



Distinguishing Short Quantum Computations

Bill Rosgen

► **To cite this version:**

Bill Rosgen. Distinguishing Short Quantum Computations. Susanne Albers, Pascal Weil. STACS 2008, Feb 2008, Bordeaux, France. IBFI Schloss Dagstuhl, pp.597-608, 2008. <hal-00255825>

HAL Id: hal-00255825

<https://hal.archives-ouvertes.fr/hal-00255825>

Submitted on 14 Feb 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

DISTINGUISHING SHORT QUANTUM COMPUTATIONS

BILL ROSGEN ¹

¹ Institute for Quantum Computing and School of Computer Science, University of Waterloo
E-mail address: wrosgen@iqc.ca

ABSTRACT. Distinguishing logarithmic depth quantum circuits on mixed states is shown to be complete for QIP, the class of problems having quantum interactive proof systems. Circuits in this model can represent arbitrary quantum processes, and thus this result has implications for the verification of implementations of quantum algorithms. The distinguishability problem is also complete for QIP on constant depth circuits containing the unbounded fan-out gate. These results are shown by reducing a QIP-complete problem to a logarithmic depth version of itself using a parallelization technique.

1. Introduction

Much of the difficulty in implementing quantum algorithms in practice is that qubits quickly decohere upon interacting with the environment. This entanglement destroying process limits the length of the computations that can be realized by experiment. Implementing quantum algorithms as circuits of low depth can provide a way to perform as much computation as possible within the limited time available, and for this reason there is significant interest in finding short quantum circuits for important problems.

Log-depth quantum circuits have been found for several significant problems including the approximate quantum Fourier transform [3] and the encoding and decoding operations for many quantum error correcting codes [10]. In addition to these applications, a procedure for parallelizing to log-depth a large class of quantum circuits has recently been discovered [2]. These examples demonstrate the surprising power of short quantum circuits.

Much of the work on quantum circuits is done in the standard model of unitary quantum circuits on pure states. In this paper a slightly different model of computation is considered: the model of mixed state quantum circuits, introduced by Aharonov, Kitaev, and Nisan [1]. While much of the previous complexity-theoretic work on short quantum circuits has been in the unitary model [4, 6], there has also been work outside of this model [13]. There are several advantages to considering the more general model of mixed state circuits. The primary advantage is that the mixed state model is able to capture any process allowed by quantum mechanics, so that results on this model may have implications for experimental work in quantum computing. The problem of distinguishing circuits may thus be thought of as the problem of distinguishing potentially noisy physical processes. As an example,

Key words and phrases: quantum information, computational complexity, quantum circuits, quantum interactive proof systems.

finding an error in an implementation of a quantum algorithm is simply the problem of distinguishing the constructed circuit from one that is known to be correct.

Unfortunately, in this paper it is shown that the apparent power of short quantum computations comes with a price: logarithmic depth quantum circuits are *exactly* as difficult to distinguish as polynomial depth quantum circuits. This equivalence implies the surprising result that distinguishing log-depth quantum computations is complete for the class QIP, the set of all problems that have quantum interactive proof systems. As $\text{PSPACE} \subseteq \text{QIP} \subseteq \text{EXP}$ [8], this result also implies that the problem is PSPACE-hard.

The result on circuit distinguishability is shown using the closely related problem of determining if two circuits can be made to output states that are close together. This problem was introduced by Kitaev and Watrous [8] who show it to be both complete for QIP and contained in EXP. The main result of the present paper is obtained by reducing an instance of this problem of polynomial depth to an equivalent instance of logarithmic depth. This demonstrates that the problem of close images remains complete for QIP even under a logarithmic depth restriction. The hardness of distinguishing short quantum circuits is then demonstrated by a modification to the argument in [12] to show that the equivalence of close images problem and the distinguishability holds even for log-depth circuits.

