \{1, 2\} \\ x_1 + \dots + x_r \equiv 1 \pmod{2}}} Q_{x_1} \otimes \dots \otimes Q_{x_r}$$

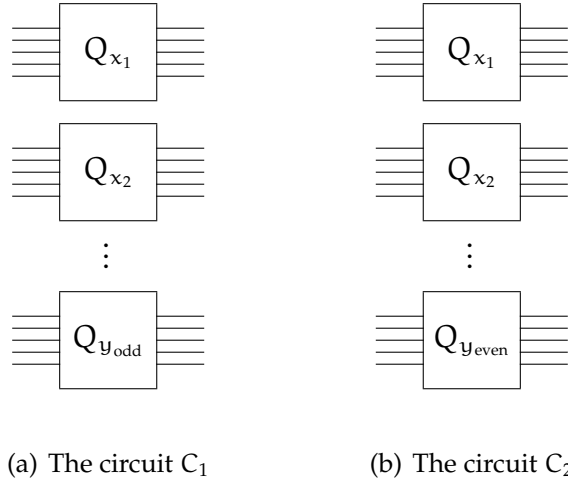


Figure 3.2: Circuits C_1 and C_2 output by the construction in Lemma 3.28. The circuit C_i consists of $r - 1$ independent circuits Q_{x_i} , each chosen randomly at run-time from $\{Q_1, Q_2\}$, and one final circuit chosen so that the parity of the indices of the chosen circuits is odd in the case $i = 1$ and even in the case that $i = 2$.

while C_2 performs a similar transformation defined as

$$C_2 = \frac{1}{2^{r-1}} \sum_{\substack{x_1, \dots, x_r \in \{0,1\} \\ x_1 + \dots + x_r \equiv 0 \pmod{2}}} Q_{x_1} \otimes \dots \otimes Q_{x_r}.$$

These circuits run r copies of Q_1 and/or Q_2 in parallel, where the choice of Q_1 or Q_2 determined uniformly at random subject to the constraint that C_1 applies an odd number of copies of Q_1 while C_2 applies an even number. Such circuits may be constructed in time polynomial in the sizes of Q_1 and Q_2 by using ancillary qubits with Hadamard and dephasing gates to generate the randomness.

A proof by induction based on Proposition 3.27 establishes the desired equality. This proof is included here for completeness. The base case, $r = 1$, leaves nothing to prove. Let $r > 1$, and let D_1, D_2 be the channels C_1 and C_2 for the case $r - 1$. Notice that

$$\|C_1 - C_2\|_{\diamond} = \frac{1}{2} \|Q_1 \otimes D_2 + Q_2 \otimes D_1 - Q_1 \otimes D_1 - Q_2 \otimes D_2\|_{\diamond},$$

which can be observed from the construction of C_1 and C_2 by considering the case that the first transformation is Q_1 or Q_2 and applying the parity conditions. Applying Proposition 3.27 to this, we have

$$\|C_1 - C_2\|_{\diamond} = \frac{1}{2} \|Q_1 - Q_2\|_{\diamond} \|D_1 - D_2\|_{\diamond} = 2 \left(\frac{\|Q_1 - Q_2\|_{\diamond}}{2} \right)^r,$$

where the last equality is by the induction hypothesis on $\|D_1 - D_2\|_\diamond$. \square

These two constructions, taken together, suffice to prove the polarization theorem. The proof consists of an application of Lemma 3.28, followed by an application of Lemma 3.26, followed by one more application of Lemma 3.28. The proof is intuitively simple, though it is quite technical: the value of r used in each transformation must be chosen very carefully.

Theorem 3.29. *Let the constants $a, b \in (0, 2)$ satisfy $2b < a^2$. There exists a deterministic polynomial-time procedure that, given input $(Q_1, Q_2, 1^n)$, where Q_1 and Q_2 are mixed-state quantum circuits, outputs quantum circuits (C_1, C_2) such that*

$$\begin{aligned}\|Q_1 - Q_2\|_\diamond \leq b &\implies \|C_1 - C_2\|_\diamond < 2^{-n} \\ \|Q_1 - Q_2\|_\diamond \geq a &\implies \|C_1 - C_2\|_\diamond > 2 - 2^{-n}.\end{aligned}$$

Proof. First, we apply Lemma 3.28 to the input $(Q_1, Q_2, 1^r)$, with

$$r = \lceil \log(16n) / \log(a^2/(2b)) \rceil.$$

This result in circuits (Q'_1, Q'_2) such that

$$\begin{aligned}\|Q_1 - Q_2\|_\diamond \leq b &\implies \|Q'_1 - Q'_2\|_\diamond \leq 2(b/2)^r, \\ \|Q_1 - Q_2\|_\diamond \geq a &\implies \|Q'_1 - Q'_2\|_\diamond \geq 2(a/2)^r.\end{aligned}$$

Next, we apply Lemma 3.26 to the input $(Q'_1, Q'_2, 1^s)$, where

$$s = \lfloor (b/2)^{-r} / 4 \rfloor.$$

This procedure produces circuits (Q''_1, Q''_2) satisfying

$$\begin{aligned}\|Q_1 - Q_2\|_\diamond \leq b &\implies \|Q''_1 - Q''_2\|_\diamond \leq 2(b/2)^r (b/2)^{-r} / 4 = 1/2, \\ \|Q_1 - Q_2\|_\diamond \geq a &\implies \|Q''_1 - Q''_2\|_\diamond > 2 - 2 \exp(-\frac{s}{2}(a/2)^{2r}) \geq 2 - 2e^{-2n+1}.\end{aligned}$$

The last inequality is due to the fact that

$$\frac{s}{2}(a/2)^{2r} + 1 \geq \frac{1}{8}(b/2)^{-r}(a/2)^{2r} \geq \frac{1}{8} \left(\frac{a^2}{2b} \right)^r,$$

where the $+1$ term on the left is due to the floor in the definition of s . Taking logarithms of both sides, this is

$$\log \left(\frac{s}{2}(a/2)^{2r} + 1 \right) \geq \log \frac{1}{8} \left(\frac{a^2}{2b} \right)^r \geq r \log \frac{a^2}{2b} - 3 \geq \frac{\log 16n}{\log a^2/(2b)} \log \frac{a^2}{2b} - 3 = \log 2n.$$

This implies that $2 - 2 \exp(-\frac{s}{2}(\alpha/2)^{2r}) \geq 2 - 2e^{-2n+1}$, as required.

Finally, applying Lemma 3.28 once more, this time to $(Q_1'', Q_2'', 1^t)$, where

$$t = \lceil (n+1)/2 \rceil,$$

we obtain circuits (C_1, C_2) such that

$$\|Q_1 - Q_2\|_{\diamond} \leq b \implies \|C_1 - C_2\|_{\diamond} \leq (1/2)^{(n+1)/2} (1/2)^{(n-1)/2} = 2^{-n}$$

$$\|Q_1 - Q_2\|_{\diamond} \geq a \implies \|C_1 - C_2\|_{\diamond} > (2 - 2e^{-2n+1})^{\lceil (n+1)/2 \rceil} (1/2)^{\lceil (n+1)/2 \rceil - 1} > 2 - 2^{-n}.$$

The final inequality is due to the fact that by Bernoulli's inequality

$$2(1 - e^{-2n+1})^{\lceil (n+1)/2 \rceil} \geq 2(1 - \lceil (n+1)/2 \rceil e^{-2n+1}) > 2 - 2^{-n}.$$

The circuits (C_1, C_2) have size rst times the size of the original circuits (Q_1, Q_2) . By inspecting these quantities we find that $r, t \in O(n)$ and $s \in O(n^c)$ for c a constant depending on the constants a, b . This implies that the construction can be implemented in time polynomial in n and the size of the original circuits. \square

3.8 Conclusion

In this chapter several different measures on quantum states and channels have been introduced. Many of the important properties of these measures have also been defined. This forms the basis for the remainder of the thesis: the quantities described here, and their properties, will find use throughout the problems studied later. For this reason it is hoped that this chapter will stand as a useful reference for these concepts.

There are two new results contained within this chapter. The first of these is the theorem in Section 3.5.1 that demonstrates that the maximum in the diamond norm on the difference of two channels is achieved on a pure quantum state, as opposed to a general linear operator. This property will be extremely useful when the diamond norm is later used as a way to quantify the distinguishability of two quantum channels, which is the central problem considered in this thesis. The second new result is the polarization technique for the diamond norm in Theorem 3.29. Part of this result is used in Chapter 6 to reduce the error in the reductions of the distinguishability problem to the degradable and the antidegradable channels. Both of these results are joint work with John Watrous [RW05].

Chapter 4

The Close Images Problem

Given two quantum channels it is natural to ask how close the outputs of the two channels can be. When these channels are given as mixed-state quantum circuits this becomes the computational problem `CLOSE IMAGES` that is considered in this chapter. This problem is **QIP**-complete, as it is a restatement of the definition the complexity class. This result is due to Kitaev and Watrous [KW00], but it is included here because it is the result that all of the other hardness results in the thesis depend on.

The main result of this chapter is that restricting this problem to input circuits of logarithmic depth does not reduce the computational difficulty. This is shown by constructing a reduction from an instance of `CLOSE IMAGES` to an instance on log-depth circuits. This provides further evidence for the computational power of log-depth circuits. The reduction that proves this result involves a simulation of the two input circuits by log-depth circuits. The maximum output fidelity of the constructed circuits is related to the maximum output fidelity for the original two circuits, so that this reduction preserves the structure of the close images problem.

The results in this chapter on log-depth circuits have been published in [Ros08b].

Contents

4.1	Log-depth mixed-state quantum circuits	78
4.2	QIP completeness of close images	79
4.3	The swap test	83
4.4	Reduction to logarithmic depth	87
4.5	Correctness of the reduction	93
4.6	Conclusion	98

4.1 Log-depth mixed-state quantum circuits

A significant practical problem in quantum information is that quantum systems quickly decohere when allowed to interact with the environment. This process severely limits the length of quantum computations that can be experimentally realized. Short quantum circuits provide a model of computation that can capture the kinds of computation that we can perform under this type of time limit. For this reason it is of significant interest to find short quantum circuits for important problems.

Log-depth quantum circuits have been found for several significant problems including the approximate quantum Fourier transform [CW00] and the encoding and decoding operations for many quantum error correcting codes [MN02]. In addition to these applications, a procedure for parallelizing to log-depth a large class of quantum circuits is known [BK09]. These examples demonstrate the surprising power of short quantum circuits. It has been conjectured by Jozsa that any quantum algorithm can be performed with logarithmic quantum circuit depth interspersed with polynomial time classical computation [Joz06].

The standard circuit model of quantum computation is the unitary circuit model applied to pure quantum states. In this thesis we consider the more general model of mixed-state quantum computation introduced in Section 2.1. While much of the previous work on short quantum circuits has been in the unitary circuit model [FGHZ05, GHMP02], there has also been work outside of this model [TD04]. The primary advantage of considering this more general model is that the mixed state model is able to capture any physically realizable quantum operation, and so results on this model may have implications for experimental quantum information.

In this chapter it is shown that the apparent power of short quantum computations comes with a price: the close images problem on logarithmic depth quantum circuits is *exactly* as difficult as the general problem on polynomial depth circuits. This result will be used, in Chapter 5, to show that the problem of distinguishing mixed state circuits is also no easier when restricted to log-depth circuits.

The remainder of this chapter is organized as follows. In the next section, the close images problem is discussed, and the result due to Kitaev and Watrous [KW00] that the close images problem is complete for **QIP** is detailed. In Section 4.3 a key component of the reduction to log-depth circuits is considered: the Swap Test. This procedure can be used to ensure that two pure quantum states are close together, and as such is a key component of many quantum algorithms. In Section 4.4 a Karp reduction from the polynomial depth to logarithmic depth versions of the close images problem is

presented in detail. The correctness of this reduction is shown formally in Section 4.5.

4.2 QIP completeness of close images

In this section an overview is given of the close images problem as it relates to the complexity class **QIP**. **CLOSE IMAGES** is essentially a restatement of the acceptance condition for the verifier in a quantum interactive proof system, and so it will be important to review this connection, as this connects all of the other computational problems studied in the thesis to the class **QIP**.

In order to model the hardness of the class of problems having quantum interactive proof systems, Kitaev and Watrous introduced the close images problem [KW00]. This problem can be given the following formal definition.

Problem 4.1 (Close Images). For constants $0 < b < a \leq 1$, the input consists of quantum circuits Q_1 and Q_2 that implement transformations in $\mathbf{T}(\mathcal{H}, \mathcal{K})$. The promise problem is to distinguish the two cases:

Yes: $F(Q_1(\rho), Q_2(\xi)) \geq a$ for some $\rho, \xi \in \mathbf{D}(\mathcal{H})$,

No: $F(Q_1(\rho), Q_2(\xi)) \leq b$ for all $\rho, \xi \in \mathbf{D}(\mathcal{H})$.

This is simply the problem of determining if there are inputs to Q_1 and Q_2 that cause them to output states that are nearly the same. It will be helpful to abbreviate this problem as $\text{CI}_{a,b}$ when the constants a and b will be significant. It is the aim of the present chapter to prove that this problem remains complete for **QIP** when restricted to circuits Q_1 and Q_2 that are of depth logarithmic in the number of input qubits. This will be achieved in the case of perfect soundness error, i.e. $a = 1$ in the above problem definition. As discussed below, this problem remains complete for **QIP** in this case. This restriction serves only to simplify the problem, as distinguishing the two cases for a weaker promise can only be more difficult, so a hardness result on this case will also imply the hardness of the more general problem. For the sake of brevity, the log-depth version of this problem will be referred to as $\text{LOG-DEPTH CI}_{a,b}$ and since this problem is a restriction of a problem in **QIP**, as argued below, it is clear that it is also in **QIP**. Similarly, the abbreviation $\text{CONST-DEPTH CI}_{a,b}$ will be used to denote the version of this problem on constant-depth circuits.

Although this problem was introduced by Kitaev and Watrous to show that $\mathbf{QIP} \subseteq \mathbf{EXP}$, it was not explicitly defined in [KW00]. For this reason the reduction from the




Figure 4.1: The operations and Hilbert spaces corresponding to a three message quantum interactive proof system.

value of a quantum interactive proof system to the close images problem is repeated here. The hardness of this problem is significant for the thesis: it is from this reduction that all of the other **QIP**-hardness results in the thesis follow.

Recall that, due to the results of Kitaev and Watrous [KW00], any quantum interactive proof system can be parallelized to three messages. For any input x , this results in unitary circuits $P_1, V_1, P_2,$ and V_2 acting on systems $\mathcal{V}, \mathcal{M}, \mathcal{P}$ for the verifier's private space, the message space, and the prover's private space respectively. These spaces and transformations are illustrated in Figure 4.1. Recall from Section 2.2 that these circuits depend on the input, but the verifier's circuits V_1 and V_2 must be generated in polynomial time in the length of x . For this reason the input string does not appear in the description of the protocol: it is "hard-coded" into the circuits of the two parties. As the verifier accepts depending on the result of a measurement on one of the message qubits at the end of the protocol, the value of the quantum interactive proof system for fixed transformations P_i is given by

$$\text{tr} [\Pi (V_2 \circ P_2 \circ V_1 \circ P_1(|0\rangle\langle 0|))],$$

where Π is the projector onto the verifier's accepting subspace. The prover can make the verifier accept if there exist circuits P_1 and P_2 such that this probability is large.

To reduce this problem to an instance of CI we must find transformations Q_1 and Q_2 that have close images if and only if the verifier can be made to accept. The construction of these two transformations is outlined in Figure 4.2. The transformation Q_1 is the first half of the protocol, consisting of the unitary circuit V_1 applied to the prover's first message. The transformation Q_2 represents the second half of the protocol run in reverse: starting from an accepting state $|1\rangle\langle 1| \otimes \sigma$ and performing V_2^* , which is the inverse of the unitary circuit V_2 . These transformations are given in [KW00] more




Figure 4.2: Construction of the circuits Q_1 and Q_2 in the reduction from three message quantum interactive proof system to an instance of **CLOSE IMAGES**. The space \mathcal{P} of the prover's private system is not shown and ρ represents the prover's first message.

formally as

$$\begin{aligned} Q_1(\rho) &= \text{tr}_{\mathcal{M}} V_1 (|0\rangle\langle 0| \otimes \rho) V_1^*, \\ Q_2(\sigma) &= \text{tr}_{\mathcal{M}} V_2^* (|1\rangle\langle 1| \otimes \sigma) V_2. \end{aligned} \tag{4.1}$$

These transformations do not take inputs of the same dimension, but this can easily be fixed by padding the input space of Q_1 with qubits that will later be traced out. The idea is that the verifier will accept in this protocol if and only if there are states ρ and σ that are consistent with a transcript of the **QIP** protocol where the verifier accepts. In this case, these states exist if and only if the verifier's private qubits do not change between V_1 and V_2 , which happens exactly when there are ρ and σ such that $F(Q_1(\rho), Q_2(\sigma))$ is large. To see this more formally, if the verifier accepts with certainty, then the output of Q_1 on the input state the proof system is exactly the output of Q_2 on some accepting configuration of the proof system. This is because there is a strategy for the prover that both causes the verifier to accept with probability one and does not change the state of the verifier's private space (because this is not allowed in the **QIP** model). Therefore, if the verifier accepts with certainty, there exist ρ, σ such that

$$Q_1(\rho) = Q_2(\sigma),$$

so that we have a valid instance of $\text{CI}_{1,b}$ in this case.

To see this more formally, the following lemma of Kitaev and Watrous [KW00] characterizes the probability that the verifier can be made to accept in a three-message quantum interactive proof system.

Lemma 4.2 (Kitaev and Watrous [KW00]). *Let V_1 and V_2 describe a verifier in a quantum interactive proof system, as shown in Figure 4.1, and let Q_1 and Q_2 be as given in Equation (4.1).*

The maximum probability that the verifier can be made to accept is

$$F_{\max}(Q_1, Q_2)^2.$$

Proof. By the definition of the model, the acceptance probability is given by a projector Π_{acc} onto the subspace of the verifier's private qubits with the first qubit in the state $|1\rangle$. By re-adding the prover's private space \mathcal{P} , we may assume that all states during the protocol are pure.. The maximum acceptance probability is defined as a maximum over the prover's strategy P and the initial state $|\phi\rangle$ corresponding to the prover's first message, which is in the $|0\rangle$ state for all but the message qubits and the prover's private qubits. Doing so, this quantity is

$$\begin{aligned} \max_{P, |\phi\rangle} \text{tr}(\Pi_{\text{acc}} V_2 P V_1 |\phi\rangle \langle \phi| V_1^* P^* V_2^*) &= \max_{P, |\phi\rangle, |\psi\rangle} |(\langle \psi | \Pi_{\text{acc}}) V_2^* P V_1 |\phi\rangle|^2 \\ &= \max_{P, |\phi\rangle, |\nu\rangle} |\text{tr } P(V_1 |\phi\rangle \langle \nu | V_2)|^2 \end{aligned} \quad (4.2)$$

where $|\nu\rangle$ is restricted to be an accepting state, i.e. the first of the verifier's private qubits is in the $|1\rangle$ state, and $|\phi\rangle$ is restricted to be an initial state, i.e. the verifier's private qubits are in the $|0\rangle$ state. If we first trace out the space \mathcal{V} in this equation, then by Lemma 3.7 the resulting quantity is equal to

$$\max_{P, |\phi\rangle, |\nu\rangle} |\text{tr } \text{tr}_{\mathcal{V}} P(V_1 |\phi\rangle \langle \nu | V_2)|^2 = \max_{|\phi\rangle, |\nu\rangle} \|\text{tr}_{\mathcal{V}} V_1 |\phi\rangle \langle \nu | V_2\|_{\text{tr}}^2 = \max_{|\phi\rangle, |\nu\rangle} \|\text{tr}_{\mathcal{M} \otimes \mathcal{P}} V_1 |\phi\rangle \langle \nu | V_2\|_{\text{tr}}^2,$$

where the final equality follows from the fact that the complementary reduced states of a pure state have the same singular values. Combining this with Equation (4.2) implies that the verifier can be made to accept with probability

$$\max_{|\phi\rangle, |\nu\rangle} \|\text{tr}_{\mathcal{M} \otimes \mathcal{P}} V_1 |\phi\rangle \langle \nu | V_2\|_{\text{tr}}^2 = \max_{\rho, \sigma} F(Q_1(\rho), Q_2(\sigma))^2 = F_{\max}(Q_1, Q_2)^2$$

where the first equality is by Lemma 3.21. Recall that the state $|\phi\rangle = |0\rangle \otimes |\phi'\rangle$ is a valid initial state and that $|\nu\rangle = |1\rangle \otimes |\nu'\rangle$ is a valid accepting state for the verifier: these conditions on the two pure states conform exactly to the states in the definition of Q_1 and Q_2 in Equation (4.1). \square

This lemma implies directly that $\text{CI}_{a,b}$ is **QIP**-hard for any probabilities a, b that suffice for the definition of **QIP** in Section 2.2. As it is known that these parameters may be any values such that $0 < b < a \leq 1$ with at least an inverse polynomial gap between a and b , this implies that $\text{CI}_{a,b}$ is hard for these same values.

To see that this problem is in **QIP**, consider the following protocol for the verifier, due to Kitaev and Watrous¹. The verifier starts with two circuits implementing

¹John Watrous, private communication

transformations Q_1 and Q_2 with

$$Q_i(\rho) = \text{tr}_{\mathcal{B}} U_i(\rho \otimes |0\rangle\langle 0|)U_i^*.$$

As a first step, the prover sends a state ρ , promised to be a such $Q_1(\rho)$ is close to a state in the image of Q_2 . The verifier computes

$$U_1(\rho \otimes |0\rangle\langle 0|)U_1^*$$

and sends the part of the state in \mathcal{B} to the prover. If there is a state σ such that $Q_2(\sigma) = Q_1(\rho)$, then the prover and verifier together hold a purification of this state. In this case, the prover can apply a unitary to his portion of the system to obtain a state corresponding to the purification that would have been obtained had the verifier instead evaluated

$$U_2(\sigma \otimes |0\rangle\langle 0|)U_2^*.$$

The prover performs such a computation, and sends the state in \mathcal{B} back to the verifier, who applies U_2^* and checks to see that the result is a valid initial state (i.e. his private qubits are in the $|0\rangle$ state).

The above argument implies that when $Q_1(\rho) = Q_2(\sigma)$ the prover can succeed with certainty. In the general case, the maximum probability that the verifier can be made to accept is given by Lemma 4.2, which in the case of this proof system is exactly $F_{\max}(Q_1, Q_2)^2$. This argument shows that $CI_{a,b}$ is in **QIP** for all $0 < b < a \leq 1$ with at least an inverse polynomial gap between a and b .

The preceding arguments imply that problem is complete for **QIP**. This argument appears implicitly in [KW00], where it is used to show that **QIP** \subseteq **EXP**.

Theorem 4.3 (Kitaev and Watrous [KW00]). *For any $0 < b < a \leq 1$, the problem $CI_{a,b}$ is **QIP**-complete.*

4.3 The swap test

The swap test provides a simple way to detect if two pure quantum states are the same. It was introduced by Buhrman et al. in the context of quantum communication complexity [BCWdW01], but it has also found applications in error correction [BBD⁺97] and in the estimation of various properties of quantum states [EAO⁺02]. Generalizations of the swap test to more than two inputs have also been considered [KNY08], though they will not be needed here.




Figure 4.3: A Constant-depth implementation of the W gate.




Figure 4.4: A circuit implementing the swap test.

An essential component of the swap test is the operator $W \in \mathbf{U}(\mathcal{H} \otimes \mathcal{H})$ that swaps the states in the two spaces, i.e. $W|a\rangle|b\rangle = |b\rangle|a\rangle$ for all $|a\rangle, |b\rangle \in \mathcal{H}$. Expressing W in the computational basis gives

$$W = \sum_{i,j} |j\rangle\langle i| \otimes |i\rangle\langle j|,$$

from which it is clear that W is both Hermitian and unitary.

As circuit depth is one of the primary considerations of this chapter, notice that as a circuit on $2n$ qubits, the operation W can be implemented in constant depth. Such an implementation can be given by n independent two-qubit swaps, as shown in Figure 4.3. These two qubit swaps can each be implemented in the usual basis of quantum gates using three controlled-not gates, as shown in Figure 2.3 of Section 2.1.

The swap test is built using a controlled- W operator to determine how close two states are to each other. A circuit performing the swap test is given in Figure 4.4.

An alternate characterization of the swap test is simply as a projective measurement onto the symmetric and antisymmetric subspaces of the systems it is applied to. Let $\{|i\rangle : 1 \leq i \leq d\}$ be a basis for \mathcal{H} . On $\mathcal{H} \otimes \mathcal{H}$, the symmetric and antisymmetric

subspaces are defined as

$$\begin{aligned}
(\mathcal{H} \otimes \mathcal{H})_{\text{sym}} &= \{|\psi\rangle \in \mathcal{H} \otimes \mathcal{H} : W|\psi\rangle = |\psi\rangle\} \\
&= \text{span}\{|i\rangle|j\rangle + |j\rangle|i\rangle : i \leq j\} \\
(\mathcal{H} \otimes \mathcal{H})_{\text{asym}} &= \{|\psi\rangle \in \mathcal{H} \otimes \mathcal{H} : W|\psi\rangle = -|\psi\rangle\} \\
&= \text{span}\{|i\rangle|j\rangle - |j\rangle|i\rangle : i < j\},
\end{aligned}$$

where these two subspaces represent the $\binom{d+1}{2}$ dimensional subspace corresponding to the +1 eigenvalues of W and the $\binom{d}{2}$ dimensional subspace corresponding to the -1 eigenvalues of W . From this representation it is easy to see that these two subspaces make up the whole space, i.e. that

$$\mathcal{H} \otimes \mathcal{H} = (\mathcal{H} \otimes \mathcal{H})_{\text{sym}} \oplus (\mathcal{H} \otimes \mathcal{H})_{\text{asym}},$$

and that the projections onto these subspaces are given by $(\mathbb{1} + W)/2$ and $(\mathbb{1} - W)/2$.

To see that this formulation of the swap test is equivalent to the circuit presented in Figure 4.4, consider the result of the measurement on the control qubit and work through the circuit in reverse. If this measurement result is $|0\rangle$, then the state of the control qubit after the controlled- W operation is $|0\rangle + |1\rangle$, up to normalization. As this is also the state of this qubit before the controlled- W operation, then applying W did not change the phase of the system, which implies that the input has been projected onto the symmetric subspace. On the other hand, if the measurement result is $|1\rangle$, then the state of the control qubit after the controlled- W operation is $|0\rangle - |1\rangle$. In this case, the system has been projected onto the subspace where applying W results in a phase of -1 , which is exactly the antisymmetric subspace of $\mathcal{H} \otimes \mathcal{H}$. Thus the circuit in Figure 4.4 applies the projective measurement given by $(\mathbb{1} + W)/2$ and $(\mathbb{1} - W)/2$, exactly as required.

This characterization immediately gives the probability that the swap test returns the antisymmetric outcome when applied to two pure states. To see this, observe that on pure states $|\psi\rangle$ and $|\phi\rangle$ this occurs with probability

$$\begin{aligned}
\frac{1}{2} \text{tr}((\mathbb{1} - W)|\psi\rangle\langle\psi| \otimes |\phi\rangle\langle\phi|) &= \frac{1}{2} (\text{tr} |\psi\rangle\langle\psi| \otimes |\phi\rangle\langle\phi| - \text{tr} |\phi\rangle\langle\psi| \otimes |\psi\rangle\langle\phi|) \\
&= \frac{1}{2} (1 - |\langle\phi|\psi\rangle|^2). \tag{4.3}
\end{aligned}$$

This result can be found in [BCWdW01].

The results that follow will make use of a generalization of this equation to the case of two (potentially entangled) mixed states. In the process of making this generalization the square in Equation (4.3) is lost, and so we will only show a lower bound.

Notice also that the requirement in the lemma that the two input states be reduced states of each other can be made in full generality, by applying the theorem to $\rho = \sigma \otimes \xi$ in the case that the input states are not entangled.

Lemma 4.4. *If $\rho \in \mathbf{D}(\mathcal{A} \otimes \mathcal{B})$ then a swap test on $\mathcal{A} \otimes \mathcal{B}$ returns the antisymmetric outcome with probability at least*

$$\frac{1}{2} - \frac{1}{2} F(\text{tr}_{\mathcal{A}} \rho, \text{tr}_{\mathcal{B}} \rho).$$

Proof. Let $|\psi\rangle \in \mathcal{A} \otimes \mathcal{B} \otimes \mathcal{C}$ be a purification of ρ , where \mathcal{C} is an arbitrary space with $\dim \mathcal{C} \geq \dim \mathcal{A} \dim \mathcal{B}$ to allow such a purification. The swap test measures the state on $\mathcal{A} \otimes \mathcal{B}$ with the projectors $\frac{1}{2}(\mathbb{1} - W)$ and $\frac{1}{2}(\mathbb{1} + W)$. As W is Hermitian and $W^2 = \mathbb{1}$, the antisymmetric outcome occurs with probability given by

$$\frac{1}{2} \text{tr} [(\mathbb{1}_{\mathcal{A} \otimes \mathcal{B} \otimes \mathcal{C}} - W \otimes \mathbb{1}_{\mathcal{C}})|\psi\rangle\langle\psi|] = \frac{1}{2} - \frac{1}{2} \langle\psi|W \otimes \mathbb{1}_{\mathcal{C}}|\psi\rangle. \quad (4.4)$$

The operator W is also unitary and so the states $|\psi\rangle$ and $(W \otimes \mathbb{1}_{\mathcal{C}})|\psi\rangle$ each purify both $\text{tr}_{\mathcal{A} \otimes \mathcal{C}} |\psi\rangle\langle\psi|$ and $\text{tr}_{\mathcal{B} \otimes \mathcal{C}} |\psi\rangle\langle\psi|$, and so by Uhlmann's Theorem (Theorem 3.18) Equation (4.4) implies

$$\frac{1}{2} - \frac{1}{2} \langle\psi|W \otimes \mathbb{1}_{\mathcal{C}}|\psi\rangle \geq \frac{1}{2} - \frac{1}{2} F(\text{tr}_{\mathcal{A} \otimes \mathcal{C}} |\psi\rangle\langle\psi|, \text{tr}_{\mathcal{B} \otimes \mathcal{C}} |\psi\rangle\langle\psi|). \quad (4.5)$$

Finally, by observing that

$$\text{tr}_{\mathcal{A} \otimes \mathcal{C}} |\psi\rangle\langle\psi| = \text{tr}_{\mathcal{A}} (\text{tr}_{\mathcal{C}} |\psi\rangle\langle\psi|) = \text{tr}_{\mathcal{A}} \rho$$

and

$$\text{tr}_{\mathcal{B} \otimes \mathcal{C}} |\psi\rangle\langle\psi| = \text{tr}_{\mathcal{B}} (\text{tr}_{\mathcal{C}} |\psi\rangle\langle\psi|) = \text{tr}_{\mathcal{B}} \rho,$$

Equation (4.5) is the lower bound in the statement of the lemma. \square

To see that this lemma generalizes Equation (4.3) up to a square, consider the input state $\rho = |\phi\rangle\langle\phi| \otimes |\psi\rangle\langle\psi|$ and apply the definition of the Fidelity.

