QUANTUM INTERACTIVE PROOFS AND THE COMPLEXITY OF ENTANGLEMENT DETECTION

KEVIN MILNER
SCHOOL OF COMPUTER SCIENCE
MCGILL UNIVERSITY
MONTREAL

A thesis submitted to McGill University in partial fulfilment of the requirements of the degree of M.Sc.

August 2013

Abstract

This thesis identifies a formal connection between physical problems related to entanglement detection and complexity classes in theoretical computer science. In particular, we prove that to nearly every quantum interactive proof complexity class (including BQP, QMA, QMA(2), QSZK, and QIP), there corresponds a natural entanglement or correlation detection problem that is complete for that class. In this sense, we can say that an entanglement or correlation detection problem captures the expressive power of each quantum interactive proof complexity class, and the contrast between such problems gives intuition to the differences between classes of quantum interactive proofs. It is shown that the difficulty of entanglement detection also depends on whether the distance measure used is the trace distance or the one-way LOCC distance. We also provide analysis for another problem of this flavour, which we show is decidable by a two-message quantum interactive proof system while being hard for both NP and QSZK, the first nontrivial example of such a problem.

Sommaire

Ce mémoire met en évidence un lien formel entre les problèmes physiques de détection d'intrication et les classes de complexité de l'informatique théorique. Plus particulièrement, nous établissons une correspondance entre la plupart des classes de complexité naturelles issues de preuves interactives quantiques (incluant BQP, QMA, QMA(2), QSZK, et QIP), et une intrication ou un problème de détection de corrélation qui est complet pour cette classe. En ce sens, nous pouvons dire que l'intrication, ou le problème de détection de corrélation, capture la puissance expressive de chaque classe de complexité de preuve interactive quantique et que le contraste entre de tels problèmes donne une idée sur les différences entre les classes de preuves interactives quantiques. Il est démontré que la difficulté de la détection d'intrication varie considérablement du fait que la mesure de distance utilisée soit la distance de trace ou LOCC unidirectionnel. Nous fournissons également l'analyse d'un problème similaire, et montrons que celui-ci est décidable par un système de preuve interactive quantique (de deux messages) tout en étant NP-dur ainsi que QSZK-dur, le premier exemple non trivial d'un tel problème.

Acknowledgements

Foremost, I would like to thank my supervisor, Patrick Hayden, for his instruction, patience, guidance, and vast expertise. It is because of his wisdom that I have the opportunity to write on such a fascinating topic.

I would also like to thank Mark Wilde for his dedication and motivation, as well as his unwavering support in the face of my continuous and occasionally foolish inquiries. He has been my co-supervisor in all but name, and I have learned an incredible amount from him. I cannot imagine a greater resource for a student to have.

Further, my thanks to others that aided with thesis: Gus Gutoski for the fruitful discussions that contributed significantly to the extent of this thesis, Claude Crépeau for his examination and helpful corrections, and Teri Drummond for her dedicated proofreading and constructive comments.

It has been an honour and a pleasure to work with so many brilliant and accomplished people throughout my research.

More personally, I feel it is of great importance to acknowledge others that have had significant impact on the course of my academic career: Jim Hoover, for his encouragement and mentorship from the very beginning of my computing science education; Bob Beck, for his boisterous stories and advice; Lorna Stewart, who introduced me to so many aspects of theoretical computer science; and the late Piotr Rudnicki, who represents to me some of the greatest moments and challenges in my education to date.

Finally, I would like to recognize the many friends and family – near and far – who have supported me directly or indirectly over the last two years, the names of whom I cannot hope to contain within this page. They are all greatly appreciated, likely more than they would suspect. In particular though, I would like to thank my parents, without whom this thesis most certainly would not have been possible for a great many reasons, and who continue to fortify me through my many endeavours.

Preface

This thesis represents the combination of three manuscripts created in the course of research, and as such, much of the following text has been, or will be, submitted for publication in one or more of those manuscripts [HMW13, MGHW13]. Nonetheless, this thesis is presented in its current form because it is hoped that the combination of these manuscripts into a single coherent document will be of aid to those approaching the topic in the future. In light of this, however, it is especially important to recognize the contributions of the collaborators which were instrumental in the research for this thesis, and whose names have appeared – or will appear – on the manuscripts comprising this work.

Mark Wilde has contributed all figures within this document with the exception of the class hierarchy diagram, as well as some text (most notably a majority of the QSEP-CIRCUIT and QSEP-CHANNEL sections, along with the discussion of geometric measures of entanglement) and significant time aiding in editing. A large portion of the work represented within these pages is due to our discussions.

Patrick Hayden sparked many of the ideas that led to this work, and most importantly his intuition during our discussions led to the basic element used within the proofs of hardness for many of these problems as well as the protocol to place $\mathsf{QPROD\text{-}CIRCUIT}_{\mathrm{Tr}}$ in QSZK. He also aided in the editing of this document.

Gus Gutoski prompted the writing of our third manuscript, and it is because of this and our helpful discussions that we considered further additions to our work, which now appear in this thesis.

Finally, the author contributed the majority of analysis and text contained in this thesis, as well as significant time and effort towards the research which this document represents.

Contents

List of Figures

1

Introduction

Certain families of decision problems have proven to be particularly versatile and expressive in complexity theory, in the sense that slightly varying their formulation can tune the difficulty of the problems through a wide range of complexity classes. Adding quantifiers to the problem of evaluating a Boolean formula, for example, brings the venerable satisfiability problem up through the levels of the polynomial hierarchy [Sto76] all the way up to PSPACE [Sip96], at each level providing a decision problem complete for the associated complexity class. Moreover, adding limitations to the format of the Boolean satisfiability problem gives decision problems complete for a variety of more limited classes.[1] Likewise, in the domain of interactive proofs [Bab85, GMR85, BM88, GMR89, Wat03, KW00, Wat09a], problems based on distinguishing probability distributions or quantum states, depending on the setting, arise very naturally.

In the domain of quantum information theory, quantum mechanical entanglement is responsible for many of the most surprising and, not coincidentally, useful potential applications of quantum information [HHHH09], including quantum teleportation [BBC+93], super-dense coding [BW92], enhanced communication capacities [BSST99, BSST02, CLMW10], device-independent quantum key distribution [Eke91, VV12, BB84], and communication complexity [CB97]. Thus, deciding whether the correlations in a given state represent true quantum entanglement is a prominent and long-standing question that frequently resurfaces in different forms. The complexity of determining whether a given mixed quantum state is separable or entangled therefore arose early and was resolved: the problem is NP-complete with respect to Cook reductions when the state is specified as a density matrix and one demands an error tolerance no smaller than an inverse polynomial in the dimension of the matrix [Gur03, Gha10].

1 For example, it is known that if clauses of the Boolean satisfiability problem are limited to two variables each, the resulting problem (2SAT) is NL-complete [Pap94], while if one allows only Horn clauses the resulting problem (HORNSAT) is P-complete [CN10], and if one removes any such limitations on clauses the problem (SAT) is NP-complete [Coo71].

From a physics or engineering perspective, however, it is often more natural to specify a quantum state as arising from a sequence of specified operations or the application of a local Hamiltonian. In this thesis we explore several variations on the complexity of determining whether or not a state specified by a quantum circuit is entangled. We vary, for example, whether we allow general mixed states or restrict to pure states. We also compare the difficulty of deciding whether entanglement is present (separable versus entangled states) with the difficulty of identifying any correlation whatsoever (product versus non-product states). One of the most subtle and interesting variations is to alter the metric used to measure distance between quantum states: we show that the complexity of detecting entanglement produced by an isometry is either QMA(2)-complete or QMA-complete according to whether one measures distance using the familiar trace distance or the more forgiving "one-way LOCC" distance of Ref. [MWW09].

We consider all problems in terms of general multipartite states, though only bipartite states are required for the hardness results—this indicates that in general, detecting multipartite entanglement or correlation may be no more difficult than the detection of bipartite entanglement or correlation. We also consider these problems for quantum channels, asking whether there exists an input to the channel with the specified properties. In most cases, the resulting problem proves to be complete for a complexity class of quantum interactive proofs, providing characterizations of BQP, QMA, QMA(2), and QSZK [Wat09a]. (We will define these classes below for those unfamiliar with them.)

2

Summary of Results

Figure 1 provides a concise summary of our results. Refer to this table for a brief description of the promise problems. Below we give more details of our results along with their relation to prior results in the literature:

- QPROD-PURE-STATE is complete for the class BQP for any inverse polynomial gap in completeness and soundness parameters, as stated in Theorem 12. We demonstrate that this problem is in BQP by employing the "product test" introduced in [MCKB05] along with the analysis of its success probability in [HM10]. We also provide a simple reduction of a general BQP quantum circuit to the problem of pure-state entanglement detection.

- QSEP-ISOMETRY$_{1,1\text{-LOCC}}$ is complete for the class QMA for any inverse polynomial gap in completeness and soundness parameters. We show that this problem is in QMA by building upon prior work of Brandão *et al.* [BCY11a] and the notion of k-extendibility [Wer89a, DPS04]. The QMA proof system requires the prover to provide: 1) a quantum input to the isometry such that the output is close to some product state $|\psi\rangle_A \otimes |\phi\rangle_B$ and 2) k copies of $|\phi\rangle_B$. The verifier then checks whether the prover is being honest by performing phase estimation over the symmetric group on all of the B systems [Kit95] (also called the "permutation test" [BBD$^+$97, KNY08]). This proof system extends naturally to the multipartite case as well. We prove QMA-hardness of QSEP-ISOMETRY$_{1,1\text{-LOCC}}$ by reusing the BQP reduction technique mentioned above.

- QSEP-STATE$_{1,1\text{-LOCC}}$ is decidable by a two-message quantum interactive proof system for a wide range of parameters, is QSZK-hard, and is NP-hard with respect to Cook reductions. Section 4.3.1 builds upon the approach of Brandão *et al.* [BCY11b] and the notion of k-extendibility [DPS02, DPS04] in order to provide a two-message quantum interactive proof to decide the problem. Section

Problem	Summary	Complexity	Circuit
QPROD-PURE-STATE	Is the state generated by the pure-state quantum circuit close to a product state?	BQP-complete	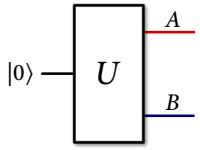
QSEP-ISOMETRY$_{1,1\text{-LOCC}}$	Is there an input to the isometry such that the output is close to a separable state in the trace distance, or does every input lead to an output that is far from separable in 1-LOCC distance?	QMA-complete	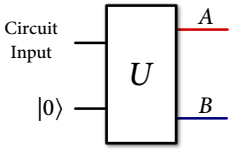
QPROD-ISOMETRY QSEP-ISOMETRY	Is there an input to the isometry such that the output is close to a product/separable state?	QMA(2)-complete	
QPROD-STATE	Is the state generated by the mixed-state circuit close to a product state?	QSZK-complete	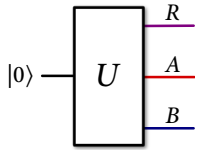
QSEP-STATE$_{1,1\text{-LOCC}}$	Is the state generated by the mixed-state circuit close to a separable state?	In QIP(2), QSZK-hard, NP-hard	
QSEP-CHANNEL$_{1,1\text{-LOCC}}$	Is there an input to the channel such that the output is close to a separable state in trace distance or does every input lead to an output that is far from separable in 1-LOCC distance?	QIP-complete	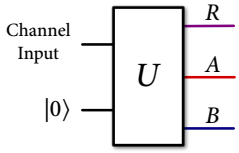

Figure 1: The collected results of entanglement detection problems and their complexity. The leftmost column gives the name of the promise problem. Problem names subscripted with "$1, 1$-LOCC" indicate that the distance measure for yes-instances is the trace distance while the distance measure for no-instances is the one-way LOCC distance. The second column gives a question to which the problem corresponds. The third column states our complexity-theoretic characterization of the problem. The final column depicts a quantum circuit corresponding to the promise problems.

4.3.2 gives a reduction from the QSZK-complete promise problem QUANTUM-STATE-DISTINGUISHABILITY to QSEP-STATE$_{1,1\text{-LOCC}}$ that is somewhat similar to a previous reduction of Rosgen and Watrous [RW05], but the setting here is different and thus requires a different analysis. Section 4.3.3 shows that NP-hardness follows from the fact that the matrix version of the quantum separability problem is NP-hard, though we require some results of Knill which show that one can encode a quantum state efficiently by a unitary circuit if given a matrix description of the state [Kni95].

- QSEP-CHANNEL$_{1,1\text{-LOCC}}$ is QIP-complete for a wide range of completeness and soundness parameters. The reasoning for this proof is similar to that of our earlier results for QSEP-STATE$_{1,1\text{-LOCC}}$, and we again exploit the results of Rosgen and Watrous [RW05] (in particular, the fact that QUANTUM-CIRCUIT-DISTINGUISHABILITY is QIP-complete).

- QPROD-ISOMETRY and QSEP-ISOMETRY are both complete for the class QMA(2) for any inverse polynomial gap in completeness and soundness parameters. We give a QMA(2) proof system in which the verifier performs the product test mentioned above, and we can employ the QMA(k)-amplification results of Harrow and Montanaro to reduce the completeness and soundness errors to become negligible [HM10]. We then show that these problems are QMA(2)-hard by reducing a general QMA(2) verifier circuit to one for which the output can be made product if and only if there exist two product inputs to the original general QMA(2) circuit that would cause the verifier to accept.

- QPROD-STATE is complete for the class QSZK for a wide range of completeness and soundness parameters. This problem differs from the BQP-complete problem QPROD-PURE-STATE in that it allows for a mixed-state quantum circuit to generate the state rather than a unitary circuit. We show that QPROD-STATE is in QSZK by specifying a statistical zero-knowledge proof system that decides it, and we show QSZK-hardness by giving a reduction from the canonical QSZK-complete promise problem QUANTUM-STATE-DISTINGUISHABILITY (QSD)

[Wat02, Wat09b] to QPROD-STATE using a similar reduction to the one used in the analysis of $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}$.

This thesis thus gives a variety of entanglement or correlation detection promise problems that are complete for BQP, QMA, QIP, QMA(2) and QSZK, along with a problem that is in QIP(2). The thesis is structured around this set of of correlation detection problems, beginning with preliminary concepts related to quantum information and quantum interactive proof classes. In the subsequent chapter, we give detailed definitions and justify our aforementioned claims that these various entanglement detection problems are in one-to-one correspondence with BQP, QMA, QIP, QMA(2) and QSZK, as well as prove the stated properties of $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}$. In Section 4.7, we briefly mention how our various proof systems provide operational interpretations for several geometric measures of entanglement (see Refs. [WG03, CAH13] and references therein). Finally, we conclude in Chapter 5 with a summary of our results and a discussion of directions for future research.

3

Background

In this section, we review concepts and complexity classes that will be used throughout the paper, though general background knowledge of both quantum information theory and quantum computational complexity theory is assumed. For more in depth overviews of these fields, consult [NC00, Wil11, Wil13] and [Wat09a, Aar13], respectively.

3.1 Quantum states and channels

A quantum state is a positive semidefinite, unit-trace operator (referred to as the density operator) acting on some Hilbert space $\mathcal{H}$. Let $\mathcal{D}(\mathcal{H})$ denote the set of density operators acting on a Hilbert space $\mathcal{H}$. An *extension* of a quantum state $\rho \in \mathcal{D}(\mathcal{H}_A)$ is some state $\omega \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$ (on a larger Hilbert space) such that $\rho = \mathrm{Tr}_{\mathcal{H}_B}\{\omega\}$. A quantum state is pure if its density operator is unit rank, in which case it has an equivalent representation as a unit vector $|\psi\rangle \in \mathcal{H}$. A *purification* of a density operator $\rho \in \mathcal{D}(\mathcal{H})$ is a pure extension of ρ. Throughout this work, we restrict ourselves to finite-dimensional Hilbert spaces, so that a d-dimensional Hilbert space is isomorphic to $\mathbb{C}^d$. A quantum channel is a linear, completely positive, trace-preserving (CPTP) map $\mathcal{N} : \mathcal{D}(\mathcal{H}_{\mathrm{in}}) \to \mathcal{D}(\mathcal{H}_{\mathrm{out}})$. The Stinespring representation theorem states that every CPTP map can be realized by tensoring its input with an ancillary environment system in some fiducial state $|0\rangle_E \in \mathcal{H}_E$ where $\dim(\mathcal{H}_E) \leq \dim(\mathcal{H}_{\mathrm{in}}) \dim(\mathcal{H}_{\mathrm{out}})$, performing some unitary operation on the joint Hilbert space $\mathcal{H}_{\mathrm{in}} \otimes \mathcal{H}_E$, factoring the unitary's output Hilbert space as $\mathcal{H}_{\mathrm{out}} \otimes \mathcal{H}_{E'}$, and finally tracing over the Hilbert space $\mathcal{H}_{E'}$ [Sti55]. That is, for every CPTP map $\mathcal{N}$, there exists some unitary U such that the following relation holds for all $\rho \in \mathcal{D}(\mathcal{H}_{\mathrm{in}})$:

$$\mathcal{N}(\rho) = \mathrm{Tr}_{E'}\left\{U\left(\rho \otimes |0\rangle\langle 0|_E\right) U^\dagger\right\}.$$

This theorem is the essential reason for the equivalence in computational power between the unitary and mixed-state circuit models of quantum computation [AKN98].

In this section, we review concepts and complexity classes that will be used throughout the paper, though general background knowledge of both quantum information theory and quantum computational complexity theory is assumed. For more in depth overviews of these fields, consult [NC00, Wil11, Wil13] and [Wat09a, Aar13], respectively.