The remainder of this paper is organized as follows. In the next section, some of the notation and results that will be needed are summarized. This is followed by Section 3, where the complete problems for QIP are discussed. In Section 4 the reduction from the polynomial depth to logarithmic depth versions of the close images problem is given, and the correctness of this construction is shown in Section 5. The equivalence between the log-depth close images problem and the problem of distinguishing log-depth computations is discussed in Section 6.

2. Preliminaries

This section outlines some of the definitions and results that will be used throughout the paper. For a more thorough treatment of the concepts introduced here see [9] and [11].

Throughout the paper scripted letters such as \mathcal{H} will refer to finite dimensional Hilbert spaces, $\mathbf{D}(\mathcal{H})$ will denote the set of all density matrices on \mathcal{H} , and $\mathbf{U}(\mathcal{H}, \mathcal{K})$ will denote the norm-preserving linear operators from \mathcal{H} to \mathcal{K} . The proof of the main result will make extensive use of two notions of distance between quantum states. The first of these is the fidelity. The *fidelity* between two positive semidefinite operators X and Y on a space \mathcal{H} can be defined as

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

This definition is known as Uhlmann's Theorem, and it is used here as it is more directly applicable to the task at hand than the usual definition. As any purification of a state necessarily purifies the partial trace of that state, this equation implies that the fidelity is nondecreasing under the partial trace. This property is known as *monotonicity* and can be stated more formally as $F(X, Y) \leq F(\text{tr}_{\mathcal{K}} X, \text{tr}_{\mathcal{K}} Y)$ where X, Y are positive semidefinite operators on $\mathcal{H} \otimes \mathcal{K}$. The final property of the fidelity that will be needed is the result that the maximum fidelity of any outputs of two transformations is multiplicative with respect to the tensor product. This result can be found in [9] (see Problem 11.10 and apply the multiplicativity of the diamond norm with respect to the tensor product).

Theorem 2.1 (Kitaev, Shen, and Vyalıi [9]). *For any completely positive transformations $\Phi_1, \Phi_2, \Psi_1, \Psi_2$ on states in \mathcal{H}*

$$\max_{\rho, \xi \in \mathbf{D}(\mathcal{H} \otimes \mathcal{H})} F((\Phi_1 \otimes \Phi_2)(\rho), (\Psi_1 \otimes \Psi_2)(\xi)) = \prod_{i=1,2} \max_{\rho, \xi \in \mathbf{D}(\mathcal{H})} F(\Phi_i(\rho), \Psi_i(\xi))$$

The second notion of distance that will be used is the *trace norm*, which can be defined for any linear operator X by $\|X\|_{\text{tr}} = \text{tr} \sqrt{X^* X}$, or equivalently as the sum of the singular values of X . This quantity is a norm, and so in particular it satisfies the triangle inequality. Similar to the fidelity, the trace norm is monotone under the partial trace, though in this case the trace norm is non-increasing under this operation. The proofs that follow will make essential use of the Fuchs-van de Graaf Inequalities [5] that relate the trace norm and the fidelity. For any density operators ρ and ξ on the same space, these inequalities are

$$1 - F(\rho, \xi) \leq \frac{1}{2} \|\rho - \xi\|_{\text{tr}} \leq \sqrt{1 - F(\rho, \xi)}.$$

In addition to these measures on quantum states, it will be helpful to have a distance measure on quantum transformations. One such measure is the *diamond norm*, which for a completely positive transformation Φ on density operators on \mathcal{H} is given by

$$\|\Phi\|_{\diamond} = \sup_{\rho \in \mathbf{D}(\mathcal{H} \otimes \mathcal{H})} \|(\Phi \otimes I_{\mathbf{L}(\mathcal{H})})(\rho)\|_{\text{tr}}.$$

This norm is essential when considering transformations as it represents the distinguishability of two transformations when a reference system is taken into account. The simple supremum of the trace norm over all inputs to the channel is not stable under the addition of a reference system, and so the diamond norm is used in place of the simpler one. More properties and a more thorough definition of this norm can be found in [9].