The main result of this chapter concerns log-depth circuits, and so it is important to note that the swap test can be performed in log-depth. As discussed in Proposition 2.1, controlled operations on n qubits can be implemented by adding only log-depth overhead. This implies that the swap test can be implemented with a log-depth circuit, as with the exception of the overhead for the controlled operation, Figure 4.4 provides a constant depth circuit. Additionally, this implies that if the unbounded fan-out gate is allowed into the model of computation, this overhead can be reduced to constant-depth, and so in this case the swap test can be performed with a constant depth circuit. As it is not at all clear that including this gate produces a reasonable circuit model, any results that depend on the addition of this gate to the circuit model are clearly marked with this requirement.

4.4 Reduction to logarithmic depth

In this section the reduction from the general close images problem to the log-depth restriction of the problem is described. This is done in the case of one-sided error, i.e. the problem $CI_{1,b}$, where the two circuits either have intersecting images or the largest fidelity between any two outputs is at most b . The fact that this restricted version of the problem is hard implies that we obtain the desired hardness result even when we assume that the input instance has $\alpha = 0$.

The general idea behind the construction is to simply slice the circuits of an instance of $CI_{1,b}$ into constant-depth pieces and run them in parallel. These circuits will have much larger input spaces than the original circuit, but they are able to simulate the original circuit. This is due to the fact that if for each of the constant depth pieces, the input to one piece of the circuit is identical to the output of the previous piece, then the output of the final piece of the circuit will be equal to the output of the original circuit. This need not be the case if the intermediate inputs are not the outputs of the previous pieces, and so additional tests that ensure these inputs are at least close to the desired states are required. The swap test will be used extensively to perform these tests, though care must be taken to ensure that the resulting circuits have logarithmic depth.

This construction is similar to an idea of Gottesman and Chuang [GC99] in which a circuit is sliced into constant depth pieces with teleportation used to transfer the information between the pieces. Conditioned on all of the teleportations not requiring a Pauli correction, this process produces a constant depth simulation (as a mixed state circuit) of the original circuit. This process, however, does not perform any verification that the teleportation operations were successful, and so the resulting simulation is only accurate with exponentially small probability. This technique was used by Terhal and DiVincenzo [TD04] to show that exactly simulating these circuits in polynomial time leads to unexpected complexity-theoretic results ($\mathbf{P} = \mathbf{PP}$). The circuit used in the paper does not conform to our circuit model, however, as an infinite set of gates is allowed into the model. This difficulty was eliminated by Fenner et. al [FGHZ05] who implemented the construction in a circuit model equivalent to the one used here. This paper also provides an approximate simulation of the constant depth circuits in classical polynomial time, which suggests that extremely short quantum circuits are not interesting from a computational perspective.

For the reduction of CLOSE IMAGES to circuits of logarithmic depth we use a similar technique to that of Gottesman and Chuang [GC99] of slicing the circuit into pieces.




Figure 4.5: The original circuit Q_i decomposed into constant depth unitary circuits.

In the place of teleportation, however, we demand that the inputs to each of the pieces are provided before the start of the computation. These intermediate inputs are then verified using a verification procedure to ensure that the constructed circuit faithfully simulates the original circuit.

To describe the reduction, let Q_1 and Q_2 be the circuits from an instance of $CI_{1,b}$, and let n be the size of Q_1 and Q_2 , by padding the smaller circuit, if necessary. In order to slice the circuits into pieces it is assumed that Q_1 and Q_2 first introduce any necessary ancillary qubits, then apply local unitary gates, and finally trace out any qubits that are not part of the output. This form for a circuit is shown in Figure 2.7, and, as discussed in Section 2.1, this can be assumed with no loss of generality with only polynomial overhead, by delaying any partial trace operations until the end of the computation and introducing any needed ancillary qubits at the start of the computation.

A simple way to decompose Q_1 into constant depth pieces is to simply let each gate of Q_1 be a single piece in the decomposition. Let U_1, U_2, \dots, U_n be these pieces, with the additional complication that the operation U_1 both adds the ancillary qubits and performs the first gate of the circuit. In a similar way, Q_2 can be decomposed into constant depth pieces V_1, V_2, \dots, V_n . Such a decomposition is shown in Figure 4.5. If the circuits Q_1 and Q_2 implement transformations in $\mathbf{T}(\mathcal{H}, \mathcal{K})$, then as we have assumed that they are in Stinespring form, these circuits first introduce ancillary qubits in some space \mathcal{A} , apply some unitary in $\mathbf{U}(\mathcal{H} \otimes \mathcal{A}, \mathcal{B} \otimes \mathcal{K})$, and finally trace out the space \mathcal{B} . It can be assumed that the spaces \mathcal{A} and \mathcal{B} are of the same dimension for both Q_1 and Q_2 , once again by padding the smaller circuit with unused ancillary qubits that are later traced out. This implies that the spaces $\mathcal{H} \otimes \mathcal{A}$ and $\mathcal{B} \otimes \mathcal{K}$ are isomorphic. Using these spaces, and implicitly this isomorphism, we have

$$\begin{aligned} U_1, V_1 &\in \mathbf{U}(\mathcal{H}_1, \mathcal{B}_1 \otimes \mathcal{K}_1) \\ U_i, V_i &\in \mathbf{U}(\mathcal{H}_i \otimes \mathcal{A}_i, \mathcal{B}_i \otimes \mathcal{K}_i) \quad \text{for } 2 \leq i \leq n, \end{aligned}$$

where the subscripted spaces are copies of the non-subscripted spaces that hold the




Figure 4.6: Testing that the output of U_i is close to the input of U_{i+1} . The inputs $|\psi_j\rangle$ are the ideal inputs to U_j , and are labelled for clarity only – no assumptions are made about these states. Qubits that do not reach the right edge are traced out, but for clarity these operations are not shown in the figure.

input or output of one of the pieces of the original circuits. As an example of this notation, if $\rho \in \mathbf{D}(\mathcal{H})$, then the output of the circuit Q_1 on ρ is given by

$$\text{tr}_{\mathcal{B}_n} U_n U_{n-1} \cdots U_1 \rho U_1^* U_2^* \cdots U_n^*, \quad (4.6)$$

and the output of Q_2 is given by the same expression with the V_i in place of the unitaries U_i .

This decomposition of Q_1 and Q_2 will be used to construct circuits C_1 and C_2 that have logarithmic depth and still, in some sense, faithfully implement Q_1 and Q_2 . This is done by placing the circuits corresponding to U_1, \dots, U_n in parallel, and tracing out all the qubits that are not in the output of U_n . Such a circuit is constant depth, but does not necessarily output a state in the image of Q_1 , as the input to U_{i+1} is not necessarily close to the output from U_i . This problem is solved by comparing the output of U_i to the input to U_{i+1} using the swap test. The swap test will fail to detect the case that the two inputs are different with some probability, but in Section 4.5 it is shown that this probability can be upper bounded by an expression involving the trace norm of the two states.

In order for this comparison procedure to be done in log-depth an auxiliary input is first compared against the input to U_{i+1} and then held in reserve to compare to the output of U_i . This strategy avoids the comparing the input to U_{i+1} directly to the output of U_i , which leads to a circuit of linear depth. This depth reduction comes at a cost, however, as the two states are always compared through an intermediary state, which can at worst halve the probability of detecting when these two states differ, since one test is replaced with two. This constant loss will not affect the main result in




Figure 4.7: The outputs of C_1 and C_2 . The dummy $|0\rangle$ qubits of one circuit line up with the outputs of the swap tests of the other.

a significant way. An example of the construction used to ensure that the output of U_i agrees with the input to U_{i+1} is given in Figure 4.6.

To simplify the analysis of the constructed circuits these two tests are controlled so that exactly one of the two tests is performed. This will increase the failure probability by another factor of two, but allows the analysis of each swap test to ignore the effect of the other. To implement this scheme a control qubit is used so that either the first or the second test is performed between every two pieces U_i, U_{i+1} of the circuit. If a test is not performed, then the value of the output qubit of the swap test is left unchanged, and so the result of the test is a qubit in the $|0\rangle$ state. In the case that a test is performed, the output is either $|0\rangle$ for the symmetric subspace (i.e. the two states are the same) or $|1\rangle$ for the antisymmetric subspace (i.e. the two states differ). These outputs are classical values, but they are treated as the two orthogonal quantum states $|0\rangle$ and $|1\rangle$ for convenience. Controlled application of these swap tests can be implemented in log-depth using the techniques described in Proposition 2.1.

After adding these two tests between each piece of the circuit there is one final modification to obtain the circuits C_1 and C_2 . If any of the swap tests fail, i.e. detect states in the antisymmetric subspace, then they will output qubits in the $|1\rangle$ state. As yes instances of $CI_{1,b}$ have outputs that are close together, we can ensure that if any of the swap tests fail then the outputs of the constructed circuits are far apart by adding dummy qubits in the $|0\rangle$ state to be compared to the outputs of the swap tests in the other circuit. The arrangement of these dummy qubits is shown in Figure 4.7.

The constructed circuits C_1 and C_2 are obtained by decomposing Q_1 and Q_2 into constant depth pieces, inserting the swap tests shown in Figure 4.6, and adding dummy qubits to ensure that the swap tests in the other circuit do not fail. The final circuit C_1 constructed from Q_1 , including these dummy qubits, is shown in Figure 4.8, the circuit C_2 is similar, with the exception that the qubits corresponding to the swap test

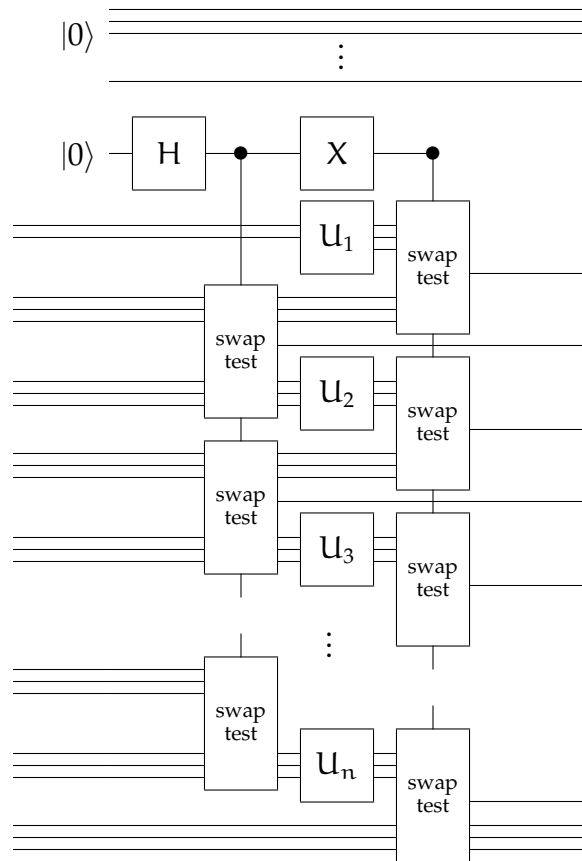


Figure 4.8: The constructed circuit C_1 . In the circuit C_2 the dummy zero output qubits are swapped with the qubits containing the results of the swap tests. All qubits that do not reach the right edge of the figure are traced out, but this is for notational convenience only: the constructed circuits are in Stinespring form.

outputs and dummy qubits have been swapped, as shown in 4.7. At the end of these circuits, all qubits are traced out except the output (in the space \mathcal{K}_n) of U_n or V_n , the output of the swap tests, and the dummy zero qubits. Notice that the circuit C_i can be computed from Q_i in polynomial time, as it is simply a rearrangement of the gates of the original circuit with the addition of a linear number of extra gates.

If the outputs of the circuits C_1 and C_2 are close together, then at an intuitive level the output of the swap tests in each circuit must be close to zero and the output of U_n and V_n must also be close. If the swap tests do not fail with high probability (i.e. the outputs are close to zero), then these circuits will more or less faithfully reproduce the output of Q_1 and Q_2 . Thus, in the case that the outputs of C_1 and C_2 can be made close, we will be able to argue that the output of Q_1 and Q_2 can also be made close. Proving that this picture is accurate forms the content of Section 4.5.

In the other direction, it is much simpler to argue that if there are states $\rho, \xi \in \mathbf{D}(\mathcal{H})$ such that $Q_1(\rho) = Q_2(\xi)$, then there are similar states for the constructed circuits C_1 and C_2 . This is the content of the following proposition.

Proposition 4.5. *If there exist states ρ, ξ such that $Q_1(\rho) = Q_2(\xi)$, then there exist states ρ', ξ' such that $C_1(\rho') = C_2(\xi')$.*

Proof. To prove the proposition, states ρ' and ξ' are constructed so that

$$C_1(\rho') = |0\rangle\langle 0| \otimes Q_1(\rho) = |0\rangle\langle 0| \otimes Q_2(\xi) = C_2(\xi'). \quad (4.7)$$

To find these states, notice that both the output fidelity and construction of the circuits C_i do not change if additional ancillary qubits are added to the circuits Q_i to allow purification of the input states, so long as these extra qubits are traced out at the end of the circuit. These purifications are pure states and all operations performed during the circuit Q_i are unitary, which implies that the intermediate states of the circuits are also pure.

If a purification of the state ρ is $|\psi\rangle$, then by providing the pure state

$$|\gamma\rangle = |\psi\rangle \otimes (U_1|\psi\rangle)^{\otimes 2} \otimes (U_2U_1|\psi\rangle)^{\otimes 2} \otimes \cdots \otimes (U_{n-1}U_{n-2}\cdots U_1|\psi\rangle)^{\otimes 2} \quad (4.8)$$

as input to C_1 , the output of each block of the circuit will be identical to the input to the next block, by construction. All but the first piece of this state is repeated twice: this is to provide the correct intermediate inputs that are used by the swap tests to compare the output of one block to the input of the next, as shown in Figure 4.6. This ensures that all the swap tests will succeed with probability one, which can be seen from Equation (4.3). Let this constructed state ρ' be given by $\rho' = |\gamma\rangle\langle\gamma|$.

It remains to check that on ρ' that C_1 simulates Q_1 on ρ . By the construction of C_1 and ρ' , the output is exactly

$$C_1(\rho') = |0\rangle\langle 0| \otimes \text{tr}_{\mathcal{B}_n} U_n U_{n-1} \cdots U_1 \rho U_1^* U_2^* \cdots U_n^*,$$

which is equal to the output of Q_1 on ρ , up to a number of qubits in the $|0\rangle$ state, which correspond to the dummy qubits and the outputs of the swap tests. By the symmetry of the construction, a state ξ' for the circuit C_2 can be constructed from ξ in the same way, and for these constructed ρ' and ξ' Equation 4.7 is satisfied, which completes the proof. \square

This proposition implies that the reduction presented in this section maps yes instances of $CI_{1,b}$ to yes instances of $\text{LOG-DEPTH } CI_{1,b}$. The remaining direction is considerably more intricate, and forms the content of the next section.

4.5 Correctness of the reduction

In this section it is argued formally that the reduction presented in the previous section maps negative instances (Q_1, Q_2) of the close images problem, i.e. those for which $Q_1(\rho)$ and $Q_2(\sigma)$ are far apart for all states ρ and σ , to negative instances of log-depth close images. Less formally, it is shown that if the images of the original circuits Q_1 and Q_2 are far apart then so must be the images of the constructed circuits C_1 and C_2 . This argument is technical but fairly straightforward: the basic idea is to transform the fidelity to the trace norm using the Fuchs-van de Graaf Inequalities (Theorem 3.22), apply the triangle inequality to reduce the problem to individual blocks of the circuits C_1 and C_2 , and finally return to the fidelity with another application of the Fuchs-van de Graaf Inequalities. As might be expected, this technique results in poor error bounds: the value b , which is the maximum output fidelity allowed in a no instance of the close images problem, ends up polynomially close to 1. This is dealt with using a parallelization technique due to Kitaev and Watrous [KW00] that improves the value of b to any constant $b > 0$.

As the constructed circuits C_1 and C_2 can be used to simulate Q_1 and Q_2 , by Proposition 4.5, the result is obtained by arguing that either the outputs of C_1 and C_2 correspond to the outputs of Q_1 and Q_2 , respectively, or the outputs of C_1 and C_2 are far apart. In the case that this simulation is not faithful it is shown that some swap test fails with non-negligible probability. This implies that outputs of the constructed circuits are far apart, as the failing swap test produces a state of the form $(1 - p)|0\rangle\langle 0| + p|1\rangle\langle 1|$ that has low fidelity with the corresponding dummy zero qubit of the other circuit.

Lemma 4.4 which describes the behaviour of the swap test on mixed states does not immediately apply to the circuits C_1 and C_2 . This is because in these circuits the output of one block of the circuit is not directly compared to the input to the next block, but instead each of these states are with probability $1/2$ compared to some intermediate value. In order to deal with this difficulty, we use the Fuchs-van de Graaf Inequalities (Theorem 3.22) to translate the fidelity to a relation involving the trace norm, which we can then apply the triangle inequality to. The triangle inequality shows that at least one of the two swap tests fails with probability bounded below by an expression involving the fidelity, which is lower bounded by Lemma 4.4. In the proof of the following corollary the reduced states of various parts of the input to either of the circuits C_1 or C_2 are used, but no assumption is made on the form of the input state, i.e. it is not assumed that the input is separable across the block boundaries of the circuit. For instance, the density matrices ρ_i , σ_i , and ξ_i that appear in the lemma may be part of some larger entangled pure state, so that the failure probabilities of the two swap tests need not be independent. To clarify the notation used below, the state ξ_i is the output of the $(i - 1)$ st block (i.e. the output of U_{i-1} in Figure 4.2), the state ρ_i is the input to the i th block, and the state σ_i is the intermediate state used to indirectly compare ρ_i and ξ_i .

Corollary 4.6. *If $|\psi\rangle$ is input to the circuit C_a for $a \in \{1, 2\}$, with ρ_i the reduced state of $|\psi\rangle\langle\psi|$ on $\mathcal{H}_i \otimes \mathcal{A}_i$, then at least one of the swap tests on the i th block of C_a fails with probability at least*

$$\frac{1}{128} \|U_i \rho_{i-1} U_i^* - \rho_i\|_{\text{tr}}^2.$$

Proof. In the i th block of C_a there are two inputs to the first swap test: let the reduced density operators of these inputs be ρ_i and σ_i , as discussed above. The inputs to the second swap test are then given by σ_i and $U_i \rho_{i-1} U_i^* = \xi_i$. As exactly one of these tests is performed we do not need to consider the effect of the first test on the state when considering the second test, and so the same input state σ_i is used for both swap tests.

By Lemma 4.4, the failure probability of the first and second tests, when performed, are at least $\frac{1}{2}(1 - F(\rho_i, \sigma_i))$ and $\frac{1}{2}(1 - F(\sigma_i, \xi_i))$, respectively. Thus, the probability p that at least one of these tests fails, given that each of them is performed with probability $1/2$, is at least

$$p \geq \frac{1}{2} \max \left\{ \frac{1}{2}(1 - F(\sigma_i, \xi_i)), \frac{1}{2}(1 - F(\rho_i, \sigma_i)) \right\} = \frac{1}{4} (1 - \min\{F(\sigma_i, \xi_i), F(\rho_i, \sigma_i)\}).$$

By the Fuchs-van de Graaf inequalities (Theorem 3.22), this fidelity may be replaced

by the trace norm. Doing so, we obtain

$$p \geq \frac{1}{32} \max(\|\sigma_i - \xi_i\|_{\text{tr}}^2, \|\rho_i - \sigma_i\|_{\text{tr}}^2),$$

where we have made use of Bernoulli's inequality to show that $\sqrt{1-x} \leq 1-x/2$ when simplifying this equation. Finally, as this maximum must be at least the average of the two values,

$$p \geq \frac{1}{32} \left(\frac{\|\sigma_i - \xi_i\|_{\text{tr}}}{2} + \frac{\|\rho_i - \sigma_i\|_{\text{tr}}}{2} \right)^2 \geq \frac{1}{128} \|\rho_i - \xi_i\|_{\text{tr}}^2,$$

where the last inequality is the triangle inequality. \square

By repeatedly applying some of the properties of the trace norm discussed in Chapter 3 it is somewhat tedious but not difficult to use Corollary 4.6 to bound the distance between the images of the constructed circuits in terms of the distance between the images of the original circuits. In the following theorem n is the size of the circuits Q_1 and Q_2 , as in Section 4.4. Informally, this theorem states that “no” instances of the problem $\text{CI}_{1,b}$ are mapped to “no” instances of the problem $\text{LOG-DEPTH CI}_{1,b'}$ with the resulting value b' only polynomially closer to 1 than the value b , which shows the **QIP-completeness** of log-depth close images for these large values of b .

Theorem 4.7. *If $F(Q_1(\rho_0), Q_2(\xi_0)) < 1 - c$ for all $\rho_0, \xi_0 \in \mathbf{D}(\mathcal{H})$ then*

$$F(C_1(\rho), C_2(\xi)) < 1 - \frac{c^2}{576n^2}$$

for all $\rho, \xi \in \mathbf{D}(\mathcal{H} \otimes \bigotimes_{i=2}^{2n} \mathcal{H}_i \otimes \mathcal{A}_i)$

Proof. Let ρ and ξ be inputs to C_1 and C_2 , and let ρ_i, ξ_i be the reduced states of these inputs on $\mathcal{H}_i \otimes \mathcal{A}_i$ for $1 \leq i \leq 2n$, where the states for $i > n$ are the inputs that are only used by the swap tests, which will not be referred to explicitly. That is, ρ_i and ξ_i for $1 \leq i \leq n$ are the portions of the state that are input to the unitaries U_i and V_i that make up the circuits Q_1 and Q_2 , as shown in Figure 4.2. The output of the circuits C_1 and C_2 is given by the output qubits corresponding to the swap tests as well as the states $\text{tr}_{\mathcal{B}_n} \rho_n$ and $\text{tr}_{\mathcal{B}_n} \xi_n$, where \mathcal{B}_n is simply the space that is traced out to obtain the output from the unitary representations of the original circuits. In this notation, ρ_1 and ξ_1 are the inputs to the first pieces U_1 and V_1 of the constructed circuits C_1 and C_2 . These two states are density matrices in $\mathbf{D}(\mathcal{H}_1) \cong \mathbf{D}(\mathcal{H})$, and we can also consider them as potential inputs to the original circuits Q_1 and Q_2 .

By the condition on the fidelity of Q_1 and Q_2 in the statement of the theorem, as well as the Fuchs-van de Graaf inequalities (Theorem 3.22), we have

$$2c < \|Q_1(\rho_1) - Q_2(\xi_1)\|_{\text{tr}}.$$

Using the triangle inequality we can relate this to the distance between the constructed circuits. By adding terms and simplifying, we obtain

$$\begin{aligned} 2c &< \|Q_1(\rho_1) - \text{tr}_{\mathcal{B}_n} \rho_n + \text{tr}_{\mathcal{B}_n} \xi_n - Q_2(\xi_1) + \text{tr}_{\mathcal{B}_n} \rho_n - \text{tr}_{\mathcal{B}_n} \xi_n\|_{\text{tr}} \\ &\leq \|Q_1(\rho_1) - \text{tr}_{\mathcal{B}_n} \rho_n\|_{\text{tr}} + \|\text{tr}_{\mathcal{B}_n} \xi_n - Q_2(\xi_1)\|_{\text{tr}} + \|\text{tr}_{\mathcal{B}_n} \rho_n - \text{tr}_{\mathcal{B}_n} \xi_n\|_{\text{tr}}. \end{aligned}$$

We now observe that $\|\text{tr}_{\mathcal{B}_n} \rho_n - \text{tr}_{\mathcal{B}_n} \xi_n\|_{\text{tr}} \leq \|C_1(\rho) - C_2(\xi)\|_{\text{tr}}$ by the monotonicity of the trace norm under quantum operations (Theorem 3.8), since the former can be obtained from the later by applying the parital trace on the appropriate space. Using this we have

$$2c < \|Q_1(\rho_1) - \text{tr}_{\mathcal{B}_n} \rho_n\|_{\text{tr}} + \|\text{tr}_{\mathcal{B}_n} \xi_n - Q_2(\xi_1)\|_{\text{tr}} + \|C_1(\rho) - C_2(\xi)\|_{\text{tr}} \quad (4.9)$$

As the three terms on the right are nonnegative, at least one of them must be larger than the average $2c/3$. If $\|C_1(\rho) - C_2(\xi)\|_{\text{tr}} > 2c/3$ then $F(C_1(\rho), C_2(\xi)) < 1 - c^2/18$, by Theorem 3.22, and there is nothing left to prove.

The cases where one of the first two terms of (4.9) exceeds $2c/3$ are symmetric, and so we can consider only the first term. Expanding $Q_1(\rho_1)$ in terms of the U_i , we obtain

$$\begin{aligned} \frac{2c}{3} &< \|Q_1(\rho_1) - \text{tr}_{\mathcal{B}_n} \rho_n\|_{\text{tr}} \\ &= \|\text{tr}_{\mathcal{B}_n} U_n U_{n-1} \cdots U_1 \rho_1 U_1^* U_2^* \cdots U_n^* - \text{tr}_{\mathcal{B}_n} \rho_n\|_{\text{tr}} \\ &\leq \|U_n U_{n-1} \cdots U_1 \rho_1 U_1^* U_2^* \cdots U_n^* - \rho_n\|_{\text{tr}}, \end{aligned}$$

where once again the monotonicity of the trace norm under the partial trace (Theorem 3.8) has been used. By adding and subtracting the term $U_n \cdots U_2 \rho_1 U_2^* \cdots U_n^*$ inside the norm, and then applying both the triangle inequality and the unitary invariance of the trace norm, we have

$$\frac{2c}{3} < \|U_1 \rho_1 U_1^* - \rho_2\|_{\text{tr}} + \|U_n U_{n-1} \cdots U_2 \rho_2 U_2^* U_3^* \cdots U_n^* - \rho_n\|_{\text{tr}}.$$

Here the unitary invariance of the trace norm has been used to discard the operators U_2, \dots, U_n from the first term. Repeating this strategy, by adding and subtracting the term $U_n \cdots U_3 \rho_3 U_3^* \cdots U_n^*$ and once again applying the triangle inequality results in

$$\frac{2c}{3} < \|U_1 \rho_1 U_1^* - \rho_2\|_{\text{tr}} + \|U_2 \rho_2 U_2^* - \rho_3\|_{\text{tr}} + \|U_n U_{n-1} \cdots U_3 \rho_3 U_3^* U_4^* \cdots U_n^* - \rho_n\|_{\text{tr}}.$$

Continuing in this fashion we have

$$\frac{2c}{3} < \sum_{i=1}^{n-1} \|U_i \rho_i U_i^* - \rho_{i+1}\|_{\text{tr}}.$$

As all terms in this sum are nonnegative, there must be at least one term in the sum that exceeds $2c/(3n)$, as this is a lower bound on the average of all terms. Thus, for some value of i , we have $\|U_i \rho_i U_i^* - \rho_{i+1}\|_{\text{tr}} > 2c/(3n)$, and so by Corollary 4.6 one of the corresponding swap tests fails with probability $p > c^2/(288n^2)$. The qubit representing the output value of this swap test is then of the form $(1-p)|0\rangle\langle 0| + p|1\rangle\langle 1|$, and so, by the monotonicity of the fidelity under the partial trace (Theorem 3.20),

$$F(C_1(\rho), C_2(\xi)) \leq F((1-p)|0\rangle\langle 0| + p|1\rangle\langle 1|, |0\rangle\langle 0|) = \sqrt{1-p} < 1 - \frac{c^2}{576n^2},$$

as in the statement of the theorem. \square

By combining Theorem 4.7 with Proposition 4.5, and the multiplicativity of the maximum output fidelity of two transformations, given as Theorem 3.24, we obtain the main result of this chapter.

Corollary 4.8. $\text{LOG-DEPTH CI}_{a,b}$ is **QIP-complete** for any constants $0 < b < a \leq 1$.

Proof. Theorem 4.7 together with Proposition 4.5 establish the result that $\text{CI}_{1,b}$ reduces to $\text{LOG-DEPTH CI}_{1,b'}$ for $b' \geq 1 - (1-b)^2/(576n^2)$, where n is an upper bound on the size of the circuits.

The value of b' can be improved using Theorem 3.24 of Kitaev, Shen, and Vayli [KSV02], which shows that if the circuits C_1 and C_2 are repeated r times in parallel, then the maximum output fidelity is

$$\max_{\rho, \xi} F(C_1^{\otimes r}(\rho), C_2^{\otimes r}(\xi)) = \max_{\rho, \xi} F(C_1(\rho), C_2(\xi))^r.$$

This implies that $\text{CI}_{1,b}$ reduces to $\text{LOG-DEPTH CI}_{1,b'}$ for all

$$b' \geq \left(1 - \frac{(1-b)^2}{576n^2}\right)^r,$$

and so, by taking r polynomially large in n and b , we may take $b' \leq b$, which implies that $\text{CI}_{1,b}$ reduces to $\text{LOG-DEPTH CI}_{1,b}$.

This shows that $\text{LOG-DEPTH CI}_{1,b}$ is then **QIP-complete** for any constant $0 < b < 1$, as by Theorem 4.3 due to Kitaev and Watrous [KW00], $\text{CI}_{a,b}$ is **QIP-complete** for all

$0 < b < a \leq 1$. Generalizing the log depth close images problem for all values of a gives the problem LOG-DEPTH $CI_{a,b}$ for $0 < b < a \leq 1$, which is also complete for **QIP** as it can be obtained by weakening the promise. This more general problem is in **QIP** as it is a restriction of $CI_{a,b}$ to log-depth circuits. \square

As the circuits constructed by the reduction only make use of logarithmic depth when performing swap tests, and the controlled swap operations performed by these tests can be accomplished in constant depth using unbounded fan-out gates, as described in Proposition 2.2, the following Corollary follows immediately from the previous one.