3.1.1 Distance measures

One distance measure often used in quantum information theory to quantify the distance between quantum states is the *trace distance*, induced by the *trace norm*. The trace norm of an operator A is $\|A\|_1 \equiv \mathrm{Tr}\{\sqrt{A^\dagger A}\}$. The trace distance has an important operational interpretation as the bias in distinguishing two states ρ and σ, each elements of $\mathcal{D}(\mathcal{H})$, so that the maximum probability p_{succ} of successfully discriminating them is given by

$$p_{\mathrm{succ}} = \frac{1}{2}\left(1 + \frac{1}{2}\|\rho - \sigma\|_1\right).$$

A variational characterization of the trace distance is as follows:

$$\|\rho - \sigma\|_1 = 2 \max_{0 \leq \Lambda \leq I} \mathrm{Tr}\{\Lambda(\rho - \sigma)\},$$

where the measurement $\{\Lambda, I - \Lambda\}$ that achieves this maximum is known as the Helstrom measurement [Hel69, Hol72, Hel76]. This also leads to the following useful inequality that holds for all Γ such that $0 \leq \Gamma \leq I$:

$$\mathrm{Tr}\{\Gamma\rho\} \geq \mathrm{Tr}\{\Gamma\sigma\} - \|\rho - \sigma\|_1. \tag{1}$$

The quantum fidelity $F(\rho, \sigma)$ between two quantum states ρ and σ is another measure of distinguishability, defined as follows:

$$F(\rho, \sigma) \equiv \left\| \sqrt{\rho}\sqrt{\sigma} \right\|_1^2 . \tag{2}$$

Uhlmann proved that the fidelity is the optimal squared overlap between any two purifications of ρ and σ [Uhl76]:

$$F(\rho, \sigma) = \max_{|\phi_\rho\rangle, |\phi_\sigma\rangle} |\langle \phi_\rho | \phi_\sigma \rangle|^2 .$$

Uhlmann's characterization gives the fidelity an operational interpretation as the optimal probability with which a purification of ρ would pass a test for being a purification of σ. Since all purifications are related by unitary transformations acting on the purifying system, it follows that

$$F(\rho, \sigma) = \max_U |\langle \phi_\rho | (U \otimes I_{\mathcal{H}}) | \phi_\sigma \rangle|^2 \tag{3}$$

for any fixed purifications $|\phi_\rho\rangle$ and $|\phi_\sigma\rangle$ of ρ and σ, respectively. The equivalence between (2) and (3) is commonly known as Uhlmann's theorem. The fidelity and trace distance for general states are related by the Fuchs-van-de-Graaf inequalities [FvdG99]:

$$1 - \sqrt{F(\rho, \sigma)} \leq \tfrac{1}{2} \|\rho - \sigma\|_1 \leq \sqrt{1 - F(\rho, \sigma)}, \tag{4}$$

with the following equality holding for pure states

$$\tfrac{1}{2} \|\psi - \phi\|_1 = \sqrt{1 - F(\psi, \phi)}. \tag{5}$$

The final relevant distance measure that we require is based on the maximum distinguishability of two states when restricting to local operations and one-way classical communication between the systems of the two states. This distance measure

is known as the one-way LOCC distance (1-LOCC), induced by a 1-LOCC norm [MWW09]:

$$\|\rho_{AB} - \sigma_{AB}\|_{1\text{-LOCC}} \equiv \max_{\Lambda_{B\to X}} \|(I_A \otimes \Lambda_{B\to X})(\rho_{AB} - \sigma_{AB})\|_1 ,$$

for two bipartite states $\rho_{AB}, \sigma_{AB} \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and where the maximization on the RHS is over all quantum-to-classical channels

$$\Lambda_{B\to X}(\omega) \equiv \sum_{x\in\mathcal{X}} \text{Tr}\{\Lambda_x \omega\} |x\rangle \langle x| ,$$

with $\Lambda_x \geq 0$ for all $x \in \mathcal{X}$, $\sum_{x\in\mathcal{X}} \Lambda_x = I$, and $\{|x\rangle\}$ some orthonormal basis. (Note that we could also define the 1-LOCC distance with respect to measurement maps on the A systems, which generally gives a different value). This distance is the natural distance measure in the setting of Bell experiments [Bel64] or quantum teleportation. It is also clear from the definitions that

$$\|\rho - \sigma\|_{1\text{-LOCC}} \leq \|\rho - \sigma\|_1 , \tag{6}$$

because a 1-LOCC protocol to distinguish states cannot do any better than a general protocol.

The 1-LOCC distance measure has been extended to multipartite quantum states in [LW12, BaC12, BaH13]. On an l-partite system, the l-partite 1-LOCC distance is given by

$$\|\rho_{A_1\cdots A_l} - \sigma_{A_1\cdots A_l}\|_{1\text{-LOCC}} \equiv \max_{\Lambda_{A_2},\ldots,\Lambda_{A_l}} \|(I_{A_1} \otimes \Lambda_{A_2} \otimes \cdots \otimes \Lambda_{A_l})(\rho_{A_1\cdots A_l} - \sigma_{A_1\cdots A_l})\|_1 , \tag{7}$$

where each of $\Lambda_{A_2}, \ldots, \Lambda_{A_l}$ are quantum-to-classical channels. The interpretation here is that the last $l-1$ parties each perform a local measurement on their system and communicate the results to the first party, who then does her best to distinguish the two states.

3.1.2 Separability and k-extendibility

A multipartite state $\rho_{A_1 \cdots A_l} \in \mathcal{H}_{A_1} \otimes \cdots \otimes \mathcal{H}_{A_l}$ is said to be *separable* if it admits a decomposition of the following form:

$$\rho_{A_1 \cdots A_l} = \sum_{y \in \mathcal{Y}} p_Y(y)\, \sigma^{1,y}_{A_1} \otimes \cdots \otimes \sigma^{l,y}_{A_l}, \tag{8}$$

for collections $\{\sigma^{1,y}_{A_1}\}, \ldots, \{\sigma^{l,y}_{A_l}\}$ of quantum states and some probability distribution $p_Y(y)$ over an alphabet $\mathcal{Y}$ [Wer89b]. By applying the spectral theorem to each density operator, we can always find a decomposition of any separable state in terms of pure product states:

$$\rho_{A_1 \cdots A_l} = \sum_{z \in \mathcal{Z}} p_Z(z)\, |\psi^{1,z}\rangle\langle\psi^{1,z}|_{A_1} \otimes \cdots \otimes |\psi^{l,z}\rangle\langle\psi^{l,z}|_{A_l}. \tag{9}$$

States which cannot be written in this form are *entangled*. Let $\mathcal{S}$ denote the set of separable states. Throughout this work we refer to states in $\mathcal{S}$ as states that are separable across all named systems unless a specific cut is indicated.

States of the form $\sigma^1_{A_1} \otimes \cdots \otimes \sigma^l_{A_l}$ (such that the distribution $p_Y(y)$ in (8) is degenerate) are known as *product states*. Let $\mathcal{P}$ denote the set of product states. $\mathcal{P}$ is not a convex set, and the convex closure of $\mathcal{P}$ is the set $\mathcal{S}$. Operationally, product states are those that are completely uncorrelated between systems and so can be prepared on systems in complete isolation, while separable states can be prepared by means of classical communication between the systems. Furthermore, the correlation exhibited in separable states may be simulated by classical systems in a non-locality, Bell-like test [Wer89b].

Separability has a close connection with the notion of k-extendibility. A bipartite state $\rho_{AB} \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$ is k-extendible [Wer89a, DPS02, DPS04] if there exists a state $\omega_{AB_1 \cdots B_k} \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_{B_1} \otimes \cdots \otimes \mathcal{H}_{B_k})$ such that

1. Each Hilbert space $\mathcal{H}_{B_i}$ is isomorphic to $\mathcal{H}_B$ for all $i \in \{1, \ldots, k\}$.

2. The state $\omega_{AB_1\cdots B_k}$ is invariant under permutations of the systems B_1 through B_k. That is,

$$\forall \pi \in S_k : \omega_{AB_1\cdots B_k} = \left(I_A \otimes W^{\pi}_{B_1\cdots B_k}\right) \omega_{AB_1\cdots B_k} \left(I_A \otimes W^{\pi}_{B_1\cdots B_k}\right)^{\dagger}, \tag{10}$$

where S_k is the symmetric group on k elements and $W^{\pi}_{B_1\cdots B_k}$ is a unitary transformation that implements the permutation π on the B systems.

3. The state $\omega_{AB_1\cdots B_k}$ is an extension of ρ_{AB}:

$$\rho_{AB} = \mathrm{Tr}_{B_2\cdots B_k} \left\{\omega_{AB_1\cdots B_k}\right\}.$$

Let $\mathcal{E}_k$ denote the set of k-extendible states. Every separable state is k-extendible for all $k \geq 2$, because the following state,

$$\sum_{x\in\mathcal{X}} p_X(x) \left|\psi_x\right\rangle\left\langle\psi_x\right|_A \otimes \left|\phi_x\right\rangle\left\langle\phi_x\right|_{B_1} \otimes \cdots \otimes \left|\phi_x\right\rangle\left\langle\phi_x\right|_{B_k},$$

is a suitable k-extension of any separable state $\sum_{x\in\mathcal{X}} p_X(x) \left|\psi_x\right\rangle\left\langle\psi_x\right|_A \otimes \left|\phi_x\right\rangle\left\langle\phi_x\right|_B$. On the other hand, if a state is not separable, there always exists some k for which the state is not k-extendible, and furthermore, for every $l > k$, the state is also not l-extendible [DPS02, DPS04]. This forms a hierarchy of approximations to the set of separable states, becoming exact in the limit as $k \to \infty$.

The bipartite notion of k-extendibility has been further expanded in [DPS05, BaH13] to multipartite states, which requires that every system is extendible according to the conditions above. Specifically, a multipartite state $\rho_C \in \mathcal{D}(\mathcal{H}_{A_1} \otimes \cdots \otimes \mathcal{H}_{A_l})$ (we abbreviate the combined systems $A_1 \cdots A_l$ as C for simplicity) is k-extendible if there exists a state $\omega_{C_1\cdots C_k} \in \mathcal{D}(\mathcal{H}_{C_1} \otimes \cdots \otimes \mathcal{H}_{C_k})$ such that

1. Each Hilbert space $\mathcal{H}_{C_{i,j}}$ is isomorphic to $\mathcal{H}_{A_j}$ for all $i \in \{1,\ldots,k\}$ and $j \in \{1,\ldots,l\}$, where the notation $C_{i,j}$ refers to the j^{th} subsystem of C_i.

2. For all parties $j \in \{1,\ldots,l\}$, the state $\omega_{CC_2\cdots C_k}$ is invariant under permutations of the systems $C_{1,j}$ through $C_{k,j}$. Note that there are $l \cdot k!$ such permutations.

3. The state $\omega_{CC_2\cdots C_k}$ is an extension of ρ_C:

$$\rho_C = \mathrm{Tr}_{C_2\cdots C_k}\left\{\omega_{CC_2\cdots C_k}\right\}.$$

Let $\mathcal{E}_k$ denote the set of k-extendible states for l parties (we suppress the dependence of $\mathcal{E}_k$ on l as it should be clear from context). A fully separable state $\sigma_{A_1:\cdots:A_l}$ of the form in 9 has a k-extension of the following form for all k:

$$\sum_{x\in\mathcal{X}} p_X(x)\left(\left|\psi_x^1\right\rangle\left\langle\psi_x^1\right|_{A_1}\right)^{\otimes k}\otimes\cdots\otimes\left(\left|\psi_x^l\right\rangle\left\langle\psi_x^l\right|_{A_l}\right)^{\otimes k}. \tag{11}$$

The following lemma is essential for some of our quantum interactive proof systems and expands Theorem 2 of [BaH13] to establish a notion of approximate k-extendibility. The proof of Lemma 1 is straightforward and can be found in Appendix A.

Lemma 1 *Let $\rho_{A_1\cdots A_l}$ be ε-far in one-way LOCC distance from the set of fully separable states, for some $\varepsilon > 0$:*

$$\min_{\sigma_{A_1\cdots A_l}\in\mathcal{S}} \left\|\rho_{A_1\cdots A_l} - \sigma_{A_1\cdots A_l}\right\|_{1\text{-LOCC}} \geq \varepsilon.$$

Then the state $\rho_{A_1\cdots A_l}$ is δ-far in trace distance from the set of k-extendible states:

$$\min_{\sigma_{A_1\cdots A_l}\in\mathcal{E}_k} \left\|\rho_{A_1\cdots A_l} - \sigma_{A_1\cdots A_l}\right\|_1 \geq \delta,$$

for $\delta < \varepsilon$ and where

$$k = \left\lceil l + \frac{4l^2\log|C|}{(\varepsilon-\delta)^2}\right\rceil.$$

3.1.3 Quantum interactive proofs

We now formally introduce the quantum interactive proof complexity classes that are relevant to this work. Quantum interactive proof systems involve multiple parties who exchange quantum information: a verifier who has access to a computationally bounded quantum computer and one or more untrustworthy provers who have access

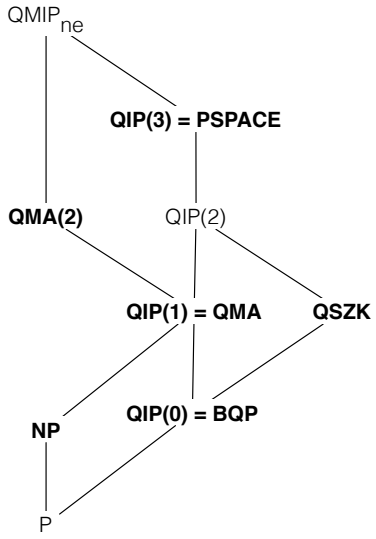

Figure 2: The quantum interactive proof hierarchy and related classes referenced in this paper. A line denotes inclusion of the lower class in the higher class, for example P is a subset of NP. Classes for which there is an entanglement detection problem proven to be complete are in bold type.

to powerful quantum computers bounded only by the laws of quantum mechanics (these provers can perform any unitary operation). The verifier aims to decide whether one of two promises is true—he can receive help from the provers by exchanging quantum messages with them, but he must perform tests to make sure that the provers are not trying to fool him. QMA(2) is the only multi-prover quantum interactive proof complexity classes that we consider in this work. All others that we consider (BQP, QMA, QIP(2), QIP(3), and QSZK) have just one prover.

3.1.4 BQP

The least powerful class within the quantum interactive proof hierarchy consists of a verifier who does not exchange any quantum messages with a prover. Bounded error quantum polynomial time (BQP) includes all promise problems that can be decided by a quantum verifier in polynomial time, and it is the most natural quantum extension of BPP and P, the classical probabilistic and deterministic verifier regimes,

respectively. (The term verifier is used for consistency with what follows. However, in this case, there is no proof being verified—the verifier is simply working on his own.)

Definition 2 (BQP) *Let $A = (A_{yes}, A_{no})$ be a promise problem, and let $c, s : \mathbb{N} \to [0, 1]$ be polynomial-time computable functions such that the gap $c - s$ is at least an inverse polynomial in the input length. Then $A \in \mathsf{BQP}(c, s)$ if there exists a polynomial-time generated family of circuits $U = \{U_n : n \in \mathbb{N}\}$ that satisfies the following properties:*

1. *Completeness: For all input strings $x \in A_{yes}$, the probability of acceptance is at least $c(|x|)$.*
2. *Soundness: For all input strings $x \in A_{no}$, the probability of acceptance is at most $s(|x|)$.*

We define $\mathsf{BQP} = \mathsf{BQP}(2/3, 1/3)$, though note that one can amplify the gap between c and s such that they become exponentially close to their extremes by employing parallel repetition. Thus $\mathsf{BQP} = \mathsf{BQP}(1 - 2^{-p(n)}, 2^{-p(n)})$ for any polynomial function $p(n)$.

3.1.5 QMA

Giving the verifier access to a quantum proof, also called a witness state, seems to greatly expand the set of problems that the verifier can decide in polynomial time. This class is known as Quantum Merlin-Arthur (QMA) [Kit99, Wat00], after the analogous probabilistic verifier class, Merlin-Arthur (MA) [Bab85], in which a computationally bounded verifier (Arthur) wishes to solve a problem with the help of a computationally unbounded but potentially dishonest prover (Merlin). This class is the most natural fully quantum extension of the famous deterministic class NP.

Definition 3 (QMA) *Let $A = (A_{yes}, A_{no})$ be a promise problem, and let $c, s : \mathbb{N} \to [0, 1]$ be polynomial-time computable functions such that the gap $c - s$ is at least an inverse polynomial in the input length. Then $A \in \mathsf{QMA}(c, s)$ if there exists a polynomial-time generated family of circuits $U = \{U_n : n \in \mathbb{N}\}$ that satisfies the following properties:*

1. *Completeness: For all input strings $x \in A_{yes}$, there exists a witness state on a polynomial number of qubits such that the probability of acceptance is at least $c(|x|)$.*

2. *Soundness: For all input strings $x \in A_{no}$ and all witness states, the probability of acceptance is at most $s(|x|)$.*

Note that it suffices for the prover to provide a pure quantum witness state rather than a mixed one. By a simple convexity argument, one can see that for every mixed quantum witness state there exists a pure quantum witness state which has an acceptance probability that is only larger than or equal to that of the mixed witness state.

It is conventional to define $\mathsf{QMA} = \mathsf{QMA}(2/3, 1/3)$, but note that, as in the case of BQP, one can amplify the gap between c and s such that they become exponentially close to their extremes and thus $\mathsf{QMA} = \mathsf{QMA}(1 - 2^{-p(n)}, 2^{-p(n)})$ for any polynomial function $p(n)$. To obtain this result, one can exploit the QMA amplification technique of Marriott and Watrous in [MW05] or the more recent fast amplification procedure of Nagaj *et al.* in [NWZ09].

3.1.6 $\mathsf{QIP}(m)$

We now formally define the namesake family of quantum interactive proof classes. The class $\mathsf{QIP}(m)$ is defined as the class of problems that a verifier can decide if he is allowed to exchange at most m messages with the prover, and it is analogous to the class $\mathsf{IP}(m)$ [GMR89] in the classical probabilistic verifier regime.

Definition 4 (QIP) *Let $A = (A_{yes}, A_{no})$ be a promise problem, and let $c, s : \mathbb{N} \to [0, 1]$ be polynomial-time computable functions such that the gap $c - s$ is at least an inverse polynomial in the input length. Let m be a positive integer no larger than a polynomial in the input length. Then $A \in \mathsf{QIP}(m, c, s)$ if there exists an m-message quantum interactive proof system with the following properties:*

1. *Completeness: For all input strings $x \in A_{yes}$, there exists a prover that causes the verifier to accept with probability at least $c(|x|)$.*

2. *Soundness: For all input strings $x \in A_{no}$, every prover causes the verifier to accept with probability at most $s(|x|)$.*

We define $\mathsf{QIP}(m) = \mathsf{QIP}(m, 2/3, 1/3)$, though Kitaev and Watrous proved in [KW00] that $\mathsf{QIP}(m) = \mathsf{QIP}(3, 1 - 2^{-p(n)}, 2^{-p(n)})$ for all $m \geq 3$ no larger than a polynomial in the input length. As such, we refer to the class $\mathsf{QIP}(3)$ as QIP. For the case that $m = 2$, it is known that $\mathsf{QIP}(2) = \mathsf{QIP}(2, 1 - 2^{-p(n)}, 2^{-p(n)})$ using the error reduction procedure of Jail *et al.* in [JUW09]. A recent breakthrough result in quantum computational complexity theory is that $\mathsf{QIP} = \mathsf{PSPACE}$ [JJUW10].

We also note that any promise problem in BQP and QMA can be decided by a quantum interactive proof system, as $\mathsf{QIP}(0) = \mathsf{BQP}$ and $\mathsf{QIP}(1) = \mathsf{QMA}$. This gives rise to a four-level quantum interactive proof hierarchy, ranging from the verifier alone to a verifier who exchanges no more than three messages with the prover. This hierarchy is shown in Figure 2 along with related classes.