The circuit model that will be used in this paper is the mixed state model introduced by Aharonov, Kitaev, and Nisan [1]. Circuits in this model are composed of qubits that are acted upon by arbitrary trace preserving and completely positive operations. This model allows for non-unitary operations, such as measurement or the introduction of ancillary qubits, to occur in the middle of the circuit. It is important to note that this model captures any physical process that quantum mechanics allows, and so in particular, any computation that can be done on mixed states with measurements can be represented in this model. Fortunately this model is polynomially equivalent to the standard model of unitary quantum circuits (with ancilla) followed by measurement, as shown in [1]. This will allow us to consider only circuits composed of unitary gates from some finite basis of one and two qubit gates with the additional operations of introducing qubits in the $|0\rangle$ state and measuring in the computational basis. This restriction can be strengthened, again with no loss of generality, to assume that all ancillary qubits are introduced at the start of the circuit and that all measurements are performed at the end.

We will often add to this circuit model one additional gate: the unbounded fan-out gate. This gate, in constant depth, applies a controlled-not operation from one qubit to an arbitrary number of output qubits. It is not clear that this gate is a reasonable choice in a standard basis of gates, and so it will be clearly marked when this gate is allowed into the circuit model under consideration. As an example of the power of this gate it can be used to build a constant depth circuit for the approximate quantum Fourier transform [7]. This gate is considered here for the sole reason that if it is included in the standard set of gates, the main result will also hold for constant depth circuits.




Figure 1: A circuit implementing the swap test.

For spaces \mathcal{H} and \mathcal{K} of the same dimension, we use $W \in \mathbf{U}(\mathcal{H} \otimes \mathcal{K}, \mathcal{H} \otimes \mathcal{K})$ to represent the operation that swaps the states in the two spaces. As W is a permutation matrix when expressed in the computational basis, and the permutation that it encodes is composed exclusively of transpositions, the swap operation is both hermitian and unitary. Furthermore, W can easily be implemented in constant depth, as all of the required transpositions can be performed in parallel. This operator is the essential component of the *swap test*, where a controlled W operation is used to determine how close two states are to each other. A circuit performing the swap test is given in Figure 1, where the measurement is performed in the computational basis. Another way to view the swap test is as a projective measurement onto the symmetric and antisymmetric subspaces. The projections in this measurement are given by $(I + W)/2$ and $(I - W)/2$. This formulation of the swap test is equivalent to the circuit presented in Figure 1.

It is not immediately clear how a controlled operation on n qubits, such as the controlled-swap operation used in the swap test, can be performed in depth logarithmic in n . The straightforward implementation requires using one control qubit to control each of the gates in the operation. However, Moore and Nilsson [10] give a simple construction that allows such an operation to be performed in log-depth.

Proposition 2.2 (Moore and Nilsson). *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 constant depth case, but the method of proof that they use immediately extends to the log depth case. They prove this proposition by using a tree of $\log n$ controlled-not operations to ‘duplicate’ the control qubit. These copies can then be used to control the remaining operations, with each control qubit used at most a logarithmic number of times. This proposition, as an example, implies that the swap test circuit on n qubits shown in Figure 1 can be implemented in depth $O(\log n)$.

If the fan-out gate is allowed into the standard basis of gates, then controlled operations can be performed with only constant depth overhead. A circuit that performs this can be obtained by simply using one fan-out gate to make n copies (in the computational basis) of the control qubit onto ancillary qubits. These ‘copies’ may then be used to control each of the n operations, with a final application of the fan-out gate to restore the ancillary qubits to the $|0\rangle$ state. As controlled operations will be the only place that the circuits constructed here exceed constant depth, this will allow the proof of the main result for constant depth circuits with fan-out.

3. Complete Problems for QIP

The **Close Images** problem, defined and shown to be complete for QIP in [8] can be stated as follows.