Corollary 4.9. *The problem CONST-DEPTH $CI_{a,b}$ on circuits with the unbounded fan-out gate is **QIP**-complete for any constants $0 < b < a \leq 1$.*

4.6 Conclusion

In this chapter, the problem CLOSE IMAGES has been introduced. This problem asks: given two quantum channels, as mixed state quantum circuits, how close are the images of the two channels? More concretely, how large is the minimum distance of any two outputs of the channels, where the fidelity is used as the notion of distance. This problem is complete for the class **QIP** [KW00].

The main result of the chapter is a reduction of this problem to the case of logarithmic depth circuits. This reduction works only for the case that the two circuits are promised to either have intersecting images or images that are far apart, but a hardness result on this special case also implies the hardness of the general problem. This restriction is necessary to the proof that the reduction is correct as it enables the use of a parallel repetition technique to strengthen the promise of the class of instances that is shown to be hard.

This hardness result is the base for the main result of the next chapter, which is that the computational problem of distinguishing quantum circuits is also **QIP**-hard. The result of this chapter enables the hardness of this distinguishability problem to be extended even to the case of channels implemented by log-depth circuits.

Chapter 5

Distinguishability of Quantum Computations

A natural problem in quantum information is to discriminate between two quantum channels. In the model where channels are represented as quantum circuits, this is the computational distinguishability problem on channels, though the difficulty of the problem does not change if the circuits are replaced by black boxes that can be performed, but not inspected. The main result of this chapter is that this problem is computationally very difficult, as it is complete for the class **QIP** of problems having quantum interactive proof systems. This also implies [JJUW09] that this problem is complete for **PSPACE**, which gives a new quantum characterization for a classical complexity class.

The majority of the results in this chapter are in collaboration with John Watrous, and have been published in [RW05]. The results in Section 5.6 have appeared in [Ros08b].

Contents

5.1	Overview of distinguishability problems	100
5.2	Quantum circuit distinguishability	102
5.3	QIP protocol	104
5.4	Reduction from Close Images	107
5.5	Correctness of the reduction	110
5.6	Distinguishing log-depth computations	115
5.7	Conclusion	116

5.1 Overview of distinguishability problems

The problem of distinguishing two computations is central to computer science. This is the problem that asks, given two computations represented in some way, do the two computations always produce results that are close together, or are there inputs on which they act differently? This is an important problem both theoretically and practically: determining if some process has been correctly implemented is one of the most important tasks in experimental quantum computing.

The problem of estimating an unknown quantum channel is known as *process tomography* [CN97, PCZ97]. All known approaches for approximate process tomography unfortunately require exponential time. This is not a surprise, however, as the complete characterization of a quantum channel on n qubits requires an exponential number of parameters. The main result of this chapter is that even the simpler task of *distinguishing* two quantum channels, given as mixed-state circuits, is computationally intractable. That the distinguishability problem reduces to process tomography is clear: one way to solve the problem is characterize the two channels with enough accuracy to detect the case that they are far apart.

Returning to classical complexity theory, the most basic distinguishability problem asks: given two classical deterministic circuits, is there an input on which they produce different outputs? This problem is in **NP**, as given such an input a verifier can both simulate the two circuits and check to see if they agree, all in polynomial time. This problem is also complete for **NP**. To see this, notice that a circuit is satisfiable if and only if it is distinguishable from the circuit that always outputs false. The satisfiability problem is the original **NP**-complete problem [Coo71], and so the problem of distinguishing classical deterministic computations must also be **NP**-complete.

Adding randomness to the circuit model, in the form of gates that produce unbiased coin flips, appears to increase the difficulty of the problem. Averaged over the values of the coin flips, the two circuits in the distinguishability problem produce probability distributions – distinguishing these distributions is also a nontrivial problem. To avoid the problem of distinguishing randomized circuits being artificially difficult, the additional promise is given that the two probabilistic circuits to be distinguished either produce output distributions that are very close for any input, or that there exists some input on which the distributions produced are far apart. The usual measure of distance on probability distributions is the difference in the ℓ_1 norm, that is defined, for distributions p, q , as $\|p - q\|_{\ell_1} = \sum_x |p(x) - q(x)|$. This norm is simply the classical analogue of the trace norm of the difference of two density operators. Even when given




Figure 5.1: Complexity classes and distinguishability problems.

an input on which they produce maximally distant output distributions, distinguishing two randomized circuits is complete for the class **SZK** of problems with statistical zero-knowledge proof systems [SV03]. This is in sharp contrast to the deterministic case where a verifier can check, given an input, whether or not two deterministic circuits produce the same output. The best complexity theoretic upper bound known for the randomized circuit distinguishability problem is **AM**, by simply having the prover first produce an optimal distinguishing input for the two circuits and then performing the **SZK** protocol due to Sahai and Vadhan [SV03] for the statistical difference problem that remains. This problem is not known to be complete for **AM**.

Extending this problem in the direction of quantum information leads first to the natural problem of distinguishing unitary circuits. As in the case of randomized classical circuits, the outputs of unitary quantum circuits are nontrivial to compare. Once again, a promise is used to keep the problem from being artificially difficult. In this case, the distance measure can be either the diamond norm or the trace norm, as it is known that they agree on unitary transformations [AKN98, CPR00], and the promise is that the distance between the two given transformations, which is given by

$$\max_{|\psi\rangle \in \mathcal{H}} \|U|\psi\rangle\langle\psi|U^* - V|\psi\rangle\langle\psi|V^*\|_{\text{tr}},$$

is either close to zero or close to two. The input in this equation may be restricted to pure states by Theorem 3.17. This problem has been shown to be complete for **QMA** by Janzing, Wocjan, and Beth [JWB05]. This result implies that, given an optimal input state, a verifier with access to a quantum computer can solve the distinguishability problem on unitary circuits. This is not unexpected, as unitary circuits are similar to deterministic quantum computation.

Adding both the elements of randomness and quantumness to the computations being distinguished results in a significantly more difficult problem than adding just

one of the two elements. This is the distinguishability problem for mixed-state quantum circuits, as defined in Section 2.1. Once again we require a promise to avoid an artificially difficult problem: the circuits to be distinguished can be assumed to have a diamond norm difference that is either close to zero or close to two. It is surprising that this problem is **QIP**-complete. The class **QIP** is believed to be much larger than the classes **QMA** and **AM**. As evidence of this, the polynomial hierarchy collapses if **QIP** is contained in either of these classes. These complexity classes, known inclusions among them, and the distinguishability problems related to them, are summarized in Figure 5.1. Removing either the quantum computing, leaving randomized circuits, or randomness, leaving unitary quantum circuits, seems to change the essential character of the problem: the hardness appears to lie in the combination of both of these ingredients.

The main result of this section is showing the **QIP**-completeness for the mixed-state circuit distinguishability problem using a Karp reduction from the problem **CLOSE IMAGES** of Chapter 4. The main technique is a scheme for using two transformations that are close together for some inputs and producing from them two transformations that act very differently on a specific input state, whereas when the scheme is applied to transformations that are far apart, the resulting transformations are very close together. In some sense this reduction inverts the distance between two circuits: circuits that are close together are mapped to circuits that are far apart, and vice versa, though the definitions of distance used in the close images and distinguishability problems are not the same.

5.2 Quantum circuit distinguishability

The problem of distinguishing mixed state quantum circuits can be stated more intuitively in the following way: given a black box that is promised to implement one of two known quantum channels, with what probability can the channel be identified with only a single use of the black box? As was discussed in Section 3.5, the maximum probability that the correct channel can be identified is given by

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

where the channels Φ_i represent the two known channels. This implies that the problem of estimating the diamond norm of the difference of two channels is equivalent to estimating the probability that the black box can be correctly identified with a single use, given descriptions of the two channels Φ_1, Φ_2 .

To obtain a computational problem, let these channels be given by mixed-state quantum circuits: this results in the quantum circuit distinguishability problem that is the focus of this chapter. The main result of the chapter is to show that this problem is complete for **QIP**. As a by-product, the definition of this problem implies that simply deciding if the two channels are close together for all inputs or far apart on some input state is equivalent to determining with what probability the correct channel is identified in the black box problem. The formal definition of the distinguishability (promise) problem is given below.

Problem 5.1 (Quantum Circuit Distinguishability). For constants $0 \leq b < a \leq 2$, the input consists of quantum circuits Q_1 and Q_2 that implement transformations in $\mathbf{T}(\mathcal{H}, \mathcal{K})$. The promise problem is to distinguish the two cases:

Yes: $\|Q_1 - Q_2\|_{\diamond} \geq a$,

No: $\|Q_1 - Q_2\|_{\diamond} \leq b$.

Less formally, this problem asks: is there an input density matrix ρ on which the circuits Q_1 and Q_2 can be made to act differently? Theorem 3.17 implies that this problem can be stated in terms of pure state inputs. For notational convenience, this problem will be referred to as $\text{QCD}_{a,b}$, with the logarithmic and constant-depth variants referred to as $\text{LOG-DEPTH QCD}_{a,b}$ and $\text{CONST-DEPTH QCD}_{a,b}$, though they will not be encountered until Section 5.6.

The Quantum Circuit Distinguishability problem appears on the surface to be very similar to the Close Images problem considered in Chapter 4, but a closer inspection reveals that this is not the case. Given two circuits Q_1 and Q_2 , the close images problem asks if there are two inputs ρ and σ on which the two circuits act the same, i.e. $Q_1(\rho) \approx Q_2(\sigma)$. On the other hand, the circuit distinguishability problem asks if there is one input ρ for which the states $Q_1(\rho)$ and $Q_2(\rho)$ are nearly orthogonal. The two problems ask for the two channels to have significantly different properties, though the problems do not appear to be dual to each other in any real sense.

In addition to this, the circuit distinguishability problem has an operational meaning in terms of how well an unknown quantum process chosen from a set of two known channels can be identified. An alternate characterization of the problem is, given two quantum channels, are they almost the same, or are there inputs on which they differ significantly. This is a simplification of quantum process tomography [CN97, PCZ97], which has many applications in quantum information, but is unfortunately intractable in the computational sense. In contrast, the Close Images problem, as discussed in

Section 4.2, is quite closely related to whether or not the verifier can be made to accept in a quantum interactive proof system. This makes the circuit distinguishability problem interesting for the study of the class **QIP**, as it gives a quantum information theoretic characterization of the class that is not a restatement of the definition.

5.3 QIP protocol

The aim of this section is to present and analyze a protocol that puts the circuit distinguishability problem inside of **QIP**. This is an essential step in showing that this problem is complete for this class.

The basic idea of the protocol used to achieve this is to have the prover send a state on which the two circuits are maximally distinguishable, apply one of the two circuits at random, and then ask the prover to determine which circuit has been applied. It is not hard to see that by playing honestly, the prover will be able to succeed with probability related to the diamond norm of the difference of the two circuits. It is only slightly more difficult to see that this is also the optimal strategy for a dishonest prover. A more complete description of the protocol follows.

Protocol 5.2 (Quantum Circuit Distinguishability). As input, both the prover P and the verifier V receive circuits $Q_1, Q_2 \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ of size at most n . The three steps of the protocol are:

1. V receives from P a state $\rho \in \mathbf{D}(\mathcal{H})$.
2. V chooses $i \in \{1, 2\}$ uniformly at random and sends $Q_i(\rho)$ back to P .
3. V receives from P some $j \in \{1, 2\}$, accepts if $i = j$, and rejects otherwise.

The idea behind this protocol is that if the two circuits are far apart the prover can find an input state on which they are distinguishable. The prover is then asked to perform this distinguishing task. Both choosing the state ρ and performing the measurement to distinguish the output states $Q_1(\rho)$ and $Q_2(\rho)$ may be computationally intractable and so in the protocol the prover is required to perform them: the verifier only needs to flip a coin and apply one of the two circuits, which can be done in polynomial time, given circuit descriptions.

Step 3 of this protocol does not strictly fit into the model of quantum interactive proof systems. This is because the prover sends classical information to the verifier,

and not a quantum message. This difficulty can be avoided by allowing the prover to send a qubit to the verifier, who immediately measures it in the computational basis. Such a modification does not change the protocol, as any state sent during this step gives the prover no way to do better than simply sending either $|1\rangle$ or $|2\rangle$.

To show that this protocol puts the distinguishability problem into **QIP**, it remains to show that the prover can succeed with probability p on yes instances and with probability at most q on no instances, for some values of p, q that are at least polynomially far apart. Error reduction for **QIP** allows this to be amplified to any constant gap, as discussed in Section 2.2. This is the content of the following theorem.

Theorem 5.3. *For any constants a, b with $0 \leq b < a \leq 2$, $\text{QCD}_{a,b} \in \text{QIP}$.*

Proof. To show that the verifier of Protocol 5.2 forms a quantum interactive proof system for $\text{QCD}_{a,b}$ bounds must be placed on the error probability of the protocol for both positive and negative instances of the problem. This will be done by showing that in either case, the maximum acceptance probability of the verifier is given by

$$\frac{1}{2} + \frac{1}{4} \|Q_1 - Q_2\|_{\diamond},$$

which is simply the optimal probability that a black box can be identified as either Q_1 or Q_2 with only a single use, as discussed in Corollary 3.16. It is not hard to see that this is exactly the task faced by the prover in the protocol.

By Theorem 3.17 there exists a Hilbert space \mathcal{F} and a pure state $|\psi\rangle \in \mathcal{H} \otimes \mathcal{F}$ such that

$$\|Q_1 - Q_2\|_{\diamond} = \|(Q_1 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|) - (Q_2 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|)\|_{\text{tr}}.$$

For this state $|\psi\rangle$, let

$$\begin{aligned} \rho_1 &= (Q_1 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|), \\ \rho_2 &= (Q_2 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|). \end{aligned}$$

Let Π_1 and $\Pi_2 = \mathbb{1} - \Pi_1$ be projection operators on $\mathcal{K} \otimes \mathcal{F}$ that specify an optimal projective measurement for distinguishing ρ_1 from ρ_2 . These projection operators form the Helstrom measurement, which is discussed in Theorem 3.9. Such a measurement satisfies

$$\text{tr} \Pi_1(\rho_1 - \rho_2) = \text{tr} \Pi_2(\rho_2 - \rho_1) = \frac{1}{2} \|\rho_1 - \rho_2\|_{\text{tr}},$$

as Π_1 is the projector onto the positive eigenvalues of $\rho_1 - \rho_2$ and $\text{tr}(\rho_1 - \rho_2) = 0$.

A strategy for the prover that convinces the verifier to accept with probability

$$\frac{1}{2} + \frac{1}{4} \|Q_1 - Q_2\|_{\diamond}$$

is as follows. The prover prepares the state $|\psi\rangle$ and sends the reduced state $\rho = \text{tr}_{\mathcal{F}} |\psi\rangle\langle\psi|$ to the verifier, keeping the portion of the state on \mathcal{F} in reserve.

Upon receiving $\sigma = Q_i(\rho)$ from the verifier, the prover measures the state on $\mathcal{K} \otimes \mathcal{F}$ with the measurement $\{\Pi_1, \Pi_2\}$ and returns the result to the verifier. By Theorem 3.9, this measurement correctly determines i with probability

$$\frac{1}{2} + \frac{1}{4} \|\rho_1 - \rho_2\|_{\text{tr}} = \frac{1}{2} + \frac{1}{4} \|Q_1 - Q_2\|_{\diamond}.$$

That this strategy is optimal can be argued as follows. Let $\xi \in \mathbf{D}(\mathcal{H} \otimes \mathcal{F})$ be the state of the system immediately after the first message is sent, where the space \mathcal{F} represents the private space of the prover (which need not be the same size as the space \mathcal{F} considered above). The verifier applies either Q_1 or Q_2 to the space \mathcal{H} , which results in the global state $(Q_1 \otimes I_{\mathcal{F}})(\xi)$ with probability $1/2$ and $(Q_2 \otimes I_{\mathcal{F}})(\xi)$ with probability $1/2$. This state is, after step 2 of the protocol, in the possession of the prover. The prover's final message to the verifier is immediately measured by the verifier, resulting in a single bit. This process may be viewed as a two-valued measurement on $\mathcal{K} \otimes \mathcal{F}$. The probability that the optimal measurement of this type is correct is given by Theorem 3.9, and so

$$\frac{1}{2} + \frac{1}{4} \|(Q_1 \otimes I_{\mathcal{F}})(\xi) - (Q_2 \otimes I_{\mathcal{F}})(\xi)\|_{\text{tr}} \leq \frac{1}{2} + \frac{1}{4} \|Q_1 - Q_2\|_{\diamond}$$

is an upper bound on the success probability of the prover, as required.

This gives a quantum interactive proof system for $\text{QCD}_{a,b}$ that accepts yes instances with probability $1/2 + a/4$ and accepts no instances with probability at most $1/2 + b/4$. This proves that $\text{QCD}_{a,b} \in \mathbf{QIP}$ as $b < a$ with at least a polynomial gap between them, by the definition of the distinguishability problem. \square

As discussed in Section 5.1, the version of this problem defined on classical randomized circuits is contained in the complexity class \mathbf{AM} . One way to see this is to consider Protocol 5.2 with all of the quantum information removed. The prover can still send an optimal distinguishing input in Step 1 and decide which distribution the sample is from in Step 3. The analysis of this classical protocol is virtually identical to the quantum one: this generalizes a result of Sahai and Vadhan [SV03] on the statistical difference problem to the case of circuits that take input states.

To complete the proof that $\text{QCD}_{a,b}$ is \mathbf{QIP} -complete the Close Images problem is reduced to it. The next section contains a description of this reduction.

5.4 Reduction from Close Images

This section presents a reduction from the close images problem to the circuit distinguishability problem that will be used to show that $\text{QCD}_{a,b}$ is **QIP**-hard for any constants a and b such that $0 < b < a \leq 2$. This is done using a standard polynomial-time Karp reduction: a polynomial time procedure that transforms instances of one problem to another that outputs a yes instance of QCD if and only if the input instance of CI was also a yes instance. The analysis of the reduction presented in this section appears in Section 5.5.

The reduction takes as input an instance of the CI problem, which is given by a pair (Q_1, Q_2) of mixed-state quantum circuits implementing channels in $\mathbf{T}(\mathcal{H}, \mathcal{K})$. The reduction produces as output a pair of circuits (C_1, C_2) that form an instance of the QCD problem.

As in the case of the reduction in Chapter 4, we may assume that the input circuits are given in Stinespring form. In this form the circuit consists of three parts. First is the introduction of any ancillary qubits in the $|0\rangle$ state, second is a unitary circuit applied to the input and ancillary qubits, and finally, the third part is the tracing out of any qubits that are not a part of the output space. A general mixed-state quantum circuit can be put into this form in polynomial time, and this assumption can be made without loss of generality that the input circuits are of this form. This is discussed in more detail in Section 2.1. As the reduction will modify the circuits Q_1 and Q_2 , it is helpful to identify the names of the various Hilbert spaces associated with them. As mentioned above, the circuit Q_i implements an operation in $\mathbf{T}(\mathcal{H}, \mathcal{K})$. The spaces \mathcal{H} and \mathcal{K} will be referred to as the “input” and “output” spaces of the circuit, respectively. Given in Stinespring form the circuits Q_i makes use of ancillary qubits. Let the space \mathcal{A} represent the space these ancillary qubits are added in, and let \mathcal{B} represent the space that is traced out after the unitary is applied. The space \mathcal{A} will be called the “ancillary” space of Q_i and \mathcal{B} will be called the “environment” space. Furthermore, let U_i be the unitary operation in $\mathbf{U}(\mathcal{H} \otimes \mathcal{A}, \mathcal{K} \otimes \mathcal{B})$ that is applied as part of the circuit Q_i . As we may assume without loss of generality that each circuit uses the same number of ancillary qubits by padding one of the circuit with ancillary qubits that are left unused and later traced out, we take the four Hilbert spaces $\mathcal{H}, \mathcal{K}, \mathcal{A}, \mathcal{B}$ to be the same for each of the input circuits. Notice also that since U_i is unitary the spaces $\mathcal{H} \otimes \mathcal{A}$ and $\mathcal{K} \otimes \mathcal{B}$ have the same dimension, and so are isomorphic. The various Hilbert spaces associated with the circuit Q_i are summarized in Figure 5.2.

An important piece of the reduction will be a circuit that, based on the value of




Figure 5.2: The circuit Q_i in Stinespring form, with the Hilbert spaces labelled.

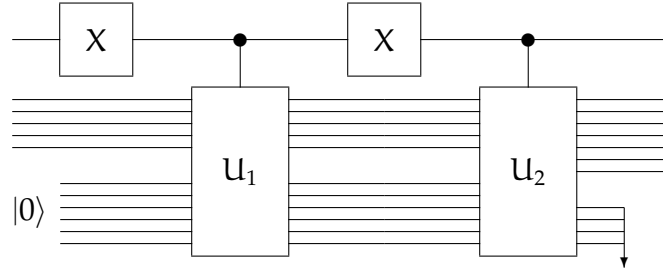


Figure 5.3: A Circuit to apply Q_1 or Q_2 based on the value of a control qubit.

a control qubit, applies either Q_1 or Q_2 to the input state. Such a circuit is easily constructed in polynomial time by simply replacing each gate of Q_1 and Q_2 with gates that are controlled by the value of the control qubit. These controlled gates need not be in the family of gates in the circuit model, but since the number of gates in the model is finite, decompositions of these controlled gates in terms of gates in the model can be constructed efficiently. Let, for concreteness, the circuit that implements one of the two input circuits implement Q_1 when the control qubit is $|0\rangle$ and Q_2 when the control qubit is $|1\rangle$. One construction for such a circuit is given in Figure 5.3. Let \mathcal{Q} be the Hilbert space containing the control qubit, so that the constructed transformation is a channel in $\mathbf{T}(\mathcal{Q} \otimes \mathcal{H}, \mathcal{Q} \otimes \mathcal{K})$. There is some ambiguity in the constructed transformation: the controlled U_1 operation takes an input in $\mathcal{Q} \otimes \mathcal{H} \otimes \mathcal{K}$ and produces output in $\mathcal{Q} \otimes \mathcal{K} \otimes \mathcal{B}$ and this is followed by a controlled U_2 operation. This operation also expects an input in the space $\mathcal{Q} \otimes \mathcal{H} \otimes \mathcal{K}$, not the space $\mathcal{Q} \otimes \mathcal{K} \otimes \mathcal{B}$. Fortunately both of these spaces have the same dimension, and so by implicitly making use of the isomorphism between them, this potential difficulty is avoided.

The circuit shown in Figure 5.3 is very close to the circuits that will be the output of the reduction. To obtain these circuits one critical modification is made: instead of tracing out the environment space \mathcal{B} , the “output” space \mathcal{K} is traced out instead. This reversal of the purposes of the output and environment spaces is essential to the reduction. Taking a Stinespring representation of a channel and tracing out the output




Figure 5.4: Circuits output by the reduction.

instead of the environment leads to what has been called a *conjugate* or *complementary* channel [DS05, Hol07, KMNR07]. Viewed in this light, this circuit simply applies a fixed conjugate of either Q_1 or Q_2 , depending on the value of a control qubit. This is how the circuit C_1 , one of the two circuits output by the reduction is constructed. This circuit is demonstrated in Figure 5.4(a).

The circuit C_2 is constructed in the same way as the circuit C_1 , with one difference. This is a Pauli Z operation is applied to the control qubit after the controlled operations. The circuit C_2 is shown in Figure 5.4(b). This Z gate will make a substantial difference in the output of the two circuits when the control qubit in the space \mathcal{Q} has not been decohered by the other operations of the circuit C_2 . The output of the reduction is an instance of QCD given by the pair of circuits (C_1, C_2) .

The key to this reduction is that when an input is given to either C_1 or C_2 with the control qubit in a superposition of $|0\rangle$ and $|1\rangle$, then both of the circuits Q_1 and Q_2 are run. By tracing out the “output” space \mathcal{K} the idea is that if the outputs of C_1 and C_2 are sufficiently far apart, then tracing out the space \mathcal{K} is akin to measuring the control qubit but forgetting the result. Intuitively, if there is enough information in \mathcal{K} to identify which of the two circuits Q_1 or Q_2 has been performed, then the control qubit will be subject to decoherence. In this case the Pauli Z gate in C_2 has no effect: the control qubit has decohered to a mixture of the form $p|0\rangle\langle 0| + (1 - p)|1\rangle\langle 1|$, and

applying Z to such a state has no effect, since

$$Z(p|0\rangle\langle 0| + (1-p)|1\rangle\langle 1|)Z = pZ|0\rangle\langle 0|Z + (1-p)Z|1\rangle\langle 1|Z = p|0\rangle\langle 0| + (1-p)|1\rangle\langle 1|$$

On the other hand, if the outputs of Q_1 and Q_2 are sufficiently close, then there should be no information about the control qubit in the space \mathcal{K} , so tracing it out will not have any effect. In this case the control qubit remains in a pure state such as $(|0\rangle + |1\rangle)/\sqrt{2}$, so that applying the Pauli Z operation in the circuit C_2 results in the state

$$Z|+\rangle = Z\frac{|0\rangle + |1\rangle}{\sqrt{2}} = \frac{|0\rangle - |1\rangle}{\sqrt{2}} = |-\rangle,$$

which is orthogonal to the control qubit output by C_1 . In this way the reduction effectively inverts the closeness of the circuits: if the circuits Q_1 and Q_2 can be made to output states that are close together, then the output of the circuits C_1 and C_2 can be made far apart. In the other direction, if the original circuits always output states that are distinguishable, then the constructed circuits will always be close together, as the control qubit is left in an incoherent mixture after tracing out the space \mathcal{K} , so that the Z operation in C_2 has little effect.

This argument is not complete: the circuits C_1 and C_2 also output the environment space \mathcal{B} of the original circuits, which has been ignored. The notions of closeness used in the problems CI and QCD are also not the same. Significant care must be taken to formalize this intuitive picture. This is the content of the next section, which contains a formal proof that this reduction implies the **QIP**-hardness of the QCD problem.

5.5 Correctness of the reduction

This section contains the formal proof that the reduction presented in Section 5.4 implies that the QCD problem is **QIP**-hard. Proving that this problem is **QIP**-hard implies that it is also **QIP**-complete, as it is argued in Section 5.3 that these problems belong to **QIP**.

This section is quite technical. The reader uninterested in the details of the proof of the main result is invited to skip the proofs of the lemmas found here: the proofs are not overly difficult but much of the intuition has already been presented in the previous section, so it is unlikely that a detailed study of these proofs will provide a clearer picture of the results.

As in the previous section, let (Q_1, Q_2) be the instance of CI provided as input to the reduction, where these circuits implement transformations in $\mathbf{T}(\mathcal{H}, \mathcal{K})$ given by

$$\begin{aligned} Q_1(\rho) &= \text{tr}_{\mathcal{B}} U_1(\rho \otimes |0\rangle\langle 0|)U_1^*, \\ Q_2(\rho) &= \text{tr}_{\mathcal{B}} U_2(\rho \otimes |0\rangle\langle 0|)U_2^*, \end{aligned}$$

where $U_1, U_2 \in \mathbf{U}(\mathcal{H} \otimes \mathcal{A}, \mathcal{K} \otimes \mathcal{B})$. This is summarized in Figure 5.2.

From these circuits the reduction described in the previous section produces as output (C_1, C_2) , a pair of circuits that form an instance of QCD. Let, for notational convenience, the operator V be the operator that applies the operator U_{i+1} when a control qubit is in the $|i\rangle$ state, i.e. let the operation V be defined by

$$\begin{aligned} V(|0\rangle \otimes |\psi\rangle) &= |0\rangle \otimes U_1|\psi\rangle \\ V(|1\rangle \otimes |\psi\rangle) &= |1\rangle \otimes U_2|\psi\rangle, \end{aligned}$$

for all states $|\psi\rangle$. One implementation for the operation V is given by Figure 5.3, with the exception that the operation V does not trace out the qubits in the space \mathcal{B} .

Using this notation, the circuits in Figure 5.4 implement the operations in the space $\mathbf{T}(\mathcal{Q} \otimes \mathcal{H}, \mathcal{Q} \otimes \mathcal{B})$ given by

$$\begin{aligned} C_1(\rho) &= \text{tr}_{\mathcal{K}} V(\rho \otimes |0\rangle\langle 0|)V^* \\ C_2(\rho) &= \text{tr}_{\mathcal{K}} Z_{\mathcal{Q}} V(\rho \otimes |0\rangle\langle 0|)V^* Z_{\mathcal{Q}}. \end{aligned} \tag{5.1}$$

The operation $Z_{\mathcal{Q}}$ in this equation is simply shorthand for the application of the Pauli Z gate to the qubit represented by \mathcal{Q} , i.e. $Z_{\mathcal{Q}} = Z \otimes \mathbb{1}_{\mathcal{K}} \otimes \mathbb{1}_{\mathcal{B}}$. This characterization of the circuits produced by the reduction will be used to show that C_1 and C_2 are distinguishable if and only if Q_1 and Q_2 have close images.

The main result of this section is that the maximum output fidelity of Q_1 and Q_2 is equal to the diamond norm of the difference of C_1 and C_2 . The proof of this is presented in two steps. The first, and simplest, of these steps is to show that the diamond norm provides a lower bound on the maximum output fidelity. This is argued directly from the properties of the diamond norm and the constructed circuits.