There exist trivially complete promise problems for $\mathsf{QIP}(2)$ and $\mathsf{QIP}(3)$ called CLOSE-IMAGE and CLOSE-IMAGES respectively, which amount to rewriting the definitions of the classes [RW05, Ros09], and which we define below:

Problem 5 (CLOSE-IMAGE) *Fix two constants $c, s \in [0, 1]$ such that $c > s$. Given are a mixed-state quantum circuit to generate the m-qubit state ρ_0 and a mixed-state quantum circuit Q_1, with an n-qubit input state and an m-qubit output state. Decide whether*

1. *Yes: There exists an n-qubit state ρ_1 such that*

$$\max_{\rho_1} F\left(\rho_0, Q_1\left(\rho_1\right)\right) \geq c.$$

2. *No: For all n-qubit states ρ_1, it holds that*

$$\max_{\rho_1} F\left(\rho_0, Q_1\left(\rho_1\right)\right) \leq s.$$

Problem 6 (CLOSE-IMAGES) *Fix two constants $c, s \in [0, 1]$ such that $c > s$. Given are two mixed-state quantum circuits Q_0 and Q_1, each accepting n-qubit inputs and having m-qubit outputs. Decide whether*

1. *Yes: There exist n-qubit states ρ_0 and ρ_1 such that*

$$\max_{\rho_0,\rho_1} F\left(Q_0\left(\rho_0\right), Q_1\left(\rho_1\right)\right) \geq c.$$

2. *No: For all n-qubit states ρ_0 and ρ_1, it holds that*

$$\max_{\rho_0,\rho_1} F\left(Q_0\left(\rho_0\right), Q_1\left(\rho_1\right)\right) \leq s.$$

The following promise problem is also complete for QIP(3), but proving so requires more than a trivial rewriting of the definition of QIP(3) [RW05, Ros09]:

Problem 7 (QUANTUM-CIRCUIT-DISTINGUISHABILITY) *Fix a constant $\varepsilon \in [0,1)$. Given are two mixed-state quantum circuits Q_0 and Q_1, each with n-qubit inputs and m-qubit outputs. Decide whether*

1. *Yes: There is a quantum input for which the circuits are distinguishable:*

$$\max_{\rho \in \mathcal{D}(\mathcal{H}_R \otimes \mathcal{H}_{in})} \left\| \left(I_R \otimes Q_0\right)(\rho) - \left(I_R \otimes Q_1\right)(\rho) \right\|_1 \geq 2 - \varepsilon.$$

2. *No: No quantum input can distinguish the circuits:*

$$\max_{\rho \in \mathcal{D}(\mathcal{H}_R \otimes \mathcal{H}_{in})} \left\| \left(I_R \otimes Q_0\right)(\rho) - \left(I_R \otimes Q_1\right)(\rho) \right\|_1 \leq \varepsilon.$$

In what follows, we abbreviate QUANTUM-CIRCUIT-DISTINGUISHABILITY as QCD.

3.1.7 QMA(2)

Although we have considered only single-prover classes so far, we can also consider a natural extension of QMA in which the verifier has access to unentangled quantum proofs from multiple quantum provers. It is clear that entanglement is a powerful tool in quantum information, and the ways in which the prover can fool the verifier in QMA are directly related to their ability to entangle the witness state. The class QMA(k)

consists of all promise problems that can be decided with the help of k unentangled quantum witness states.

Definition 8 (QMA(k)) *Let $A = (A_{yes}, A_{no})$ be a promise problem, and let $c, s : \mathbb{N} \to [0,1], k \in \mathbb{N}$ be polynomial-time computable functions such that the gap $c - s$ is at least an inverse polynomial in the input length. Then $A \in \mathsf{QMA}(k, c, s)$ if there exists a polynomial-time generated family of circuits $U = \{U_n : n \in \mathbb{N}\}$ that satisfies the following properties:*

1. *Completeness: For all input strings $x \in A_{yes}$, there exist k unentangled quantum witness states on a polynomial number of qubits each, such that the probability of acceptance is at least $c(|x|)$.*

2. *Soundness: For all input strings $x \in A_{no}$ and all possible k unentangled witness states, the probability of acceptance is at most $s(|x|)$.*

Note that allowing classical communication between the provers does not change this complexity class. Indeed, by coordinating with classical communication, they could prepare a separable state to send to the verifier. However, it suffices for the provers to provide a pure, product quantum witness state rather than a mixed separable state. Again, by a simple convexity argument and the decomposition in (9), one can see that for every separable quantum witness state there exists a pure product quantum witness state which has an acceptance probability that is only larger than or equal to that for the separable state; classical communication does not help them cheat.

This family of classes was originally defined in [KMY01]. We define $\mathsf{QMA}(k)$ as $\mathsf{QMA}(k, 2/3, 1/3)$, though Harrow and Montanaro recently showed in [HM10] that $\mathsf{QMA}(k) = \mathsf{QMA}(2)$ for k no larger than a polynomial in the length of the input x, and further that $\mathsf{QMA}(2) = \mathsf{QMA}(2, 1 - 2^{-p(n)}, 2^{-p(n)})$ for any polynomial function $p(n)$. It remains unclear exactly how powerful $\mathsf{QMA}(2)$ is in relation to other quantum interactive proof classes, but there is evidence that the guarantee of unentangled proofs is a very powerful resource [ABD+09]. Estimating the minimum energy of a sparse Hamiltonian over all bipartite product states is a non-trivial promise problem that is complete for $\mathsf{QMA}(2)$ [CS12]. This thesis gives another non-trivial promise problem that is complete for $\mathsf{QMA}(2)$.

3.1.8 QSZK

Classical zero-knowledge proof systems were first considered by Goldwasser *et al.* in the same paper that introduced the classical interactive proof hierarchy [GMR89]. In their work they also introduced *knowledge complexity* as a measure of the amount of knowledge that the prover must transfer to the verifier in order to convince him of the truth of some statement. An interactive proof system for a language is said to be *zero-knowledge* if for every $x \in A_{yes}$, the prover can convince the verifier to accept without the verifier learning anything that he could not have computed himself. In statistical zero knowledge, it is required that in a YES instance, the interaction with the prover must be below some constant in trace distance (traditionally 1/10) to a distribution corresponding to a computation that the verifier can perform themselves.

Quantum statistical zero-knowledge extends this definition to apply to a quantum interactive proof system instead [Wat02, Wat09b], with the requirement being that in a YES instance a computationally bounded quantum computer could simulate the verifier's state at any point to within some constant trace distance.

Definition 9 (QSZK) *A promise problem $A = (A_{yes}, A_{no})$ is in $\mathsf{QSZK}(c, s)$ if there exists a statistical zero-knowledge quantum interactive proof system that satisfies the following properties:*

1. *Completeness: For all input strings $x \in A_{yes}$, the prover can convince the verifier to accept with probability at least $c(|x|)$.*

2. *Soundness: For all input strings $x \in A_{no}$, the prover can convince the verifier to accept with probability at most $s(|x|)$.*

The traditional definition of QSZK is $\mathsf{QSZK}(2/3, 1/3)$, though [Wat02] proved that $\mathsf{QSZK} = \mathsf{QSZK}(1 - 2^{-p(n)}, 2^{-p(n)})$ for any polynomial function $p(n)$. Several facts are known about QSZK: it is closed under complement, any QSZK proof system can be parallelized to two messages, and honest-verifier QSZK is equal to QSZK with a potentially cheating verifier [Wat09b]. The canonical QSZK-complete problem is QUANTUM-STATE-DISTINGUISHABILITY, commonly abbreviated QSD, and defined as follows:

Problem 10 (QUANTUM-STATE-DISTINGUISHABILITY) *Fix a constant $\varepsilon \in [0, 1)$. Given is a mixed-state quantum circuit to generate the n-qubit states ρ_0 and ρ_1. Decide whether*

1. *Yes:* $\|\rho_0 - \rho_1\|_1 \geq 2 - \varepsilon$.
2. *No:* $\|\rho_0 - \rho_1\|_1 \leq \varepsilon$.

4

Results

4.1 QPROD-PURE-STATE is BQP-complete

In order to set the stage for problems in the upcoming sections, we begin with the simplest of our entanglement detection promise problems: determining if a quantum circuit generates a state close to a product state. Unlike the problems in the subsequent sections, the analysis of QPROD-PURE-STATE does not require the help of a prover; it is a straightforward application of the prior results of Harrow and Montanaro [HM10] combined with a reduction from a general BQP circuit.

Problem 11 (QPROD-PURE-STATE(δ_c, δ_s)) *Given is a description of a quantum circuit to generate the n-qubit pure state $|\psi\rangle_{A_1 \cdots A_l}$, along with a labeling of the output qubits for systems $A_1, \ldots, A_l$. Decide whether*

1. *Yes: There is a product state $|\phi_1\rangle_{A_1} \otimes \cdots \otimes |\phi_l\rangle_{A_l}$ that is δ_c-close to $|\psi\rangle_{A_1 \cdots A_l}$ in trace distance:*

$$\min_{|\phi_1\rangle, \ldots, |\phi_l\rangle} \left\| |\psi\rangle\langle\psi|_{A_1 \cdots A_l} - |\phi_1\rangle\langle\phi_1|_{A_1} \otimes \cdots \otimes |\phi_l\rangle\langle\phi_l|_{A_l} \right\|_1 \leq \delta_c. \tag{12}$$

2. *No: Every product state is at least δ_s-far from $|\psi\rangle_{A_1 \cdots A_l}$ in trace distance:*

$$\min_{|\phi_1\rangle, \ldots, |\phi_l\rangle} \left\| |\psi\rangle\langle\psi|_{A_1 \cdots A_l} - |\phi_1\rangle\langle\phi_1|_{A_1} \otimes \cdots \otimes |\phi_l\rangle\langle\phi_l|_{A_l} \right\|_1 \geq \delta_s. \tag{13}$$

Theorem 12 *QPROD-PURE-STATE(δ_c, δ_s) is BQP-complete if there exist polynomial-time computable functions $\delta_c, \delta_s : \mathbb{N} \to [0, 1]$ such that the difference $\frac{11}{2048}\delta_s^2 - \frac{1}{2}\delta_c^2$ is larger than an inverse polynomial in the circuit size.*

Proof. We first show that QPROD-PURE-STATE$(\delta_c, \delta_s) \in$ BQP. The BQP algorithm for deciding QPROD-PURE-STATE(δ_c, δ_s) is to generate two copies of the state $|\psi\rangle_{A_1 \cdots A_l}$

by running the circuit twice, then to perform SWAP tests over each of the pairs of l systems separately, and to accept if and only if all SWAP tests pass. This procedure is known as the product test [MCKB05, HM10].

The promise in (12) implies that

$$\max_{\phi_1,\ldots,\phi_l} |\langle\psi|\phi_1 \otimes \cdots \otimes \phi_l\rangle|^2 \geq 1 - \frac{\delta_c^2}{4},$$

by employing the Fuchs-van-de-Graaf equality in (5). The promise in (13) likewise implies that

$$\max_{\phi_1,\ldots,\phi_l} |\langle\psi|\phi_1 \otimes \cdots \otimes \phi_l\rangle|^2 \leq 1 - \frac{\delta_s^2}{4}.$$

Harrow and Montanaro have determined bounds on the success probability of the product test in Theorem 1 of [HM10]. The verifier accepts if every swap test passes, the probability of which is no smaller than $1 - \frac{\delta_c^2}{2}$ in a YES case, while in a NO case the probability of every swap test passing is no larger than $1 - \frac{11\delta_s^2}{2048}$. Thus, so long as

$$\frac{11}{2048}\delta_s^2 - \frac{1}{2}\delta_c^2$$

is larger than an inverse polynomial in the circuit size, repetition of this procedure no more than a polynomial number of times is sufficient to place the problem in BQP.

We now show that QPROD-PURE-STATE is BQP-hard. Let U denote a quantum circuit for an arbitrary promise problem in BQP acting on $p(n)$ qubits with completeness and soundness error each less than ε, where the decision to accept or reject is based on a measurement of one of the output qubits (the decision qubit) in the computational basis.

We reduce this circuit to QPROD-PURE-STATE by appending three qubits in the state $|0\rangle_{A_1} |\Phi^+\rangle_{A_2A_3}$ to the output of the BQP circuit U. We perform a bit flip on

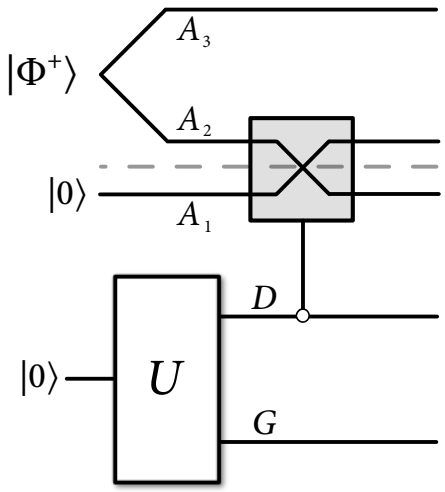

Figure 3: Reduction from a general BQP circuit U to QPROD-PURE-STATE. The constructed circuit works by initializing a Bell state across the A_2 and A_3 systems and swapping it with $|0\rangle$ on A_1 controlled on the value of the decision qubit. This causes the output of the circuit to be product across the $DGA_1 : A_2A_3$ cut if the decision qubit is equal to $|1\rangle$ (and otherwise entangled if the decision qubit is equal to $|0\rangle$).

the decision qubit and a controlled-SWAP from the decision qubit to the qubits in systems A_1 and A_2. The resulting state is as follows:

$$|\psi\rangle_{DGA_1A_2A_3} \equiv (|1\rangle\langle 1|_D \otimes I_G)\,|\phi\rangle_{DG}\,|0\rangle_{A_1}\left|\Phi^+\right\rangle_{A_2A_3} + (|0\rangle\langle 0|_D \otimes I_G)\,|\phi\rangle_{DG}\,|0\rangle_{A_2}\left|\Phi^+\right\rangle_{A_1A_3},$$

where $|\phi\rangle_{DG}$ denotes the state $U\,|0\rangle^{\otimes p(n)}$. This reduction is shown in Figure 3.

We could then feed the result of this computation into an instance of QPROD-PURE-STATE and use an algorithm that decides QPROD-PURE-STATE to determine whether the state is product (or close to product) with respect to the bipartite cut $DGA_1 : A_2A_3$. Given an arbitrary problem in BQP with completeness and soundness error ε, then in a YES instance the following acceptance probability is high:

$$\|(|1\rangle\langle 1|_D \otimes I_G)\,|\phi\rangle_{DG}\|_2^2 \geq 1 - \varepsilon.$$

Thus, after performing the additional steps mentioned above, the resulting state $|\psi\rangle_{DGA_1A_2A_3}$ has a high fidelity with

$$|\phi\rangle_{DG} \otimes |0\rangle_{A_1} \otimes \left|\Phi^+\right\rangle_{A_2A_3}$$

because

$$\begin{aligned}\left|\langle\psi|_{DGA_1A_2A_3} \left(|\phi\rangle_{DG} \otimes |0\rangle_{A_1} \otimes \left|\Phi^+\right\rangle_{A_2A_3}\right)\right|^2 &\geq |\langle\phi|_{DG} \left(|1\rangle\langle 1|_D \otimes I_G\right) |\phi\rangle_{DG}|^2 \\ &= \left\|\left(|1\rangle\langle 1|_D \otimes I_G\right) |\phi\rangle_{DG}\right\|_2^2 \\ &\geq 1 - \varepsilon.\end{aligned}$$

By the Fuchs-van-de-Graaf inequalities, it then follows that

$$\left\||\psi\rangle\langle\psi|_{DGA_1A_2A_3} - |\phi\rangle\langle\phi|_{DG} \otimes |0\rangle\langle 0|_{A_1} \otimes \Phi^+_{A_2A_3}\right\|_1 \leq 2\sqrt{2\varepsilon},$$

so that the state is approximately product with respect to the bipartite cut $DGA_1 : A_2A_3$, and thus

$$\min_{|\zeta\rangle,|\theta\rangle} \left\||\psi\rangle\langle\psi|_{DGA_1A_2A_3} - |\zeta\rangle\langle\zeta|_{DGA_1} \otimes |\theta\rangle\langle\theta|_{A_2A_3}\right\|_1 \leq 2\sqrt{2\varepsilon}. \tag{14}$$

So a YES instance of any promise problem in BQP reduces to a YES instance of QPROD-PURE-STATE.

On the other hand, in the case of a NO instance, the following rejection probability is high:

$$\left\|\left(|0\rangle\langle 0|_D \otimes I_G\right) |\phi\rangle_{DG}\right\|_2^2 \geq 1 - \varepsilon.$$

Thus, after performing the additional steps mentioned above, the resulting state $|\psi\rangle_{DGA_1A_2A_3}$ has a high overlap with

$$|\phi\rangle_{DG} \otimes |0\rangle_{A_2} \otimes \left|\Phi^+\right\rangle_{A_1A_3},$$

because

$$\begin{aligned}\left|\langle\psi|_{DGA_1A_2A_3}\left(|\phi\rangle_{DG}\otimes|0\rangle_{A_2}\otimes\left|\Phi^+\right\rangle_{A_1A_3}\right)\right|^2 &\geq \left|\langle\phi|_{DG}\left(|0\rangle\langle 0|_D\otimes I_G\right)|\phi\rangle_{DG}\right|^2\\ &= \left\|\left(|0\rangle\langle 0|_D\otimes I_G\right)|\phi\rangle_{DG}\right\|_2^2\\ &\geq 1-\varepsilon.\end{aligned}$$

Now, we consider the maximum separable fidelity [Wat04, CAH13] of $|\phi\rangle\langle\phi|_{DG}\otimes|0\rangle\langle 0|_{A_2}\otimes\Phi^+_{A_1A_3}$ with respect to the cut $DGA_1 : A_2A_3$

$$\max_{\sigma_{DGA_1:A_2A_3}\in\mathcal{S}} F\left(|\phi\rangle\langle\phi|_{DG}\otimes|0\rangle\langle 0|_{A_2}\otimes\Phi^+_{A_1A_3},\sigma_{DGA_1:A_2A_3}\right). \tag{15}$$

Since the first state is pure, the fidelity takes the special form

$$\langle\phi|_{DG}\langle 0|_{A_2}\left\langle\Phi^+\right|_{A_1A_3}\sigma_{DGA_1:A_2A_3}|\phi\rangle_{DG}|0\rangle_{A_2}\left|\Phi^+\right\rangle_{A_1A_3},$$

and it is clear that a pure product state optimizes (15). Furthermore, it is clear that we can take the state σ on the systems DG and A_2 to be $|\phi\rangle_{DG}|0\rangle_{A_2}$ since this state is product with respect to the cut $DGA_1 : A_2A_3$. We find that (15) is equal to

$$\max_{|\zeta\rangle,|\theta\rangle}\left|\left\langle\Phi^+\right|_{A_1A_3}|\zeta\rangle_{A_1}|\theta\rangle_{A_3}\right|^2=\frac{1}{2},$$

where the equality follows from [Wat04]. Exploiting the Fuchs-van-de-Graaf inequalities once again, we find that

$$\begin{aligned}\left\||\psi\rangle\langle\psi|_{DGA_1A_2A_3}-|\phi\rangle\langle\phi|_{DG}\otimes|0\rangle\langle 0|_{A_2}\otimes\Phi^+_{A_1A_3}\right\|_1 &\leq 2\sqrt{2\varepsilon},\\ \min_{|\zeta\rangle,|\theta\rangle}\left\||\phi\rangle\langle\phi|_{DG}\otimes|0\rangle\langle 0|_{A_2}\otimes\Phi^+_{A_1A_3}-|\zeta\rangle\langle\zeta|_{DGA_1}\otimes|\theta\rangle\langle\theta|_{A_2A_3}\right\|_1 &\geq 2\left(1-\frac{1}{\sqrt{2}}\right)\\ &\geq\frac{1}{2}.\end{aligned}$$

Using the triangle inequality, we end up with

$$\min_{|\zeta\rangle,|\theta\rangle}\left\||\psi\rangle\langle\psi|_{DGA_1A_2A_3}-|\zeta\rangle\langle\zeta|_{DGA_1}\otimes|\theta\rangle\langle\theta|_{A_2A_3}\right\|_1\geq\frac{1}{2}-2\sqrt{2\varepsilon}, \tag{16}$$

so that a NO instance of any promise problem in BQP reduces to a NO instance of QPROD-PURE-STATE. Since the gap between the lower bound in (16) and the upper bound in (14) is equal to a positive constant $1/2 - 4\sqrt{2\varepsilon}$ for small enough ε, it follows that QPROD-PURE-STATE is BQP-hard. ■

4.2 QSEP-ISOMETRY$_{1,1\text{-LOCC}}$ is QMA-complete

In this section, we provide a proof of the QMA-completeness of QSEP-ISOMETRY$_{1,1\text{-LOCC}}$, the problem of determining if an isometry can generate a state close to some separable state in the trace distance or if all inputs to the isometry lead to a state far from all separable states in the one-way LOCC distance. There are many other problems known to be QMA-complete, including the problem of testing whether a quantum channel (specified by a mixed-state quantum circuit) is not close to an isometry [Ros11], estimating the ground state of a k-local Hamiltonian [KKR06, KSV02] and many more (see [Boo12] for an overview). Nonetheless, it is of interest to note that this problem is QMA-complete when the soundness condition is defined in terms of the one-way LOCC distance, in comparison to the result in the subsequent section that QSEP-ISOMETRY is QMA(2)-complete.