Close Images. *For constants $0 < b < a \leq 1$, the input consists of quantum circuits Q_1 and Q_2 that implement transformations from \mathcal{H} to \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 the name of this problem as $\mathbf{Cl}_{a,b}$.

A closely related problem is that of distinguishing two quantum circuits. This problem was introduced and shown complete for QIP in [12].

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 from \mathcal{H} to \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? This problem will be referred to as $\mathbf{QCD}_{a,b}$.

It is our aim to prove that these problems remain 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. $a = 1, 2$ in the above problem definitions. Both of these problem remain complete for QIP in this case. This restriction serves only to make these problems easier, as distinguishing the two cases for a weaker promise can only be more difficult, so the results of this paper will also imply the hardness of the more general problems. The log-depth versions of these problems will be referred to as **Log-depth $\mathbf{Cl}_{1,b}$** and **Log-depth $\mathbf{QCD}_{2,b}$** , and since these are restrictions of QIP-complete problems it is clear that they are also in QIP. Similarly, the abbreviations **Const-depth $\mathbf{Cl}_{1,b}$** and **Const-depth $\mathbf{QCD}_{2,b}$** for the versions of these problems on constant-depth circuits will be convenient.

4. Log-Depth Construction

In this section the reduction from the general $\mathbf{Cl}_{1,b}$ problem to the log-depth restriction of the problem is described. The general idea behind the construction is to simply slice the circuits of an instance of $\mathbf{Cl}_{1,b}$ into logarithmic-depth pieces and run them in parallel. These circuits will require more input, but if each piece of the circuit is given as input the same state output by the previous piece, then the output of the last piece of the circuit will be equal to the output of the original circuit. This may 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.

To describe the reduction, let Q_1 and Q_2 be the circuits from an instance of $\mathbf{Cl}_{1,b}$, and let n be the size (number of gates) of Q_1 and Q_2 (by padding the smaller circuit, if necessary). In order to perform the slicing of the circuit into pieces it is assumed that Q_1




Figure 2: The original circuits Q_1 and Q_2 decomposed into constant depth unitary circuits.

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 input. This restriction can be made with no loss in generality, as any quantum circuit, even one that incorporates measurements and other non-unitary operations, can be approximated by such a circuit, and furthermore, this circuit uses a number of gates that is a polynomial in the size of the original circuit [1].

A simple way to decompose Q_1 into constant depth pieces is to simply let each gate of Q_1 be a 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 . These pieces are shown in Figure 2. If Q_1 and Q_2 implement transformations from \mathcal{H} to \mathcal{K} , using ancillary qubits that fit into \mathcal{A} , and trace out the qubits in \mathcal{B} , then the spaces $\mathcal{H} \otimes \mathcal{A}$ and $\mathcal{B} \otimes \mathcal{K}$ are isomorphic, since by assumption Q_1 and Q_2 first introduce any needed ancilla and only trace qubits out at the end of the computation. Using these spaces, and implicitly this isomorphism, we have

$$U_1, V_1 \in \mathbf{U}(\mathcal{H}_1, \mathcal{B}_1 \otimes \mathcal{K}_1)$$

$$U_i, V_i \in \mathbf{U}(\mathcal{H}_i \otimes \mathcal{A}_i, \mathcal{B}_i \otimes \mathcal{K}_i) \quad \text{for } 2 \leq i \leq n,$$

where the subscripted spaces are copies of the non-subscripted spaces that hold the 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 Q_1 on ρ is given by

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

and the output of Q_2 is given by the same expression using the V_i operators.