Lemma 5.4. *Given circuits Q_1 and Q_2 , and the circuits C_1 and C_2 constructed from them given by (5.1)*

$$\frac{1}{2} \|C_1 - C_2\|_{\diamond} \leq \max_{\rho, \sigma \in \mathbf{D}(\mathcal{H})} F(Q_1(\rho), Q_2(\sigma)).$$

Proof. By Theorem 3.17 the diamond norm of the difference of two channels is achieved on a pure state in some larger system. Let this state be $|\psi\rangle \in \mathcal{Q} \otimes \mathcal{H} \otimes \mathcal{F}$, where \mathcal{F} is the reference system implied by the theorem. For this state, we have

$$\|C_1 - C_2\|_{\diamond} = \|(C_1 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|) - (C_2 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|)\|_{\text{tr}} \quad (5.2)$$

As $|\psi\rangle$ is a unit vector, it may be written in terms of components on the two subspaces where the qubit in the space \mathcal{Q} is either $|0\rangle$ or $|1\rangle$. More formally, there is some $p \in [0, 1]$ and states $|\psi_0\rangle, |\psi_1\rangle \in \mathcal{H} \otimes \mathcal{F}$ such that

$$|\psi\rangle = \sqrt{p}|0\rangle|\psi_0\rangle + \sqrt{1-p}|1\rangle|\psi_1\rangle.$$

Using this decomposition we can evaluate the circuits C_1 and C_2 on the state $|\psi\rangle$. The input state can be decomposed as

$$\begin{aligned} p|0\rangle\langle 0| \otimes |\psi_0\rangle\langle\psi_0| + (1-p)|1\rangle\langle 1| \otimes |\psi_1\rangle\langle\psi_1| \\ + \sqrt{p(1-p)}|0\rangle\langle 1| \otimes |\psi_0\rangle\langle\psi_1| \\ + \sqrt{p(1-p)}|1\rangle\langle 0| \otimes |\psi_1\rangle\langle\psi_0|. \end{aligned} \quad (5.3)$$

For the sake of brevity, further notation is introduced. Let $|\phi_i\rangle = (U_{i+1} \otimes \mathbf{1}_{\mathcal{F}})|\psi_i\rangle$. This notation suffices as the unitary U_1 is only applied to the state $|\psi_0\rangle$, and likewise with U_2 with the state $|\psi_1\rangle$. Making use of this notation, we can consider the behaviour of the circuits C_1 and C_2 on the terms in Equation (5.3). The output of C_1 is given by

$$(C_1 \otimes I_{\mathcal{F}})(|i\rangle\langle j| \otimes |\psi_i\rangle\langle\psi_j|) = |i\rangle\langle j| \otimes \text{tr}_{\mathcal{X}} |\phi_i\rangle\langle\phi_j|, \quad (5.4)$$

for all $i, j \in \{0, 1\}$. The output of C_2 differs only slightly, being given by

$$(C_2 \otimes I_{\mathcal{F}})(|i\rangle\langle j| \otimes |\psi_i\rangle\langle\psi_j|) = (-1)^{i+j}|i\rangle\langle j| \otimes \text{tr}_{\mathcal{X}} |\phi_i\rangle\langle\phi_j|, \quad (5.5)$$

where the $(-1)^{i+j}$ factor is due to the Pauli Z gate in the circuit C_2 . Notice that when $i = j$, as in the first two terms of Equation (5.3), the two circuits produce identical output. The difference between the two circuits lies in the behaviour on the final two terms of this equation. On these two terms the circuits agree, up to a multiplicative factor of -1 , as can be seen from Equations (5.4) and (5.5). Using this observation, the difference between the outputs of the two circuits is given by

$$(C_1 \otimes I_{\mathcal{F}} - C_2 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|) = 2\sqrt{p(1-p)}(|0\rangle\langle 1| \otimes \text{tr}_{\mathcal{X}} |\phi_0\rangle\langle\phi_1| + |1\rangle\langle 0| \otimes \text{tr}_{\mathcal{X}} |\phi_1\rangle\langle\phi_0|).$$

Combining this with Equation (5.2) yields

$$\begin{aligned} \|C_1 - C_2\|_{\diamond} &= \|(C_1 \otimes I_{\mathcal{F}} - C_2 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|)\|_{\text{tr}} \\ &= 2\sqrt{p(1-p)}\| |0\rangle\langle 1| \otimes \text{tr}_{\mathcal{X}} |\phi_0\rangle\langle\phi_1| + |1\rangle\langle 0| \otimes \text{tr}_{\mathcal{X}} |\phi_1\rangle\langle\phi_0| \|_{\text{tr}}. \end{aligned}$$

From this equation, as well as Lemmas 3.10 and 3.21, we find that

$$\begin{aligned}
\|C_1 - C_2\|_\diamond &= 4 \sqrt{p(1-p)} \|\text{tr}_{\mathcal{X}} |\phi_0\rangle\langle\phi_1|\|_{\text{tr}} \\
&= 4 \sqrt{p(1-p)} F(\text{tr}_{\mathcal{B}\otimes\mathcal{F}} |\phi_0\rangle\langle\phi_0|, \text{tr}_{\mathcal{B}\otimes\mathcal{F}} |\phi_1\rangle\langle\phi_1|) \\
&= 4 \sqrt{p(1-p)} F(Q_1(\text{tr}_{\mathcal{F}} |\psi_0\rangle\langle\psi_0|), Q_2(\text{tr}_{\mathcal{F}} |\psi_1\rangle\langle\psi_1|)) \\
&\leq 2 \max_{\rho, \sigma} F(Q_1(\rho), Q_2(\sigma)),
\end{aligned}$$

which completes the proof of the lemma. \square

The following lemma formalizes the intuitive picture presented in Section 5.4 that the constructed circuits C_1 and C_2 are distinguishable if the original circuits Q_1 and Q_2 have output states with high fidelity. This is the second direction of the proof of the main result of this chapter.

Lemma 5.5. *Given circuits Q_1 and Q_2 , and the circuits C_1 and C_2 constructed from them given by (5.1),*

$$\frac{1}{2} \|C_1 - C_2\|_\diamond \geq \max_{\rho, \sigma \in \mathbf{D}(\mathcal{H})} F(Q_1(\rho), Q_2(\sigma)).$$

Proof. Let $\rho_1, \rho_2 \in \mathbf{D}(\mathcal{H})$ be two arbitrary states. For these states, we will show that

$$\|C_1 - C_2\|_\diamond \geq 2F(Q_1(\rho_1), Q_2(\rho_2)).$$

Let the states $|\psi_0\rangle, |\psi_1\rangle \in \mathcal{H} \otimes \mathcal{F}$ be purifications of ρ_1 and ρ_2 , where \mathcal{F} is any Hilbert space large enough to admit these purifications. These states will play a similar role to the states of the same name used in the proof of Lemma 5.4. Following the notation in this lemma further, let $|\phi_i\rangle = (U_{i+1} \otimes \mathbb{1}_{\mathcal{F}})|\psi_i\rangle$ be the states produced by applying the “appropriate” unitary to these states. Using this notation, consider the input state $|\psi\rangle \in \mathcal{Q} \otimes \mathcal{H}$ to C_1 and C_2 given by

$$|\psi\rangle = \frac{1}{\sqrt{2}}|0\rangle|\psi_0\rangle + \frac{1}{\sqrt{2}}|1\rangle|\psi_1\rangle.$$

On this state the output of the two channels is exactly as discussed in Lemma 5.4 with $p = 1/2$. In particular, the channel C_1 produces the output

$$(C_1 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|) = \frac{1}{2} \sum_{i,j \in \{0,1\}} |i\rangle\langle j| \otimes \text{tr}_{\mathcal{X}} |\phi_i\rangle\langle\phi_j| \quad (5.6)$$

while the circuit C_2 produces the output

$$(C_2 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|) = \frac{1}{2} \sum_{i,j \in \{0,1\}} (-1)^{i+j} |i\rangle\langle j| \otimes \text{tr}_{\mathcal{X}} |\phi_i\rangle\langle\phi_j|. \quad (5.7)$$

These equations follow from identical reasoning to the derivation of equations (5.4) and (5.5) of the previous lemma.

Notice that the pure states $|\phi_0\rangle, |\phi_1\rangle \in \mathcal{K} \otimes \mathcal{B} \otimes \mathcal{F}$ are purifications of $Q_1(\rho_1)$ and $Q_2(\rho_2)$, respectively. This allows us to use Lemma 3.21 to transform the trace norm of $\text{tr}_{\mathcal{K}} |\phi_0\rangle\langle\phi_1|$ into the fidelity of the input states ρ_1 and ρ_2 to the original circuits. This trace norm will be essential, as the difference between the outputs of the circuits C_i consists only of terms of this form, which can be seen from Equations (5.6) and (5.7). This, along with Theorem 3.14 and Lemma 3.10 show that

$$\begin{aligned} \|C_1 - C_2\|_{\diamond} &\geq \| (C_1 \otimes I_{\mathcal{F}} - C_2 \otimes I_{\mathcal{F}})(|\psi\rangle\langle\psi|) \|_{\text{tr}} \\ &= \| |0\rangle\langle 1| \otimes \text{tr}_{\mathcal{K}} |\phi_0\rangle\langle\phi_1| + |1\rangle\langle 0| \otimes \text{tr}_{\mathcal{K}} |\phi_1\rangle\langle\phi_0| \|_{\text{tr}} \\ &= 2 \| \text{tr}_{\mathcal{K}} |\phi_0\rangle\langle\phi_1| \|_{\text{tr}} \\ &= 2 F(Q_1(\rho_1), Q_2(\rho_2)). \end{aligned}$$

This completes the proof of the lemma. \square

With these two lemmas, most of the work in proving the main result has been completed. We have so far shown that the diamond norm difference of the constructed instance (C_1, C_2) of QCD and the maximum output fidelity of the input instance (Q_1, Q_2) of CI are equal. This fact is stated as the following theorem for easy reference. This also proves that the reduction correctly produces “yes” instances of QCD if and only if it is given “yes” instances of CI.

Theorem 5.6. *Given circuits Q_1 and Q_2 , and the circuits C_1 and C_2 constructed from them given by (5.1),*

$$\frac{1}{2} \|C_1 - C_2\|_{\diamond} = \max_{\rho, \sigma \in \mathcal{D}(\mathcal{H})} F(Q_1(\rho), Q_2(\sigma)).$$

Proof. Lemma 5.4 provides the upper bound and Lemma 5.5 provides the lower bound. Taken together they prove the desired equation. \square

This theorem immediately implies the main result of the chapter: the **QIP** hardness of the distinguishability problem for mixed-state quantum circuits.

Corollary 5.7. *For any $0 < b < a \leq 2$ the problem $\text{QCD}_{a,b}$ is **QIP**-complete.*

Proof. Theorem 5.6 and the construction in Section 5.4 imply the reduction

$$\text{CI}_{a/2, b/2} \leq_m^p \text{QCD}_{a,b}.$$

As $CI_{a/2,b/2}$ is **QIP**-hard for any $0 < b < a \leq 2$ by Theorem 4.3, which is due to Kitaev and Watrous [KW00], this reduction implies that the distinguishability problem is **QIP**-hard for the values of a, b specified in the theorem.

The distinguishability problem is also complete for **QIP** for these values of a and b , as it is in **QIP** by Theorem 5.3. \square

5.6 Distinguishing log-depth computations

In this section it is discussed how to extend the hardness of the QCD problem to the case of input circuits that have logarithmic depth. This can be done by simply noting that the reduction of Section 5.4 can be modified to produce output circuits of logarithmic depth, and the hardness will follow from the hardness of the log-depth close images problem.

To see that the **QIP**-hardness of the problem $\text{LOG-DEPTH } CI_{a/2,b/2}$ can be extended to the problem $\text{LOG-DEPTH } QCD_{a,b}$, observe that the reduction in Section 5.4 simply takes the input circuits and produces circuits that apply them based on the value of a control qubit. These controlled operations can be implemented in logarithmic depth using a tree structure with copies of the control qubit made in the computational basis – this is discussed in Proposition 2.1. If this more careful implementation of the controlled U_1 and U_2 operations is made, then the output circuits of Figure 5.4 have logarithmic depth if and only if the input circuits do. This requires that the circuits for the operations U_i that implement the input circuits Q_i by the equation

$$Q_i(\rho) = \text{tr}_B U_i(\rho \otimes |0\rangle\langle 0|)U_i^*$$

can be assumed to have logarithmic depth when the mixed-state circuits Q_i do, but these circuits can be constructed by simply delaying any partial trace operations that are performed during the circuit. These circuits have the same depth as the original mixed state circuits and they can be constructed in polynomial time. This implies that

$$\text{LOG-DEPTH } CI_{a/2,b/2} \leq_m^p \text{LOG-DEPTH } QCD_{a,b}$$

which, by Corollary 4.8 immediately implies the following corollary, since $\text{LOG-DEPTH } QCD$ is in **QIP** by Theorem 5.3 as it is a restriction of the general problem.

Corollary 5.8. *For any $0 < b < a \leq 2$ the problem $\text{LOG-DEPTH } QCD_{a,b}$ is **QIP**-complete.*

As in Chapter 4, the only place that this construction requires logarithmic depth circuits are the controlled operations. If the unbounded fan-out gate is allowed into

the basis of computational gates, then the circuits can be reduced to constant depth, as discussed in Proposition 2.2. This implies the following result.

Corollary 5.9. *For any $0 < b < a \leq 2$, the problem $\text{CONST-DEPTH QCD}_{a,b}$ on circuits with the unbounded fan-out gate is **QIP**-complete.*

5.7 Conclusion

In this chapter, the problem $\text{QUANTUM CIRCUIT DISTINGUISHABILITY}$ has been introduced and studied. This is the problem of determining if two quantum channels, given as mixed-state quantum circuits, is there an input on which they produce nearly orthogonal output states, or are they effectively the same on all inputs?

The main result of the chapter is that this problem is complete for the class **QIP** of problems that have quantum interactive proof systems, when the phrases “nearly orthogonal” and “effectively the same” are formalized as large and small diamond norm distance, respectively. This result requires many of the results on the diamond norm from Chapter 3, such as Theorem 3.17, which proves that the diamond norm of the difference of two channels is achieved on a pure state input. This result can also be extended to the case of channels specified by logarithmic depth circuits, or even constant depth circuits if the unbounded fan-out gate is included in the circuit model.

The main result of this chapter is essential for the result in the next two chapters that this distinguishability problem remains hard when restricted to circuits that implement convex mixtures of unitary channels and when restricted to the degradable or antidegradable channels. These results will be shown by reducing the problem considered here to restricted versions.

Chapter 6

Degradable and Antidegradable Channels

The degradable and antidegradable channels are two of the most interesting classes of quantum channels. The degradable channels are, informally, the channels where the output space contains more information about the input than the environment, in the sense that the output state can be used to reconstruct the state of the environment. The antidegradable channels can be similarly thought of as those channels whose output contains less information about the input than the environment does. These channels have many nice properties when considered for the transmission of quantum information. This is interesting as these channels can be otherwise awkward to work with: as an example, the set of degradable channels is not even convex!

The main result of this chapter is that the quantum circuit distinguishability problem considered in Chapter 5 remains **QIP**-complete on both the degradable and antidegradable channels. This lends evidence to the notion that the difficulty of distinguishing quantum channels has little to do with how well they preserve information.

Contents

6.1 Degradable and antidegradable channels	118
6.2 Simulation by a degradable channel	119
6.3 Distinguishing degradable channels	121
6.4 Simulation by an antidegradable channel	123
6.5 Distinguishing antidegradable channels	126
6.6 Conclusion	127

6.1 Degradable and antidegradable channels

As defined in Chapter 1, a channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ is degradable if there exists a second channel Δ that maps the output state of Φ to the state of the environment. More precisely, if

$$\Phi(\rho) = \text{tr}_{\mathcal{B}} \mathbf{U}(\rho \otimes |0\rangle\langle 0|)\mathbf{U}^*,$$

then Φ is called degradable if there exists a channel $\Delta \in \mathbf{T}(\mathcal{K}, \mathcal{B})$ such that

$$(\Delta \circ \Phi)(\rho) = \text{tr}_{\mathcal{K}} \mathbf{U}(\rho \otimes |0\rangle\langle 0|)\mathbf{U}^* = \Phi^{\mathcal{C}}(\rho).$$

The channel $\Phi^{\mathcal{C}}$ 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 viewed as a part of the degrading map Δ . Loosely, these are the channels whose output contains more information than the environment, because the output can be degraded to give the state of the environment. These channels were introduced by Shor and Devetak [DS05] 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 Proposition 7.2).

A channel is called antidegradable if the complementary channel is degradable. Alternately, a channel Φ is antidegradable if there exists a map A such that $A \circ \Phi^{\mathcal{C}} = \Phi$, where once again the channel $\Phi^{\mathcal{C}}$ is only defined up to an isometry, but this isometry can also be part of the map A , so that the antidegradable channels are also well-defined. This class of channels has been introduced by Wolf and Pérez-García [WPG07], and can be informally thought of as the class of those 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 [CRS08], where it is shown that, unlike the degradable channels, the set of antidegradable channels is convex.

The degradable and antidegradable channels are very interesting from a quantum information perspective. A simple no-cloning argument implies that the antidegradable channels have zero capacity for the transmission of quantum information. This argument was first presented for erasure channels in [BDS97], extended to lossy bosonic channels in [GLMS03], and finally applied to antidegradable channels in [GF05]. 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. that the formula for the quantum capacity does not require regularization [DS05].

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 [CRS08] that the additivity of this quantity on degradable channels is equivalent to the general case, making use of a result from [FW07]. As it is also known that this additivity problem is equivalent on the complementary class of channels [Hol07, KMNR07], this implies that the additivity of the antidegradable channels is also equivalent to the general case. Finally, using the recent result of Hastings [Has09], there are degradable and antidegradable channels that do not have additive Holevo capacity.

Interestingly, we can adapt the same construction used by Cubitt, Ruskai, and Smith [CRS08] to show that the quantum circuit distinguishability problem restricted to either the degradable or antidegradable channels remains **QIP**-complete. These results are the focus of the remainder of this chapter.

6.2 Simulation by a degradable channel

Given a circuit Q implementing a transformation in $\mathbf{T}(\mathcal{H}, \mathcal{K})$, the goal is to efficiently construct a circuit C implementing a degradable channel in $\mathbf{T}(\mathcal{H}, \mathcal{K})$ that is closely related to the original circuit Q . This reduction will make use of the results used in the case of the minimum output entropy [CRS08]: the construction presented here, as well as the proof that the resulting channel is degradable, can both be found in this work.

To describe the channel, we assume that $\dim \mathcal{H} = \dim \mathcal{K}$, i.e. that the circuit Q has identical input and output dimension. This may be assumed without loss of generality by padding the smaller space with unused $|0\rangle$ qubits, since these qubits will not affect the diamond norm used in the definition of the distinguishability problem. Once this padding has been completed, we may view Q as an implementation of some channel in $\mathbf{T}(\mathcal{H}, \mathcal{H})$. The channel C constructed from Q will make use of an additional output qubit in the space \mathcal{C} of dimension 2, so that $Q \in \mathbf{T}(\mathcal{H}, \mathcal{C} \otimes \mathcal{H})$.

The basic idea is to implement the channel

$$C(\rho) = \frac{1}{2}|0\rangle\langle 0| \otimes \rho + \frac{1}{2}|1\rangle\langle 1| \otimes Q(\rho). \quad (6.1)$$




Figure 6.1: The degradable channel C constructed from Q .

This is just the channel that applies the circuit Q with probability $1/2$, and does nothing to the input with probability $1/2$. If Q is given in Stinespring form with unitary U , so that

$$Q(\rho) = \text{tr}_B U(\rho \otimes |0\rangle\langle 0|)U^*,$$

then the channel C can be implemented as shown in Figure 6.1. The idea in this implementation is that the top ancillary qubit (which is the qubit in the space \mathcal{C}) is placed in the $|+\rangle$ state, which results in the circuit for Q being applied with probability one-half, as the value of this qubit is ‘copied’ onto one of the environment qubits by the controlled-not gate. This results in the mixture in Equation (6.1).

To see that the circuit C implements a degradable operation, the degrading map that takes the output state to the environment state can be explicitly constructed. As the complementary channel is defined only up to an isometry, we may construct this map for *any* of the complementary channels C^C defined by C , as this isometry can be added to the degrading map as required. For this reason, we consider the complementary channel defined by the implementation in Figure 6.1, i.e. the channel from the input to all those qubits that are traced out. This results in the complementary channel

$$C^C(\rho) = \frac{1}{2}|0\rangle\langle 0| \otimes |0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1| \otimes Q^C(\rho), \quad (6.2)$$

where Q^C has implementation

$$Q^C(\rho) = \text{tr}_{\mathcal{C}} U(\rho \otimes |0\rangle\langle 0|)U^*,$$

which is obtained by tracing out the ‘output’ space of the original circuit.

It is not hard to see how to implement the degrading map Δ_C for this channel. Starting with the output state of C

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




Figure 6.2: The degrading channel Δ_C corresponding to the channel in C in Figure 6.1.

as given by Equation (6.1), this channel can, based on the flag state in the space \mathcal{C} either output $|0\rangle\langle 0|$ or $Q^C(\rho)$. More formally, when the flag state is $|0\rangle$, the state in \mathcal{H} is the original input ρ , so that the channel Q^C can be applied to it by performing the unitary U from the circuit Q and tracing out the appropriate space. On the other hand, when this flag state is $|1\rangle$, the degrading map needs to output $|0\rangle\langle 0|$, which can be done by producing the correct number of untouched ancillary qubits as output. All that remains is to invert the flag qubit to get exactly the output of C^C . A circuit implementation of the channel Δ_C is presented in Figure 6.2. We can formally verify that this map performs the required operation by observing that

$$\begin{aligned} \Delta_C(C(\rho)) &= \frac{1}{2}\Delta_C(|0\rangle\langle 0| \otimes \rho + |1\rangle\langle 1| \otimes Q(\rho)) \\ &= \frac{1}{2}|1\rangle\langle 1| \otimes Q^C(\rho) + \frac{1}{2}|0\rangle\langle 0| \otimes |0\rangle\langle 0| \\ &= C^C(\rho), \end{aligned}$$

where the final equality is Equation (6.2). This argument, due to Cubitt, Ruskai, and Smith [CRS08] proves that the channel C is degradable. In the next section we consider the implications of this construction for the computational hardness of the problem of distinguishing quantum circuits.

6.3 Distinguishing degradable channels

The construction in the previous section essentially embeds any channel into a degradable channel. This construction can be used to show that distinguishing degradable channels is no easier than distinguishing general channels.

As a first step towards this, a formal definition of the circuit distinguishability problem of Chapter 5 is presented. This is simply the general problem with the extra restriction that the input circuits implement channels that are degradable.

Problem 6.1 (Degradable Quantum Circuit Distinguishability). For constants $0 \leq b < a \leq 2$, the input consists of quantum circuits C_1 and C_2 that implement degradable transformations in $\mathbf{T}(\mathcal{H}, \mathcal{K})$. The promise problem is to distinguish the two cases:

Yes: $\|C_1 - C_2\|_{\diamond} \geq a$,

No: $\|C_1 - C_2\|_{\diamond} \leq b$.

The primary ingredient in the proof that this problem is **QIP**-complete, is the result that the construction in the previous section does not significantly affect the diamond norm of the difference of two channels. This is not difficult to see from the output of the construction, given by Equation (6.1), but for completeness it is argued formally in the following lemma.

Lemma 6.2. *Let Q_1, Q_2 be quantum circuits implementing transformations in $\mathbf{T}(\mathcal{H}, \mathcal{K})$. If $C_1, C_2 \in \mathbf{T}(\mathcal{H}, \mathcal{C} \otimes \mathcal{K})$ are given by*

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

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

$$\|C_1 - C_2\|_{\diamond} = \frac{1}{2} \|Q_1 - Q_2\|_{\diamond}.$$

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

$$\begin{aligned} \|(C_1 \otimes I_{\mathcal{F}} - C_2 \otimes I_{\mathcal{F}})(\rho)\|_{\text{tr}} &= \frac{1}{2} \||0\rangle\langle 0| \otimes (\rho - \rho) + |1\rangle\langle 1| \otimes ([Q_1 \otimes I_{\mathcal{F}} - Q_2 \otimes I_{\mathcal{F}}](\rho))\|_{\text{tr}} \\ &= \frac{1}{2} \||1\rangle\langle 1| \otimes ([Q_1 \otimes I_{\mathcal{F}} - Q_2 \otimes I_{\mathcal{F}}](\rho))\|_{\text{tr}} \\ &= \frac{1}{2} \|(Q_1 \otimes I_{\mathcal{F}} - Q_2 \otimes I_{\mathcal{F}})(\rho)\|_{\text{tr}}. \end{aligned}$$

Since the diamond norm is defined as the maximization over all states ρ , this implies the statement of the lemma. \square

Let (Q_1, Q_2) be an instance of $\text{QCD}_{a,b}$. As it is demonstrated in the previous section how to efficiently construct the channels C_i from the channels Q_i , this lemma implies the following reduction

$$\text{QCD}_{a,b} \leq_m \text{DEGRADABLE QCD}_{a/2, b/2}.$$

This implies that distinguishing degradable circuits is hard for all $0 < b < a \leq 1$, using the hardness result for general circuits (Corollary 5.9).

This result can be strengthened using a result from Section 3.7 on the polarization of the diamond norm. The general construction does not preserve degradability, but for the special case of interest using only a portion of the polarization construction will suffice. The strategy is to take an instance (C_1, C_2) of $\text{DEGRADABLE QCD}_{1,\epsilon}$ and construct the instance $(C_1^{\otimes k}, C_2^{\otimes k})$. 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 to a value close to 2, but it will also have the property that the norm of ‘no’ instances is not made too large. This is a straightforward consequence of Lemma 3.25, which appears as part of the procedure for polarizing the diamond norm.

Corollary 6.3. *For any constants $0 < b < a < 2$, the problem $\text{DEGRADABLE QCD}_{a,b}$ is QIP-complete.*

Proof. This problem is in QIP as it is a restriction of the general problem, which is in QIP by Theorem 5.3. To see that it is QIP-hard, take an instance (Q_1, Q_2) of the QIP-complete problem $\text{QCD}_{2,\epsilon}$, for $\epsilon > 0$ a constant.

Applying the construction of Section 6.2 to (Q_1, Q_2) results in the instance (C_1, C_2) of $\text{DEGRADABLE QCD}_{1,\epsilon/2}$, by Lemma 6.2. As the degradable channels are closed under tensor products, $(C_1^{\otimes k}, C_2^{\otimes k})$ is a pair of circuits implementing degradable channels. By Lemma 3.25, we have the following implications

$$\begin{aligned} \|C_1 - C_2\|_{\diamond} \geq 1 &\implies \|C_1^{\otimes k} - C_2^{\otimes k}\|_{\diamond} \geq 2 - 2^{-k/8}, \\ \|C_1 - C_2\|_{\diamond} \leq \frac{\epsilon}{2} &\implies \|C_1^{\otimes k} - C_2^{\otimes k}\|_{\diamond} \leq \frac{k\epsilon}{2}. \end{aligned}$$

These equations imply that for any constants $0 < b < a < 2$, there are choices of the constants k, ϵ so that

$$\begin{aligned} \|Q_1 - Q_2\|_{\diamond} = 2 &\implies \|C_1^{\otimes k} - C_2^{\otimes k}\|_{\diamond} \geq a, \\ \|Q_1 - Q_2\|_{\diamond} \leq \epsilon &\implies \|C_1^{\otimes k} - C_2^{\otimes k}\|_{\diamond} \leq b, \end{aligned}$$

which implies the QIP hardness of $\text{DEGRADABLE QCD}_{a,b}$ □

6.4 Simulation by an antidegradable channel

In this section a construction very similar to that used in Section 6.2 is presented that takes any circuit Q to a circuit C implementing an antidegradable channel. The idea is to (with probability one-half) send the input state to the environment, so that




Figure 6.3: The antidegradable channel C constructed from Q .

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 a construction used in [CRS08] for degradable channels.

Once again we may assume that Q implements a channel in $\mathbf{T}(\mathcal{H}, \mathcal{H})$, i.e. that Q has the same input and output dimension, by embedding the smaller space into the larger, if necessary. As in Section 6.2, the constructed circuit C will use one additional output qubit, implementing an antidegradable transformation in $\mathbf{T}(\mathcal{H}, \mathcal{C} \otimes \mathcal{H})$.

Let Q implement the transformation given by

$$Q(\rho) = \text{tr}_B U(\rho \otimes |0\rangle\langle 0|)U^*,$$

where, as usual, since Q is assumed (without loss of generality) to be in Stinespring form, the input specifies a circuit for computing the unitary U . The channel C will be constructed as

$$C(\rho) = \frac{1}{2}|0\rangle\langle 0| \otimes |0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1| \otimes Q(\rho). \quad (6.3)$$

This is just the channel that applies Q with probability one-half, outputs $|0\rangle$ with probability one-half, and outputs a flag qubit in the space \mathcal{C} to indicate which case has occurred. In a way very similar to the construction in Section 6.2, this channel can be implemented using a controlled- U operation. In this case, however, we will also need the operation W that swaps the states in two spaces (i.e. $W|a\rangle|b\rangle = |b\rangle|a\rangle$). An implementation of the channel C is given in Figure 6.3. This circuit will, depending on the value of the control qubit in the space \mathcal{C} either apply Q or output the pure state $|0\rangle$, as required.

To show that the circuit C implements an antidegradable channel, we explicitly construct the map A_C that maps the environment state of C to the output state. The




Figure 6.4: The anti-degrading channel corresponding to the channel in C in Figure 6.3.

environment state of C is once again simply the state produced by C^C , the complementary channel to C. As before, this channel is only defined up to an isometry, but this is not significant as this isometry can be absorbed into the definition of A_C . One implementation of the channel C^C is obtained by considering the channel mapping the input of C to the space traced out by the circuit in Figure 6.3. This channel is given by

$$C^C(\rho) = \frac{1}{2}|0\rangle\langle 0| \otimes \rho + \frac{1}{2}|1\rangle\langle 1| \otimes Q^C(\rho), \quad (6.4)$$

where once again the channel Q^C is given by

$$Q^C(\rho) = \text{tr}_{\mathcal{C}'} U(\rho \otimes |0\rangle\langle 0|)U^*.$$

Given the state in Equation (6.4) it is not hard to see how to map it to the state in Equation (6.3). This can be done by implementing one of two operations, depending on the value of the flag qubit in the space \mathcal{C}' , which is the ‘copy’ of the control qubit traced out in Figure 6.3. If this qubit is in the state $|0\rangle$, then the other portion of the input state is ρ , the original input to C, so that applying the circuit for Q produces the state $Q(\rho)$. If the control qubit is in the $|1\rangle$ state, however, the remainder of the input state is $Q^C(\rho)$. This state can be discarded (i.e. traced out) and ancillary qubits in the state $|0\rangle$ can be used as the output, using the swap operation W. As before, the value of the qubit in \mathcal{C}' needs to be flipped with a Pauli X gate so that the state is exactly correct. A circuit implementing this is shown in Figure 6.4.