Problem 13 (QSEP-ISOMETRY$_{1,1\text{-LOCC}}(\delta_c, \delta_s)$) *Given is a description of a quantum circuit to implement a unitary U acting on an n-qubit input and m ancilla qubits, as well as a labeling of the systems $A_1, \ldots, A_l$. Decide whether*

1. *Yes: There is an input ρ_S such that the output of U is δ_c-close in trace distance to a separable state:*

$$\min_{\rho, \sigma_{A_1 \cdots A_l} \in \mathcal{S}} \left\| U(\rho_S \otimes |0\rangle\langle 0|^{\otimes m}) U^\dagger - \sigma_{A_1 \cdots A_l} \right\|_1 \leq \delta_c. \tag{17}$$

2. *No: For all inputs ρ_S, the output of U is at least δ_s-far in 1-LOCC distance from a separable state:*

$$\min_{\rho, \sigma_{A_1 \cdots A_l} \in \mathcal{S}} \left\| U(\rho_S \otimes |0\rangle\langle 0|^{\otimes m}) U^\dagger - \sigma_{A_1 \cdots A_l} \right\|_{1\text{-LOCC}} \geq \delta_s. \tag{18}$$

Theorem 14 *QSEP-ISOMETRY$_{1,1\text{-LOCC}}(\delta_c, \delta_s)$ is QMA-complete if there are polynomial-time computable functions $\delta_c, \delta_s : \mathbb{N} \to [0,1]$ such that the difference $\delta_s^2/8 - 4\sqrt{\delta_c}$ is larger than an inverse polynomial in the circuit size.*

Proof. We first show that QSEP-ISOMETRY$_{1,1\text{-LOCC}}(\delta_c, \delta_s) \in$ QMA. Note that by Lemma 27 in Appendix B, the condition in (17) implies the existence of pure states $|\psi\rangle, |\phi_1\rangle, \ldots, |\phi_l\rangle$ such that

$$\left\| U(|\psi\rangle\langle\psi|_S \otimes |0\rangle\langle 0|^{\otimes m})U^\dagger - |\phi_1\rangle\langle\phi_1|_{A_1} \otimes \cdots \otimes |\phi_l\rangle\langle\phi_l|_{A_l} \right\|_1 \leq 4\sqrt{\delta_c}. \tag{19}$$

In a YES instance, the prover can provide the state $|\psi\rangle$ and k copies of the states $|\phi_1\rangle$,...,$|\phi_l\rangle$ to the verifier. The verifier then runs U on $|\psi\rangle_S \otimes |0\rangle^{\otimes m}$ to generate a state close to $|\phi_1\rangle_{A_1} \otimes \cdots \otimes |\phi_l\rangle_{A_l}$ and performs a permutation test over all copies on each of the systems (see Section 4.3.1 for details of the permutation test).

The promise in (17) implies that the permutation test will succeed with probability at least $1 - 4\sqrt{\delta_c}$ for any k. This follows from applying (1) to (19).

In a NO instance, we can employ the promise in (18) and Lemma 1 by requiring k to be larger than

$$\left\lceil l + \frac{4l^2(n+m)}{(\delta_s - \delta_s')^2} \right\rceil,$$

in order to guarantee that

$$\min_{\sigma_{A_1\cdots A_l} \in \mathcal{E}_k} \left\| U(\rho_S \otimes |0\rangle\langle 0|^{\otimes m})U^\dagger - \sigma_{A_1\cdots A_l} \right\|_1 \geq \delta_s',$$

for δ_s' strictly less than δ_s, which can be enforced by setting $\delta_s' = \delta_s/\sqrt{2}$. This gives the following bound on the probability that the permutation test succeeds (for the full analysis of the permutation test see Section 4.3.1):

$$\max_{\sigma_{A_1\cdots A_l} \in \mathcal{E}_k} F(U(\rho_S \otimes |0\rangle\langle 0|^{\otimes m})U^\dagger, \sigma_{A_1\cdots A_l}) \leq 1 - \delta_s^2/8.$$

Note that l cannot be larger than the total number of qubits acted upon, and thus the promise that $\delta_s^2/8 - 4\sqrt{\delta_c}$ is larger than an inverse polynomial is sufficient to place the problem in QMA.

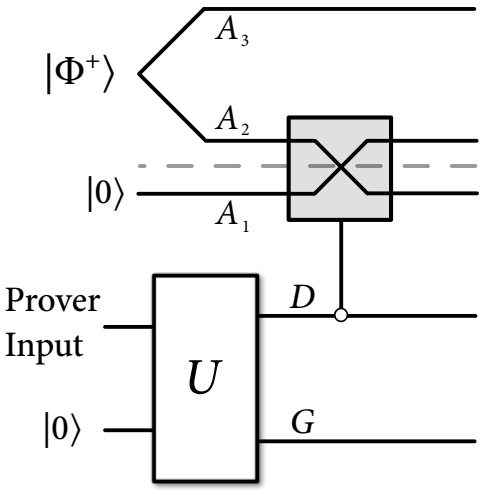

Figure 4: Similar to the BQP reduction, the constructed circuit works by initializing a Bell state across the A_2 and A_3 systems and swapping it with $|0\rangle$ on A_1 controlled on the decision qubit being equal to $|0\rangle$. This causes the input to be separable across the $DGA_1 : A_2A_3$ cut if the decision qubit is close to $|1\rangle$.

The QMA-hardness of $\mathsf{QSEP\text{-}ISOMETRY}_{1,1\text{-LOCC}}$ follows similarly to how we proved BQP-hardness of QPROD-PURE-STATE, by swapping a Bell state across the output systems controlled on the decision qubit being equal to $|0\rangle$. This reduction is shown in Figure 4. This reduction creates a unitary for which the analysis in a YES instance proceeds identically to the analysis in Section 4.1, so that a YES instance of a general QMA problem becomes a YES instance of $\mathsf{QSEP\text{-}ISOMETRY}_{1,1\text{-LOCC}}$ with $\delta_c = 2\sqrt{2\varepsilon}$.

In a NO instance, we wish to show that the one-way LOCC distance from the output of the circuit in Figure 4 to the nearest separable state is larger than an appropriate constant. To show this, we proceed by using the A_1 and A_3 systems in the CHSH game (a reformulation of a Bell experiment as a nonlocal game [CyTW04]), so that we can distinguish the output of the circuit from all separable states by means of a one-way LOCC protocol. In such a protocol, we imagine that Alice has system A_1 and flips a coin x to choose one of two binary-outcome measurements to perform on her qubit. She sends both x and the measurement outcome a to Bob who we imagine has system A_3. Bob then flips a coin y and performs one of two binary outcome

measurements on his qubit, naming the measurement result b. Bob declares the state to be a Bell state in the case that $x \wedge y = a \oplus b$ and otherwise declares that it is not. It is well known that the probability of winning such a game with a Bell state is equal to $\cos^2(\pi/8) \approx 0.85$, while the maximum probability of winning such a game with any separable state is equal to 0.75 [CyTW04]. From this, we can easily lower bound the one-way LOCC distance of the final state from the reduction:

$$\begin{aligned}
\min_{\rho, \sigma_{DGA_1:A_2A_3} \in \mathcal{S}} & \left\| V(\rho_S \otimes |0\rangle\langle 0|^{\otimes m})V^\dagger - \sigma_{DGA_1:A_2A_3} \right\|_{1\text{-LOCC}} \\
& \geq \left\| \Phi^+_{A_1A_3} - \sigma_{A_1:A_3} \right\|_{1\text{-LOCC}} \\
& \qquad - \left\| \Phi^+_{A_1A_3} - \mathrm{Tr}_{DGA_2}\{V(\rho_S \otimes |0\rangle\langle 0|^{\otimes m})V^\dagger\} \right\|_{1\text{-LOCC}} \\
& \geq 0.2 - 2\sqrt{\varepsilon},
\end{aligned}$$

where V denotes the transformation realized by U and the controlled-SWAP and $1-\varepsilon$ is a lower bound on the fidelity between the state of the decision qubit and $|0\rangle$ as in the BQP reduction. The second inequality follows from the fact that the fidelity of $V(\rho_S \otimes |0\rangle\langle 0|^{\otimes m})V^\dagger$ with Φ^+ over the $A_1 : A_3$ system is equal to the fidelity of the decision qubit with $|0\rangle$ (due to the controlled swap). After the reduction, then, this becomes an instance of QSEP-ISOMETRY$_{1,1\text{-LOCC}}$ with $\delta_s = 0.2 - 2\sqrt{\varepsilon}$.

Thus, as long as ε is small enough (so that $0.2 - (2 + 2\sqrt{2})\sqrt{\varepsilon} > 0$), there is an appropriate gap between the completeness and soundness errors, and as such QSEP-ISOMETRY$_{1,1\text{-LOCC}}$ is QMA-hard. With this, we conclude the proof that QSEP-ISOMETRY$_{1,1\text{-LOCC}}$ is QMA-complete. ■

4.3 QSEP-STATE$_{1,1\text{-LOCC}}$ and QIP(2)

Given the many applications of entanglement, it is clearly important to be able to decide if a particular bipartite state is separable or entangled. When the state is specified as the rational entries of a density matrix acting on a finite-dimensional Hilbert space $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$, one can formulate several variations of the problem, all of them being known collectively as the *quantum separability problem*, and characterize

their computational complexity [Gur03, Gha10, BCY11b] (also see Ref. [Ioa07] for a useful, though now somewhat outdated review). Gurvits proved that it is NP-hard (with respect to Cook reductions) to decide if a state $\rho_{AB} \in \mathcal{S}$ or if

$$\min_{\sigma_{AB} \in \mathcal{S}} \|\rho_{AB} - \sigma_{AB}\|_2 \geq \varepsilon,$$

where $\|A\|_2 \equiv \sqrt{\mathrm{Tr}\{A^\dagger A\}}$ is the Hilbert-Schmidt norm and ε is some positive number larger than an inverse exponential in $\dim(\mathcal{H})$. Gharibian later improved upon this result by showing that this formulation of the quantum separability problem is strongly NP-hard with respect to Cook reductions: it is still NP-hard even if ε is promised to be larger than an inverse polynomial in $\dim(\mathcal{H})$. Brandão, Christandl, and Yard then offered a quasi-polynomial time algorithm that decides the quantum separability problem if it is promised that ε is a positive constant [BCY11b], by appealing to their Pinsker-inequality-like lower bound on the squashed entanglement [BCY11a] and to the k-extendibility separability test of Doherty *et al.* [DPS02, DPS04]. They also considered a variant of the promise problem where the Hilbert-Schmidt distance is replaced by the one-way LOCC distance [MWW09], which characterizes the distinguishability of ρ_{AB} and $\mathcal{S}$ if Alice and Bob are allowed to perform local operations and to send one message of classical communication (from either Alice to Bob or Bob to Alice). See [Bei10, HM13] for other results related to separability testing and computational complexity.

In the circuit model of quantum computation, quantum states are generated by unitary circuits acting on some number of qubits (with some of them being traced out in the mixed-state circuit model [AKN98]), and we measure the complexity of a quantum computation by how the circuit size (number of gates and wires) scales with the length of the input [Wat09a]. (Note that if the circuit size is polynomial in the input length, then the number of qubits on which the circuit acts is likewise polynomial in the input length.) Thus, from the perspective of quantum computational complexity theory [Wat09a], one might consider the prior formulations of the quantum separability problem to be somewhat restrictive. The reason is the same as that given in [Ros09]: the mathematical description of a bipartite quantum state is polynomial in the dimension of the Hilbert space, but this Hilbert space is exponential in the number

of qubits in the state. Thus, the matrix representation is exponentially larger than it needs to be when we are in the setting of the circuit model of quantum computation. Also, the circuit model is natural physically, as the evolution induced by a time-varying two-body Hamiltonian can be efficiently described by a quantum circuit [BACS07].

In this section—with this dual computational and physical motivation in mind—we take an approach to the quantum separability problem along the above lines. We define QSEP-STATE$_{1,1\text{-LOCC}}$ and give a protocol for deciding QSEP-STATE$_{1,1\text{-LOCC}}$ in QIP(2), followed by a proof of QSZK-hardness, and finally a proof of NP-hardness with respect to Cook reductions

4.3.1 QSEP-STATE$_{1,1\text{-LOCC}}$ is in QIP(2)

Problem 15 (QSEP-STATE$_{1,1\text{-LOCC}}(\delta_c, \delta_s)$) *Given is a mixed-state quantum circuit to generate the n-qubit state ρ_C, along with a labeling of the qubits in the reference system R and the output qubits for each system $A_1, \ldots, A_l \in C$. Decide whether*

1. *Yes: There is a fully separable state $\sigma_C \in \mathcal{S}$ that is δ_c-close to ρ_C in trace distance:*
$$\min_{\sigma_C \in \mathcal{S}} \|\rho_C - \sigma_C\|_1 \leq \delta_c.$$
2. *No: All fully separable states are at least δ_s-far from ρ_C in 1-LOCC distance:*
$$\min_{\sigma_C \in \mathcal{S}} \|\rho_C - \sigma_C\|_{1-LOCC} \geq \delta_s.$$

Theorem 16 *QSEP-STATE$_{1,1\text{-LOCC}}(\delta_c, \delta_s) \in$ QIP(2) if there are polynomial-time computable functions $\delta_c, \delta_s : \mathbb{N} \to [0,1]$, such that the difference $\delta_s^2/8 - 2\sqrt{\delta_c}$ is larger than an inverse polynomial in the circuit size.*

Proof. Figure 5 depicts a two-message quantum interactive proof system for QSEP-STATE$_{1,1\text{-LOCC}}$. The protocol begins with the verifier preparing the state $|\psi_\rho\rangle_{RC}$, a particular purification of ρ_C, by running the quantum circuit U_ρ as given in the problem instance. The verifier transmits the reference system to the prover, who then acts on R and some ancillary qubits with a unitary P_1 that has output systems

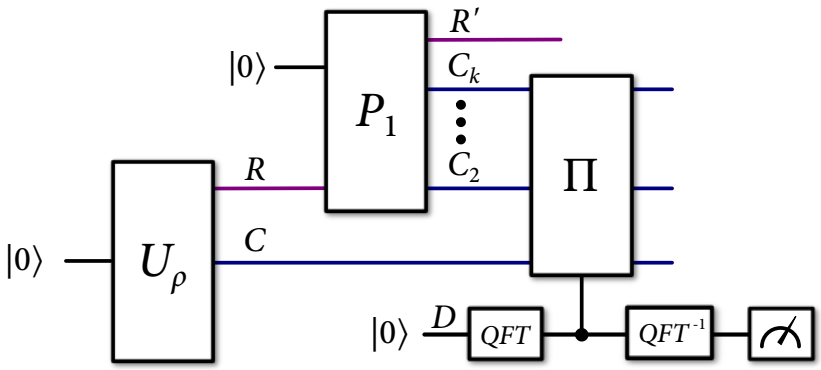

Figure 5: A two-message quantum interactive proof system for $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}$. The proof system begins with the verifier executing the circuit U_ρ that generates the state ρ_C. He sends the reference system to the prover. In the case that ρ_C is fully separable, the prover should be able to act with a unitary on the reference system and some ancillas in order to generate a k-extension of ρ_C to the systems C_2 through C_k. The prover sends all of the extension systems back to the verifier, who then performs the permutation test in order to test if the state sent by the prover is a valid k-extension.

R', C_2, …, C_k. The prover transmits systems C_2, …, C_k to the verifier. The verifier then performs phase estimation over the symmetric group [Kit95, BBD+97] (also known as the "permutation test" [KNY08]) on the registers C, C_2, …, C_k, using the qubits in system D as the control. This control register requires $\lceil \log(l \cdot k!) \rceil$ qubits because the permutations included in the test are those from the definition of multipartite k-extendibility. The verifier performs a computational basis measurement on all of the qubits in the control register D and accepts if and only if the measurement outcome is all zeros.

This protocol is just an implementation of a k-extendibility test on a quantum computer. We can build intuition for why it works on YES instances by examining the exact case, when ρ_C is actually a separable state. In this case, we know that ρ_C has a decomposition of the form given in (9), and as such, it has an extension of the form in (11) for all $k > 1$. Thus, the following state is a purification of ρ_C:

$$|\phi_{\sigma,k}\rangle_{R'CC_2\cdots C_k} \equiv \sum_{x\in\mathcal{X}} \sqrt{p_X(x)}|x\rangle_{R'} \left(|\psi_x^1\rangle_{A_1}\right)^{\otimes k} \otimes \cdots \otimes \left(|\psi_x^l\rangle_{A_l}\right)^{\otimes k}, \tag{20}$$

where $\{|x\rangle_{R'}\}$ is some orthonormal basis for the reference system. Since all purifications are related by unitaries on the reference system, the prover can append ancilla qubits to the R system received from the verifier and perform a unitary P_1 that takes $|\psi_\rho\rangle_{RC}|0\rangle$ to $|\phi_{k,\rho}\rangle_{R'CC_2\cdots C_k}$. The prover then sends the systems C_2, ..., C_k to the verifier. The verifier performs a permutation test on the systems C, ..., C_k. Since the state $|\phi_{k,\rho}\rangle_{R'RC\cdots C_k}$ is invariant under permutations of the systems C, ..., C_k, the qubits in the control register D do not acquire a phase. Thus, after the final quantum Fourier transform is applied, the qubits in the control register D are in the all-zero state with certainty.