Using this decomposition of Q_1 and Q_2 , circuits C_1 and C_2 are constructed that are logarithmic in depth and still in some sense faithfully implement Q_1 and Q_2 . This is done by running 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 is not necessarily close to the output from U_{i-1} . This problem can be dealt with by comparing the output of U_{i-1} to the input to U_i . In order to do this in logarithmic depth an auxiliary input that is first compared against the input to U_i and then held in reserve to compare to the output of U_{i-1} is needed. To compare these quantum states the swap test can be used. This test will fail with some probability depending on the distance between the two states. An example of the construction used to ensure that the output of U_{i-1} agrees with the input to U_i is given in Figure 3. To simplify the analysis of the constructed circuits these tests are controlled so that either one or the other is performed. This will affect the failure probability by a




Figure 3: 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.

factor of at most two, but will allow the analysis of each swap test to ignore the effect of the other. To implement this 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. These controlled operations can be implemented in logarithmic depth using the technique of Moore and Nilsson [10].

After adding these tests between each piece of the circuit there is one final modification required. If any of the swap tests fail, i.e. detect states that are not the same, then they will output qubits in the $|1\rangle$ state. As yes instances of $\mathbf{Cl}_{1,b}$ have outputs that are close together, we can ensure that no outputs of the constructed circuits can be close if any swap tests fail by adding dummy qubits in the $|0\rangle$ state to be compared to the outputs of the swap tests in the other circuit. These dummy qubits are shown in Figure 4.

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 3, and adding dummy qubits to ensure that the swap tests in the other circuit do not fail. 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. If the outputs of C_1 and C_2 are close together, then intuitively 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 intuitive picture is accurate forms the content of the next section.

In the other direction, it is not hard to see 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 . To do this, notice that the circuit construction does not change if additional qubits are added to the circuits to allow purification of the states ρ and ξ to be used as inputs to C_1 and C_2 . These additional qubits are traced out with the other qubits at the end of the circuit, so that the output state of the circuit are not changed. As these purifications are pure states and all operations performed during the circuit are unitary, the intermediate states of the




Figure 4: The outputs of C_1 and C_2 .

circuits must also be pure states. If the input state to C_1 is $|\psi\rangle$, then by providing the state

$$|\psi\rangle \otimes U_1|\psi\rangle \otimes \cdots \otimes U_{n-1}U_{n-2} \cdots U_1|\psi\rangle$$

as input to C_1 , the output of each block of the circuit will be identical to the input to the next block, ensuring that all the swap tests will succeed with probability one. It remains only to check on such input states that C_1 produces the same output as Q_1 on ρ . This can be observed by noting that the output of the circuit is exactly

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

which by construction is equal to the output of Q_1 on ρ . Thus if the circuits Q_1 and Q_2 have intersecting images then so do the circuits C_1 and C_2 . This observation proves the completeness of the construction. Soundness is considerably more intricate, and is the focus of the next section.

5. Soundness of the Construction

In this section it is demonstrated 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 . As the constructed circuits essentially simulate Q_1 and Q_2 the desired result can be obtained by arguing that either the outputs of C_1 and C_2 are far apart or the input to at least one of the constructed circuits is not a faithful simulation of the corresponding original circuit. In the case that this simulation is not faithful it will be shown that there is some swap test that fails with reasonable probability. This implies that outputs of the constructed circuits must also be distant, 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.

As a first step, we place a lower bound on the failure probability of a swap test in terms of the fidelity of the two states being compared. In the following lemma the swap test is viewed as a measurement of the symmetric and antisymmetric projectors, with the outcome that produces a qubit in the state $|1\rangle$ corresponding to the antisymmetric case.