To see that A_C correctly implements the anti-degrading map for C, we may compute

$$\begin{aligned} A_C(C^C(\rho)) &= \frac{1}{2}A_C(|0\rangle\langle 0| \otimes \rho + |1\rangle\langle 1| \otimes Q^C(\rho)) \\ &= \frac{1}{2}|1\rangle\langle 1| \otimes Q(\rho) + \frac{1}{2}|0\rangle\langle 0| \otimes |0\rangle\langle 0| \\ &= C(\rho), \end{aligned}$$

where the final equality is Equation 6.3. This demonstrates that the channel C constructed from Q is antidegradable. In the following section the implications of this construction for the hardness of computationally distinguishing antidegradable channels is considered.

6.5 Distinguishing antidegradable channels

In a very similar way to the degradable case, the construction in the previous section embeds any channel into an antidegradable one. In exactly the same manner as Section 6.3, this can be used to show the hardness of distinguishing circuits that implement antidegradable transformations.

As in the degradable case, the distinguishability problem in the antidegradable case is simply the restriction of the problem to a smaller class of channels.

Problem 6.4 (Antidegradable Quantum Circuit Distinguishability). For constants $0 \leq b < a \leq 2$, the input consists of quantum circuits C_1 and C_2 that implement antidegradable transformations in $\mathbf{T}(\mathcal{H}, \mathcal{K})$. The promise problem is to distinguish the two cases:

Yes: $\|C_1 - C_2\|_{\diamond} \geq a$,

No: $\|C_1 - C_2\|_{\diamond} \leq b$.

Once again the key technique to proving that the problem is **QIP**-complete is to place bounds on the diamond norm of the difference of two channels that have had the construction of the previous section applied to them. The proof of this lemma is identical to the proof of Lemma 6.2.

Lemma 6.5. *Let Q_1, Q_2 be quantum circuits implementing transformations in $\mathbf{T}(\mathcal{H}, \mathcal{K})$. If $C_1, C_2 \in \mathbf{T}(\mathcal{H}, \mathcal{C} \otimes \mathcal{K})$ be given by*

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

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

$$\|C_1 - C_2\|_{\diamond} = \frac{1}{2} \|Q_1 - Q_2\|_{\diamond}.$$

Proof. Let $\rho \in \mathbf{D}(\mathcal{H} \otimes \mathcal{F})$ be arbitrary. Then

$$\|(C_1 \otimes I_{\mathcal{F}} - C_2 \otimes I_{\mathcal{F}})(\rho)\|_{\text{tr}} = \frac{1}{2} \|(Q_1 \otimes I_{\mathcal{F}} - Q_2 \otimes I_{\mathcal{F}})(\rho)\|_{\text{tr}},$$

as in the proof of Lemma 6.2. This implies the statement of the lemma. \square

Exactly as in the degradable case, this implies the **QIP**-hardness of distinguishing antidegradable channels for constants $0 < b < a \leq 1$. Once again, this can be strengthened to any constants $0 < b < a < 2$ using the polarization techniques of Section 3.7, using the property that the antidegradable channels are closed under tensor products. The next corollary follows from Lemma 6.5 in exactly the same manner that Corollary 6.3 follows from Lemma 6.2, so the proof has been omitted.

Corollary 6.6. *For any constants $0 < b < a < 2$, the problem $\text{ANTIDEGRADABLE QCD}_{a,b}$ is **QIP**-complete.*

6.6 Conclusion

This chapter has presented a construction for embedding an arbitrary channel into a degradable channel due to Cubitt, Ruskai, and Smith [CRS08], 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 of the chapter 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. The proof of this result makes use of the diamond norm polarization techniques of Section 3.7.

Chapter 7

Mixed-Unitary Channels

The mixed-unitary channels are an interesting class of quantum operations. These are the channels that probabilistically apply one of a set of unitary operations. These channels have several interesting properties and many of the common transformations used in quantum information are mixed-unitary. For these reasons the problems of determining the additivity of the classical capacity of a mixed-unitary channel and distinguishing circuits that implement mixed-unitary operations are important steps toward understanding these problems.

In the distinguishability case it is shown that distinguishing mixed-unitary channels is exactly as computationally difficult as general channels, using a reduction that essentially simulates a general channel with a mixed unitary one. In the case of additivity, a similar reduction is used to show that given a channel, there is a mixed-unitary channel that is approximately additive if and only if the original channel is additive. By sending the approximation error to zero this produces a sequence of mixed-unitary channels with the property that the original channel is additive if and only if the tail of the sequence consists of additive mixed-unitary channels.

The results in this chapter have been published in [Ros08a].

Contents

7.1	Mixed-unitary channels	130
7.2	Unital channels	132
7.3	Mixed-unitary approximation	134
7.3.1	Simulating the partial trace	135
7.3.2	Simulating the ancillary space	137
7.3.3	Mixed-unitary approximation of a general channel	141

7.4	Properties of the constructed channel	142
7.5	Multiplicativity of mixed-unitary transformations	145
7.6	Mixed-unitaries and minimum output entropy	147
7.7	Circuit constructions	149
7.8	QIP-completeness of distinguishing mixed-unitary circuits	156
7.9	Conclusion	160

7.1 Mixed-unitary channels

As defined in Chapter 1, a quantum channel Φ is mixed-unitary if there exist unitary operators U_1, \dots, U_n and a probability distribution p_1, \dots, p_n such that

$$\Phi(X) = \sum_{i=1}^n p_i U_i X U_i^*. \quad (7.1)$$

These channels have many interesting properties. These channels are commonly known as the *random unitary* channels, but they will be referred to as the mixed-unitary channels here to avoid confusion with unitary operators drawn chosen at random according to some measure. This notational choice was suggested by Watrous in [Wat09a].

It has been shown by Gregoratti and Werner [GW03] that the mixed-unitary channels describe exactly the noise processes that can be corrected using classical information obtained by measuring the environment. One way to see that this correction is possible is to consider a Stinespring representation for the channel in Equation (7.1). One such representation can be constructed using the operations V and W given by

$$\begin{aligned} V|\psi\rangle|i\rangle &= (U_i|\psi\rangle)|i\rangle, \\ W|0\rangle &= \sum_i \sqrt{p_i}|i\rangle. \end{aligned}$$

The operation V is a unitary operation in $\mathbf{U}(\mathcal{H} \otimes \mathcal{A}, \mathcal{K} \otimes \mathcal{B})$ and the operator W can be extended to a unitary operation in $\mathbf{U}(\mathcal{A})$ in an arbitrary way. A Stinespring representation for Φ is then given by

$$\Phi(\rho) = \text{tr}_{\mathcal{B}} V(I_{\mathcal{H}} \otimes W)(\rho \otimes |0\rangle\langle 0|)(I_{\mathcal{H}} \otimes W^*)V^*,$$

where the operator W prepares a weighted superposition of the ancillary space, the operator V applies the corresponding unitary operator from Equation (7.1), and finally

the partial trace over \mathcal{B} produces the desired mixture. To see that this can be perfectly reversed with a measurement of \mathcal{B} , notice that when applied to the state ρ , before the partial trace the system is in the state

$$\sigma = \sum_{i,j} \sqrt{p_i p_j} (U_i \rho U_j^*) \otimes |i\rangle\langle j|.$$

Measuring the second system in the computational basis gives an outcome a with probability p_a , leaving the system in the state

$$\frac{(\mathbb{1}_{\mathcal{H}} \otimes \langle a|) \sigma (\mathbb{1}_{\mathcal{H}} \otimes |a\rangle)}{p_a} = U_a \rho U_a^*.$$

The original state can then be recovered by simply applying the U_a^* selected by the outcome of the measurement. This example describes in principle any mixed-unitary channel, due to the uniqueness of the Stinespring representation, up to an isometry on the space \mathcal{B} which corresponds to a different measurement. The fact that the mixed-unitary channels are the only channels that can be corrected using the strategy of measuring the environment and applying a correction is more complicated and can be found in [GW03].

One question that this correction scheme raises is how much classical information must be recovered from the environment to correct a mixed-unitary channel? This corresponds to minimizing the number of operators U_i in Equation (7.1). A simple bound on this quantity is given by Buscemi [Bus06], who shows that the number n of unitary operators in Equation (7.1) is at most the square of the minimum number of Kraus operators in a Kraus representation of Φ .

Audenaert and Scheel have also provided a characterization of the mixed-unitary channels [AS08], and used it to construct a measure of the distance from a quantum channel to the set of mixed-unitary channels.

The remainder of this chapter provides an answer to the question: are the problems of the additivity of the classical capacity and the distinguishability of quantum channels simplified when restricted to mixed-unitary channels? This is answered in the negative, using a method to approximate an arbitrary quantum channel by a mixed-unitary one. This approximation will only faithfully implement the channel on low-entropy outputs, as the technique used will be able to decide when the approximation would fail and instead produce a highly mixed state. This suffices to produce a channel with the same minimum output entropy or maximum output p -norm, however, as these quantities are defined only by the low-entropy outputs of the channel.

These results extends the work of Fukuda [Fuk07] on unital channels to the mixed-unitary case. The unital case is discussed in the next section. In the Section 7.3 the mixed-unitary case is described. Section 7.4 proves some properties of the construction that will be used in Sections 7.5 and 7.6 that consider the multiplicativity of the p -norm and the additivity of the minimum output entropy (respectively). An efficient circuit construction for this reduction is then presented in Section 7.7, which is used in Section 7.8 to reduce the circuit distinguishability problem from general channels to the mixed-unitary channels.

7.2 Unital channels

Recall from Chapter 1 that a channel $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ is doubly stochastic if $\Phi(\mathbb{1}_{\mathcal{H}}) = \mathbb{1}_{\mathcal{K}}$, and unital if it is also the case that $\mathcal{H} = \mathcal{K}$. This section provides an overview of a result related to the main results of the chapter. This result is Fukuda's proof that the additivity of the minimum output entropy or the multiplicativity of the maximum output p -norm of an arbitrary channel Φ is equivalent to the same problem on a related unital channel Φ' [Fuk07].

The unital channels can be characterized in a similar way to the mixed-unitary channels. Mendl and Wolf have shown that any unital channel Φ can be represented in the form

$$\Phi(\rho) = \sum_i \lambda_i U_i \rho U_i^*,$$

where the U_i are unitary operators and $\sum_i \lambda_i = 1$, with $\lambda_i \in \mathbb{R}$ for all i [MW09]. It is also of note that for channels on qubits the mixed-unitary channels are exactly the unital channels [Tre86, KM87]. This is no longer true for channels on systems of larger dimension, as shown by Landau and Streater [LS93]. An interesting fact is that unital but not mixed-unitary channel that they use is, up to unitary conjugation, the same channel used by Werner and Holevo to find the first example of the super-multiplicativity of the p -norm of a quantum channel [WH02]. This might suggest that these channels are the key to this property, but the results of this chapter imply that the mixed-unitary channels do not hold a special place with respect to this property of super-multiplicativity. Indeed, Hayden and Winter [HW08] have shown that mixed-unitary channels also exhibit this property.

The key ingredient in Fukuda's reduction to the unital case is the addition of an extra input system that allows for an input determined selection of one of the discrete Weyl operators $W_{i,j}$ introduced in Chapter 1 to be applied to the output of the channel.

Letting $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$, with $d = \dim \mathcal{K}$, the channel Φ' is constructed as

$$\Phi'(\rho \otimes |i, j\rangle\langle i, j|) = W_{i,j} \Phi(\rho) W_{i,j}^*, \quad (7.2)$$

for any $1 \leq i, j \leq d$. This defines a channel Φ' on $\mathbf{T}(\mathcal{H} \otimes \mathcal{K} \otimes \mathcal{K}, \mathcal{K})$ by linearity. Such a channel can be implemented by measuring the input space $\mathcal{K} \otimes \mathcal{K}$ in the computational basis to decide which of the Weyl operators to apply and then tracing out the result. This is the same construction used by Shor to prove that the additivity of the minimum output entropy of a channel Φ implies the additivity of the Holevo χ -capacity of the channel Φ' [Sho04]. This construction was discussed in Section 3.3.1.

To see that this channel is doubly stochastic, notice that on any input of the form $\rho \otimes \tilde{\mathbb{1}}_{\mathcal{K} \otimes \mathcal{K}}$

$$\Phi'(\rho \otimes \tilde{\mathbb{1}}_{\mathcal{K} \otimes \mathcal{K}}) = \frac{1}{d^2} \sum_{i,j=1}^d W_{i,j} \Phi(\rho) W_{i,j}^* = \tilde{\mathbb{1}}_{\mathcal{K}}.$$

The fact that this mixture of discrete Weyl operators mixes states in this way is shown in Proposition 7.2 as part of the proof that this operation is mixed-unitary. This equation implies that, on the particular input $\mathbb{1}_{\mathcal{H} \otimes \mathcal{K} \otimes \mathcal{K}}$ the output of Φ' is given by $\mathbb{1}_{\mathcal{K}}$, as required. This channel is not unital, but it can be made so by adding an additional output space of the correct dimension in which the output state is always completely mixed. This extra mixed state will affect the minimum output entropy or the maximum output p-norm by a constant depending on d , and so it will have no effect on additivity or multiplicativity.

Fukuda proves the following result about this construction.

Theorem 7.1 (Fukuda [Fuk07]). *Let $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K}), \Psi \in \mathbf{T}(\mathcal{X}, \mathcal{Y})$ and let Φ' be the doubly stochastic channel constructed from Φ as in Equation (7.2). For these channels, and any $p \in [1, \infty]$,*

$$\begin{aligned} S_{\min}(\Psi \otimes \Phi) &= S_{\min}(\Psi \otimes \Phi') \\ \|\Psi \otimes \Phi\|_p &= \|\Psi \otimes \Phi'\|_p. \end{aligned}$$

Proof. Only a proof of the minimum output entropy case is provided, as the proof for the maximum output p-norm is identical, with the concavity of the entropy replaced by the triangle inequality.

To see that $S_{\min}(\Psi \otimes \Phi) \geq S_{\min}(\Psi \otimes \Phi')$, notice that since $W_{0,0} = \mathbb{1}$, if the input in the space that controls the Weyl operations is given as $|0, 0\rangle$, then

$$(\Psi \otimes \Phi')(\rho \otimes |0, 0\rangle\langle 0, 0|) = \mathbb{1}_{\mathcal{X}}(\Psi \otimes \Phi)(\rho) \mathbb{1}_{\mathcal{K}} = (\Psi \otimes \Phi)(\rho),$$

from which the desired inequality follows immediately.

In the other direction, notice that since the channel Φ' can be assumed to immediately measure the control input, the channel $\Psi \otimes \Phi'$ can be written as the probabilistic application of one of the discrete Weyl operators to the output of $\Psi \otimes \Phi$. To this end, let ρ be an input minimizing $S((\Psi \otimes \Phi')(\rho))$ and let the result of measuring the control input in the computational basis be $|i, j\rangle$ with probability $p_{i,j}$, and let $\rho_{i,j}$ be the state after the measurement has produced outcome i, j . In this notation we have

$$\begin{aligned}
S_{\min}(\Psi \otimes \Phi') &= S((\Psi \otimes \Phi')(\rho)) \\
&= S\left(\sum_{i,j} p_{i,j} (\mathbb{1} \otimes W_{i,j})(\Psi \otimes \Phi)(\rho_{i,j})(\mathbb{1} \otimes W_{i,j}^*)\right) \\
&\geq \sum_{i,j} p_{i,j} S((\mathbb{1} \otimes W_{i,j})(\Psi \otimes \Phi)(\rho_{i,j})(\mathbb{1} \otimes W_{i,j}^*)) \\
&= \sum_{i,j} p_{i,j} S((\Psi \otimes \Phi)(\rho_{i,j})) \\
&\geq S_{\min}(\Psi \otimes \Phi),
\end{aligned}$$

where the concavity of the entropy and the unitary invariance of the entropy have been used. \square

In Sections 7.6 and 7.5 similar results are shown for the mixed-unitary channels, though the techniques used to prove them do not seem to be directly related to Fukuda's construction in the unital case.

7.3 Mixed-unitary approximation

Given a representation of a channel Φ in Stinespring form, that is, an implementation of the form

$$\Phi(X) = \text{tr}_{\mathcal{B}} U(|0\rangle\langle 0| \otimes X)U^*, \quad (7.3)$$

there are only two operations that are not mixed-unitary. These are the partial trace over the system \mathcal{B} and the introduction of the ancillary system in the state $|0\rangle$. The goal of this section is to describe a method for approximating these two operations with mixed-unitaries, so that when combined with the circuit for the operation U in Equation (7.3), the result is a mixed-unitary approximation of Φ .

Though this approximation does have an efficient circuit implementation, the discussion of mixed-state quantum circuits is postponed to Section 7.7 as this allows the




Figure 7.1: The channel Φ to be approximated by a mixed-unitary, in Stinespring form with labelled spaces.

construction to be described in simpler terms. Additionally, some of the applications of this simulation technique do not depend on efficient circuit implementations, so this simplified exposition is useful for readers not interested in computational complexity. Despite this avoidance of the quantum circuit model, the figures in this will use circuit diagrams, but this is done only for clarity: no assumptions are made regarding implementations of the channels depicted.

To describe this approximation we fix notation throughout the next three sections. To this end let $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ be a quantum channel, and let a Stinespring representation for Φ be as given in Equation 7.3. In this representation let \mathcal{A} be the space containing the ancillary space starting in the $|0\rangle$ state, and let \mathcal{B} be the space that is traced out. This implies that the operator U is a unitary map from $\mathcal{A} \otimes \mathcal{H}$ to $\mathcal{K} \otimes \mathcal{B}$. These spaces are summarized in Figure 7.1.

7.3.1 Simulating the partial trace

Of the two operations in Equation (7.3) that are not mixed-unitary, the partial trace is the easiest to simulate with a mixed-unitary channel. At an intuitive level, the partial trace represents the loss of information to the environment in a quantum channel, but this operation is not mixed-unitary as it changes the dimension of the system being considered. The direct approach to simulating this with a mixed-unitary is to model the loss of information with a completely depolarizing channel, which avoids the issue of the change in dimensionality. It is not hard to prove that this approach works, nor is it hard to see that the completely depolarizing channel is mixed-unitary.

This depolarizing channel can be implemented as a mixture of the discrete Weyl operators, which are also known as the generalized Pauli operators [AMTdW00, BR03, HLSW04]. These unitary operators, as discussed in Chapter 1, form an orthogonal basis for the space $\mathbf{L}(\mathcal{A})$ of linear operators on a Hilbert space \mathcal{A} of dimension d .

Proposition 7.2. *The completely depolarizing channel on \mathcal{A} has implementation as a mixed-unitary channel given by*

$$\mathbf{N}(\rho) = \frac{1}{d^2} \sum_{a,b=0}^{d-1} W_{a,b} \rho W_{a,b}^* = \tilde{\mathbf{1}}_{\mathcal{A}}$$

Proof. Let $\rho \in \mathbf{D}(\mathcal{A})$ be a density matrix and let $d = \dim \mathcal{A}$. By Equation 1.7 the operators $W_{a,b}$ form a basis of $\mathbf{L}(\mathcal{A})$, so that ρ can be decomposed as

$$\rho = \sum_{e,f=1}^d \lambda_{e,f} W_{e,f}, \quad (7.4)$$

for some coefficients $\lambda_{e,f} \in \mathbb{C}$. Notice also that since $W_{a,b}$ has trace zero unless $a = b = 1$ and $W_{0,0} = \mathbf{1}_{\mathcal{A}}$, it is the case that

$$\lambda_{0,0} = \frac{\text{tr } \rho}{\text{tr } \mathbf{1}_{\mathcal{A}}} = \frac{1}{d}.$$

Putting this decomposition into the proposed implementation, we obtain

$$\frac{1}{d^2} \sum_{a,b} W_{a,b} \rho W_{a,b}^* = \frac{1}{d^2} \sum_{a,b,e,f} \lambda_{e,f} W_{a,b} W_{e,f} W_{a,b}^*.$$

Using Equation 1.6 and the unitarity of the discrete Weyl operators to manipulate this sum gives

$$\frac{1}{d^2} \sum_{a,b,e,f} \lambda_{e,f} \omega^{be-af} W_{e,f} = \sum_{e,f} \lambda_{e,f} \left(\sum_{a,b} \omega^{be-af} \right) W_{e,f}.$$

Since ω is a primitive d th root of unity, this inner summation is zero unless $e = f = 0$, and so we have shown that

$$\frac{1}{d^2} \sum_{a,b} W_{a,b} \rho W_{a,b}^* = \frac{1}{d^2} \lambda_{0,0} \sum_{a,b=0}^{d-1} W_{0,0} = \frac{1}{d^3} \sum_{a,b=0}^{d-1} \mathbf{1}_{\mathcal{A}} = \tilde{\mathbf{1}}_{\mathcal{A}} = \mathbf{N}(\rho),$$

as required. \square

This proves that the channel $\mathbf{N}_{\mathcal{B}}$ that completely depolarizes the space \mathcal{B} can be implemented as a mixed-unitary channel. To see that this channel can be used to replace the partial trace observe that one implementation of this channel simply traces out the state in \mathcal{B} and replaces it with the state $\tilde{\mathbf{1}}_{\mathcal{B}}$ that has been separately prepared. From this implementation it is clear that for a state $\rho \in \mathbf{D}(\mathcal{A} \otimes \mathcal{B})$ it holds that

$$\mathbf{N}_{\mathcal{B}}(\rho) = (\text{tr}_{\mathcal{B}} \rho) \otimes \tilde{\mathbf{1}}_{\mathcal{B}}, \quad (7.5)$$

and this property will hold no matter how $N_{\mathcal{B}}$ is implemented. This implies that if the system to be traced out instead has $N_{\mathcal{B}}$ applied to it, the resulting state is the same, up to a tensor factor of a maximally mixed state in the space \mathcal{B} . By replacing the partial trace over \mathcal{B} in Equation 7.3 with this channel, the result is

$$N_{\mathcal{B}}(\mathbf{U}(|0\rangle\langle 0| \otimes X)\mathbf{U}^*) = \Phi(X) \otimes \tilde{\mathbb{1}}_{\mathcal{B}}.$$

This is the best that can be hoped for, as a mixed-unitary transformation cannot change the dimension of the system it acts on.

7.3.2 Simulating the ancillary space

Replacing the introduction of the ancillary space \mathcal{A} with a mixed-unitary operation is more complicated than replacing the partial trace. In order to do this the input space of the transformation is to be expanded to include the space \mathcal{A} . The input state of this system will not, in general, be the desired state $|0\rangle$, so additional operations are needed to ensure that this is the case for any input state that either maximizes distinguishability or minimizes the output entropy of the resulting channel.

Because these quantities of interest, the minimum output entropy and the maximum output p -norm, involve optimizing over input states, the channel can be constructed so that those inputs that achieve the optimal value have the desired property: the input state in the space \mathcal{A} is (close to) the state $|0\rangle$. Given this property, the values of these optimizations will stay approximately the same when taken over the mixed-unitary simulations of the original channels.

To this end, the ideal operation Λ to ensure this condition does not alter any input state of the form $|0\rangle\langle 0| \otimes \sigma$, but takes any orthogonal state to the completely mixed state $\tilde{\mathbb{1}}_{\mathcal{A} \otimes \mathcal{H}}$. This operation is, unfortunately, not mixed-unitary, as it is not unital, since

$$\Lambda(\mathbb{1}_{\mathcal{A} \otimes \mathcal{H}}) = \frac{1}{\dim \mathcal{A}} |0\rangle\langle 0| \otimes \mathbb{1}_{\mathcal{H}} + \left(1 - \frac{1}{\dim \mathcal{A}}\right) \mathbb{1}_{\mathcal{A} \otimes \mathcal{H}}. \quad (7.6)$$

Notice, however, that this channel deviates from unitality with additive error $1/\dim \mathcal{A}$: there is a very good unital, and indeed mixed-unitary, approximation to this ideal channel, which is described in the remainder of this section.

This closely related mixed-unitary channel first projects the input state either onto the subspace $S_0 = |0\rangle \otimes \mathcal{H}$ or the orthogonal subspace $S_0^\perp = |0\rangle^\perp \otimes \mathcal{H}$. This projection is then followed by a completely depolarizing channel on the subspace S_0^\perp . These operations can be implemented using mixed-unitary channels, and the distance from

the ideal channel will go as $O(1/\dim \mathcal{A})$, which allows the error to be made arbitrarily small by padding \mathcal{A} with an unused ancillary space.

The mixing process on S_0^\perp is be introduced first. It is given by the channel M that does not affect the subspace S_0 but completely depolarizes the space S_0^\perp . More concretely, on a state $\rho = q\rho_{S_0} + (1 - q)\rho_{S_0^\perp}$ where $\rho_{S_0} = |0\rangle\langle 0| \otimes \sigma$ is a density operator on S_0 and $\rho_{S_0^\perp}$ a density operator on S_0^\perp , the output of M is given by

$$\begin{aligned} M(\rho) &= qM(\rho_{S_0}) + (1 - q)M(\rho_{S_0^\perp}) = q\rho_{S_0} + (1 - q)\tilde{\mathbb{1}}_{S_0^\perp} \\ &= q|0\rangle\langle 0| \otimes \sigma + (1 - q)\frac{\mathbb{1}_{\mathcal{A}} - |0\rangle\langle 0|}{\dim \mathcal{A} - 1} \otimes \tilde{\mathbb{1}}_{\mathcal{H}}. \end{aligned} \quad (7.7)$$

Here notation has been abused somewhat: S_0 and S_0^\perp are Hilbert spaces, but the whole space $\mathcal{A} \otimes \mathcal{H}$ is *not* the tensor product of these two spaces.

The channel M can be implemented as a mixed-unitary channel in the same way as the completely depolarizing channel: a uniform mixture of the discrete Weyl operators, except here these operators are taken over the subspace S_0^\perp . These operators exist, and the whole construction is very similar to the one given in Proposition 7.2. More concretely, where $W_{a,b}$ for $a, b \in \mathbb{Z}_d$ are the discrete Weyl operators on the space S_0^\perp and $\mathbb{1}_{S_0}$ is the identity on the space S_0 , the channel M can be implemented as

$$M(\rho) = \frac{1}{d^2} \sum_{a,b=0}^{d-1} (\mathbb{1}_{S_0} \oplus W_{a,b})\rho(\mathbb{1}_{S_0} \oplus W_{a,b})^*,$$

where all of the operators $\mathbb{1}_{S_0} \oplus W_{a,b}$ are mixed-unitary by construction.

As previously mentioned, this channel does not implement the ideal transformation. If the output of M on $\rho_{S_0^\perp}$ in Equation (7.7) were the completely mixed state on $\mathcal{A} \otimes \mathcal{H}$ and not the subspace S_0^\perp then this process would create an essentially error-free mixed-unitary approximation of the original channel (for the purpose of minimizing the output entropy or maximizing distinguishability). Fortunately, the error involved at this step can be shown, in Lemma 7.7, to be $O(1/\dim \mathcal{A})$, which, by Equation 7.6, is as close as a mixed-unitary channel can come to the ideal case. Fortunately this error can be made arbitrarily small by taking the space \mathcal{A} large enough, and so this construction can be used to approximate the ideal case.

It will be helpful for the analysis of this construction to remove the coherences between the subspaces S_0 and S_0^\perp . The channel M does perform this operation. This is the operation commonly known as dephasing that, applied to a density matrix expressed in some basis, removes the off-diagonal terms. This aids the analysis of the construction, because once this dephasing is applied, an equation similar to Equation (7.7)

would hold for all input states ρ , not just those states that have no entanglement between the subspaces S_0 and S_0^\perp .

While we will only need to apply dephasing between the two subspaces S_0 and S_0^\perp , a mixed-unitary construction for the general case is provided below. This is the channel D that completely decoheres all information not stored in the computational basis. More specifically, this channel implements

$$D(|i\rangle\langle j|) = \delta_{ij}|i\rangle\langle j| = \begin{cases} |i\rangle\langle i| & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases} \quad (7.8)$$

That this channel is mixed-unitary is simple to prove, using a construction based on the discrete Weyl operators, similar to that used for the complete depolarizing channel in Proposition 7.2. That this channel can be implemented in this way has been observed in [DFH06].

Proposition 7.3. *The completely dephasing channel on \mathcal{A} defined in Equation (7.8) has implementation as a mixed-unitary channel given by*

$$D(\rho) = \frac{1}{d} \sum_{b=0}^{d-1} W_{0,b} \rho W_{0,b}^*$$

Proof. Recall that $W_{0,b}|j\rangle = Z^b|j\rangle = \omega^{bj}|j\rangle$ as introduced in Chapter 1, where ω is a d th primitive root of unity, with $d = \dim \mathcal{A}$. To see that this channel has the desired effect, let $\rho = \sum_{i,j=0}^{d-1} a_{ij}|i\rangle\langle j|$, so that

$$\frac{1}{d} \sum_{b=0}^{d-1} W_{0,b} \rho W_{0,b}^* = \frac{1}{d} \sum_{b=0}^{d-1} \sum_{i,j=0}^{d-1} a_{i,j} Z^b |i\rangle\langle j| Z^{-b} = \frac{1}{d} \sum_{b=0}^{d-1} \sum_{i,j=0}^{d-1} a_{i,j} \omega^{(i-j)b} |i\rangle\langle j|. \quad (7.9)$$

Then, since ω is a d th root of unity

$$\sum_{b=0}^{d-1} \omega^{(i-j)b} = d\delta_{ij} = \begin{cases} d & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

Combining this property with Equation (7.9) gives

$$\frac{1}{d} \sum_{b=0}^{d-1} W_{0,b} \rho W_{0,b}^* = \frac{1}{d} \sum_{i=0}^{d-1} d a_{i,i} |i\rangle\langle i| = \sum_{i=0}^{d-1} a_{i,i} |i\rangle\langle i| = D(\rho),$$

which is exactly the channel defined by Equation (7.8). \square

For the specific case that we will use here, a simpler construction suffices: instead of applying this dephasing channel to the whole of the ancillary space \mathcal{A} it only needs to be applied to remove coherences between the two orthogonal subspaces S_0 and S_0^\perp . If these two subspaces are viewed as a two-dimensional Hilbert space, the construction in Proposition 7.3 can be reduced to the application of a specific unitary operator V with probability one-half. The action of this unitary V on basis states is given by

$$V|i\rangle = \begin{cases} |i\rangle & \text{if } |i\rangle \in S_0, \\ -|i\rangle & \text{if } |i\rangle \in S_0^\perp. \end{cases} \quad (7.10)$$

In other words, V applies a phase of -1 to states in S_0^\perp and does not change states in S_0 . When V is applied with probability one half the result is complete dephasing between the two subspaces. This can be seen by restricting the construction in Proposition 7.3 to the case of a two-dimensional system with orthogonal states that represent the subspaces S_0 and S_0^\perp . More concretely, when this is applied to a density matrix expressed in the computational basis, the result is, by a simple calculation, the zeroing of the off-diagonal elements of the first row and column. Let this simplified dephasing channel be given by

$$D_{S_0}(\rho) = \frac{1}{2} [V\rho V^* + \rho],$$

where the operator V is given in Equation (7.10). When this operation is applied to a density operator $\rho \in \mathbf{D}(\mathcal{A} \otimes \mathcal{H})$, the result is

$$D_{S_0}(\rho) = q\rho_{S_0} + (1 - q)\rho_{S_0^\perp} = q|0\rangle\langle 0| \otimes \sigma + (1 - q)\rho_{S_0^\perp}, \quad (7.11)$$

where $\rho_{S_0} = |0\rangle\langle 0| \otimes \sigma$ is a density operator on the subspace $S_0 = |0\rangle \otimes \mathcal{H}$, $\rho_{S_0^\perp}$ is a density operator on the orthogonal subspace S_0^\perp , and $0 \leq q \leq 1$ is a probability.