The analysis for a YES instance follows the above intuition closely. In this case, there is some state $\sigma_C \in \mathcal{S}$ that is δ_c-close in trace distance to ρ_C. By Uhlmann's theorem and the Fuchs-van-de-Graaf inequalities in (4), there is a purification $|\psi_\sigma\rangle_{RC}$ of σ_C such that

$$\left\| |\psi_\rho\rangle\langle\psi_\rho|_{RC} - |\psi_\sigma\rangle\langle\psi_\sigma|_{RC} \right\|_1 \leq 2\sqrt{\delta_c}. \tag{21}$$

Thus the prover can just operate as above, but choosing his unitary P_1 to correspond to the state $|\psi_\sigma\rangle_{RC}$ instead. Writing as U the unitary corresponding to P_1 followed by the permutation test, we obtain the following lower bound on the probability with which the verifier accepts:

$$\begin{aligned}
&\operatorname{Tr}\left\{ |0\rangle\langle 0|_D U \left(|\psi_\rho\rangle\langle\psi_\rho|_{RC} \right) U^\dagger \right\} \\
&= \operatorname{Tr}\left\{ U^\dagger |0\rangle\langle 0|_D U \left(|\psi_\rho\rangle\langle\psi_\rho|_{RC} \right) \right\} \\
&\geq \operatorname{Tr}\left\{ U^\dagger |0\rangle\langle 0|_D U \left(|\psi_\sigma\rangle\langle\psi_\sigma|_{RC} \right) \right\} - \left\| |\psi_\rho\rangle\langle\psi_\rho|_{RC} - |\psi_\sigma\rangle\langle\psi_\sigma|_{RC} \right\|_1 \\
&\geq 1 - 2\sqrt{\delta_c},
\end{aligned} \tag{22}$$

where the first inequality follows from (1), and the second inequality follows by applying (21) and because the protocol accepts with probability one for a separable state.

The analysis for a NO instance has two components:

1. demonstrating that the maximum k-extendible fidelity is an upper bound on the maximum acceptance probability

2. using Lemma 1 regarding approximate k-extendibility and the first item above to specify how large k should be in order to obtain a good upper bound on the maximum acceptance probability.[1]

For the protocol in Figure 5, the state generated by the verifier's first circuit is as follows:

$$\rho_C \otimes |\text{perm}\rangle \langle \text{perm}|_D ,$$

where $|\text{perm}\rangle_D$ is a superposition over all possible permutations of k elements over all l systems resulting from an application of the quantum Fourier transform [NC00] to the state $|0\rangle_D$:

$$|\text{perm}\rangle_D \equiv \frac{1}{\sqrt{l \cdot k!}} \sum_{\pi \in S_k^{\otimes l}} |\pi\rangle_D , \tag{23}$$

so that the D register requires $\lceil \log_2 (l \cdot k!) \rceil$ qubits. (Note that Figure 5 depicts the verifier generating $|\text{perm}\rangle_D$ later in the protocol, but we could just as easily reorder things so that he generates this state in the first step.) The channel generated by the inverse of the verifier's circuit conditional on accepting is

$$\begin{aligned} &\mathcal{M}_{C,C_2 \cdots C_k \to CD} (\sigma_{AB B_2 \cdots B_k}) \\ &\quad \equiv \text{Tr}_{C_2 \cdots C_k} \left\{ (U_\Pi)_{C C_2 \cdots C_k D} (\sigma_{C C_2 \cdots C_k} \otimes |\text{perm}\rangle \langle \text{perm}|_D) \left(U_\Pi^\dagger \right)_{C C_2 \cdots C_k D} \right\}, \end{aligned} \tag{24}$$

where $(U_\Pi)_{C C_2 \cdots C_k D}$ is a controlled-permutation operation:

$$(U_\Pi)_{C C_2 \cdots C_k D} \equiv \sum_{\pi \in S_k} W^\pi_{C C_2 \cdots C_k} \otimes |\pi\rangle \langle \pi|_D , \tag{25}$$

and $W^\pi_{C C_2 \cdots C_k}$ is a unitary operation corresponding to permutation π. Note that we can write the verifiers maximum acceptance probability as the maximum fidelity of the channel generated by the inverse of the verifier's circuit with the initial output state [KW00, RW05, Ros09], and thus the maximum acceptance probability is equal to

$$\max_{\sigma_{C C_2 \cdots C_k}} F \left(\rho_C \otimes |\text{perm}\rangle \langle \text{perm}|_D , \mathcal{M}_{C C_2 \cdots C_k \to CD} (\sigma_{C C_2 \cdots C_k}) \right) .$$

1 For a YES instance, the value of k does not matter because the lower bound on the maximum acceptance probability is always as given above.

Since the fidelity can only increase under the discarding of the control register D,[2] the maximum acceptance probability is upper bounded by the following quantity:

$$\max_{\sigma_{CC_2\cdots C_k}\in\mathcal{S}} F\left(\rho_C, \mathcal{M}_{CC_2\cdots C_k\to C}\left(\sigma_{CC_2\cdots C_k}\right)\right), \tag{26}$$

where

$$\begin{aligned}
&\mathcal{M}_{CC_2\cdots C_k\to C}\left(\sigma_{CC_2\cdots C_k}\right) \\
&\qquad = \operatorname{Tr}_D\left\{\mathcal{M}_{CC_2\cdots C_k\to CD}\left(\sigma_{CC_2\cdots C_k}\right)\right\} \\
&\qquad = \frac{1}{l\cdot k!}\sum_{\pi\in S_k}\operatorname{Tr}_{C_2\cdots C_k}\left\{\left(I_A\otimes W^{\pi}_{CC_2\cdots C_k}\right)\sigma_{CC_2\cdots C_k}\left(I_A\otimes W^{\pi}_{CC_2\cdots C_k}\right)^{\dagger}\right\},
\end{aligned}$$

which is just the channel that applies a random permutation over the l parts of the C systems and discards the last $k-1$ systems $C_2, \ldots, C_k$. Clearly, since the channel $\mathcal{M}_{CC_2\cdots C_k\to C}$ symmetrizes the state of the systems $CC_2\cdots C_k$, the maximum in (26) is achieved by a state $\sigma_{CC_2\cdots C_k}$ for which systems $CC_2\cdots C_k$ are permutation symmetric. Thus, by recalling the definition of k-extendibility, we can rewrite (26) as the maximum k-extendible fidelity of ρ_C:

$$\max_{\sigma_{CC_2\cdots C_k}} F\left(\rho_C, \mathcal{M}_{CC_2\cdots C_k\to C}\left(\sigma_{CC_2\cdots C_k}\right)\right) = \max_{\sigma_C\in\mathcal{E}_k} F\left(\rho_C,\sigma_C\right). \tag{27}$$

This demonstrates that the maximum k-extendible fidelity is an upper bound on the maximum acceptance probability and completes our proof of the first item above.

The second part of the analysis of a NO instance involves determining how large k needs to be. Suppose that

$$\min_{\sigma_C\in\mathcal{S}} \left\|\rho_C-\sigma_C\right\|_{1\text{-LOCC}} \geq \delta_s.$$

2 We can interpret discarding the control register as actually giving it to the prover, so that the resulting fidelity corresponds to the maximum acceptance probability in a modified protocol in which the prover controls the inputs to D.

According to Lemma 1, if we take k to be larger than

$$\left\lceil l + \frac{4l^2 \log |C|}{(\delta_s - \delta_s')^2} \right\rceil,$$

then we can guarantee that

$$\min_{\sigma_C \in \mathcal{E}_k} \|\rho_C - \sigma_C\|_1 \geq \delta_s',$$

for δ_s' strictly less than δ_s. We can enforce this latter condition by setting $\delta_s' = \delta_s/\sqrt{2}$. Then, using the following manipulation of the Fuchs-van-de-Graaf inequalities in (4):

$$F(\rho, \sigma) \leq 1 - \frac{1}{4} \|\rho - \sigma\|_1^2,$$

we have that

$$\begin{aligned} \max_{\sigma_C \in \mathcal{E}_k} F(\rho_C, \sigma_C) &\leq 1 - \frac{1}{4} \min_{\sigma_C \in \mathcal{E}_k} \|\rho_C - \sigma_C\|_1^2 && (28) \\ &\leq 1 - \frac{1}{4} (\delta_s')^2 && (29) \\ &= 1 - \delta_s^2/8. && (30) \end{aligned}$$

In the above, we have separated the probability of accepting and the probability of rejecting by an inverse polynomial in the number of qubits (namely, from the promise that the difference $\delta_s^2/8 - 2\sqrt{\delta_c}$ is at least an inverse polynomial in the circuit size), and it is known that an inverse polynomial gap is sufficient to place this protocol in QIP(2) (see Section 3.2 of Ref. [JUW09] for how to amplify an inverse polynomial gap). Thus, we have given a two-message quantum interactive proof system that decides the quantum separability problem. ■

4.3.2 QSEP-STATE$_{1,1\text{-LOCC}}$ is QSZK-hard

Having placed an upper bound on the difficulty of solving QSEP-STATE$_{1,1\text{-LOCC}}$, we now move on to lower bounds, beginning in this section with a proof that it is QSZK-hard. Our approach is to exhibit a Karp reduction from the QSZK-complete promise problem QSD to QSEP-STATE$_{1,1\text{-LOCC}}$. The essential idea behind the reduction is similar to Rosgen and Watrous's reduction of CLOSE-IMAGES to QCD [RW05, Ros09].

In order to demonstrate this reduction, we have to show that there is a polynomial-time algorithm that encodes YES instances of QSD into YES instances of QSEP-STATE$_{1,1\text{-LOCC}}$ and the same for the NO instances. Recall that for QSD, we are given a description of circuits U_{ρ_0} and U_{ρ_1} that generate mixed states ρ_0 and ρ_1. The output qubits of the circuit are divided into two sets: qubits in a reference system R that are traced out and qubits in a system S which contains ρ_i. For $i \in \{0, 1\}$, let

$$|\psi_{\rho_i}\rangle_{RS} \equiv U_{\rho_i} |0\rangle ,$$

so that

$$\rho_i = \mathrm{Tr}_R \left\{ |\psi_{\rho_i}\rangle \langle\psi_{\rho_i}|_{RS} \right\} .$$

Figure 6 depicts a circuit that accomplishes the reduction. From the description of the circuits U_{ρ_0} and U_{ρ_1}, one can generate a description of the circuit in Figure 6 in polynomial time, and furthermore, the resulting circuit runs efficiently on a quantum computer [Ros09]. The circuit takes as input a Bell state

$$\left|\Phi^+\right\rangle_{AB} \equiv \frac{1}{\sqrt{2}} \left(|00\rangle_{AB} + |11\rangle_{AB} \right) ,$$

and performs the following controlled unitary from the qubit B to the ancilla qubits:

$$|0\rangle \langle 0|_B \otimes U_{\rho_0} + |1\rangle \langle 1|_B \otimes U_{\rho_1} .$$

The resulting state is as follows:

$$|\varphi\rangle_{ABRS} \equiv \frac{1}{\sqrt{2}} \left(|0\rangle_A |0\rangle_B |\psi_{\rho_0}\rangle_{RS} + |1\rangle_A |1\rangle_B |\psi_{\rho_1}\rangle_{RS} \right) .$$

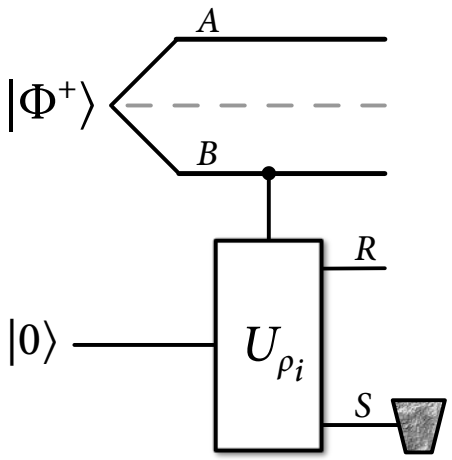

Figure 6: Given respective circuit descriptions U_{ρ_0} and U_{ρ_1} for generating the states ρ_0 and ρ_1 on the output system S, one can compute a description for the above circuit in polynomial time, and furthermore, the above circuit can be run efficiently on a quantum computer. This serves as a reduction from QUANTUM-STATE-DISTINGUISHABILITY to QSEP-STATE$_{1,1\text{-LOCC}}$, i.e., where one should decide if the state on systems A and BR is separable with respect to this cut.

The output qubits are divided into three sets: environment qubits in the system S that are traced out, a single qubit in system A, and qubits in systems BR. Thus, the state resulting from applying the circuit in Figure 6 is as follows:

$$\omega_{A:BR} \equiv \mathrm{Tr}_S \left\{ |\varphi\rangle \langle\varphi|_{ABRS} \right\}. \tag{31}$$

The task is to decide whether the state on systems A and BR is separable across this cut, subject to the promise in Problem 15. Our claim is that YES instances of QSD map to YES instances of QSEP-STATE$_{1,1\text{-LOCC}}$, with the same holding true for NO instances.

The intuition for why this reduction works is as follows. In the case of a YES instance of QSD, the states ρ_0 and ρ_1 are approximately orthogonal, so that tracing

out the S system of the circuit in Figure 6 decoheres the Bell state, leaving a state on A and BR close to the following state:

$$\omega^{\text{sep}}_{A:BR} \equiv \frac{1}{2}\left(|0\rangle\langle 0|_A \otimes |0\rangle\langle 0|_B \otimes (\psi_{\rho_0})_R + |1\rangle\langle 1|_A \otimes |1\rangle\langle 1|_B \otimes (\psi_{\rho_1})_R\right). \tag{32}$$

The above state is clearly separable with respect to the bipartite cut $A : BR$. In the case of a NO instance of QSD, the states ρ_0 and ρ_1 are approximately indistinguishable, and tracing over the S system of the circuit in Figure 6 does little to decohere the entanglement shared between A and BR. Thus, Bob can perform a local unitary operation on systems B and R to distill out a pure Bell state shared between A and B. After this, Alice and Bob can perform a Bell experiment on the distilled Bell state to determine if they indeed share a Bell state. Since these two operations can be performed with local operations and one message of classical communication, the resulting state is 1-LOCC distinguishable from the set of separable states.

We now give a formal proof to justify this reduction:

Theorem 17 *QSEP-STATE$_{1,1\text{-LOCC}}$ with constant promise gap is QSZK-hard.*

Proof. We first prove that the circuit in Figure 6 maps YES instances of QSD to YES instances of QSEP-STATE$_{1,1\text{-LOCC}}$. So we begin by assuming that

$$\|\rho_0 - \rho_1\|_1 \geq 2 - \varepsilon, \tag{33}$$

and we will use this condition to show that the fidelity between $\omega^{\text{sep}}_{A:BR}$ in (32) and the reduced state $\omega_{A:BR}$ in (31) is close to one. So, recall from Uhlmann's theorem that the fidelity between $\omega^{\text{sep}}_{A:BR}$ and $\omega_{A:BR}$ is the maximum squared overlap between any purifications of these states. Thus, if we can show that the squared overlap between two *particular* purifications of $\omega^{\text{sep}}_{A:BR}$ and $\omega_{A:BR}$ is large, then this implies a lower bound on the fidelity between these two states. Consider the following particular purification of $\omega^{\text{sep}}_{A:BR}$:

$$|\omega^{\text{sep}}_{ABB'RS}\rangle \equiv \frac{1}{\sqrt{2}}\left(|0\rangle_A |0\rangle_B |0\rangle_{B'} |\psi_{\rho_0}\rangle_{RS} + |1\rangle_A |1\rangle_B |1\rangle_{B'} |\psi_{\rho_1}\rangle_{RS}\right).$$

Recall that the trace distance bound in (33) implies the existence of a two-outcome projective measurement $\{\Pi_0, \Pi_1\}$ (known as a Helstrom measurement [Hel69, Hol72, Hel76]) that has the following success probability in discriminating ρ_0 from ρ_1 if they are chosen uniformly at random:

$$\begin{aligned}\frac{1}{2}\mathrm{Tr}\left\{\Pi_0\rho_0\right\} + \frac{1}{2}\mathrm{Tr}\left\{\Pi_1\rho_1\right\} &= \frac{1}{2}\left(1 + \frac{1}{2}\left\|\rho_0 - \rho_1\right\|_1\right) \\ &\geq 1 - \frac{\varepsilon}{2}. \end{aligned} \tag{34}$$

Performing the following "Helstrom isometry"

$$U^H_{S\to SB'} \equiv (\Pi_0)_S \otimes |0\rangle_{B'} + (\Pi_1)_S \otimes |1\rangle_{B'}$$

on the S system of $|\varphi\rangle_{ABRS}$ produces a particular purification of the state $\omega_{A:BR}$:

$$U^H_{S\to SB'}|\varphi\rangle_{ABRS} = \frac{1}{\sqrt{2}}\sum_{i,j\in\{0,1\}} |i\rangle_A |i\rangle_B \otimes |j\rangle_{B'} \otimes (\Pi_j)_S |\psi_{\rho_i}\rangle_{RS}.$$

The overlap between these purifications is

$$\begin{aligned}&\langle\omega^{\text{sep}}_{ABB'RS}|U^H_{S\to SB'}|\varphi\rangle_{ABRS} \\ &= \frac{1}{2}\left(\sum_{k\in\{0,1\}} \langle k|_A \langle k|_B \langle k|_{B'} \langle\psi_{\rho_k}|_{RS}\right)\left(\sum_{i,j\in\{0,1\}} |i\rangle_A |i\rangle_B \otimes |j\rangle_{B'} \otimes (\Pi_j)_S |\psi_{\rho_i}\rangle_{RS}\right) \\ &= \frac{1}{2}\sum_{i,j,k\in\{0,1\}} \langle k|i\rangle_A \ \langle k|i\rangle_B \ \langle k|j\rangle_{B'} \ \langle\psi_{\rho_k}|_{RS} I_R \otimes (\Pi_j)_S |\psi_{\rho_i}\rangle_{RS} \\ &= \frac{1}{2}\sum_{i\in\{0,1\}} \langle\psi_{\rho_i}|_{RS} I_R \otimes (\Pi_i)_S |\psi_{\rho_i}\rangle_{RS} \\ &= \frac{1}{2}\mathrm{Tr}\left\{\Pi_0\rho_0\right\} + \frac{1}{2}\mathrm{Tr}\left\{\Pi_1\rho_1\right\} \\ &\geq 1 - \frac{\varepsilon}{2}, \end{aligned}$$

where the inequality follows from (34). Squaring the overlap gives the following lower bound on the fidelity:

$$F(\omega^{\text{sep}}_{A:BR}, \omega_{A:BR}) \geq 1 - \varepsilon,$$

which imply by the Fuchs-van-de-Graaf inequalities in (4) that

$$\min_{\sigma_{A:BR}\in\mathcal{S}} \|\omega_{A:BR} - \sigma_{A:BR}\|_1 \leq 2\sqrt{\varepsilon}. \tag{35}$$

Thus, the circuit in Figure 6 transforms a YES instance of QSD to a YES instance of QSEP-STATE$_{1,1\text{-LOCC}}$. We note that the above argument is reminiscent of similar ones from quantum information theory [Dev05].