Lemma 5.1. *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 of sufficient dimension to allow such a purification. The swap test measures the state on $\mathcal{A} \otimes \mathcal{B}$ with the projectors $\frac{1}{2}(I - W)$ and $\frac{1}{2}(I + W)$, where W is the swap operator on $\mathcal{A} \otimes \mathcal{B}$. Thus, the antisymmetric outcome occurs with probability

$$\frac{1}{4} \text{tr}([(I - W) \otimes I]|\psi\rangle\langle\psi|[(I - W^*) \otimes I]) = \frac{1}{2}\langle\psi|I \otimes I - W \otimes I|\psi\rangle = \frac{1}{2} - \frac{1}{2}\langle\psi|W \otimes I|\psi\rangle,$$

as W is hermitian. Then as W is also unitary, the states $|\psi\rangle$ and $W|\psi\rangle$ 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

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

After tracing out the space \mathcal{C} , this is exactly the statement of the lemma. ■

This lemma cannot be immediately applied to the circuits C_1 and C_2 , as 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 to translate the fidelity to a relation involving the trace norm, which we can then apply the triangle inequality to. This application of the triangle inequality shows that at least one of the two swap tests fails with probability bounded below by an expression involving the fidelity. In 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 it is not assumed that these states are given in a separable form. 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.

Corollary 5.2. *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*

$$\frac{1}{64} \|U_i \rho_{i-1} U_i^* - \rho_i\|_{\text{tr}}^2.$$

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 . 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 in both swap tests.

By Lemma 5.1, the failure probability of 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, this fidelity may be replaced by the trace norm. Doing so, we obtain

$$p \geq \frac{1}{16} \max(\|\sigma_i - \xi_i\|_{\text{tr}}^2, \|\rho_i - \sigma_i\|_{\text{tr}}^2).$$

Finally, as this maximum must be at least the average of the two values,

$$p \geq \frac{1}{16} \left(\frac{\|\sigma_i - \xi_i\|_{\text{tr}}}{2} + \frac{\|\rho_i - \sigma_i\|_{\text{tr}}}{2} \right)^2 \geq \frac{1}{64} \|\rho_i - \xi_i\|_{\text{tr}}^2,$$

where the last inequality follows from an application of the triangle inequality. ■

By repeatedly applying some of the properties of the trace norm discussed in Section 2 it is somewhat tedious but not difficult to reduce the problem at hand to the previous Corollary. This is the content of the following theorem.

Theorem 5.3. *If $F(Q_1(\rho_0), Q_2(\xi_0)) < 1 - c$ for all $\rho_0, \xi_0 \in \mathcal{H}$ then*

$$F(C_1(\rho), C_2(\xi)) < 1 - \frac{c^2}{144n^2}$$

for all $\rho, \xi \in (\mathcal{H} \otimes \mathcal{A})^{\otimes 2n}$.

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 $0 \leq i \leq 2n$, where the states for $i > n$ are the inputs that are only used by the swap tests, which we will not need to refer to explicitly. That is, ρ_i and ξ_i for $0 \leq i \leq n$ 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 . The output of the circuits C_1 and C_2 is then given by a number of 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.

By the condition on the fidelity of Q_1 and Q_2 and the Fuchs-van de Graaf inequalities, we have $2c < \|Q_1(\rho_0) - Q_2(\xi_0)\|_{\text{tr}}$. Using the triangle inequality we can relate this to the distance between the constructed circuits. Adding terms and simplifying, we obtain

$$\begin{aligned} 2c &< \|Q_1(\rho_0) - \text{tr}_{\mathcal{B}_n} \rho_n + \text{tr}_{\mathcal{B}_n} \xi_n - Q_2(\xi_0) + \text{tr}_{\mathcal{B}_n} \rho_n - \text{tr}_{\mathcal{B}_n} \xi_n\|_{\text{tr}} \\ &\leq \|Q_1(\rho_0) - \text{tr}_{\mathcal{B}_n} \rho_n\|_{\text{tr}} + \|\text{tr}_{\mathcal{B}_n} \xi_n - Q_2(\xi_0)\|_{\text{tr}} + \|\text{tr}_{\mathcal{B}_n} \rho_n - \text{tr}_{\mathcal{B}_n} \xi_n\|_{\text{tr}}. \end{aligned}$$

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 the partial trace, since the former can be obtained from the later by tracing out the appropriate spaces. Using this we have

$$2c < \|Q_1(\rho_0) - \text{tr}_{\mathcal{B}_n} \rho_n\|_{\text{tr}} + \|\text{tr}_{\mathcal{B}_n} \xi_n - Q_2(\xi_0)\|_{\text{tr}} + \|C_1(\rho) - C_2(\xi)\|_{\text{tr}} \tag{5.1}$$

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 $F(C_1(\rho), C_2(\xi)) < 1 - c^2/144$ and there is nothing left to prove.