Combining Equations (7.7) and (7.11), the output of D_{S_0} followed by M on a density operator ρ on $\mathcal{A} \otimes \mathcal{H}$ is given by a state of the form

$$\begin{aligned} (M \circ D_{S_0})(\rho) &= qM(|0\rangle\langle 0| \otimes \sigma) + (1 - q)M(\rho_{S_0^\perp}) \\ &= q|0\rangle\langle 0| \otimes \sigma + (1 - q)\frac{\mathbb{1}_{\mathcal{A}} - |0\rangle\langle 0|}{\dim \mathcal{A} - 1} \otimes \tilde{\mathbb{1}}_{\mathcal{H}}. \end{aligned}$$

This operation, $M \circ D_{S_0}$, will be used as a way to force any input that results in a low output entropy to be close to the subspace S_0 of inputs having the ‘ancilla’ space \mathcal{A} in the desired $|0\rangle$ state. On these inputs the constructed mixed-unitary channel will behave in a similar way to the original channel that is being approximated. On inputs that are far from this subspace, the resulting state has high entropy, and so it will not be close to a state minimizing the output entropy and it will not be useful for distinguishing two channels constructed in this way.

7.3.3 Mixed-unitary approximation of a general channel

Putting these pieces together, given a channel $\Phi(\rho) = \text{tr}_{\mathcal{B}} \mathbf{U}(\rho \otimes |0\rangle\langle 0|)\mathbf{U}^*$, the mixed-unitary approximation Φ' is constructed as

$$\Phi'(\rho) = \mathbf{N}_{\mathcal{B}} (\mathbf{U} [(M \circ D_{S_0})(\rho)] \mathbf{U}^*), \quad (7.12)$$

which, more plainly, is simply the application of the ancilla simulation procedure of Section 7.3.2, the unitary operation from a Stinespring dilation of Φ , and finally the completely mixing channel to the space that would have been traced out by Φ , as discussed in Section 7.3.1. As the mixed-unitary channels are closed under composition, the channel Φ' is mixed-unitary.

It will be useful to observe that the constructed channel Φ' specified in Equation (7.12) can be used to simulate the original channel Φ . This occurs when the input $|0\rangle\langle 0| \otimes \sigma$, i.e. an input in the space S_0 , is provided to Φ' . This is argued in the following proposition.

Proposition 7.4. *Let $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$. If $\Phi' \in \mathbf{T}(\mathcal{A} \otimes \mathcal{H}, \mathcal{K} \otimes \mathcal{B})$ is the mixed-unitary channel that is constructed from Φ in Equation (7.12), then*

$$\Phi'(|0\rangle\langle 0| \otimes \sigma) = \Phi(\sigma) \otimes \tilde{\mathbf{1}}_{\mathcal{B}}.$$

Proof. Notice that both D_{S_0} and M do not affect this input: the decoherence operation D_{S_0} does not affect the state as it is in the subspace S_0 and M does not affect the state by Equation (7.7). Thus, the output of the channel Φ' is

$$\begin{aligned} \Phi'(|0\rangle\langle 0| \otimes \sigma) &= \mathbf{N}_{\mathcal{B}} (\mathbf{U} [(M \circ D_{S_0})(|0\rangle\langle 0| \otimes \sigma)] \mathbf{U}^*) \\ &= \mathbf{N}_{\mathcal{B}} (\mathbf{U} (|0\rangle\langle 0| \otimes \sigma) \mathbf{U}^*) \\ &= \text{tr}_{\mathcal{B}} (\mathbf{U} (|0\rangle\langle 0| \otimes \sigma) \mathbf{U}^*) \otimes \tilde{\mathbf{1}}_{\mathcal{B}} \\ &= \Phi(\sigma) \otimes \tilde{\mathbf{1}}_{\mathcal{B}}, \end{aligned}$$

where the penultimate equality is an application of Equation (7.5). \square

Combining this proposition with Equation (7.11) that demonstrates the effect of the $M \circ D_{S_0}$ on states not of this form, and the observation that applying $M \circ D_{S_0}$ twice has no further effect than applying it once, the output of Φ' on an arbitrary input state ρ is given by

$$\Phi'(\rho) = p\Phi'(|0\rangle\langle 0| \otimes \sigma) + (1-p)\Phi'(\rho_{S_0^\perp}) = p\Phi(\sigma) \otimes \tilde{\mathbf{1}}_{\mathcal{B}} + (1-p)\Phi'(\rho_{S_0^\perp}), \quad (7.13)$$

where as in Equation (7.11) $\rho_{S_0^\perp}$ is a density operator on the subspace S_0^\perp of inputs orthogonal to those with the state $|0\rangle$ on the space \mathcal{A} . The most significant portion of the technical results in the next section lies in bounding the distance from the maximally mixed state of the second term in this equation, from which most of the results will follow straightforwardly.

7.4 Properties of the constructed channel

This section provides the basis for the analysis of the mixed-unitary approximation constructed in the previous section. The main result is a lower bound on the output entropy when the constructed channel is applied to a state in S_0^\perp , the subspace of inputs where the ‘ancillary’ subspace \mathcal{A} is not in the desired $|0\rangle$ state. This result is not difficult to show, but it will be essential to the results that follow.

Throughout this section, and the two sections that follow, the channel Φ will represent an arbitrary transformation, and Φ' will represent the mixed-unitary transformation constructed from it, as in Equation (7.12). The names of the Hilbert spaces that Φ acts on will be consistent with the previous section: Φ maps mixed states on \mathcal{H} to \mathcal{K} , using the ancillary system \mathcal{A} and tracing out the system \mathcal{B} . The constructed channel Φ' is mixed-unitary, mapping density matrices on $\mathcal{A} \otimes \mathcal{H}$ to $\mathcal{K} \otimes \mathcal{B}$.

As a first step to showing that Φ' approximates Φ it is shown that mixed-unitary channels do not increase the distance of a state from the completely mixed state. This lemma can be interpreted as the statement that the output of a mixed-unitary channel is not more pure than the input. The Hilbert space \mathcal{B} appearing in this lemma will correspond to a reference system needed for the results in Section 7.8 – this generality will not be needed for the results on the maximum output p -norm or the minimum output entropy.

Lemma 7.5. *Let $\|\cdot\|$ be a unitarily invariant norm on $\mathbf{L}(\mathcal{A} \otimes \mathcal{B})$. If $\Psi \in \mathbf{T}(\mathcal{A})$ is mixed-unitary, then for any $\rho \in \mathbf{D}(\mathcal{A} \otimes \mathcal{B})$*

$$\|(\Psi \otimes I_{\mathcal{B}})(\rho) - \tilde{\mathbf{1}}_{\mathcal{A}} \otimes \text{tr}_{\mathcal{A}} \rho\| \leq \|\rho - \tilde{\mathbf{1}}_{\mathcal{A}} \otimes \text{tr}_{\mathcal{A}} \rho\|$$

Proof. As Ψ is mixed-unitary, let $\Psi(X) = \sum_i p_i U_i X U_i^*$ with the U_i unitary, $0 \leq p_i \leq 1$, and $\sum_i p_i = 1$. For brevity, let $\hat{U}_i = U_i \otimes \mathbf{1}_{\mathcal{B}}$ for all i . Using this notation

$$\begin{aligned} \|(\Psi \otimes I_{\mathcal{B}})(\rho) - \tilde{\mathbf{1}}_{\mathcal{A}} \otimes \text{tr}_{\mathcal{A}} \rho\| &= \left\| \sum_i p_i \hat{U}_i \rho \hat{U}_i^* - \tilde{\mathbf{1}}_{\mathcal{A}} \otimes \text{tr}_{\mathcal{A}} \rho \right\| \\ &\leq \sum_i p_i \|\hat{U}_i \rho \hat{U}_i^* - \tilde{\mathbf{1}}_{\mathcal{A}} \otimes \text{tr}_{\mathcal{A}} \rho\|. \end{aligned} \quad (7.14)$$

Notice that $U_i \tilde{\mathbf{1}}_{\mathcal{A}} U_i^* = \tilde{\mathbf{1}}_{\mathcal{A}}$, which implies that $\hat{U}_i(\tilde{\mathbf{1}}_{\mathcal{A}} \otimes \sigma) \hat{U}_i^* = \tilde{\mathbf{1}}_{\mathcal{A}} \otimes \sigma$. Using this fact, as well as the unitary invariance of the norm, Equation (7.14) becomes

$$\sum_i p_i \left\| \left\| \hat{U}_i(\rho - \tilde{\mathbf{1}}_{\mathcal{A}} \otimes \text{tr}_{\mathcal{A}} \rho) \hat{U}_i^* \right\| \right\| = \sum_i p_i \left\| \left\| \rho - \tilde{\mathbf{1}}_{\mathcal{A}} \otimes \text{tr}_{\mathcal{A}} \rho \right\| \right\| = \left\| \left\| \rho - \tilde{\mathbf{1}}_{\mathcal{A}} \otimes \text{tr}_{\mathcal{A}} \rho \right\| \right\|.$$

Combining this with Equation (7.14) yields the statement of the lemma. \square

This lemma will be used to show not only that the ancilla simulation procedure sends states in the subspace S_0^\perp of states where the ancillary space is not in the $|0\rangle$ state to states that are highly mixed, but that the channel Φ' also has this behaviour. Before doing this, however, the lemma is extended to the case of the von Neumann entropy, where the proof is essentially identical, with the exception that the triangle inequality is replaced by concavity.

Corollary 7.6. *If $\Psi \in \mathbf{T}(\mathcal{A})$ is mixed-unitary, and $\rho \in \mathbf{D}(\mathcal{A})$, then*

$$S(\rho) \leq S(\Psi(\rho)).$$

Proof. Let $\Psi(\rho) = \sum_i p_i U_i \rho U_i^*$ as in the proof of Lemma 7.5. Using this notation, and the concavity of the von Neumann entropy

$$S(\Psi(\rho)) = S\left(\sum_i p_i U_i \rho U_i^*\right) \geq \sum_i p_i S(U_i \rho U_i^*) = \sum_i p_i S(\rho) = S(\rho),$$

where the penultimate equality is due to the unitary invariance of the entropy. \square

The next lemma shows that when the input is in the subspace S_0^\perp the output of Φ' is very close to completely mixed. The distance measure used is the trace norm, but this can be applied also to the case of the maximum output p -norm due to the fact that $\|\rho\|_{\text{tr}} = \|\rho\|_1 \geq \|\rho\|_p$ for all $p \in [1, \infty]$. This is the key lemma in the proof of the results on the additivity and multiplicativity conjectures, though it is not difficult to prove.

Lemma 7.7. *On input states $\rho \in S_0^\perp$ the output of the channel Φ' given in Equation (7.12) satisfies*

$$\left\| \Phi'(\rho) - \tilde{\mathbf{1}}_{\mathcal{X} \otimes \mathcal{B}} \right\|_{\text{tr}} \leq \frac{2}{\dim \mathcal{A}}.$$

Proof. On input $\rho \in S_0^\perp$ the operation D_{S_0} that introduces decoherence between the subspaces S_0 and S_0^\perp has no effect. This implies that the output of $M \circ D_{S_0}$ on ρ is obtained by setting $q = 0$ in Equation (7.7), which is

$$M(D_{S_0}(\rho)) = \frac{\mathbf{1}_{\mathcal{A}} - |0\rangle\langle 0|}{\dim \mathcal{A} - 1} \otimes \tilde{\mathbf{1}}_{\mathcal{X}}, \quad (7.15)$$

Setting $d = \dim \mathcal{A}$, the distance from from the completely mixed state on $\mathcal{A} \otimes \mathcal{H}$ is

$$\begin{aligned} \left\| \frac{\mathbb{1}_{\mathcal{A}} - |0\rangle\langle 0|}{d-1} \otimes \tilde{\mathbb{1}}_{\mathcal{H}} - \frac{\mathbb{1}_{\mathcal{A}}}{d} \otimes \tilde{\mathbb{1}}_{\mathcal{H}} \right\|_{\text{tr}} &= \left\| \frac{\mathbb{1}_{\mathcal{A}} - d|0\rangle\langle 0|}{d(d-1)} \right\|_{\text{tr}} \\ &\leq \frac{d-1}{d(d-1)} + \frac{d-1}{d(d-1)} = \frac{2}{d}. \end{aligned} \quad (7.16)$$

Finally, by noting that the remainder of the transformation Φ' is mixed-unitary and (implicitly) using the isomorphism $\mathcal{A} \otimes \mathcal{H} \cong \mathcal{K} \otimes \mathcal{B}$, an application of Lemma 7.5 yields the desired bound. \square

Once again we can extend this result to the case of the entropy. The previous lemma on the trace norm can be extended to the case of the entropy in the standard way: using Fannes' inequality [Fan73] (see also [NC00]), but given the characterization in Equation (7.15), a better bound can be obtained by explicitly computing the entropy. This bound will require that $\dim \mathcal{A} \geq 2$, but this can be assumed without loss of generality by adding an unused ancillary space.

Corollary 7.8. *Let Φ' be given as in Equation (7.12), and let $\dim \mathcal{A} \geq 2$ and $\rho \in \mathcal{S}_0^\perp$, then*

$$S(\Phi'(\rho)) \geq S(\tilde{\mathbb{1}}_{\mathcal{K} \otimes \mathcal{B}}) - \frac{1}{\dim \mathcal{A}}.$$

Proof. In the proof of Lemma 7.7, the output of $M \circ D_{S_0}$ on ρ is given by Equation (7.15), which states that

$$M(D_{S_0}(\rho)) = \frac{\mathbb{1}_{\mathcal{A}} - |0\rangle\langle 0|}{\dim \mathcal{A} - 1} \otimes \tilde{\mathbb{1}}_{\mathcal{H}}.$$

Letting $d = \dim \mathcal{A}$, this state has $(d-1)(\dim \mathcal{H})$ eigenvalues and each with value equal to $1/((d-1)(\dim \mathcal{H}))$. Using this observation, the entropy of this state can be computed as

$$\begin{aligned} S\left(\frac{\mathbb{1}_{\mathcal{A}} - |0\rangle\langle 0|}{\dim \mathcal{A} - 1} \otimes \tilde{\mathbb{1}}_{\mathcal{H}}\right) &= \log((d-1)(\dim \mathcal{H})) \\ &= \log \dim \mathcal{H} + \log \dim \mathcal{A} + \log\left(\frac{d-1}{d}\right) \\ &= S(\tilde{\mathbb{1}}_{\mathcal{A} \otimes \mathcal{H}}) - \log\left(1 + \frac{1}{d-1}\right). \end{aligned} \quad (7.17)$$

For $d \geq 2$, the last term has Taylor expansion given by

$$\log\left(1 + \frac{1}{d-1}\right) = \frac{1}{\log e} \left[\frac{1}{d-1} - \frac{1}{2(d-1)^2} + \frac{1}{3(d-1)^3} - \dots \right] \leq \frac{\log_e 2}{(d-1) \log e} \leq \frac{1}{d}.$$

Combining this with Equation (7.17) gives the lower bound on the entropy provided in the statement of the corollary, for the state after ancilla simulation procedure. By Corollary 7.6 applying the remainder of Φ' to this state cannot decrease the entropy, as this portion of Φ' is mixed-unitary. \square

This corollary and Lemma 7.7 show that when the input state to Φ' does not have large overlap with a state in S_0 , the output state is highly mixed. This property will be used to show that any state that maximizes the output norm or minimizes the output entropy will have large overlap with the space S_0 of states of the form $|0\rangle\langle 0| \otimes \sigma$, where the simulation of the original channel is faithful.

7.5 Multiplicativity of mixed-unitary transformations

In this section the construction of Section 7.3 is used to show that the maximum output p -norm of a channel is multiplicative if and only if the mixed-unitary approximations to it are also multiplicative. This will be done for all $1 \leq p < \infty$, using the analysis of the previous section.

This property is not difficult to show once it has been established that the mixed-unitary channel Φ' constructed from Φ in Equation (7.12) is a good approximation with respect to the p -norm. This is the content of the following theorem.

Theorem 7.9. *If $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$, then the mixed-unitary $\Phi' \in \mathbf{T}(\mathcal{A} \otimes \mathcal{H}, \mathcal{K} \otimes \mathcal{B})$ satisfies*

$$\nu_p(\Phi) \leq \frac{\nu_p(\Phi')}{\|\tilde{\mathbf{1}}_{\mathcal{B}}\|_p} \leq \nu_p(\Phi) + \frac{2 \dim \mathcal{B}}{\dim \mathcal{A}}.$$

Proof. For convenience, let $d = \dim \mathcal{A}$. The first inequality is simple: Proposition 7.4 shows that

$$\Phi'(|0\rangle\langle 0| \otimes \rho) = \Phi(\rho) \otimes \tilde{\mathbf{1}}_{\mathcal{B}},$$

from which it follows immediately that

$$\nu_p(\Phi) \|\tilde{\mathbf{1}}_{\mathcal{B}}\|_p \leq \nu_p(\Phi'),$$

as ν_p is a maximization over input states, and $\|\cdot\|_p$ is multiplicative with respect to the tensor product of two states.

To prove the second inequality let $\rho \in \mathbf{D}(\mathcal{A} \otimes \mathcal{H})$ be a state such that

$$\nu_p(\Phi') = \|\Phi'(\rho)\|_p. \quad (7.18)$$

Such a state exists by the compactness of $\mathbf{D}(\mathcal{A} \otimes \mathcal{H})$. The output of the channel Φ' on ρ is given by Equation (7.13), applying the triangle inequality to this yields

$$\|\Phi'(\rho)\|_p = \left\| q\Phi(\sigma) \otimes \tilde{\mathbf{1}}_{\mathcal{B}} + (1-q)\Phi'(\rho_{S_0^\perp}) \right\|_p \leq q\|\Phi(\sigma) \otimes \tilde{\mathbf{1}}_{\mathcal{B}}\|_p + (1-q)\|\Phi'(\rho_{S_0^\perp})\|_p.$$

Lemma 7.7 provides a bound on the second term of this equation, which implies that

$$\|\Phi'(\rho)\|_p \leq q \|\Phi(\sigma) \otimes \tilde{\mathbf{1}}_{\mathcal{B}}\|_p + (1-q) \left(\|\tilde{\mathbf{1}}_{\mathcal{K} \otimes \mathcal{B}}\|_p + \frac{2}{d} \right).$$

Then, as the norm $\|\cdot\|_p$ is multiplicative with respect to the tensor product of states, and $\|\tilde{\mathbf{1}}_{\mathcal{K}}\|_p \leq \|\xi\|_p$ for any state $\xi \in \mathbf{D}(\mathcal{K})$,

$$\|\Phi'(\rho)\|_p \leq q \|\Phi(\sigma)\|_p \|\tilde{\mathbf{1}}_{\mathcal{B}}\|_p + (1-q) \left(\|\tilde{\mathbf{1}}_{\mathcal{K}}\|_p \|\tilde{\mathbf{1}}_{\mathcal{B}}\|_p + \frac{2}{d} \right) \leq \|\Phi(\sigma)\|_p \|\tilde{\mathbf{1}}_{\mathcal{B}}\|_p + \frac{2}{d}.$$

Then, by the choice of the input ρ in Equation (7.18), we have shown that

$$\nu_p(\Phi') = \|\Phi'(\rho)\|_p \leq \nu_p(\Phi) \|\tilde{\mathbf{1}}_{\mathcal{B}}\|_p + \frac{2}{d}. \quad (7.19)$$

Finally, the state $\tilde{\mathbf{1}}_{\mathcal{B}}$ has $\dim \mathcal{B}$ eigenvalues, each with value $1/\dim \mathcal{B}$, which implies that

$$\|\tilde{\mathbf{1}}_{\mathcal{B}}\|_p = \left(\sum_{i=1}^{\dim \mathcal{B}} \frac{1}{(\dim \mathcal{B})^p} \right)^{1/p} = \dim \mathcal{B}^{1/p-1} \geq \frac{1}{\dim \mathcal{B}}.$$

Combining this with Equation (7.19), and expanding $d = \dim \mathcal{A}$, implies

$$\frac{\nu_p(\Phi')}{\|\tilde{\mathbf{1}}_{\mathcal{B}}\|_p} \leq \nu_p(\Phi) + \frac{2}{\|\tilde{\mathbf{1}}_{\mathcal{B}}\|_p \dim \mathcal{A}} \leq \nu_p(\Phi) + \frac{2 \dim \mathcal{B}}{\dim \mathcal{A}},$$

which completes the proof of the second inequality. \square

With this approximation result, the main theorem on the maximum output p -norm can be shown. This extends the construction of Section 7.2 due to Fukuda [Fuk07] on unital channels to the mixed-unitary unitary case, using essentially the same method of proof.

Theorem 7.10. *If $\Phi, \Psi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$ and $p \in [1, \infty)$, then*

$$\nu_p(\Phi \otimes \Psi) = \nu_p(\Phi) \nu_p(\Psi)$$

if

$$\nu_p(\Phi'_d \otimes \Psi) = \nu_p(\Phi'_d) \nu_p(\Psi),$$

for all sufficiently large d , where Φ'_d is the mixed-unitary approximation of the channel Φ obtained by applying the construction of Section 7.3 to a Stinespring dilation of Φ using a d -dimensional ancillary space.

Proof. As adding ancillary space to Φ' increases both $\dim \mathcal{A}$ and $\dim \mathcal{B}$, by taking $d = \dim \mathcal{A}$ large enough it can be assumed that $\dim \mathcal{B} \leq 2d$. Let $\epsilon > 0$, and choose d so that $4/d^{1/p} < \epsilon$. This, along with the choice of $\dim \mathcal{B} \leq 2d$ implies that

$$2 \dim \mathcal{B}^{1-1/p}/d \leq 4/d^{1/p} < \epsilon \quad (7.20)$$

Then, as $\Phi'_d(|0\rangle\langle 0| \otimes \rho) = \Phi(\rho) \otimes \tilde{\mathbf{1}}_{\mathcal{B}}$ by Proposition 7.4,

$$\nu_p(\Phi \otimes \Psi) \leq \frac{\nu_p(\Phi'_d \otimes \Psi)}{\|\tilde{\mathbf{1}}_{\mathcal{B}}\|_p}.$$

By assumption, this second quantity is multiplicative, so that

$$\nu_p(\Phi \otimes \Psi) \leq \frac{\nu_p(\Phi'_d \otimes \Psi)}{\|\tilde{\mathbf{1}}_{\mathcal{B}}\|_p} = \frac{\nu_p(\Phi'_d)\nu_p(\Psi)}{\|\tilde{\mathbf{1}}_{\mathcal{B}}\|_p}.$$

Applying Theorem 7.9 to this quantity shows that

$$\nu_p(\Phi \otimes \Psi) \leq \left[\nu_p(\Phi) + \frac{4}{d^{1/p}} \right] \nu_p(\Psi) < \nu_p(\Phi)\nu_p(\Psi) + \epsilon,$$

where the final inequality is by the choice of d to satisfy Equation (7.20). As epsilon was chosen arbitrarily, the multiplicativity of $\nu_p(\Phi'_d)$ for all large enough d implies the multiplicativity of $\nu_p(\Phi)$. \square

This theorem shows that in order to show the multiplicativity of ν_p on a class of channels it suffices to consider a related class of mixed-unitary channels. This problem may be more tractable for channels of this type: many of the known counterexamples to multiplicativity for small values of p are mixed-unitary [HW08].

7.6 Mixed-unitaries and minimum output entropy

The results of the previous section on the multiplicativity of the maximum output p -norm can be extended directly to the additivity of the minimum output entropy. This is done using very similar proof techniques as in the previous section.

The following theorem demonstrates that the mixed-unitary Φ' constructed in Equation (7.12) forms a good approximation of the original channel Φ , from which the result on the additivity will follow directly.

Theorem 7.11. *If $\Phi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$, then the mixed-unitary $\Phi' \in \mathbf{T}(\mathcal{A} \otimes \mathcal{H}, \mathcal{K} \otimes \mathcal{B})$ satisfies*

$$S_{\min}(\Phi) \geq S_{\min}(\Phi') - \log \dim \mathcal{B} \geq S_{\min}(\Phi) - \frac{1}{\dim \mathcal{A}}.$$

Proof. Exactly as in the case of Theorem 7.9, Proposition 7.4 implies the first inequality, as Φ' on a particular state can be used to simulate Φ :

$$\Phi'(|0\rangle\langle 0| \otimes \rho) = \Phi(\rho) \otimes \tilde{\mathbf{1}}_{\mathcal{B}}.$$

Let ρ be a state minimizing $S(\Phi'(\rho))$ and for convenience let $\delta = 1/\dim \mathcal{A}$. Equation (7.13) gives the output of Φ' on ρ . Applying the concavity of the entropy (Proposition 3.3) to this, we obtain

$$S_{\min}(\Phi') = S(\Phi'(\rho)) \geq qS(\Phi(\sigma) \otimes \tilde{\mathbf{1}}_{\mathcal{B}}) + (1 - q)S(\Phi'(\rho_{S_0^\perp})).$$

Applying Corollary 7.8 this becomes

$$S_{\min}(\Phi') \geq qS(\Phi(\sigma) \otimes \tilde{\mathbf{1}}_{\mathcal{B}}) + (1 - q)(S(\tilde{\mathbf{1}}_{\mathcal{A} \otimes \mathcal{H}}) - \delta).$$

Notice that, since Φ' is mixed-unitary, $\mathcal{A} \otimes \mathcal{H}$ is isomorphic to $\mathcal{K} \otimes \mathcal{B}$. This implies that $S(\tilde{\mathbf{1}}_{\mathcal{A} \otimes \mathcal{H}}) = S(\tilde{\mathbf{1}}_{\mathcal{K} \otimes \mathcal{B}})$.

Two additional properties of the entropy introduced in Section 3.1 will be useful: the additivity of the entropy on states (Equation (3.2)), $S(\sigma \otimes \xi) = S(\sigma) + S(\xi)$, for any σ, ξ ; and the fact that the entropy is maximized on completely mixed states, $S(\xi) \leq \log \dim \mathcal{K} = S(\tilde{\mathbf{1}}_{\mathcal{K}})$ for all $\xi \in \mathbf{D}(\mathcal{K})$ (Proposition 3.2). Using these three observations, in order, we find that

$$\begin{aligned} S_{\min}(\Phi') &\geq qS(\Phi(\sigma) \otimes \tilde{\mathbf{1}}_{\mathcal{B}}) + (1 - q)(S(\tilde{\mathbf{1}}_{\mathcal{K} \otimes \mathcal{B}}) - \delta) \\ &= q(S(\Phi(\sigma)) + S(\tilde{\mathbf{1}}_{\mathcal{B}})) + (1 - q)(S(\tilde{\mathbf{1}}_{\mathcal{K}}) + S(\tilde{\mathbf{1}}_{\mathcal{B}}) - \delta) \\ &\geq q(S(\Phi(\sigma)) + S(\tilde{\mathbf{1}}_{\mathcal{B}})) + (1 - q)(S(\Phi(\sigma)) + S(\tilde{\mathbf{1}}_{\mathcal{B}}) - \delta) \\ &\geq S(\Phi(\sigma)) + S(\tilde{\mathbf{1}}_{\mathcal{B}}) - \delta. \end{aligned}$$

Finally, since $S(\tilde{\mathbf{1}}_{\mathcal{B}}) = \log \dim \mathcal{B}$ and $S_{\min}(\Phi) \leq S(\Phi(\xi))$ for any ξ , we have

$$S_{\min}(\Phi') \geq S_{\min}(\Phi) + \log \dim \mathcal{B} - \delta,$$

which completes the proof of the theorem. \square

The proof that the additivity of the minimum output entropy can be equivalently restricted to mixed-unitary channels follows from the previous theorem in a way that is identical to the proof of Theorem 7.10, with the exception that the p-norm has been replaced by the minimum output entropy. The method of proof here follows Fukuda's result for unital channels [Fuk07].

Theorem 7.12. *If $\Phi, \Psi \in \mathbf{T}(\mathcal{H}, \mathcal{K})$, then*

$$S_{\min}(\Phi \otimes \Psi) = S_{\min}(\Phi) + S_{\min}(\Psi)$$

if

$$S_{\min}(\Phi'_d \otimes \Psi) = S_{\min}(\Phi'_d) + S_{\min}(\Psi),$$

for all sufficiently large d , where Φ'_d is the mixed-unitary extension of the channel obtained by applying the construction of Section 7.3 to Stinespring dilation for Φ using an ancillary space of dimension d .

Proof. Let $\epsilon > 0$, and choose d large enough so that $1/d < \epsilon$. Then, as $\Phi'_d(|0\rangle\langle 0| \otimes \rho) = \Phi(\rho) \otimes \tilde{\mathbb{1}}_{\mathcal{B}}$,

$$S_{\min}(\Phi \otimes \Psi) \geq S_{\min}(\Phi'_d \otimes \Psi) - \log \dim \mathcal{B}.$$

By assumption, this second quantity is additive, so that

$$\begin{aligned} S_{\min}(\Phi \otimes \Psi) &\geq S_{\min}(\Phi'_d \otimes \Psi) - \log \dim \mathcal{B} \\ &= S_{\min}(\Phi'_d) + S_{\min}(\Psi) - \log \dim \mathcal{B} \\ &\geq S_{\min}(\Phi) - \frac{1}{d} + S_{\min}(\Psi) \\ &> S_{\min}(\Phi) + S_{\min}(\Psi) - \epsilon \end{aligned}$$

where the penultimate inequality is an application of Theorem 7.11. As ϵ was chosen arbitrarily, the additivity of Φ'_d for all large enough d implies the additivity of Φ . \square

This theorem implies that in order to prove the additivity of the minimum output entropy for a class of channels, the hopefully simpler class of mixed-unitary approximations can instead be considered. This may be a fruitful approach: the only channels for which S_{\min} is known not to be additive are mixed-unitary [Has09] – this property may be simpler to check for mixed-unitaries having certain properties.