We now prove that the circuit in Figure 6 transforms NO instances of QSD into NO instances of QSEP-STATE$_{1,1\text{-LOCC}}$. In this case, we have the promise that the states ρ_0 and ρ_1 are nearly indistinguishable:

$$\|\rho_0 - \rho_1\|_1 \leq \varepsilon.$$

Due to the Fuchs-van-de-Graaf inequalities, we have the following lower bound on the fidelity:

$$F(\rho_0, \rho_1) \geq 1 - \varepsilon,$$

and Uhlmann's theorem implies the existence of a unitary operation U_R acting on the reference system of $|\psi_{\rho_1}\rangle_{RS}$ such that

$$\langle\psi_{\rho_0}|_{RS} U_R \otimes I_S |\psi_{\rho_1}\rangle_{RS} \geq \sqrt{1-\varepsilon}.$$

(A global phase can be fixed for U_R such that the overlap is a real number.) Thus, Bob can apply the following controlled-unitary to the state $|\varphi\rangle_{ABRS}$:

$$C^U_{BR} \equiv |0\rangle\langle 0|_B \otimes I_R + |1\rangle\langle 1|_B \otimes U_R,$$

leading to

$$\begin{aligned}(|0\rangle\langle 0|_B \otimes I_R + |1\rangle\langle 1|_B \otimes U_R)\,|\varphi\rangle_{ABRS} \\ = \frac{1}{\sqrt{2}}\left(|0\rangle_A |0\rangle_B |\psi_{\rho_0}\rangle_{RS} + |1\rangle_A |1\rangle_B U_R \otimes I_S |\psi_{\rho_1}\rangle_{RS}\right).\end{aligned}$$

Then the overlap between $|\Phi^+\rangle_{AB} \otimes |\psi_{\rho_0}\rangle_{RS}$ and the resulting state is large:

$$\begin{aligned}
&\frac{1}{2}\left(\left(\langle 0|_A \langle 0|_B + \langle 1|_A \langle 1|_B\right) \otimes \langle \psi_{\rho_0}|_{RS}\right) \\
&\qquad\qquad \left(|0\rangle_A |0\rangle_B |\psi_{\rho_0}\rangle_{RS} + |1\rangle_A |1\rangle_B U_R \otimes I_S |\psi_{\rho_1}\rangle_{RS}\right) \\
&= \frac{1}{2} + \frac{1}{2} \langle \psi_{\rho_0}|_{RS} U_R \otimes I_S |\psi_{\rho_1}\rangle_{RS} \\
&\geq \frac{1}{2} + \frac{1}{2}\sqrt{1-\varepsilon} \\
&\geq \sqrt{1-\varepsilon},
\end{aligned}$$

implying that the fidelity is larger than $1-\varepsilon$. Thus, by a local operation, Bob can distill a state which is $2\sqrt{\varepsilon}$-close in trace distance to the product state $|\Phi^+\rangle_{AB} \otimes |\psi_{\rho_0}\rangle_{RS}$:

$$\left\| C^U_{BR} |\varphi\rangle \langle\varphi|_{ABRS} \left(C^U_{BR}\right)^\dagger - \left|\Phi^+\right\rangle \left\langle\Phi^+\right|_{AB} \otimes |\psi_{\rho_0}\rangle \langle \psi_{\rho_0}|_{RS} \right\|_1 \leq 2\sqrt{\varepsilon}.$$

(We remark that the above argument is similar to a "decoupling" argument well known in quantum information theory [Dev05, ADHW09].)

Now, we would like to argue that the one-way LOCC distance between $\omega_{A:BR}$ and the separable state $\sigma^*_{A:BR} \in S$ closest to $\omega_{A:BR}$ is larger than an appropriate constant, so that we can claim that the circuit in Figure 6 maps NO instances of QSD to NO instances of $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}$. In order to do so, Bob first performs the local unitary C^U_{BR}. This transforms the state $\left(C^U_{BR}\right)^\dagger \left(\Phi^+_{AB} \otimes (\psi_{\rho_0})_R\right) C^U_{BR}$ to $\Phi^+_{AB} \otimes (\psi_{\rho_0})_R$ and the separable state $\sigma^*_{A:BR}$ to some other separable state $(\sigma^*_{A:BR})'$. Alice and Bob then perform a Bell experiment, guessing the state to be $|\Phi^+\rangle_{AB}$ if there is a violation of a Bell inequality and guessing a separable state otherwise [Bel64]. Equivalently, Alice and Bob could proceed as in the CHSH game (a reformulation of a Bell experiment as a nonlocal game [CyTW04]). In such a protocol, Alice flips a coin x and chooses one of two binary-outcome measurements to perform on her qubit. She sends both x and the measurement outcome a to Bob. Bob then flips a coin with outcome y and performs one of two binary-outcome measurements on his qubit, naming the measurement result b. Bob declares the state to be the Bell state in the case that $x \wedge y = a \oplus b$ (when they "win the CHSH game") and otherwise declares that it is not the Bell state. It

is well known that the winning probability of the CHSH game with a Bell state is equal to $\cos^2(\pi/8) \approx 0.85$, while the maximum probability with which they can win this game with a separable state is equal to 0.75 [CyTW04]. This gives the following lower bound on the one-way LOCC distance between $\left(C^U_{BR}\right)^\dagger \left(\Phi^+_{AB} \otimes (\psi_{\rho_0})_R\right) C^U_{BR}$ and $\sigma^*_{A:BR}$:

$$\begin{aligned}
&\left\| \left(C^U_{BR}\right)^\dagger \left(\Phi^+_{AB} \otimes (\psi_{\rho_0})_R\right) C^U_{BR} - \sigma^*_{A:BR} \right\|_{1-\text{LOCC}} \\
&\qquad = \left\| \Phi^+_{AB} \otimes (\psi_{\rho_0})_R - (\sigma^*_{A:BR})' \right\|_{1-\text{LOCC}} \\
&\qquad \geq \left\| \left(\cos^2(\pi/8), \sin^2(\pi/8)\right) - (0.75, 0.25) \right\|_1 \\
&\qquad \geq 0.2.
\end{aligned}$$

Thus, by combining with the distillation argument above, we have the following lower bound on the one-way LOCC distance between $\omega_{A:BR}$ and $\sigma^*_{A:BR}$:

$$\begin{aligned}
\|\omega_{A:BR} - \sigma^*_{A:BR}\|_{1-\text{LOCC}} &\geq \left\| \left(C^U_{BR}\right)^\dagger \left(\Phi^+_{AB} \otimes (\psi_{\rho_0})_R\right) C^U_{BR} - \sigma^*_{A:BR} \right\|_{1-\text{LOCC}} \\
&\quad - \left\| \left(C^U_{BR}\right)^\dagger \left(\Phi^+_{AB} \otimes (\psi_{\rho_0})_R\right) C^U_{BR} - \omega_{A:BR} \right\|_{1-\text{LOCC}} \\
&\geq \left\| \Phi^+_{AB} \otimes (\psi_{\rho_0})_R - (\sigma^*_{A:BR})' \right\|_{1-\text{LOCC}} \\
&\quad - \left\| \left(C^U_{BR}\right)^\dagger \left(\Phi^+_{AB} \otimes (\psi_{\rho_0})_R\right) C^U_{BR} - \omega_{A:BR} \right\|_1 \\
&\geq 0.2 - 2\sqrt{\varepsilon}, \qquad (36)
\end{aligned}$$

where the second inequality follows from (6) and the fact that $\left(C^U_{BR}\right)$ is a local unitary, and the third from the argument at the end of the previous paragraph. Thus, as long as ε is small enough (so that $0.2 - 4\sqrt{\varepsilon} > 0$), there is a gap between (35) and (36). In fact, Watrous showed that it is possible to make ε exponentially small with only polynomial overhead for any instance of QSD [Wat02] by exploiting a "quantized" version of the polarization lemma in [SV97]. Thus, any protocol for deciding QSEP-STATE$_{1,1\text{-LOCC}}$ could also decide QSD, implying that QSEP-STATE$_{1,1\text{-LOCC}}$ is QSZK-hard. ■

Ideally, we would like to show that QSEP-STATE$_{1,1\text{-LOCC}}$ is a complete promise problem for QIP(2), but it is not clear to us how to do so. The obvious way to attempt

this would be to reduce CLOSE-IMAGE to QSEP-STATE$_{1,1\text{-LOCC}}$, but the problem is that CLOSE-IMAGE requires a general channel, whereas our protocol for QSEP-STATE$_{1,1\text{-LOCC}}$ has a very specific channel (one that applies a random permutation to the B systems and discards the last $k-1$ of them). Alternatively, we could attempt to find a QSZK proof system for QSEP-STATE$_{1,1\text{-LOCC}}$, but the protocol that we have given to show that QSEP-STATE$_{1,1\text{-LOCC}} \in$ QIP(2) does not satisfy the zero-knowledge property because, in the case of a YES instance, the verifier ends up with a state close to a k-extension of ρ_{AB}, which he could not have generated himself using a polynomial-time quantum circuit.

4.3.3 QSEP-STATE$_{1,1\text{-LOCC}}$ is NP-hard

We now prove NP-hardness of QSEP-STATE$_{1,1\text{-LOCC}}$, with respect to Cook reductions, by finding a reduction to it from the NP-hard matrix version of the quantum separability problem. The essence of the reduction is Knill's efficient encoding of a density matrix description of a state ρ_{AB} as a description of a quantum circuit to generate it [Kni95]. We begin by recalling the matrix version of the quantum separability problem:

Problem 18 (WMEM$_\varepsilon$ (M, N)) *Given a density matrix $\rho_{AB} \in \mathcal{D}(\mathcal{H}_M \otimes \mathcal{H}_N)$ with rational entries subject to the promise that either*

(i) $\rho_{AB} \in \mathcal{S}$ *or*

(ii) $\min_{\sigma_{AB} \in \mathcal{S}} \|\rho_{AB} - \sigma_{AB}\|_2 \geq \varepsilon$,

with ε no smaller than an inverse polynomial in MN, decide which is the case.

Gharibian showed that the above promise problem is NP-hard [Gha10], so our task is just to find a Cook reduction from WMEM$_\varepsilon$ (M, N) to QSEP-STATE$_{1,1\text{-LOCC}}$. First, consider that we can diagonalize the matrix ρ_{AB} in time polynomial in $MN \log(MN/\varepsilon_1)$, where ε_1 is an error parameter characterizing the precision of the diagonalization in the trace distance. We then compute a purification $|\phi_\rho\rangle_{RAB}$ of ρ_{AB} to a reference system with dimension no larger than MN. Knill's algorithm gives a quantum circuit running on $O(\log(MN))$ qubits that generates the state $|\phi_\rho\rangle^{RAB}$ [Kni95], and this algorithm

runs in time polynomial in MN. Knill's algorithm outputs controlled single-qubit unitary gate descriptions with arbitrary precision, so we need to invoke the Solovay-Kitaev algorithm [DN06] to approximate each gate in Knill's circuit with unitaries chosen from a finite gate set, up to precision ε_2/l where l is the number of gates in Knill's circuit. The Solovay-Kitaev algorithm runs in time polylogarithmic in l/ε_2 and produces a gate sequence with length polylogarithmic in l/ε_2. This whole procedure leads to a mixed state quantum circuit generating a state ρ'_{AB} such that $\|\rho_{AB} - \rho'_{AB}\|_1 \leq \varepsilon_1 + \varepsilon_2$. The state ρ'_{AB} will be used as the input to $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}(\delta_c, \delta_s)$.

Setting $\delta_c = \varepsilon_1 + \varepsilon_2$ implies that any instance of $\mathsf{WMEM}_\varepsilon$ for which $\rho_{AB} \in \mathcal{S}$, meaning case *(i)* of the promise, gets mapped to a YES instance of $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}(\delta_c, \delta_s)$. For case *(ii)*, we know from Matthews *et al.* [MWW09] that

$$\min_{\sigma_{AB}\in\mathcal{S}} \|\rho_{AB} - \sigma_{AB}\|_{1-\text{LOCC}} \geq \frac{1}{\sqrt{153}} \min_{\sigma_{AB}\in\mathcal{S}} \|\rho_{AB} - \sigma_{AB}\|_2 \geq \frac{\varepsilon}{\sqrt{153}}. \tag{37}$$

This in turn implies that

$$\min_{\sigma_{AB}\in\mathcal{S}} \|\rho'_{AB} - \sigma_{AB}\|_{1-\text{LOCC}} \geq \frac{\varepsilon}{\sqrt{153}} - \varepsilon_1 - \varepsilon_2,$$

so if we choose $\delta_s = \varepsilon/\sqrt{153} - \varepsilon_1 - \varepsilon_2$ then case *(ii)* gets mapped to a NO instance of $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}(\delta_c, \delta_s)$. Moreover, because $\varepsilon_1 + \varepsilon_2$ can be made to shrink exponentially with the circuit size, the gap $\delta_s - \delta_c$ remains inverse polynomial in the circuit size. In particular, the instance of $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}(\delta_c, \delta_s)$ will be in $\mathsf{QIP}(2)$ for sufficiently small ε, as determined by the promise in Theorem 16.

4.4 QSEP-CHANNEL$_{1,1\text{-LOCC}}$ is QIP-complete

There is a straightforward variation of $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}$ which is a complete promise problem for $\mathsf{QIP}(3)$ (and therefore complete for QIP [KW00]). In this variation, the input is a description of a circuit that implements a quantum channel with input system S and output systems $A_1, \ldots, A_l$. (The channel is implemented by a unitary

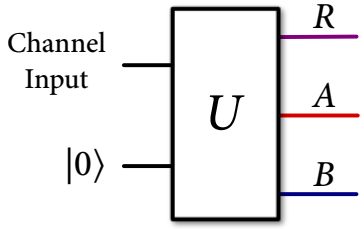

Figure 7: A quantum circuit to implement a channel. The circuit has input qubits and ancillas in the state $|0\rangle$. The circuit outputs qubits in the environment system R (which are traced out) and qubits in systems A and B.

circuit with qubits in an environment system R that are traced out.) Figure 7 depicts a circuit that implements such a channel. The task is to decide whether there is an input to the channel such that the output state on systems $A_1 \ldots A_l$ is separable.

Problem 19 (QSEP-CHANNEL$_{1,1\text{-LOCC}}(\delta_c, \delta_s)$) *Given is a mixed-state quantum circuit to generate the channel $\mathcal{N}_{S\to C}$, having an n-qubit input and an m-qubit output, along with a labeling of the qubits in the environment system R and the output qubits for each system $A_1, \ldots, A_l \in C$. Decide whether*

1. *Yes: There is an input to the channel ρ_S such that the channel output $\mathcal{N}_{S\to C}(\rho_S)$ is δ_c-close in trace distance to a separable state $\sigma_C \in \mathcal{S}$:*

$$\min_{\rho_S,\ \sigma_C \in \mathcal{S}} \|\mathcal{N}_{S\to C}(\rho_S) - \sigma_C\|_1 \leq \delta_c. \tag{38}$$

2. *No: For all channel inputs ρ_S, the channel output $\mathcal{N}_{S\to C}(\rho_S)$ is at least δ_s-far in 1-LOCC distance to a separable state:*

$$\min_{\rho_S,\ \sigma_C \in \mathcal{S}} \|\mathcal{N}_{S\to C}(\rho_S) - \sigma_C\|_{1-LOCC} \geq \delta_s.$$

Theorem 20 *QSEP-CHANNEL$_{1,1\text{-LOCC}}(\delta_c, \delta_s)$ is QIP-complete if there are polynomial-time computable functions $\delta_c, \delta_s : \mathbb{N} \to [0,1]$, such that the difference $\delta_s^2/8 - 2\sqrt{\delta_c}$ is larger than an inverse polynomial in the circuit size.*

Proof. The proof of this theorem is almost identical to the proofs of Theorems 16 and 17.

We first show that there is a three-message quantum interactive proof system for QSEP-CHANNEL$_{1,1\text{-LOCC}}$. This is just the obvious modification of the circuit in Figure 5 so that it becomes a three-message proof system. In particular, the prover first prepares a state and sends it to the verifier. The verifier inputs this state to the circuit that implements the channel $\mathcal{N}_{S\to C}$, and the rest of the proof system proceeds as in Figure 5. In the case of a positive instance, the prover can compute the states ρ_S and σ_C in (38) from the description of the channel $\mathcal{N}_{S\to C}$. He generates ρ_S with his first unitary operation and then proceeds by choosing his second unitary operation as if the state $\mathcal{N}_{S\to C}(\rho_S)$ were σ_C. Following the same analysis as in the proof of Theorem 16, the maximum probability with which the verifier accepts in this case is no smaller than $1 - 2\sqrt{\delta_c}$. In the case of a negative instance, by Lemma 1, for every state $\mathcal{N}_{S\to C}(\rho_S)$, there is some k polynomial in the circuit size such that the maximum probability with which the prover can make the verifier accept is no larger than $1 - \delta_s^2/8$. An upper bound on the maximum acceptance probability is

$$\max_{\omega_S,\ \sigma_C \in \mathcal{E}_k} F\left(\mathcal{N}_{S\to C}(\omega_S), \sigma_C\right),$$

a formula which follows from our previous analysis in the proof of Theorem 16. This leaves a gap of $\delta_s^2/8 - 2\sqrt{\delta_c}$ between completeness and soundness error (promised to be larger than an inverse polynomial) and it is known that this gap can be amplified [KW00]. Thus, QSEP-CHANNEL$_{1,1\text{-LOCC}}(\delta_c, \delta_s) \in$ QIP.

To show that QSEP-CHANNEL$_{1,1\text{-LOCC}}$ is QIP-hard, it suffices to exhibit a reduction from the QIP-complete promise problem QCD (Problem 7) to QSEP-CHANNEL$_{1,1\text{-LOCC}}$. This reduction is essentially the same as that in the proof of Theorem 17, except that the circuit in Figure 6 is modified so that the unitaries being controlled are the unitaries that generate the channels (rather than the ones that generate the states). In the case of a positive instance of QCD, there exists an input to the channels such that their outputs are nearly distinguishable, so that the output of the modified circuit is nearly separable. Also, in the case of a negative instance, the outputs of the channels for all inputs are nearly indistinguishable, so that it is possible to distill a Bell state

from the output state of the modified circuit. The CHSH game argument then applies as well. Thus, QSEP-CHANNEL$_{1,1\text{-LOCC}}$ is QIP-hard. ■

4.5 QPROD-ISOMETRY and QSEP-ISOMETRY are QMA(2)-complete

In this section we show that QPROD-ISOMETRY, the problem of determining if an isometry can produce a state close to a product state (in the trace distance), is QMA(2)-complete. We also demonstrate that it is equivalent to the problem QSEP-ISOMETRY, which is the trace distance version of the QSEP-ISOMETRY$_{1,1\text{-LOCC}}$ problem analyzed in Section 4.2. It is clear that being able to detect productness in the trace distance is of considerable use; for one, it allows a verifier to force two unentangled provers to simulate k unentangled provers as shown in the proof that QMA(2) = QMA(k) [HM10]. It seems intuitive then that these problems in the trace distance can neatly capture the power of "unentanglement," an intuition that we make precise in what follows.