The cases where one of the first two terms of (5.1) exceeds $2c/3$ are symmetric, and so we can consider only the first term. Expanding $Q_1(\rho_0)$ in terms of the U_i , we obtain

$$\begin{aligned} \frac{2c}{3} &< \|Q_1(\rho_0) - \text{tr}_{\mathcal{B}_n} \rho_n\|_{\text{tr}} \\ &= \|\text{tr}_{\mathcal{B}_n} U_n U_{n-1} \cdots U_1 \rho_0 U_1^* U_2^* \cdots U_n^* - \text{tr}_{\mathcal{B}_n} \rho_n\|_{\text{tr}} \\ &\leq \|U_n U_{n-1} \cdots U_1 \rho_0 U_1^* U_2^* \cdots U_n^* - \rho_n\|_{\text{tr}}, \end{aligned}$$

where once again the monotonicity of the trace norm under the partial trace has been used. By repeating the strategy of adding terms and then applying the triangle inequality we have

$$\frac{2c}{3} < \|U_1 \rho_0 U_1^* - \rho_1\|_{\text{tr}} + \|U_n U_{n-1} \cdots U_2 \rho_1 U_2^* U_3^* \cdots U_n^* - \rho_n\|_{\text{tr}}.$$

Here we have made use of the unitary invariance of the trace norm to discard the operators U_2, \dots, U_n from the first term. Continuing in this fashion we have

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

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 $\|U_i \rho_{i-1} U_i^* - \rho_i\|_{\text{tr}} > 2c/(3n)$, and so by Corollary 5.2 one of the corresponding swap tests fails with probability $p > c^2/(144n^2)$. 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,

$$F(C_1(\rho), C_2(\xi)) \leq F((1-p)|0\rangle\langle 0| + p|1\rangle\langle 1|, |0\rangle\langle 0|) = 1-p < 1 - \frac{c^2}{144n^2},$$

as in the statement of the theorem. ■

By combining Theorem 5.3 with the observation in Section 4 and the multiplicativity of the maximum output fidelity of two transformations, we obtain the following result.

Corollary 5.4. *The problem Log-depth $\text{Cl}_{1,b}$ is QIP-complete for any constant $0 < b < 1$.*

Proof. Theorem 5.3 establishes the completeness of the problem for any $b \geq 1 - c^2/(144n^2)$, where n is an upper bound on the size of the circuits. Using Theorem 2.1 of Kitaev, Shen, and Valyi [9] we can repeat each of the circuits r times in parallel to obtain the completeness of the problem for $b \geq (1 - c^2/(144n^2))^r$, which can be made smaller than any constant for r some polynomial in n . ■

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, the following Corollary follows immediately from the previous one.

Corollary 5.5. *The problem Const-depth $\text{Cl}_{1,b}$ on circuits with the unbounded fan-out gate is QIP-complete for any constant $0 < b < 1$.*

6. Distinguishing Log-Depth Computations

The hardness of Log-depth $\text{Cl}_{1,b}$ can be extended to Log-depth $\text{QCD}_{2,b}$ by observing that the reduction for the polynomial depth version of this problem in [12] can be made to preserve the depth of the constructed circuits. Once this observation is made, the hardness of the log-depth (and constant-depth with fan-out) versions of the circuit distinguishability problem is immediate.