7.7 Circuit constructions

In this section an efficient circuit construction is provided for the mixed-unitary approximation described in Section 7.3. This construction is used to extend the hardness of computationally distinguishing quantum circuits to the case of mixed-unitary circuits.

Before constructing these circuits, it will be important to specify the circuit models that are being used. The circuit model used to define the quantum circuit distinguishability problem is the *mixed state quantum circuit* model of Aharonov, Kitaev, and Nisan [AKN98], which is described in Section 2.1. As previously discussed, we may assume that circuits in this model first introduce any necessary ancillary qubits, then perform a unitary operation, and finally trace out those qubits that are not part of the output. This approach is equivalent to building a circuit for the Stinespring dilation of a channel. As all unitary transformations can be (approximately) implemented using one and two qubit gates there is no loss in generality in assuming that the unitary transformations implemented in such a circuit are composed of gates from some finite basis of one and two qubit gates.

The second model of quantum circuits we consider is the model of *mixed-unitary quantum circuits*. These circuits consist of one and two qubit gates from the usual circuit model as well as mixed-unitary gates, which implement a unitary gate with probability one half. More formally, the application of such a gate performs the operation

$$\rho \mapsto \frac{1}{2}U\rho U^* + \frac{1}{2}\rho,$$

where U is a one or two qubit unitary gate in the standard gate set.

For technical reasons, we need to assume that the Pauli X and Z gates, as well as controlled versions of these gates, are part of the standard basis. This restriction can be avoided by allowing gates that implement

$$\rho \mapsto \frac{1}{2}U_k \cdots U_2 U_1 \rho U_1^* U_2^* \cdots U_k^* + \frac{1}{2}\rho,$$

where the U_i are gates of the standard model. This allows sequences of multiple gates, such as approximations to gates not in the basis, to be applied with probability one half. When proving a hardness result, the model should be as restricted as possible, and it is not clear that this model is not more powerful than the model where each mixed-unitary gate is applied with an independent probability of one half.

The model of mixed-unitary circuits is an extremely simple model that does not appear to be universal for the class of transformations that implement mixed-unitary operations. It is not clear that this is the correct definition of the mixed-unitary circuit model, but since the aim of the model to prove a hardness result, an extremely weak definition has been chosen so that the result will apply to as large a class of circuit models as possible.

One drawback of this weak model is that the exact construction used in Section 7.3 cannot be implemented. Specifically, the operation D that decoheres the subspaces S_0

and S_0^\perp seems to require a unitary operation that cannot be decomposed into a series of one and two qubit gates, applied with probability one half. A similar situation occurs for the implementation of the completely depolarizing channel on the subspace S_0^\perp , the implementation of which uses the discrete Weyl operators on the subspace S_0^\perp . These operations can be implemented in a mixed-unitary way in a more permissive circuit model, but in order to keep the circuit model as simple as possible, a modified construction is presented here. This modified construction is built from pieces that perform similar tasks to those used in Section 7.3, but the specific building blocks are not exactly the same. The construction in this section can also be applied to the additivity and multiplicativity problems considered in Sections 7.5 and 7.6, but it is somewhat more complicated than the construction already presented.

In order to approximate a given circuit with a mixed-unitary circuit we once again make use of three main components, which are once again referred to as N, D, and M. These pieces are labelled in this way due to the fact that they play the same roles as the components of the same names used in Section 7.3, though the details differ slightly. The first two of these components, N the completely depolarizing channel and D the completely dephasing channel, are easy to implement as mixed-unitary operations in the chosen circuit model. More difficult to implement is the channel M, which performs a function similar to the channel described by Equation (7.7).

The complete dephasing channel D is the channel that sets to zero all of the off-diagonal elements of a density matrix. More formally, the action of this operator applied to the space \mathcal{A} , for an input ρ on $\mathcal{A} \otimes \mathcal{H}$ is given by

$$D_{\mathcal{A}}(\rho) = \sum_{i=0}^{\dim \mathcal{A} - 1} p_i |i\rangle\langle i| \otimes \rho_i, \quad (7.21)$$

where the p_i form a probability distribution. This operation is equivalent to measuring the space \mathcal{A} in the computational basis and forgetting the result. This channel is shown to be mixed-unitary in Proposition 7.3 where it is implemented as a mixture of generalized Pauli Z operations. To implement this as a mixed-unitary circuit, observe that restricting the construction of Proposition 7.3 to the case where \mathcal{A} is a two dimensional space results in exactly the channel that applies a Pauli Z gate with probability one-half. Notice also that applying this channel to each of n qubits is identical to applying the completely dephasing channel to the whole space. Thus, the operation $D_{\mathcal{A}}$ that applies D to the space \mathcal{A} can be implemented as a mixed-unitary circuit by applying the Pauli Z operation to each qubit of \mathcal{A} independently with probability 1/2. This construction can be found in [CY97].

The completely noisy channel N is also simple to implement as a mixed-unitary circuit. This channel can be realized on a single qubit by performing a uniform mixture of the Pauli operators on each qubit, which is a consequence of Proposition 7.2 when restricted to the case of a single qubit. This mixture can be implemented by, independently on each qubit, applying the Pauli Z operation with probability $1/2$, followed by applying the Pauli X operation with probability $1/2$, as shown in [BR03]. Intuitively, the Z operations will zero the off-diagonal elements of a density matrix (viewed in the computational basis), and the X operations will scramble the diagonal, resulting in the completely mixed state, $\mathbb{1}/2$, on each qubit. As the tensor product of two completely mixed qubits is the completely mixed state of the larger system, applying this construction to each qubit in a space \mathcal{B} will implement the completely depolarizing channel $N_{\mathcal{B}}$ on that space.

In Section 7.3 the channel M was implemented as a completely depolarizing channel on the subspace S_0^\perp of inputs not in the state $|0\rangle$ on the ‘ancillary space’ \mathcal{A} . While the same channel suffices for the circuit case, it is not at all obvious how this channel can be implemented using only two-qubit mixed-unitary gates. This difficulty is avoided by implementing a closely related channel. This construction is intuitively the same: it does not affect states in the subspace S_0 of inputs with the $|0\rangle$ state in the space \mathcal{A} , and it applies depolarizing noise to states in the space S_0^\perp . The difference is exactly how this noise is applied. The circuit that is constructed implements the operation M defined by

$$M(|i\rangle\langle i| \otimes \rho) = \begin{cases} \frac{1}{\dim \mathcal{A}} (\mathbb{1}_{\mathcal{A}} - |0\rangle\langle 0| + |\psi_i\rangle\langle \psi_i|) \otimes \tilde{\mathbb{1}}_{\mathcal{H}} & \text{if } i \neq 0, \\ |0\rangle\langle 0| \otimes \rho & \text{if } i = 0, \end{cases} \quad (7.22)$$

where $|\psi_i\rangle$ is a nonzero computational basis state that depends on i . The exact specification of this state can be extracted from the analysis of the circuit constructed for M , but this is not helpful.

It is perhaps not a surprise that the transformation M can be implemented using only controlled-mixing operations. Before describing this implementation, notice that the controlled application of the completely depolarizing channel N to a single qubit can be described by a mixed-unitary circuit. This is because the previously discussed implementation of N is given by a mixture of the single qubit gates X and Z . Adding a control qubit to each of these gates results in two qubit gates, which fit into the model of mixed-unitary circuits used here (because we have assumed that X and Z , as well as controlled versions of them, are included in the standard basis of gates – dropping this assumption requires the circuit model to be generalized slightly). It is not clear that




Figure 7.2: One stage of the mixing procedure on the ancillary qubits. The mixing operations applied to the qubits in the space \mathcal{H} are not shown.

general controlled mixed-unitary operations can be implemented as mixed-unitary circuits in this model, but the only controlled operation that will be needed for this construction is the completely depolarizing channel.

At an intuitive level, the implementation of the channel M consists of the application of controlled-depolarizing operations everywhere that this is possible. These operations will all be controlled by the qubits in the space \mathcal{A} , which ensures that in the case of a state in S_0 , with $|0\rangle$ in the space \mathcal{A} , the operation M acts trivially.

More formally, let m be the number of qubits in the space \mathcal{A} that are given as part of the input to M , i.e. the number of ancillary qubits used to represent the ancillary space used by the original channel. The implementation of M consists of m stages, with the j th stage testing that the j th qubit of the space \mathcal{A} is in the $|0\rangle$ state, and mixing the qubits if this is not the case. An example of one stage of the circuit is given in Figure 7.2. The j th stage consists first of an application of the controlled N operation from the j th qubit to each other qubit of $\mathcal{A} \otimes \mathcal{H}$. After these operations, stage j is completed by $m - 1$ further controlled N operations: each with the j th qubit as the target qubit and one of the other qubits of \mathcal{A} as the control qubit. An example of this construction with $m = 3$ is presented in Figure 7.3.

Given these circuit implementations of the three channels $D_{\mathcal{A}}$, $N_{\mathcal{B}}$, M , the mixed-unitary circuit C that approximates a given circuit Q is constructed in exactly the same way as in Equation (7.12). More concretely, let Q be a circuit implementing the operation

$$Q(\rho) = \text{tr}_{\mathcal{B}} U(|0\rangle\langle 0| \otimes \rho) U^*,$$

where the ancillary qubits are in the space \mathcal{A} . The circuit C that approximates it is then given by

$$C(\rho) = N_{\mathcal{B}} (U [(M \circ D_{\mathcal{A}})(\rho)] U^*). \quad (7.23)$$

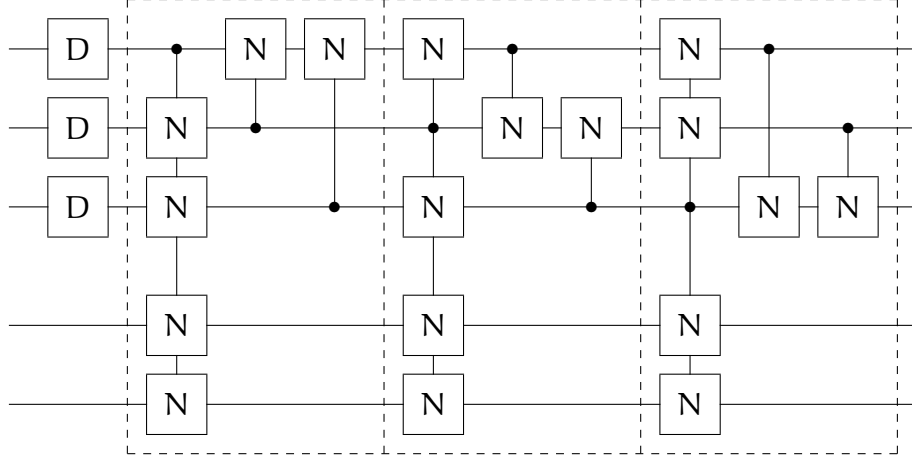


Figure 7.3: Circuit performing the ancilla simulation procedure $M \circ D_{\mathcal{A}}$. The top three qubits simulate the ancillary qubits of the original circuit in the space \mathcal{A} , and the bottom two simulate the input to the original circuit in the space \mathcal{H} . The dashed lines separate the each of the three stages of the mixing procedure.




Figure 7.4: The constructed mixed-unitary circuit C that simulates the given circuit Q , with input and output Hilbert spaces marked. The circuit U is the unitary from a the implementation of Q in Stinespring form. The circuits $D_{\mathcal{A}}$, M , and $N_{\mathcal{B}}$ are as described in the text.

This construction of the circuit C is shown in Figure 7.4. Notice that C is constructed to be a mixed-unitary circuit, as it the composition of smaller mixed-unitary circuits. Since the operations $D_{\mathcal{A}}$ and M do not affect inputs of the form $|0\rangle\langle 0| \otimes \rho$ in the space S_0 , the proof of Proposition 7.4 holds also for the circuit case, so that we have

$$C(|0\rangle\langle 0| \otimes \sigma) = Q(\sigma) \otimes \tilde{\mathbb{1}}_{\mathcal{B}}. \tag{7.24}$$

Combining this with equation (7.21) and the fact that applying $D_{\mathcal{A}}$ twice has no further effect, the output of C on an arbitrary input state ρ is of the form

$$C(\rho) = \sum_{i=0}^{\dim \mathcal{A}-1} p_i C(|i\rangle\langle i| \otimes \rho_i) = p_0 Q(\rho_0) \otimes \tilde{\mathbb{1}}_{\mathcal{B}} + \sum_{i=1}^{\dim \mathcal{A}-1} p_i C(|i\rangle\langle i| \otimes \rho_i). \tag{7.25}$$

In the remainder of the chapter it is shown that this construction does not significantly alter the distinguishability properties of quantum circuits.

As a first step towards this, it is shown that the above circuit construction correctly implements the channel M described by Equation 7.22. Much of the proof of this lemma is similar to the proof of Lemma 7.7, but the operation M considered in this section is slightly different and the proof must be extended to the case where there is an additional reference system.

This system, given by the space \mathcal{F} , is needed in the case of distinguishability, as a party attempting to distinguish two channels is permitted to use a portion of a larger entangled state as input to the channels. This is modelled by the use of the diamond norm in the definition of the computational problem QCD, from Section 5.2, the hardness of which will be extended to the mixed-unitary case.

Lemma 7.13. *On input states of the form $|k\rangle\langle k| \otimes \rho \in \mathbf{D}(\mathcal{A} \otimes \mathcal{H} \otimes \mathcal{F})$ for $|k\rangle\langle k| \in \mathbf{D}(\mathcal{A})$ with $0 < k \leq 2^m - 1$, the output of C satisfies*

$$\left\| (C \otimes \mathbf{1}_{\mathcal{F}})(|k\rangle\langle k| \otimes \rho) - \tilde{\mathbf{1}}_{\mathcal{A} \otimes \mathcal{H}} \otimes \text{tr}_{\mathcal{H}} \rho \right\|_{\text{tr}} \leq \frac{1}{2^{m-1}},$$

where m is the number of ancillary qubits used by the circuit Q .

Proof. On input of the form $|k\rangle\langle k| \otimes \rho$ the decoherence operations that are applied to the qubits in \mathcal{A} can be ignored, as they have no effect on qubits in a state of the computational basis. As $k \neq 0$ at least one qubit is in the state $|1\rangle$, and so the controlled-mixing operations in the implementation of the channel M will have an effect. Let the first nonzero qubit among the qubits of \mathcal{A} be the j th one. The first controlled N operation with nonzero control qubit that effects the j th qubit will be at the j th stage of the mixing process, where the j th qubit is the control qubit. As this qubit is not modified before this stage (as any previous qubits are in the state $|0\rangle$ by choice of j), the first $m - 1$ gates in the j th stage will mix the remaining qubits, so that the state after these gates is, using Equation (7.5),

$$|1\rangle\langle 1| \otimes \tilde{\mathbf{1}}_{\mathcal{A}'} \otimes \tilde{\mathbf{1}}_{\mathcal{H}} \otimes \text{tr}_{\mathcal{H}} \rho,$$

where for notational convenience the j th qubit has been written first, and \mathcal{A}' is the space of all but the j th qubit of \mathcal{A} . The remainder of the j th stage of the mixing process consists of $m - 1$ controlled N gates with the j th qubit as the target, each controlled by one of the $m - 1$ qubits in \mathcal{A}' . Considering the state $\mathbf{1}_{\mathcal{A}'}/2^{m-1}$ on \mathcal{A}' in the computational basis, the only term for which qubit j is not mixed by these operations is the all zero term. With this observation, the state after the j th stage is

$$\begin{aligned} & \frac{1}{2^{m-1}} \left[|1\rangle\langle 1| \otimes (|0\rangle\langle 0|)^{\otimes m-1} + \frac{|0\rangle\langle 0| + |1\rangle\langle 1|}{2} \otimes (\mathbf{1}_{\mathcal{A}'} - (|0\rangle\langle 0|)^{\otimes m-1}) \right] \otimes \tilde{\mathbf{1}}_{\mathcal{H}} \otimes \text{tr}_{\mathcal{H}} \rho \\ &= \frac{\mathbf{1}_{\mathcal{A}} + |1\rangle\langle 1| \otimes (|0\rangle\langle 0|)^{\otimes m-1} - (|0\rangle\langle 0|)^{\otimes m}}{2^m} \otimes \tilde{\mathbf{1}}_{\mathcal{H}} \otimes \text{tr}_{\mathcal{H}} \rho. \end{aligned}$$

This proves that the circuit implementing the channel M does so correctly, as this quantity is exactly the state given in Equation (7.25) with the addition of $\text{tr}_{\mathcal{H}} \rho$ in the reference system.

As in the proof of Lemma 7.7, let this state be σ . Computing the distance from this state to the desired one, we have

$$\|\sigma - \tilde{\mathbf{1}}_{\mathcal{A}} \otimes \tilde{\mathbf{1}}_{\mathcal{H}} \otimes \text{tr}_{\mathcal{H}} \rho\|_{\text{tr}} = \frac{1}{2^m} \left\| |1\rangle\langle 1| \otimes (|0\rangle\langle 0|)^{\otimes m-1} - (|0\rangle\langle 0|)^{\otimes m} \right\|_{\text{tr}} = \frac{1}{2^{m-1}}.$$

Finally, by noting that the remainder of the circuit C is mixed-unitary, Lemma 7.5 implies that the rest of the circuit cannot increase the norm. \square

In the next section this lemma is used to show that the hardness of the computational problem of distinguishing mixed-state circuits does not change when restricted to the mixed-unitary circuits.

7.8 QIP-completeness of distinguishing mixed-unitary circuits

The construction outlined in the previous section can be used to find mixed-unitary approximations to general quantum circuits, with the property that the diamond norm of the difference of two such circuits is approximately preserved. This property leads immediately to a proof that the problem of distinguishing mixed-unitary quantum circuits is **QIP**-complete, which is exactly as hard as the problem of distinguishing general quantum circuits. This will be done by taking the instance (Q_1, Q_2) of the general quantum circuit distinguishability problem (Problem 5.1), and constructing the instance (C_1, C_2) with C_1 and C_2 mixed-unitary, by applying the construction of Section 7.7 to each of these circuits.

This technique produces an instance of the mixed-unitary quantum circuit distinguishability problem, which is hereafter referred to as **MIXED-UNITARY QCD**. This problem is identical to **QCD** with the exception that the input circuits are required to be mixed-unitary circuits, in the model defined in Section 7.7. This problem is more formally defined as

Problem 7.14 (Mixed-unitary Quantum Circuit Distinguishability). For constants $0 \leq b < a \leq 2$, the input consists of mixed-unitary quantum circuits C_1 and C_2 that implement transformations in $\mathbf{T}(\mathcal{H}, \mathcal{K})$. The promise problem is to distinguish the two cases:

Yes: $\|C_1 - C_2\|_\diamond \geq a$,

No: $\|C_1 - C_2\|_\diamond \leq b$.

As this problem is a restriction of the more general circuit distinguishability problem, the protocol of Section 5.3 shows that it too is **QIP**. The remainder of this section is devoted to showing that MIXED-UNITARY QCD is **QIP**-complete for all $0 < b < a \leq 2$

The most important step in the proof that this restricted distinguishability problem is **QIP**-complete is to show that the construction in Section 7.7 does not significantly alter the diamond norm of the difference of the two circuits. This is the content of the following theorem.

Theorem 7.15. *Let Q_1 and Q_2 be arbitrary circuits implementing transformations in $\mathbf{T}(\mathcal{H}, \mathcal{K})$, and let C_i be the mixed-unitary circuit constructed from Q_i as in Equation (7.23). For any $\epsilon > 0$,*

$$\|Q_1 - Q_2\|_\diamond \leq \|C_1 - C_2\|_\diamond \leq \|Q_1 - Q_2\|_\diamond + \epsilon,$$

where the circuits C_1 and C_2 use $O(\log 1/\epsilon)$ extra qubits in the space \mathcal{A} .

Proof. The first inequality is not hard to show. Once again, the idea is sending the input state $(|0\rangle\langle 0|)^{\otimes m} \otimes \rho$ to the circuit C_i results in a simulation of Q_i , by Equation 7.24. This will imply that the distinguishability of Q_1 and Q_2 cannot be greater than the distinguishability of C_1 and C_2 . To formalize this argument, note that by the definition of the diamond norm

$$\|Q_1 - Q_2\|_\diamond = \sup_{\rho \in \mathbf{D}(\mathcal{H} \otimes \mathcal{F})} \|(Q_1 \otimes \mathbf{1}_{\mathcal{F}})(\rho) - (Q_2 \otimes \mathbf{1}_{\mathcal{F}})(\rho)\|_{\text{tr}},$$

and fix $\delta > 0$ and ρ as a state achieving a value within δ of this supremum. By Equation 7.24, if the state $(|0\rangle\langle 0|)^{\otimes m} \otimes \rho$ is given as input to the circuit C_i , then the output is $(Q_i \otimes \mathbf{1}_{\mathcal{F}})(\rho)$. Using this property we have

$$\begin{aligned} \|C_1 - C_2\|_\diamond &\geq \|(C_1 \otimes \mathbf{1}_{\mathcal{F}})((|0\rangle\langle 0|)^{\otimes m} \otimes \rho) - (C_2 \otimes \mathbf{1}_{\mathcal{F}})((|0\rangle\langle 0|)^{\otimes m} \otimes \rho)\|_{\text{tr}} \\ &= \|(Q_1 \otimes \mathbf{1}_{\mathcal{F}})(\rho) - (Q_2 \otimes \mathbf{1}_{\mathcal{F}})(\rho)\|_{\text{tr}} \\ &\geq \|Q_1 - Q_2\|_\diamond - \delta. \end{aligned}$$

Since this is true for any $\delta > 0$, it must be the case that $\|Q_1 - Q_2\|_\diamond \leq \|C_1 - C_2\|_\diamond$.

The second inequality requires somewhat more work. The idea is to once again break the input space into two subspaces: the one on which the circuits C_i simulate the circuits Q_i , and the orthogonal subspace. We will then use Lemma 7.13 to show that

on this orthogonal subspace the output states of the circuits C_i are almost completely mixed. This will in turn imply that the diamond norm on this input subspace is exponentially small, in the number of ancillary qubits added in the construction of the circuits C_i . An appeal to the decoherence operation applied as part of the circuit construction will validate the approach of treating the input state as a mixture of states from these two orthogonal subspaces.

More formally, let m be the number of ancillary qubits (the space \mathcal{A}) and let n be the number of input qubits (the space \mathcal{H}) used by the circuits Q_i . It can be assumed that the circuits Q_1 and Q_2 both use the same number of ancillary qubits by padding one of the circuits with unused qubits that are later traced out. The values of n and m may also be expressed as $m = \lceil \log \dim \mathcal{A} \rceil$ and $n = \lceil \log \dim \mathcal{H} \rceil$. By adding at most $3 + \log(1/\epsilon)$ extra ancillary qubits to the space \mathcal{A} , we may assume that

$$2^{-(m-3)} < \epsilon \quad (7.26)$$

Let $\rho \in \mathbf{D}(\mathcal{A} \otimes \mathcal{H} \otimes \mathcal{F})$ be a state such that

$$\|C_1 - C_2\|_{\diamond} - \epsilon/2 \leq \|(C_1 \otimes \mathbf{1}_{\mathcal{F}})(\rho) - (C_2 \otimes \mathbf{1}_{\mathcal{F}})(\rho)\|_{\text{tr}}, \quad (7.27)$$

and note that the reference system \mathcal{F} need not have the same dimension as the space of the same name considered in the proof of the previous inequality. The first gates applied in the circuit C_i are the decoherence gates applied to \mathcal{A} . These gates produce a state of the form $\sum_{i=0}^{2^m-1} p_i |i\rangle\langle i| \otimes \sigma_i$, for $\{p_i\}$ a probability distribution. Since applying these decoherence operations twice has no further effect, the output of the circuits C_1 and C_2 is the same on ρ as it is on this state. Applying this property and triangle inequality to Equation (7.27), the quantity of interest becomes

$$\|C_1 - C_2\|_{\diamond} - \epsilon/2 \leq \sum_{i=0}^{2^m-1} p_i \|(C_1 \otimes \mathbf{1}_{\mathcal{F}})(|i\rangle\langle i| \otimes \sigma_i) - (C_2 \otimes \mathbf{1}_{\mathcal{F}})(|i\rangle\langle i| \otimes \sigma_i)\|_{\text{tr}} \quad (7.28)$$

Then, by applying Lemma 7.13 to each term with $i \neq 0$, the states in the norm can be replaced with completely mixed states on $\mathcal{A} \otimes \mathcal{H}$ plus a small correction factor. Doing this for each of these terms we have

$$\begin{aligned} & p_i \|(C_1 \otimes \mathbf{1}_{\mathcal{F}})(|i\rangle\langle i| \otimes \sigma_i) - (C_2 \otimes \mathbf{1}_{\mathcal{F}})(|i\rangle\langle i| \otimes \sigma_i)\|_{\text{tr}} \\ & \leq p_i \left[\frac{2}{2^{m-1}} + \|\tilde{\mathbf{1}}_{\mathcal{A} \otimes \mathcal{H}} \otimes \text{tr}_{\mathcal{H}} \sigma_i - \tilde{\mathbf{1}}_{\mathcal{A} \otimes \mathcal{H}} \otimes \text{tr}_{\mathcal{H}} \sigma_i\|_{\text{tr}} \right] = p_i/2^{m-2} < p_i \epsilon/2. \end{aligned}$$

Applying this to Equation (7.28) results in

$$\|C_1 - C_2\|_{\diamond} - \epsilon/2 \leq p_0 \|(C_1 \otimes \mathbf{1}_{\mathcal{F}})(|0\rangle\langle 0| \otimes \sigma_0) - (C_2 \otimes \mathbf{1}_{\mathcal{F}})(|0\rangle\langle 0| \otimes \sigma_0)\|_{\text{tr}} + \sum_{i=1}^{2^m-1} p_i \epsilon/2.$$

By Equation 7.24 the output of the circuit C_i on this input can be replaced the output of the circuit Q_i and a maximally mixed state. When this is done to the previous equation, the desired bound is given by

$$\begin{aligned} \|C_1 - C_2\|_{\diamond} &\leq p_0 \left\| (Q_1 \otimes \mathbf{1}_{\mathcal{F}})(\sigma_0) \otimes \tilde{\mathbf{1}}_{\mathcal{B}} - (Q_2 \otimes \mathbf{1}_{\mathcal{F}})(\sigma_0) \otimes \tilde{\mathbf{1}}_{\mathcal{B}} \right\|_{\text{tr}} + (1 - p_0)\epsilon/2 + \epsilon/2 \\ &\leq p_0 \left\| (Q_1 \otimes \mathbf{1}_{\mathcal{F}})(\sigma_0) - (Q_2 \otimes \mathbf{1}_{\mathcal{F}})(\sigma_0) \right\|_{\text{tr}} + \epsilon \\ &\leq \|Q_1 - Q_2\|_{\diamond} + \epsilon. \end{aligned}$$

This completes the proof of the theorem, since $0 \leq p_0 \leq 1$. \square

The quantum circuit distinguishability problem is defined in terms of the diamond norm of the difference of two circuits. The bounds on this quantity provided by the previous theorem immediately imply the following corollary.

Corollary 7.16. *For any $0 < b < a \leq 2$ the problem MIXED-UNITARY QCD $_{a,b}$ is QIP-complete.*

Proof. Starting with an instance (Q_1, Q_2) of QCD $_{a,b/2}$ and applying the construction of Section 7.7 to each circuit results in a pair (C_1, C_2) . By Theorem 7.15 on positive instances of QCD we have

$$\|C_1 - C_2\|_{\diamond} \geq \|Q_1 - Q_2\|_{\diamond} \geq a,$$

and on negative instances we have

$$\|C_1 - C_2\|_{\diamond} \leq \|Q_1 - Q_2\|_{\diamond} + \epsilon \leq \frac{b}{2} + \epsilon.$$

By adding $O(\log 1/b)$ extra qubits in the construction of each of the circuits C_i , we can make $\epsilon < b/2$, so that the resulting pair (C_1, C_2) is an instance of the problem MIXED-UNITARY QCD $_{a,b}$, provided that $b > 0$.

To see that the reduction can be implemented efficiently, notice that the circuit construction can be done in polynomial time, since we have only added $O(\log 1/b)$ qubits, as well as a few operations that do not depend on the actual input circuits, just the number of qubits they act on. \square

A natural question to ask is whether this hardness result can be extended to the case of log-depth mixed-unitary circuits. With the construction of Section 7.7 this does not appear to be possible: the controlled-mixing operation in the procedure for mixing the input qubits if any ancillary qubits are not in the $|0\rangle$ state requires mixing operations to be applied to all other qubits, controlled by each of the qubits in the space \mathcal{A} . This

results in a linear depth circuit, since each input qubit is the target of at least $\log \dim \mathcal{A}$ mixing operations. Another approach might be to apply a construction similar to that used in Chapter 4 to a mixed-unitary CLOSE IMAGES problem. This approach runs into an immediate problem: all mixed-unitary circuits are unital, and so they must all output the completely mixed state when given one as input, i.e. the images of any two mixed-unitary transformations intersect at the completely mixed state, trivializing the close images problem on mixed-unitaries.

7.9 Conclusion

In this chapter a few problems are shown to be no easier when restricted to the class of mixed-unitary channels. This is done using a technique by which a channel is approximated using a mixed-unitary channel. While this approximation does not, in general, look very much like the original channel, for measures based on the behaviour of the channel on low-entropy outputs this approximation can be made arbitrarily good by padding the input channel with unused ancillary space. The approximation technique overcomes two main hurdles: viewed as a circuit, the original channel may trace out qubits and it may introduce fresh ancillary qubits in some pure state. The partial trace can be easily simulated by the mixed-unitary channel that maps any state to the completely mixed state, but a more complicated construction is required to deal with the ancillary space.