Problem 21 (QPROD-ISOMETRY(δ_c, δ_s)) *Given is a description of a quantum circuit to implement a unitary U acting on an n-qubit input and m ancilla qubits, as well as a labeling of the systems $A_1, \ldots, A_l$. Decide whether*

1. *Yes: There is an input ρ such that the output of U is δ_c-close in trace distance to a product state:*

$$\min_{\rho, \sigma_{A_1 \cdots A_l} \in \mathcal{P}} \left\| U(\rho_S \otimes |0\rangle\langle 0|^{\otimes m}) U^\dagger - \sigma_{A_1 \cdots A_l} \right\|_1 \leq \delta_c. \tag{39}$$

2. *No: For all inputs ρ, the output of U is at least δ_s-far in trace distance from a product state:*

$$\min_{\rho, \sigma_{A_1 \cdots A_l} \in \mathcal{P}} \left\| U(\rho_S \otimes |0\rangle\langle 0|^{\otimes m}) U^\dagger - \sigma_{A_1 \cdots A_l} \right\|_1 \geq \delta_s. \tag{40}$$

Theorem 22 *QPROD-ISOMETRY(δ_c, δ_s) is QMA(2)-complete if there are polynomial-time computable functions $\delta_c, \delta_s : \mathbb{N} \to [0, 1]$ such that the difference $\frac{11\delta_s^2}{4096} - 8\delta_c$ is larger than an inverse polynomial in the circuit size.*

Proof. We first show that QPROD-ISOMETRY is in $\mathsf{QMA}(l+1)$, from which it follows by [HM10] that it is in $\mathsf{QMA}(2)$. Our proof system has $l+1$ provers send the minimizing input and each part of the minimizing product state, followed by the verifier performing the unitary U on the input, then the product test on the provided product state.

Let $\omega_{A_1 \cdots A_l}$ denote the state that results after the verifier performs the unitary U on the input ρ_S received from the first prover, and let $\sigma_{A_1 \cdots A_l}$ denote the product state received from the other l provers. Lemma 2 of [HM10] establishes the following formula for the success probability of the product test:

$$P_{\text{test}}(\omega_{A_1 \cdots A_l}, \sigma_{A_1 \cdots A_l}) = \frac{1}{2^l} \sum_{S \subseteq \{A_1, \ldots, A_l\}} \operatorname{Tr}\left\{\omega_S \sigma_S\right\}.$$

In particular, it is clear by a convexity argument that it is optimal for the last l provers to send pure quantum states to the verifier. That is, for every set of mixed states that they could send, there exists a set of pure states that gives the same or higher probability of passing the product test. As such, we can assume without loss of generality that the last l provers send pure states.

We now analyze the YES instance. By Lemma 27, the condition in (39) implies that there exist pure states $\psi, \phi_1, \ldots, \phi_l$ such that

$$\left\| U\left(|\psi\rangle\langle\psi|_S \otimes |0\rangle\langle 0|^{\otimes m}\right) U^\dagger - |\phi_1\rangle\langle\phi_1|_{A_1} \otimes \cdots \otimes |\phi_l\rangle\langle\phi_l|_{A_l} \right\|_1 \leq 4\sqrt{\delta_c}.$$

Thus, in a YES instance, the $l+1$ provers can provide the states $\psi, \phi_1, \ldots, \phi_l$ respectively, so that running the product test between $U(\psi \otimes |0\rangle\langle 0|^{\otimes m})U^\dagger$ and $\phi_1 \otimes \cdots \otimes \phi_l$ will succeed with probability no smaller than $1 - 8\delta_c$.

We now analyze the NO instance. The promise in (40) gives the following upper bound on δ_s:

$$\begin{aligned}
\delta_s &\leq \min_{\rho,\sigma_{A_1\cdots A_l}\in\mathcal{P}} \left\| U(\rho_S \otimes |0\rangle\langle 0|^{\otimes m})U^\dagger - \sigma_{A_1\cdots A_l} \right\|_1 \\
&\leq 2\sqrt{1 - \max_{\rho,\sigma_{A_1\cdots A_l}\in\mathcal{P}} F(U(\rho_S \otimes |0\rangle\langle 0|^{\otimes m})U^\dagger, \sigma_{A_1\cdots A_l})} \\
&\leq 2\sqrt{1 - \max_{\rho,\sigma_{A_1\cdots A_l}\in\mathcal{P}} (\mathrm{Tr}\{U(\rho_S \otimes |0\rangle\langle 0|^{\otimes m})U^\dagger \sigma_{A_1\cdots A_l}\})^2} \\
&= 2\sqrt{1 - \max_{\psi,\phi_1,\ldots,\phi_l} |(\langle\psi_S| \otimes \langle 0|^{\otimes m})U^\dagger|\phi_1\rangle \otimes \cdots \otimes |\phi_l\rangle|^4}.
\end{aligned}$$

The second inequality is an application of the Fuchs-van-de-Graaf inequalities. The third inequality follows from $F(\rho,\sigma) = (\mathrm{Tr}\{|\sqrt{\rho}\sqrt{\sigma}|\})^2 \geq (\mathrm{Tr}\{\sqrt{\rho}\sqrt{\sigma}\})^2 \geq (\mathrm{Tr}\{\rho\sigma\})^2$. The final inequality follows from a convexity argument (for every set of mixed states, there is a set of pure states that can achieve the same or higher value, so that it suffices to maximize over pure states). We can rewrite the above bound as follows:

$$\max_{\psi,\phi_1,\ldots,\phi_l} |(\langle\psi_S| \otimes \langle 0|^{\otimes m})U^\dagger|\phi_1\rangle \otimes \cdots \otimes |\phi_l\rangle|^2 \leq \sqrt{1 - \delta_s^2/4} \leq 1 - \delta_s^2/8.$$

By Theorem 1 of [HM10], we can then conclude that the probability of the product test succeeding is no greater than $1 - \frac{11\delta_s^2}{4096}$. Thus, the promise in Theorem 22 is sufficient for the verifier to decide this instance, placing QPROD-ISOMETRY in $\mathsf{QMA}(l+1)$ and further in $\mathsf{QMA}(2)$ by applying the exponential amplification result of [HM10].

To show that QPROD-ISOMETRY is $\mathsf{QMA}(2)$-hard, consider an arbitrary $\mathsf{QMA}(2)$ circuit acting on $p(n)$ qubits with completeness and soundness error at most ε. On an input x, we describe the verifier's corresponding unitary as $V_x : ABW \to DG$, which takes two product inputs from the provers on the A and B systems respectively along with ancilla qubits in the $|0\rangle$ state on the W system, and outputs a decision qubit (labeled as D) along with a reference system G. Note that any $\mathsf{QMA}(2)$ verifier can be expressed in this way. The circuit V_x is depicted in Figure 8(a).

We reduce any $\mathsf{QMA}(2)$ proof system to QPROD-ISOMETRY by constructing a circuit $U : DGCC' \to A_1A_2$ shown in Figure 8(b) as follows:

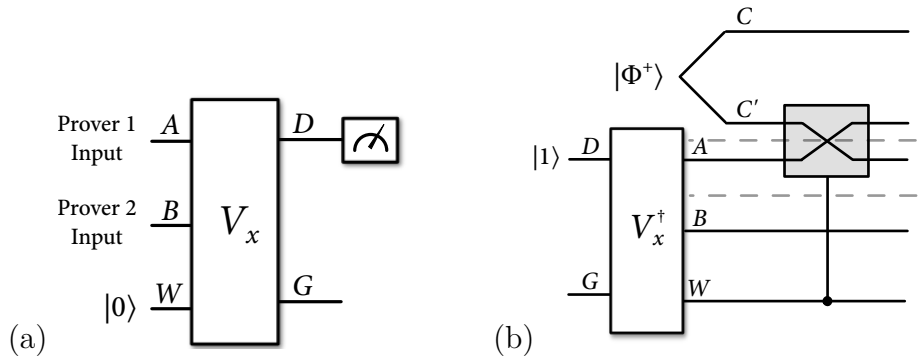

Figure 8: (a) The unitary circuit V_x for an arbitrary $\mathsf{QMA}(2)$ verifier on input x. Such a verifier has two quantum witness registers, A and B, inputs to which are provided by the two unentangled provers. The verifier initializes an ancilla register W to the all-zero state $|0\rangle$ and performs the unitary V_x. This procedure produces a single-qubit decision register D, the measurement of which indicates acceptance or rejection, and a register G which is ignored. (b) The circuit U produced by our reduction. The verifier's circuit V_x is inverted and the decision register is initialized to $|1\rangle$. After $V_x^\dagger$ is applied, the register W controls whether the Bell states are swapped in to cause the output to be entangled.

1. Prepare a Bell state across the ancilla registers C, C'.

2. Prepare the D register in the state $|1\rangle$ and perform $V_x^\dagger$ on the registers DG to obtain the registers ABW.

3. Perform the following "controlled swap" gate:

$$(|0\rangle\langle 0|^{\otimes m})_W \otimes I_{AC'} + \left(I^{\otimes m} - |0\rangle\langle 0|^{\otimes m}\right)_W \otimes \mathrm{SWAP}_{AC'}.$$

4. Label the A register as A_1 and the $CC'BW$ registers as A_2.

We begin by showing that such a circuit can produce a state close to product if there is an accepting input to V_x, and then that the circuit can only produce states that are far from separable if no such accepting input exists. (Note that the state is far from being product if it is far from being separable because $\mathcal{P} \in \mathcal{S}$.)

In the case of a YES instance of the QMA(2) problem, by a convexity argument, we know that there are pure states $|\phi\rangle_A$ and $|\psi\rangle_B$ such that

$$\langle 1|_D \operatorname{Tr}_G\{V_x(|\phi\rangle\langle\phi|_A \otimes |\psi\rangle\langle\psi|_B \otimes (|0\rangle\langle 0|^{\otimes m})_W)V_x^\dagger\}|1\rangle_D \geq 1-\varepsilon.$$

By Uhlmann's theorem, this means that there also exists a pure state $|\zeta\rangle_G$ such that

$$|\langle 1|_D \langle\zeta|_G V_x(|\phi\rangle_A \otimes |\psi\rangle_B \otimes (|0\rangle^{\otimes m})_W)|^2 \geq 1-\varepsilon.$$

Thus, there exists an input $|\zeta\rangle_G$ to the circuit U such that the output will have a large overlap with the state $|\phi\rangle_A \otimes |\psi\rangle_B \otimes |\Phi^+\rangle_{CC'} \otimes (|0\rangle^{\otimes m})_W$ (so that $1-\varepsilon$ of the weight of the state, after $V_x^\dagger$ acts, is on $(|0\rangle^{\otimes m})_W$ and the controlled-SWAP acts almost as the identity). The state $|\phi\rangle_A \otimes |\psi\rangle_B \otimes |\Phi^+\rangle_{CC'} \otimes (|0\rangle^{\otimes m})_W$ is product across the cut $A : BCC'W$, so that we map a YES instance of a general QMA(2) proof system to a YES instance of QPROD-ISOMETRY.

Now, let x be a NO instance. In such a case, we have a promise that there is no product input to V_x such that the probability of measuring $|1\rangle$ on the decision qubit is larger than ε:

$$\max_{|\phi\rangle_A,|\psi\rangle_B} \langle 1|_D \operatorname{Tr}_G\{V_x(|\phi\rangle\langle\phi|_A \otimes |\psi\rangle\langle\psi|_B \otimes (|0\rangle\langle 0|^{\otimes m})_W)V_x^\dagger\}|1\rangle_D \leq \varepsilon. \tag{41}$$

By Uhlmann's theorem, the above is equivalent to

$$\max_{|\phi\rangle_A,|\psi\rangle_B,|\zeta\rangle_G} |\langle 1|_D \langle\zeta|_G V_x(|\phi\rangle_A \otimes |\psi\rangle_B \otimes (|0\rangle^{\otimes m})_W)|^2 \leq \varepsilon. \tag{42}$$

To show that the output of the circuit U is far from a product state for any input on the system G, we will bound the following quantity:

$$\max_{|\zeta\rangle,|\sigma\rangle_{A:CC'BW}\in\mathcal{P}} \left|\langle\sigma|_{ACC'BW} U \left|\Phi^+\right\rangle_{CC'} |1\rangle_D |\zeta\rangle_G\right|, \tag{43}$$

where $|\sigma\rangle$ is product across the $A : CC'BW$ cut. To do so, note that the state $|\sigma\rangle$ can be written as the following superposition:

$$|\sigma\rangle_{ACC'BW} = \alpha_0 |0\rangle_W |\sigma_0\rangle_{ACC'B} + \alpha_1 |1\rangle_W |\sigma_1\rangle_{ACC'B}, \tag{44}$$

where $|\alpha_0|^2 + |\alpha_1|^2 = 1$. (In the above, we are now modelling the system W as a qubit in which $|0\rangle$ represents the projection onto the all-zeros state of m qubits and $|1\rangle$ represents the projection onto the complementary space. It suffices for us to do since we are only interested in these two subspaces.) Also, note that both $|\sigma_0\rangle_{ACC'B}$ and $|\sigma_1\rangle_{ACC'B}$ are product across the bipartite cut $A : CC'B$ since $|\sigma\rangle_{ACC'BW}$ is. Substituting (44) into (43), we find that

$$\begin{aligned}
&\max_{|\zeta\rangle,|\sigma\rangle_{A:CC'BW}\in\mathcal{P}} \left|\langle\sigma|_{ACC'BW} U \left|\Phi^+\right\rangle_{CC'} |1\rangle_D |\zeta\rangle_G\right| \\
=&\max_{|\zeta\rangle,|\sigma\rangle_{A:CC'BW}\in\mathcal{P}} \Big|\alpha_0 \langle 0|_W \langle\sigma_0|_{ACC'B} \left|\Phi^+\right\rangle_{CC'} V_x^\dagger(|1\rangle_D |\zeta\rangle_G) \\
&\qquad + \alpha_1 \langle 1|_W \langle\sigma_1|_{ACC'B} \mathrm{SWAP}_{AC'} \left|\Phi^+\right\rangle_{CC'} V_x^\dagger(|1\rangle_D |\zeta\rangle_G)\Big| \\
\leq&\max_{|\zeta\rangle,|\sigma\rangle_{A:CC'BW}\in\mathcal{P}} \left|\langle 0|_W \langle\sigma_0|_{ACC'B} \left|\Phi^+\right\rangle_{CC'} V_x^\dagger(|1\rangle_D |\zeta\rangle_G)\right| \\
&\quad + \max_{|\zeta\rangle,|\sigma\rangle_{A:CC'BW}\in\mathcal{P}} \left|\langle 1|_W \langle\sigma_1|_{ACC'B} \mathrm{SWAP}_{AC'} \left|\Phi^+\right\rangle_{CC'} V_x^\dagger(|1\rangle_D |\zeta\rangle_G)\right| \\
\leq&\max_{|\phi\rangle_A,|\psi\rangle_B,|\zeta\rangle_G} |\langle 0|_W \otimes \langle\phi|_A \otimes \langle\psi|_B V_x^\dagger |1\rangle_D |\zeta\rangle_G| + \max_{\sigma_{AC}\in\mathcal{S}} \sqrt{F\left(\sigma_{AC}, \Phi^+_{AC}\right)} \\
\leq&\sqrt{\varepsilon} + \frac{1}{\sqrt{2}}.
\end{aligned}$$

The first inequality follows from the triangle inequality and the fact that $|\alpha_i| \leq 1$ since $|\alpha_i|^2 \leq 1$. The second inequality follows from the monotonicity of fidelity under the tracing out of registers (C and C' for the first term and C', B, and W for the second term). Restricting the optimization in the first term to be over pure product states follows by another simple convexity argument. Performing the optimization in the second term over all separable states can achieve only the same or a higher value for the fidelity since $\mathcal{P} \in \mathcal{S}$. The final inequality follows from (42) and from the fact that the maximum fidelity of a separable state with $|\Phi^+\rangle$ is $\frac{1}{2}$ [Wat04]. The analysis

above easily generalizes to mixed states by observing that the maximization in (41) can be performed over mixed states and the rest of the analysis can proceed with the purification.

Using the Fuchs-van-de-Graaf inequality in (5), we can bound the trace distance as

$$\min_{\omega,\sigma_{A:CC'BW}\in\mathcal{P}} \left\| U(\omega \otimes |0\rangle\langle 0|^{\otimes m})U^\dagger - \sigma \right\|_1 \geq \frac{\sqrt{2}-1}{\sqrt{2}} - \sqrt{\varepsilon},$$

and so for a sufficiently small ε there is an appropriate gap between the completeness and soundness errors. Thus, we have shown that under the stated promise QPROD-ISOMETRY is both contained in QMA(2) as well as QMA(2)-hard, and thus QMA(2)-complete. ■

Corollary 23 *QSEP-ISOMETRY*(δ_c, δ_s) *is* *QMA*(2)*-complete if there are polynomial-time computable functions* $\delta_c, \delta_s : \mathbb{N} \to [0,1]$ *such that the difference* $\frac{11\delta_s^2}{2048} - 8\delta_c$ *is larger than an inverse polynomial in the circuit size.*

Proof. The main reason that this result follows easily is that allowing classical communication between the provers does not change QMA(2). QSEP-ISOMETRY is defined identically to QPROD-ISOMETRY except that the minimizations in conditions (39) and (40) are over the set of separable states rather than the set of product states. To see that this problem is also in QMA(2), note that the analysis for the inclusion of QPROD-ISOMETRY in a YES instance applies to QSEP-ISOMETRY as well, since Lemma 27 holds for separable states. The analysis of the NO instance proceeds identically.

Indeed, QSEP-ISOMETRY is QMA(2)-hard by means of a very similar reduction as we used for QPROD-ISOMETRY. We proved that in a YES instance there is an input so that the output state close to product (and thus close to separable), while in a NO instance, a very similar analysis demonstrates that all inputs will lead to an output state that is far from separable. ■

4.6 QPROD-STATE is QSZK-complete

In this section, we examine QPROD-STATE, which extends QPROD-PURE-STATE to allow for quantum circuits that output mixed states (that is, the circuit outputs a reference system and another one but traces over the reference system). This addition of a reference system thwarts our ability to use the product test since (as noted in Section 4.1) the product test fails very quickly on mixed state inputs. In contrast to the BQP-complete pure-state version and the QMA(2)-complete isometry version, we show that this problem is QSZK-complete. This new result also leads to the rather surprising conclusion that $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}$—even though it is stated with respect to the 1-LOCC distance—is at least as hard as QPROD-STATE despite the fact that QPROD-ISOMETRY is harder than $\mathsf{QSEP\text{-}ISOMETRY}_{1,1\text{-LOCC}}$ (if there is a strict separation between QMA and QMA(2)).