The reduction in [12] takes as input circuits (Q_1, Q_2) and produces circuits C_1 and C_2 . Without describing the reduction in detail, the constructed circuits C_1 and C_2 run, depending on the value of a control qubit, one of Q_1 and Q_2 , followed by a constant depth circuit. If the input circuits Q_1 and Q_2 have logarithmic depth, then the only significant difficulty is the fact that controlled versions of these circuits are needed. However, as we have already seen, if we replace the gates in Q_1 and Q_2 with controlled versions, then we can use the scheme of Moore and Nilsson [10] to implement the controlled operations in

logarithmic depth. With this modification, the reduction in [12] can be reused to show the hardness of the QCD problem on log-depth circuits.

Corollary 6.1. Log-depth $\text{QCD}_{2,b}$ is QIP-complete for any constant $0 < b < 2$.

Once again these controlled operations can be implemented in a constant depth circuit if the unbounded fan-out gate is allowed into the set of allowed gates.

Corollary 6.2. Const-depth $\text{QCD}_{2,b}$ on circuits with the unbounded fan-out gate is QIP-complete for any constant $0 < b < 2$.

7. Conclusion

The hardness of distinguishing even log-depth mixed state quantum circuits leaves several related open problems, a few of which are listed here.

- Can this new complete problem be used to further understand QIP?
- Does this result rely in an essential way on the mixed state circuit model? How difficult is it to distinguish quantum circuits in less general models of computation?
- What is the complexity of distinguishing constant depth quantum circuits that do not use the unbounded fan-out gate?

Acknowledgement. I would like to thank John Watrous for several helpful discussions, the anonymous reviewers for constructive comments, as well as Canada's NSERC and MITACS for supporting this research.

References

- [1] D. Aharonov, A. Kitaev, N. Nisan. Quantum circuits with mixed states. In *Proc. 30th ACM Symposium on the Theory of Computing*, pp. 20–30, 1998.
- [2] A. Broadbent, E. Kashefi. Parallelizing quantum circuits. arXiv:0704.1736v1 [quant-ph].
- [3] R. Cleve, J. Watrous. Fast parallel circuits for the quantum fourier transform. In *Proc. 41st ACM Symposium on the Theory of Computing*, pp. 526–536, 2000.
- [4] S. Fenner, F. Green, S. Homer, Y. Zhang. Bounds on the power of constant-depth quantum circuits. In *Proc. 15th International Symposium on Fundamentals of Computation Theory*, pp. 44–55, 2005.
- [5] C. A. Fuchs, J. van de Graaf. Cryptographic distinguishability measures for quantum-mechanical states. *IEEE Transactions on Information Theory*, 45(4):1216–1227, 1999.
- [6] F. Green, S. Homer, C. Moore, C. Pollett. Counting, fanout and the complexity of quantum ACC. *Quantum Information and Computation*, 2(1):35–65, 2002.
- [7] P. Høyer, R. Špalek. Quantum circuits with unbounded fan-out. *Theory of Computing*, 1:81–103, 2005.
- [8] A. Kitaev, J. Watrous. Parallelization, amplification and exponential time simulation of quantum interactive proof systems. In *Proc. 32nd ACM Symp. Theory of Computing*, pp. 608–617, 2000.
- [9] A. Y. Kitaev, A. H. Shen, M. N. Vyalyi. *Classical and Quantum Computation*, volume 47 of *Graduate Studies in Mathematics*. American Mathematical Society, 2002.
- [10] C. Moore, M. Nilsson. Parallel quantum computation and quantum codes. *SIAM Journal on Computing*, 31(3):799–815, 2002.
- [11] M. A. Nielsen, I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [12] B. Rosgen, J. Watrous. On the hardness of distinguishing mixed-state quantum computations. In *Proc. 20th Conference on Computational Complexity*, pp. 344–354, 2005.
- [13] B. M. Terhal, D. P. DiVincenzo. Adaptive quantum computation, constant depth quantum circuits and Arthur-Merlin games. *Quantum Information and Computation*, 4(2):134–145, 2004.