This construction is applied to the maximum output p -norm and the minimum output entropy, where it is used to show that the multiplicativity or additivity of a channel is implied by the additivity of multiplicativity of a related set of mixed-unitary channels. This can be used to show that the general multiplicativity and additivity problems are equivalent when restricted to the mixed-unitary channels, but this is no longer an interesting result, since mixed-unitary channels that are not additive [Has09] and not multiplicative [HW08] have recently been discovered.

When applied to the computational problem of distinguishing two transformations, this approximation scheme proves that this problem remains QIP-hard when restricted to mixed-unitary inputs. This is perhaps surprising: mixed-unitary channels have several nice properties [GW03], but computational distinguishability does not appear to be one of them. This can be seen as evidence that, despite the nice properties enjoyed by these channels, they may be sufficiently general to be a useful model of noise in quantum systems.

Chapter 8

Conclusion

This thesis has introduced the problem **QUANTUM CIRCUIT DISTINGUISHABILITY**, which is a computational version of channel distinguishability. This problem has been shown to be hard for the class **QIP** of problems that have quantum interactive proof systems, which is equal to the more familiar class **PSPACE** [JJUW09]. This problem gives a new quantum characterization of this class that is not closely tied to quantum interactive proof systems.

The hardness of the distinguishability problem leads naturally to a study of the problem on restricted classes of channels. This can be seen as an attempt to isolate those instances of the problem that are hard, so that more is known about those instances on which the problem is tractable. This thesis has presented four important classes of channels for which the problem remains hard: the channels implemented by log-depth circuits, the degradable channels, the antidegradable channels, and the mixed-unitary channels. These special cases demonstrate that the QCD problem is hard on a wide array of classes of channels, i.e. that the hardness of the problem is very likely not tied to just a few hard instances.

These hardness results are shown by reducing the general problem to a version of the problem restricted to a special class of channels. These reductions can have applications outside of complexity theory, as they are essentially methods for the simulation of general channels by channels in a restricted class. These simulation techniques can have powerful implications throughout quantum information. One example of such an application is the result that the additivity of the Holevo capacity or the multiplicativity of the maximum output p -norm can be (approximately) restricted to a mixed-unitary channel. It is hoped that the other reductions presented in the thesis will find similar applications.

Several natural questions are left open by this thesis. Several of the more interesting of these questions are summarized below.

- There are other models of computation that involve unitary computations applied to mixed initial states (see [ASV06, KL98]). How hard is the distinguishability problem for computations in these models?
- What is the complexity of distinguishing constant depth quantum circuits that do not use the unbounded fan-out gate? This problem is hard on constant depth circuits that have access to this gate, but the reduction used in the thesis does not produce constant-depth circuits without access to this gate.
- How hard is QCD on the entanglement-breaking channels? It was argued in Section 1.3 that this problem is hard for channels that are exponentially close together, but this problem is not very interesting with such a weak promise. It is not known how hard the distinguishability problem is on this subset of the antidegradable channels with a stronger promise.
- The random Pauli channels, also known as the Pauli diagonal channels, can be expressed as convex combinations of the channels that apply the discrete Weyl (or generalized Pauli) operators. These channels are an important subclass of the mixed-unitary channels. Many of the pieces used in the reduction to mixed-unitary channels can be expressed as channels of this form, but there is one major problem: the unitary U from a Stinespring representation of a general channel needs to be converted to such a channel. Such a simulation result would imply that the QCD problem is also hard on this class.

Bibliography

- [ABO08] D. Aharonov and M. Ben-Or. Fault-tolerant quantum computation with constant error rate. *SIAM Journal on Computing*, **38**(4):1207–1282, 2008. doi: 10.1137/S0097539799359385. EPRINT arXiv:quant-ph/9906129.
- [AG04] S. Aaronson and D. Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, **70**(5):052328, 2004. doi: 10.1103/PhysRevA.70.052328. EPRINT arXiv:quant-ph/0406196.
- [AH03] G. G. Amosov and A. S. Holevo. On the multiplicativity hypothesis for quantum communication channels. *Theory of Probability and its Applications*, **47**(1):123–127, 2003. doi: 10.1137/S0040585X97979500.
- [AHW00] G. G. Amosov, A. S. Holevo, and R. F. Werner. On some additivity problems in quantum information theory. *Problems of Information Transmission*, **36**(4):305–313, 2000. EPRINT arXiv:math-ph/0003002.
- [AJL06] D. Aharonov, V. Jones, and Z. Landau. A polynomial quantum algorithm for approximating the Jones polynomial. In *Proceedings of the 38th ACM Symposium on the Theory of Computing*, pp. 427–436, 2006. doi: 10.1145/1132516.1132579. EPRINT arXiv:quant-ph/0511096.
- [AKN98] 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.
- [AMTdW00] A. Ambainis, M. Mosca, A. Tapp, and R. de Wolf. Private quantum channels. In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science*, pp. 547–553, 2000. doi: 10.1109/SFCS.2000.892142. EPRINT arXiv:quant-ph/0003101.

- [AS08] K. M. R. Audenaert and S. Scheel. On random unitary channels. *New Journal of Physics*, **10**:023011, 2008. doi: 10.1088/1367-2630/10/2/023011. EPRINT arXiv:0709.0824 [quant-ph].
- [ASV06] A. Ambainis, L. J. Schulman, and U. Vazirani. Computing with highly mixed states. *Journal of the ACM*, **53**(3):507–531, 2006. doi: 10.1145/1147954.1147962. EPRINT arXiv:quant-ph/0003136.
- [BATS09] A. Ben-Aroya and A. Ta-Shma. On the complexity of approximating the diamond norm, 2009. EPRINT arXiv:0902.3397 [quant-ph].
- [BBD⁺97] A. Barenco, A. Berthiaume, D. Deutsch, A. Ekert, R. Jozsa, and C. Macchiavello. Stabilization of quantum computations by symmetrization. *SIAM Journal on Computing*, **26**(5):1541–1557, 1997. doi: 10.1137/S0097539796302452. EPRINT arXiv:quant-ph/9604028.
- [BCWdW01] 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.
- [BDS97] C. H. Bennett, D.P. DiVincenzo, and J. A. Smolin. Capacities of quantum erasure channels. *Physical Review Letters*, **78**(16):3217–3220, 1997. doi: 10.1103/PhysRevLett.78.3217. EPRINT arXiv:quant-ph/9701015.
- [BFS97] C. H. Bennett, C. A. Fuchs, and J. A. Smolin. Entanglement-enhanced classical communication on a noisy quantum channel. In O. Hirota, A. S. Holevo, and C. M. Caves, editors, *Quantum Communication, Computing, and Measurement: Proceedings of the Third International Conference on Quantum Communication and Measurement, 1996*, pp. 79–88. Plenum Press, New York, 1997. EPRINT arXiv:quant-ph/9611006.
- [Bha97] R. Bhatia. *Matrix Analysis, Graduate Texts in Mathematics*, volume 169. Springer, 1997.
- [BK09] A. Broadbent and E. Kashefi. Parallelizing quantum circuits. *Theoretical Computer Science*, pp. 2489–2510, 2009. doi: 10.1016/j.tcs.2008.12.046. EPRINT arXiv:0704.1736 [quant-ph].
- [BMP⁺00] P. O. Boykin, T. Mor, M. Pulver, V. Roychowdhury, and F. Vatan. A new universal and fault-tolerant quantum basis. *Information Processing Letters*, **75**(3):101–107, 2000. doi: 10.1016/S0020-0190(00)00084-3. EPRINT arXiv:quant-ph/9906054.

- [BR03] 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.
- [Bus06] F. Buscemi. On the minimum number of unitaries needed to describe a random-unitary channel. *Physics Letters A*, **360**(2):256–258, 2006. doi: 10.1016/j.physleta.2006.08.038. EPRINT arXiv:quant-ph/0607034.
- [BV97] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, **26**(5):1411–1473, 1997. doi: 10.1137/S0097539796300921.
- [BŻ06] I. Bengtsson and K. Życzkowski. *Geometry of Quantum States: An Introduction to Quantum Entanglement*. Cambridge University Press, 2006.
- [Cho75] 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.
- [CN97] I. L. Chuang and M. A. Nielsen. Prescription for experimental determination of the dynamics of a quantum black box. *Journal of Modern Optics*, **44**(11-12):2455–2467, 1997. doi: 10.1080/09500349708231894. EPRINT arXiv:quant-ph/9610001.
- [Con90] John B. Conway. *A Course in Functional Analysis, Graduate Texts in Mathematics*, volume 96. Springer, 1990.
- [Coo71] S. A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the 3rd ACM Symposium on the Theory of Computing*, pp. 151–158, 1971. doi: 10.1145/800157.805047.
- [CPR00] A. M. Childs, J. Preskill, and J. Renes. Quantum information and precision measurement. *Journal of Modern Optics*, **47**(2-3):155–176, 2000. doi: 10.1080/09500340008244034. EPRINT arXiv:quant-ph/9904021.
- [CRS08] T. S. Cubitt, M. Beth 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].
- [CW00] R. Cleve and J. Watrous. Fast parallel circuits for the quantum fourier transform. In *Proceedings of the 41st ACM Symposium on the Theory of Computing*, pp. 526–536, 2000. doi: 10.1109/SFCS.2000.892140. EPRINT arXiv:quant-ph/0006004.

- [CY97] I. L. Chuang and Y. Yamamoto. Creation of a persistent quantum bit using error correction. *Physical Review A*, **55**:114–127, 1997. doi: 10.1103/PhysRevA.55.114. EPRINT arXiv:quant-ph/9604030.
- [DFH06] N. Datta, M. Fukuda, and A. S. Holevo. Complementarity and additivity for covariant channels. *Quantum Information Processing*, **5**(3):179–207, 2006. doi: 10.1007/s11128-006-0021-6.
- [DS05] 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.
- [EAO⁺02] A. K. Ekert, C. M. Alves, D. K. Oi, M. Horodecki, P. Horodecki, and L. C. Kwak. 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.
- [Fan51] K. Fan. Maximum properties and inequalities for the eigenvalues of completely continuous operators. *Proceedings of the National Academy of Sciences*, **37**(11):760–766, 1951. URL <http://www.pnas.org/content/37/11/760.full.pdf>.
- [Fan73] M. Fannes. A continuity property of the entropy density for spin lattice systems. *Communications in Mathematical Physics*, **31**(4):291–294, 1973. doi: 10.1007/BF01646490.
- [FGHZ05] S. Fenner, F. Green, S. Homer, and Y. Zhang. Bounds on the power of constant-depth quantum circuits. In *Proceedings of the 15th International Symposium on Fundamentals of Computation Theory*, pp. 44–55, 2005. doi: 10.1007/11537311_5. EPRINT arXiv:quant-ph/0312209.
- [FKW02] M. H. Freedman, A. Kitaev, and Z. Wang. Simulation of topological field theories by quantum computers. *Communications in Mathematical Physics*, **227**(3), 2002. doi: 10.1007/s002200200635. EPRINT arXiv:quant-ph/0001071.
- [FLW02] M. H. Freedman, M. Larsen, and Z. Wang. A modular functor which is universal for quantum computation. *Communications in Mathematical Physics*, **227**(3):605–622, 2002. doi: 10.1007/s002200200645. EPRINT arXiv:quant-ph/0001108.

- [Fuk07] M. Fukuda. Simplification of additivity conjecture in quantum information theory. *Quantum Information Processing*, **6**(3):179–186, 2007. DOI: 10.1007/s11128-007-0051-8. EPRINT arXiv:quant-ph/0608010.
- [FvdG99] C. A. Fuchs and J. van de Graaf. Cryptographic distinguishability measures for quantum-mechanical states. *IEEE Transactions on Information Theory*, **45**(4):1216–1227, 1999. DOI: 10.1109/18.761271. EPRINT arXiv:quant-ph/9712042.
- [FW07] 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].
- [GB02] L. Gurvits and H. Barnum. Largest separable balls around the maximally mixed bipartite quantum state. *Physical Review A*, **66**(6):062311, 2002. DOI: 10.1103/PhysRevA.66.062311. EPRINT arXiv:quant-ph/0204159.
- [GC99] D. Gottesman and I. L. Chuang. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature*, **402**(6760):390–393, 1999. DOI: 10.1038/46503. EPRINT arXiv:quant-ph/9908010.
- [GF05] 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.
- [GHMP02] F. Green, S. Homer, C. Moore, and C. Pollett. Counting, fanout, and the complexity of quantum ACC. *Quantum Information and Computation*, **2**(1):35–65, 2002. EPRINT arXiv:quant-ph/0106017.
- [GLMS03] V. Giovannetti, S. Lloyd, L. Maccone, and P. W. Shor. Entanglement assisted capacity of the broadband lossy channel. *Physical Review Letters*, **91**(4):047901, 2003. DOI: 10.1103/PhysRevLett.91.047901. EPRINT arXiv:quant-ph/0304020.
- [GLN05] A. Gilchrist, N. K. Langford, and M. A. Nielsen. Distance measures to compare real and ideal quantum processes. *Physical Review A*, **71**(6):062310, 2005. DOI: 10.1103/PhysRevA.71.062310. EPRINT arXiv:quant-ph/0408063.

- [GW03] M. Gregoratti and R. F. Werner. Quantum lost and found. *Journal of Modern Optics*, **50**(6):915–933, 2003. doi: 10.1080/09500340308234541. EPRINT arXiv:quant-ph/0209025.
- [Has09] 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].
- [Hel67] C. W. Helstrom. Detection theory and quantum mechanics. *Information and Control*, **10**(3):254–291, 1967. doi: 10.1016/S0019-9958(67)90302-6.
- [HJ85] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.
- [HJ91] R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1991.
- [HK70] K.-E. Hellwig and K. Kraus. Operations and measurements. II. *Communications in Mathematical Physics*, **16**(2):142–147, 1970. doi: 10.1007/BF01646620.
- [HLSW04] P. Hayden, D. Leung, P. W. Shor, and A. Winter. Randomizing quantum states: constructions and applications. *Communications in Mathematical Physics*, **250**:371–391, 2004. doi: 10.1007/s00220-004-1087-6. EPRINT arXiv:quant-ph/0307104.
- [Hol98] A. S. Holevo. The capacity of the quantum channel with general signal states. *IEEE Transactions on Information Theory*, **44**(1):269–273, 1998. doi: 10.1109/18.651037. EPRINT arXiv:quant-ph/9611023.
- [Hol06] A. S. Holevo. The additivity problem in quantum information theory. In *Proceedings of the International Congress of Mathematicians*, volume 3, pp. 1000–1017, 2006. URL http://www.icm2006.org/proceedings/Vol_III/contents/ICM_Vol_3_49.pdf.
- [Hol07] 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.
- [HŠ05] P. Høyer and R. Špalek. Quantum circuits with unbounded fan-out. *Theory of Computing*, **1**:81–103, 2005. doi: 10.4086/toc.2005.v001a005. EPRINT arXiv:quant-ph/0208043.

- [HSR03] M. Horodecki, P. W. Shor, and M. B. Ruskai. Entanglement breaking channels. *Reviews in Mathematical Physics*, **15**(6):629–641, 2003. DOI: 10.1142/S0129055X03001709. EPRINT arXiv:quant-ph/0302031.
- [HW08] P. Hayden and A. Winter. Counterexamples to the maximal p -norm multiplicativity conjecture for all $p > 1$. *Communications in Mathematical Physics*, **284**(1):263–280, 2008. DOI: 10.1007/s00220-008-0624-0. EPRINT arXiv:0807.4753 [quant-ph].
- [JJUW09] Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. QIP = PSPACE, 2009. EPRINT arXiv:0907.4737 [quant-ph].
- [JKP09] N. Johnston, D. W. Kribs, and V. I. Paulsen. Computing stabilized norms for quantum operations via the theory of completely bounded maps. *Quantum Information and Computation*, **9**(1&2):16–35, 2009. EPRINT arXiv:0711.3636 [quant-ph].
- [Joz94] R. Jozsa. Fidelity for mixed quantum states. *Journal of Modern Optics*, **41**(12):2315–2323, 1994. DOI: 10.1080/09500349414552171.
- [Joz06] R. Jozsa. An introduction to measurement based quantum computation. In Dimitris G. Angelakis, Matthias Christandl, Artur Ekert, Alastair Kay, and Sergei Kilik, editors, *Quantum Information Processing – From Theory to Experiment*, pp. 137–158. IOS Press, 2006. EPRINT arXiv:quant-ph/0508124.
- [JUW09] R. Jain, S. Upadhyay, and J. Watrous. Two-message quantum interactive proofs are in PSPACE. In *Proceedings of the 50th IEEE Symposium on Foundations of Computer Science*, 2009. EPRINT arXiv:0905.1300 [cs.CC].
- [JWB05] 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.
- [Kin02] C. King. Additivity for unital qubit channels. *Journal of Mathematical Physics*, **43**(10):4641–4653, 2002. DOI: 10.1063/1.1500791. EPRINT arXiv:quant-ph/0103156.
- [Kin03] C. King. Maximal p -norms of entanglement breaking channels. *Quantum Information and Computation*, **3**(2):186–190, 2003. EPRINT arXiv:quant-ph/0212057.

- [Kit97] A. Y. Kitaev. Quantum computations: algorithms and error correction. *Russian Mathematical Surveys*, **52**(6):1191–1249, 1997. DOI: 10.1070/RM1997v052n06ABEH002155.
- [Kit99] A. Y. Kitaev. Quantum NP. Talk at the 2nd Workshop on Algorithms in Quantum Information Processing (AQIP), DePaul University, 1999.
- [KKR06] 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.
- [KL98] E. Knill and R. Laflamme. Power of one bit of quantum information. *Physical Review Letters*, **81**(25):5672–5675, 1998. DOI: 10.1103/PhysRevLett.81.5672. EPRINT arXiv:quant-ph/9802037.
- [Kle31] O. Klein. Zur Quantenmechanischen Begründung des zweiten Hauptsatzes der Wärmelehre. *Zeitschrift für Physik*, **72**(11-12):767–775, 1931. DOI: 10.1007/BF01341997.
- [KLZ98] E. Knill, R. Laflamme, and W. H. Zurek. Resilient quantum computation: error models and thresholds. *Proceedings of the Royal Society A*, **454**(1969):365–384, 1998. DOI: 10.1098/rspa.1998.0166. EPRINT arXiv:quant-ph/9702058.
- [KM87] B. Kümmerer and H. Maassen. The essentially commutative dilations of dynamical semigroups on M_n . *Communications in Mathematical Physics*, **109**(1):1–22, 1987. DOI: 10.1007/BF01205670.
- [KMN07] 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.
- [Kni96] E. Knill. Quantum randomness and nondeterminism. Technical Report LAUR-96-2186, Los Alamos National Laboratory, 1996. EPRINT arXiv:quant-ph/9610012.
- [KNY08] M. Kada, H. Nishimura, and T. Yamakami. The efficiency of quantum identity testing of multiple states. *Journal of Physics A: Mathematical and General*, **41**(39):395309, 2008. DOI: 10.1088/1751-8113/41/39/395309. EPRINT arXiv:0809.2037 [quant-ph].

- [KR01] C. King and M. B. Ruskai. Minimal entropy of states emerging from noisy quantum channels. *IEEE Transactions on Information Theory*, **47**(1):192–209, 2001. DOI: 10.1109/18.904522. EPRINT arXiv:quant-ph/9911079.
- [KSV02] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation, Graduate Studies in Mathematics*, volume 47. American Mathematical Society, 2002.
- [KW00] A. Kitaev and J. Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof systems. In *Proceedings of the 32nd ACM Symposium on the Theory of Computing*, pp. 608–617, 2000. DOI: 10.1145/335305.335387.
- [LFKN92] C. Lund, L. Fortnow, H. J. Karloff, and N. Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM*, **39**(4):859–868, 1992. DOI: 10.1145/146585.146605.
- [LS93] L. J. Landau and R. F. Streater. On Birkhoff’s Theorem for doubly stochastic completely positive maps of matrix algebras. *Linear Algebra and its Applications*, **193**:107–127, 1993. DOI: 10.1016/0024-3795(93)90274-R.
- [MN02] C. Moore and M. Nilsson. Parallel quantum computation and quantum codes. *SIAM Journal on Computing*, **31**(3):799–815, 2002. DOI: 10.1137/S0097539799355053. EPRINT arXiv:quant-ph/9808027.
- [MW05] 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.
- [MW09] C. B. Mendl and M. M. Wolf. Unital quantum channels – convex structure and revivals of Birkhoff’s Theorem. *Communications in Mathematical Physics*, **289**(3):1057–1086, 2009. DOI: 10.1007/s00220-009-0824-2. EPRINT arXiv:0806.2820 [quant-ph].
- [NC00] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [Neu43] M. A. Neumark. On a representation of additive operator set functions. *Comptes rendus (Doklady) de l’Académie des sciences de l’URSS*, **41**:359–361, 1943.

- [Pau02] V. Paulsen. *Completely Bounded Maps and Operator Algebras*, Cambridge Studies in Advanced Mathematics, volume 78. Cambridge University Press, 2002.
- [PCZ97] J. F. Poyatos, J. I. Cirac, and P. Zoller. Complete characterization of a quantum process: the two-bit quantum gate. *Physical Review Letters*, **78**(2):390–393, 1997. DOI: 10.1103/PhysRevLett.78.390. EPRINT arXiv:quant-ph/9611013.
- [Rom05] Steven Roman. *Advanced Linear Algebra, Graduate Texts in Mathematics*, volume 135. Springer, 2005.
- [Ros08a] 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].
- [Ros08b] 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/>.
- [RW05] 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.
- [Sac05a] M. F. Sacchi. Entanglement can enhance the distinguishability of entanglement-breaking channels. *Physical Review A*, **72**(1):014305, 2005. DOI: 10.1103/PhysRevA.72.014305. EPRINT arXiv:quant-ph/0505174.
- [Sac05b] M. F. Sacchi. Optimal discrimination of quantum operations. *Physical Review A*, **71**(6):062340, 2005. DOI: 10.1103/PhysRevA.71.062340. EPRINT arXiv:quant-ph/0505183.
- [Sch60] R. Schatten. *Norm Ideals of Completely Continuous Operators*. Springer-Verlag, 1960.
- [Sch96] B. Schumacher. Sending entanglement through noisy quantum channels. *Physical Review A*, **54**(4):2614–2628, 1996. DOI: 10.1103/PhysRevA.54.2614. EPRINT arXiv:quant-ph/9604023.

- [Sha48] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, **27**:379–423, 623–656, 1948. URL <http://plan9.bell-labs.com/cm/ms/what/shannonday/paper.html>.
- [Sha92] A. Shamir. $IP = PSPACE$. *Journal of the ACM*, **39**(4):869–877, 1992. DOI: 10.1145/146585.146609.
- [Sho96] P. W. Shor. Fault-tolerant quantum computation. In *Proceedings of the 37th IEEE Symposium on Foundations of Computer Science*, pp. 56 – 65, 1996. DOI: 10.1109/SFCS.1996.548464. EPRINT arXiv:quant-ph/9605011.
- [Sho02] P. W. Shor. Additivity of the classical capacity of entanglement-breaking quantum channels. *Journal of Mathematical Physics*, **43**(9):4334–4340, 2002. DOI: 10.1063/1.1498000. EPRINT arXiv:quant-ph/0201149.
- [Sho04] P. W. Shor. Equivalence of additivity questions in quantum information theory. *Communications in Mathematical Physics*, **246**(3):453–472, 2004. DOI: 10.1007/s00220-003-0981-7. EPRINT arXiv:quant-ph/0305035.
- [Smi83] R. R. Smith. Completely bounded maps between C^* -algebras. *Journal of the London Mathematical Society*, **s2-27**(1):157, 1983. DOI: 10.1112/jlms/s2-27.1.157.
- [Sti55] W. F. Stinespring. Positive functions on C^* -algebras. *Proceedings of the American Mathematical Society*, **6**(2):211–216, 1955. DOI: 10.2307/2032342.
- [SV03] A. Sahai and S. Vadhan. A complete problem for statistical zero knowledge. *Journal of the ACM*, **50**(2):196–249, 2003. DOI: 10.1145/636865.636868. EPRINT Cryptology ePrint Archive: Report 2000/056.
- [SW97] B. Schumacher and M. D. Westmoreland. Sending classical information via noisy quantum channels. *Physical Review A*, **56**(1):131–138, 1997. DOI: 10.1103/PhysRevA.56.131.
- [TD04] B. M. Terhal and D. P. DiVincenzo. Adaptive quantum computation, constant depth quantum circuits and Arthur-Merlin games. *Quantum Information and Computation*, **4**(2):134–145, 2004. EPRINT arXiv:quant-ph/0205133.
- [Tre86] S. L. Tregub. Doubly stochastic operators in finite-dimensional von Neumann algebra. *Soviet Mathematics (Iz. VUZ)*, **30**(3):105–108, 1986.

- [Uhl76] A. Uhlmann. The “transition probability” in the state space of a $*$ -algebra. *Reports on Mathematical Physics*, **9**(2):273–279, 1976. doi: 10.1016/0034-4877(76)90060-4.
- [vN27] J. von Neumann. Thermodynamik quantenmechanischer Gesamtheiten. *Göttinger Nachrichten*, **1**:273–291, 1927.
- [Vya03] M. Vyalı. QMA=PP implies that PP contains PH. Technical Report 21, Electronic Colloquium on Computational Complexity, 2003. EPRINT ECCC TR03-021.
- [Wat00] 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.
- [Wat02] J. Watrous. Limits on the power of quantum statistical zero-knowledge. In *Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science*, pp. 459 – 468, 2002. doi: 10.1109/SFCS.2002.1181970. EPRINT arXiv:quant-ph/0202111.
- [Wat03] J. Watrous. PSPACE has constant-round quantum interactive proof systems. *Theoretical Computer Science*, **292**(3):575–588, 2003. doi: 10.1016/S0304-3975(01)00375-9. EPRINT arXiv:cs/9901015.
- [Wat05] J. Watrous. Notes on super-operator norms induced by Schatten norms. *Quantum Information and Computation*, **5**(1):58–68, 2005. EPRINT arXiv:quant-ph/0411077.
- [Wat08] J. Watrous. Distinguishing quantum operations having few Kraus operators. *Quantum Information and Computation*, **8**(8&9):819–833, 2008. EPRINT arXiv:0710.0902 [quant-ph].
- [Wat09a] J. Watrous. Mixing doubly stochastic quantum channels with the completely depolarizing channel. *Quantum Information and Computation*, **9**(5&6):406–413, 2009. EPRINT arXiv:0807.2668 [quant-ph].
- [Wat09b] J. Watrous. Quantum computational complexity. In *Encyclopedia of Complexity and System Science*. Springer, 2009. EPRINT arXiv:0804.3401 [quant-ph].
- [Wat09c] J. Watrous. Semidefinite programs for completely bounded norms, 2009. EPRINT arXiv:0901.4709 [quant-ph].

- [WH02] R. F. Werner and A. S. Holevo. Counterexample to an additivity conjecture for output purity of quantum channels. *Journal of Mathematical Physics*, **43**(9):4353–4357, 2002. DOI: 10.1063/1.1498491. EPRINT arXiv:quant-ph/0203003.
- [WPG07] 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.
- [Yao93] A. C. Yao. Quantum circuit complexity. In *Proceedings of the 34th IEEE Symposium on Foundations of Computer Science*, volume 34, pp. 352–361, 1993. DOI: 10.1109/SFCS.1993.366852.

Index

- H, 11
- X, 11
- Y, 11
- Z, 11
- $\mathbf{D}(\mathcal{H})$, 15
- $\mathbb{1}_{\mathcal{H}}$, 11
- $L(\mathcal{H}, \mathcal{K})$, 10
- $\tilde{\mathbb{1}}_{\mathcal{H}}$, 15
- $I_{\mathcal{H}}$, 16

- antidegradable channel, 118
- antisymmetric subspace, 84

- BQP**, 36

- channel, 16
 - antidegradable, 118
 - degradable, 118
 - mixed-unitary, 130
 - unital, 22
- circuit
 - depth, 27
 - size, 26
 - Stinespring form, 32
- Close Images
 - constant-depth, 79
 - log-depth, 79
 - problem definition, 79
- completely positive, 16
- computational basis, 8
- constant-depth controlled operations, 35

- degradable channel, 118

- density operator, 15
- diamond norm, 59
 - condition for maximization on pure states, 62
 - direct product lemma, 70
 - multiplicativity, 60
 - polarization, 75
 - relationship to distinguishability, 61
 - stabilization, 59
- discrete Weyl operators, 12

- entanglement, 9
- entropy, 42
 - additivity, 43
 - concavity, 44
 - Klein's inequality, 43
 - minimum output entropy, 45

- fidelity, 64
 - maximum output fidelity, 68
 - monotonicity, 65
 - multiplicativity, 64
 - relation to trace norm, 66
 - Uhlmann's Theorem, 65
- Fuchs-van de Graaf Inequalities, 67

- Helstrom Measurement, 55
- Hermitian, 10
- Hilbert space, 7

- inner product, 7

- Klein's inequality, *see* entropy

- Kraus operators, 17
- log-depth controlled operations, 33
- minimum error state distinguishability, *see* Helstrom Measurement
- minimum output entropy, *see* entropy
- mixed state, 15
- mixed-unitary
 - channel, 130
 - circuits, 150
 - distinguishability, 157
- norm, 7
- normal operator, 13
- p-norm, 45
 - maximum output p-norm, 47
 - multiplicativity on states, 47
 - unitary invariance, 46
- partial trace, 16
- Pauli operators, 11
 - generalized, 12
- positive, 10
- POVM, 16
- pure state, 9
- purification, 17

- QIP**, 37
- QMA**, 36
- quantum channel, *see* channel
- Quantum Circuit Distinguishability
 - QIP** protocol, 104
 - antidegradable, 126
 - degradable, 122
 - mixed-unitary, 157
 - problem definition, 103
- qubit, 9

- Schatten p-norm, *see* p-norm
- Schmidt decomposition, 9
- separable state, 9
- singular value decomposition, 14
- spectral decomposition, 13
- Stinespring representation, 18
- support, 13
- swap operation, 84
- swap test, 84
- symmetric subspace, 84

- trace, 11
- trace norm, 53
 - Helstrom measurement, 55
 - monotonicity, 54
 - of channels, 57
 - relation to fidelity, 66
- Uhlmann’s Theorem, 65
- unital, 22
 - reduction multiplicativity of v_p , 133
 - reduction to additivity of S_{\min} , 133
- unitary, 10
- W operation, *see* swap operation
- Weyl operators, *see* discrete Weyl operators