Problem 24 (QPROD-STATE(δ_c, δ_s)) *Given is a quantum circuit U to generate the state $|\psi\rangle_{RC}$, along with a labeling of the qubits in the reference system R and the output qubits for each system $A_1, A_2, \ldots, A_l \in C$. We define the n-qubit state $\rho_C = \mathrm{Tr}_R\{|\psi\rangle\langle\psi|_{RC}\}$. Decide whether*

1. *Yes: There exists a product state such that ρ_C is δ_c-close in trace distance to it:*

$$\min_{\sigma_C \in \mathcal{P}} \|\rho_C - \sigma_C\|_1 \leq \delta_c. \tag{45}$$

2. *No: Every product state is at least δ_s-far in trace distance to ρ_C:*

$$\min_{\sigma_C \in \mathcal{P}} \|\rho_C - \sigma_C\|_1 \geq \delta_s. \tag{46}$$

Theorem 25 *QPROD-STATE(δ_c, δ_s) is QSZK-complete if there are polynomial-time computable functions $\delta_c, \delta_s : \mathbb{N} \to [0,1]$ such that the difference $\delta_s^2/4 - \delta_c$ is greater than an inverse polynomial in the circuit size.*

Proof. We begin by giving a quantum statistical zero-knowledge proof system to decide QPROD-STATE. Recall that a product state is of the form $\rho_C = \rho_{A_1}^1 \otimes \cdots \otimes \rho_{A_l}^l$.

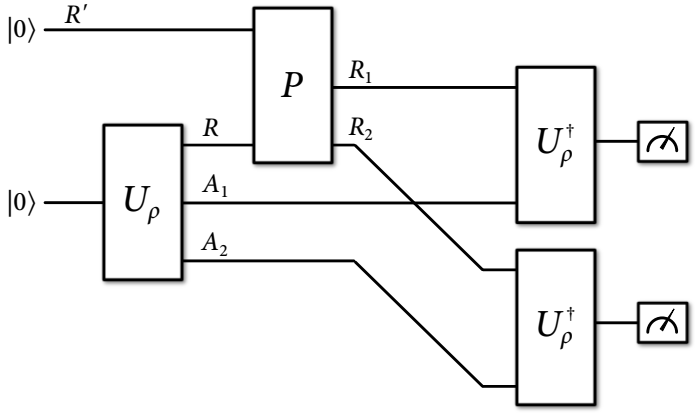

Figure 9: Our QSZK proof system for deciding QPROD-STATE. The figure depicts the case in which the task is to decide if a general bipartite mixed state is product or not, but this easily extends to l-partite states (see the main text). The proof system begins with the verifier sending the reference system R to the prover, who should be able to transform it into two separate purifications of each system of the original state. The verifier then performs the inverse of the original circuit on each pair R_1A_1 and R_2A_2 and measures to verify that the purifications sent by the prover are of the proper form. If the measurement outcomes are all zeros, then the verifier accepts that the original state is close to a product state and otherwise rejects.

In the case that a state on C is exactly product, by Uhlmann's theorem there exists an isometry that the prover can perform on the purifying system R to separate the purifications of each of the subsystems. Indeed, there exists some unitary $P_{RR'\to R_1\ldots R_l}$ acting on R and the prover's system R' such that

$$(P_{RR'\to R_1\ldots R_l}\otimes I_C)\left|\psi\right\rangle_{RC}\left|0\right\rangle_{R'} = \left|\psi\right\rangle_{R_1A_1}\otimes\left|\psi\right\rangle_{R_2A_2}\otimes\cdots\otimes\left|\psi\right\rangle_{R_lA_l},$$

where R_1 purifies A_1 and so on. Since these purifications are arbitrary, we can take them such that $R_1\cong RA_2A_3\ldots A_l$, $R_2\cong RA_1A_3\ldots A_l$ and so on.

For the proof system, the verifier need only send the reference system R to the prover, and if the state is close to product, the prover should be able to provide the purification systems $R_1,\ldots,R_l$ as above. The verifier then performs $U^\dagger$ on each system pair $R_1A_1,\ldots,R_lA_l$ and measures the output, accepting if every measurement outcome is $|0\rangle$. This proof system is depicted in Figure 9. Note also that it is statistical

zero-knowledge because, in the case of a YES instance, the verifier could have simply performed U exactly l times to create l copies of the state $|\psi\rangle_{RC}$.

For the following analysis, we first note a useful fact about product states:

$$\max_{\rho^1 \ldots \rho^l} F(\rho, \rho^1_{A_1} \otimes \cdots \otimes \rho^l_{A_l}) = \max_{P_{RR' \to R_1 \ldots R_l}} \left| \langle \psi |_{R_1 C_1} \otimes \cdots \otimes \langle \psi |_{R_l C_l} P_{RR' \to R_1 \ldots R_l} |\psi\rangle_{RC} |0\rangle_{R'} \right|^2 . \quad (47)$$

In the case of a YES instance, the fidelity is guaranteed to be at least $(1-\delta_c/2)^2 \geq 1-\delta_c$ by the condition in (45) and the Fuchs-van-de-Graaf equality in (5). Thus, the prover can perform the $P_{RR' \to R_1 \ldots R_l}$ that achieves this maximum. This gives probability at least $1 - \delta_c$ of accepting (the verifier should perform the l inverse unitaries and accept if he measures the all zero state on the output qubits).

In the case of a NO instance, the fidelity in (47) is no larger than $1 - \delta_s^2/4$ by (46). Thus, it is impossible for the prover to perform any unitary $P_{RR' \to R_1 \ldots R_l}$ that convinces the verifier to accept with probability greater than $1 - \delta_s^2/4$. So, for an inverse polynomial gap $\delta_s^2/4 - \delta_c$, there exists a QSZK proof system that decides QPROD-STATE(δ_c, δ_s).

To show QSZK-hardness, we can adapt the reduction used for QSEP-STATE$_{1,1\text{-LOCC}}$ in [HMW13] by modifying it slightly to reduce co-QSD to QPROD-STATE. Recall that for co-QSD [Wat02], we are given a description of circuits U_{ρ_0} and U_{ρ_1} that generate mixed states ρ_0 and ρ_1 on the system S as well as a reference system R that is traced out, and we are promised that either $\|\rho_0 - \rho_1\|_1 \leq \varepsilon$ in a YES instance or that $\|\rho_0 - \rho_1\|_1 \geq 2 - \varepsilon$ in a NO instance. As in the QSEP-STATE$_{1,1\text{-LOCC}}$ reduction, let

$$|\psi_{\rho_i}\rangle_{RS} \equiv U_{\rho_i} |0\rangle ,$$

for $i \in \{0, 1\}$ so that

$$\rho_i = \mathrm{Tr}_R \left\{ |\psi_{\rho_i}\rangle \langle \psi_{\rho_i}| \right\} .$$

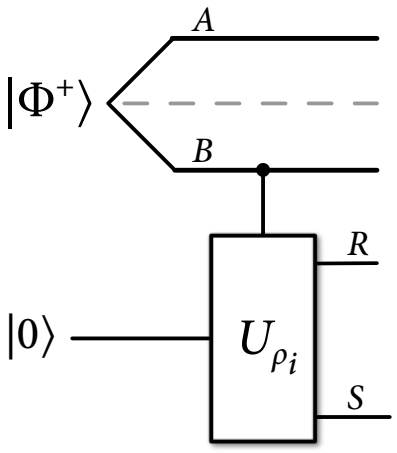

Figure 10: Given respective circuit descriptions U_{ρ_0} and U_{ρ_1} for generating the states ρ_0 and ρ_1 on the output system S, one can compute a description for the above circuit in polynomial time. This serves both as a reduction from QUANTUM-STATE-DISTINGUISHABILITY to $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}$ and from co-QUANTUM-STATE-DISTINGUISHABILITY to QPROD-STATE by tracing out different systems in each case.

From the description of the circuits U_{ρ_0} and U_{ρ_1}, one can efficiently generate a description of the circuit in Figure 10 which takes as input a Bell state across the AB systems and performs the controlled unitary

$$|0\rangle\langle 0|_B \otimes U_{\rho_0} + |1\rangle\langle 1|_B \otimes U_{\rho_1},$$

to generate the state

$$|\varphi\rangle_{ABRS} \equiv \frac{1}{\sqrt{2}}(|00\rangle_{AB}|\psi_{\rho_0}\rangle_{RS} + |11\rangle_{AB}|\psi_{\rho_1}\rangle_{RS}).$$

The output qubits are divided into three sets: the qubits in the systems BR that are traced out, the half of a Bell state on system A, and one of the states ρ_0 or ρ_1 on system S. The resulting state after tracing out BR is

$$\begin{aligned}\omega_{A:S} &\equiv \mathrm{Tr}_{BR}\left\{|\varphi\rangle\langle\varphi|_{ABRS}\right\} \\ &= \frac{1}{2}\left(|0\rangle\langle 0| \otimes \rho_0 + |1\rangle\langle 1| \otimes \rho_1\right).\end{aligned}$$

In a YES instance, we wish to show that this state is close to product by giving a product state close to $\omega_{A:S}$. To do this, we consider the state

$$\sigma = \frac{1}{2}\left(|0\rangle\langle 0| + |1\rangle\langle 1|\right) \otimes \rho_0,$$

for which the distance to $\omega_{A:S}$ is given by:

$$\begin{aligned}\left\|\omega_{A:S} - \sigma\right\|_1 &= \frac{1}{2}\left\||0\rangle\langle 0| \otimes \rho_0 + |1\rangle\langle 1| \otimes \rho_1 - |0\rangle\langle 0| \otimes \rho_0 - |1\rangle\langle 1| \otimes \rho_0\right\|_1 \\ &= \frac{1}{2}\left\||1\rangle\langle 1| \otimes \rho_1 - |1\rangle\langle 1| \otimes \rho_0\right\|_1 \\ &= \frac{1}{2}\left\|\rho_0 - \rho_1\right\|_1 \\ &\leq \frac{\varepsilon}{2}.\end{aligned}$$

Thus, in a YES instance of co-QSD, our reduction results in a YES instance of QPROD-STATE with $\delta_c = \frac{\varepsilon}{2}$.

In a NO instance, we must show that $\omega_{A:S}$ is far from any product state. Recall that trace distance is equal to the maximum probability of distinguishing states over all possible measurements [Fuc96], so we can lower bound the distance to the nearest product state by considering a particular protocol to distinguish $\omega_{A:S}$ from any product state. In this protocol, we begin by measuring the first qubit in the computational basis and by performing the Helstrom measurement $\{\Pi_0, \Pi_1\}$ on the second qubit, storing the two measurement outcomes in classical registers.

It is straightforward to calculate the state $\omega'_{A:S}$ that results after applying the protocol above to the state $\omega_{A:S}$:

$$\omega'_{A:S} = \frac{1}{2}\mathrm{Tr}\{\Pi_0\rho_0\}\left|00\right\rangle\left\langle 00\right| + \frac{1}{2}\mathrm{Tr}\{\Pi_1\rho_1\}\left|11\right\rangle\left\langle 11\right| + \frac{1}{2}\mathrm{Tr}\{\Pi_0\rho_1\}\left|10\right\rangle\left\langle 10\right| + \frac{1}{2}\mathrm{Tr}\{\Pi_1\rho_0\}\left|01\right\rangle\left\langle 01\right|. \quad (48)$$

Recall that the Helstrom measurement distinguishes two states ρ_0 and ρ_1 with the following success probability:

$$\frac{1}{2}\mathrm{Tr}\{\Pi_0\rho_0\} + \frac{1}{2}\mathrm{Tr}\{\Pi_1\rho_1\} = \frac{1}{2}\left(1 + \frac{1}{2}\left\|\rho_0 - \rho_1\right\|_1\right),$$

and the following error probability:

$$\frac{1}{2}\mathrm{Tr}\{\Pi_0\rho_1\} + \frac{1}{2}\mathrm{Tr}\{\Pi_1\rho_0\} = \frac{1}{2}\left(1 - \frac{1}{2}\left\|\rho_0 - \rho_1\right\|_1\right).$$

Using this fact, it is straightforward to establish that the trace distance between $\omega'_{A:S}$ and the perfectly correlated state $\overline{\Phi}_{A:S}$, defined as

$$\overline{\Phi}_{A:S} \equiv \frac{1}{2}(\left|00\right\rangle\left\langle 00\right| + \left|11\right\rangle\left\langle 11\right|),$$

is no larger than

$$1 - \frac{1}{2}\left\|\rho_0 - \rho_1\right\|_1 \leq \frac{\varepsilon}{2}.$$

In the case of a product state, the two measurement outcomes must be uncorrelated, and so we can write the result of applying the above protocol to any product state using the probability p of measuring $\left|0\right\rangle\left\langle 0\right|$ and the probability q of measuring Π_0:

$$\sigma_{p,q} = pq\left|00\right\rangle\left\langle 00\right| + p(1-q)\left|01\right\rangle\left\langle 01\right| + q(1-p)\left|10\right\rangle\left\langle 10\right| + (1-p)(1-q)\left|11\right\rangle\left\langle 11\right|.$$

From the monotonocity of trace distance under quantum operations, it follows that

$$\min_{\sigma_0,\sigma_1}\left\|\sigma_0 \otimes \sigma_1 - \omega_{A:S}\right\|_1 \geq \min_{p,q}\left\|\sigma_{p,q} - \omega'_{A:S}\right\|_1$$

Due to symmetry, we can take $p \leq \frac{1}{2}$ without loss of generality. We can then bound the minimum distance of $\sigma_{p,q}$ to $\omega'_{A:S}$:

$$\begin{aligned}
\min_{p,q} \|\sigma_{p,q} - \omega'_{A:S}\|_1 &\geq \min_{p,q} \left\|\sigma_{p,q} - \overline{\Phi}_{A:S}\right\|_1 - \left\|\overline{\Phi}_{A:S} - \omega'_{A:S}\right\|_1 \\
&\geq \left\|\sigma_{p,q} - \overline{\Phi}_{A:S}\right\|_1 - \frac{\varepsilon}{2} \\
&= \left|\frac{1}{2} - pq\right| + \left|\frac{1}{2} - (1-p)(1-q)\right| + |p(1-q)| + |q(1-p)| - \frac{\varepsilon}{2} \\
&= \frac{1}{2} - pq + \left|\frac{1}{2} - (1-p)(1-q)\right| + p(1-q) + q(1-p) - \frac{\varepsilon}{2} \\
&\geq \frac{1}{2} - pq + p(1-q) + q(1-p) - \frac{\varepsilon}{2} \\
&\geq \frac{1}{2} + p(1-q) - \frac{\varepsilon}{2} \\
&\geq \frac{1-\varepsilon}{2},
\end{aligned}$$

where the first line follows from the triangle inequality, and the fourth through last lines follow from the fact that $0 \leq p \leq \frac{1}{2}$ and $0 \leq q \leq 1$. Thus in a NO instance of co-QSD, our reduction results in a NO instance of QPROD-STATE with $\delta_s \geq (1-\varepsilon)/2$.

We have given a QSZK proof system to decide QPROD-STATE, as well as a reduction from the QSZK-hard problem co-QSD. This completes the proof. ■

4.7 Geometric measures of entanglement

Our work has a close connection to several entanglement measures known collectively as the geometric measure of entanglement (see [WG03, CAH13] and references therein). This is also the case with the work in [HM10], and we comment on this connection briefly.

The original definition of the geometric measure of entanglement was for a pure bipartite state $|\psi\rangle_{AB}$ and defined in terms of the following quantity:

$$\max_{|\phi\rangle_A, |\varphi\rangle_B} |\langle \phi \otimes \varphi | \psi \rangle_{AB}|^2. \tag{49}$$

Clearly, this quantity has an operational interpretation as the maximum probability with which the state $|\psi\rangle_{AB}$ would pass a test for being a pure product state. By taking the negative logarithm of this quantity, one recovers an entropic-like quantity that is equal to the geometric measure of entanglement and satisfies a list of desirable requirements that should hold for an entanglement measure. It is straightforward to extend the above definition and any of the ones below to the multipartite case.

If one has a promise that the quantity in (49) is larger or smaller than $1-\varepsilon$ or ε, respectively, (as in our specification of QPROD-PURE-STATE) then the product test and analysis of Harrow and Montanaro [HM10] demonstrate that it is easy to decide which is the case if one has access to a quantum computer. However, this does not directly give an operational intepretation to the quantity in (49). Rather, it is our QSZK proof system for QPROD-STATE that has its maximum acceptance probability equal to the quantity in (49). More generally, this QSZK proof system has its maximum acceptance probability equal to a generalization of the quantity in (49) defined as follows for mixed states:

$$\max_{\sigma_A,\omega_B} F(\rho_{AB}, \sigma_A \otimes \omega_B). \tag{50}$$

As such, it gives a direct operational interpretation to the above quantity.

In Section 4.3.1 demonstrated a QIP(2) proof system which had the following tight upper bound on its maximum acceptance probability:

$$\max_{\sigma_{AB}\in\mathcal{S}} F(\rho_{AB}, \sigma_{AB}), \tag{51}$$

which holds in the limit of large k, where k is the number of systems sent by the prover in a purported k-extension of the state ρ_{AB}. The above quantity is again related to a geometric measure of entanglement defined in prior work (see [CAH13] and references therein). Thus, the QIP(2)-proof system for $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}$ gives an operational interpretation to the quantity in (51) as the maximum probability with which a prover could convince a verifier that a state ρ_{AB} is separable if the verifier sends a purification of ρ_{AB} to the prover and then performs a check on what the prover sends back.

Finally, our work has unveiled and provided operational interpretations for other quantifiers of entanglement that fall within the geometric class. Indeed, the maximum acceptance probability for our proof system for $\mathsf{QSEP\text{-}ISOMETRY}_{1,1\text{-LOCC}}$ is upper bounded by

$$\max_{\rho,\sigma_{AB}\in\mathcal{S}} F(U(\rho_S \otimes |0\rangle\langle 0|)U^\dagger, \sigma_{AB}),$$

again a bound that holds in the large k limit. Clearly, this quantity is related to the so-called "entangling power" of the unitary U [ZZF00], that is, its ability to take a product state input to an entangled output no matter what the input is. Furthermore, the proof system for $\mathsf{QSEP\text{-}CHANNEL}_{1,1\text{-LOCC}}$ given in Section 4.4 has the following upper bound on its maximum acceptance probability:

$$\max_{\rho,\sigma_{AB}\in\mathcal{S}} F(\mathcal{N}_{S\to AB}(\rho_S), \sigma_{AB}),$$

where $\mathcal{N}_{S\to AB}$ is a quantum channel with input system S and output systems AB. Again, this bound holds in the limit of large k. The above measure is related to the entangling capabilities of a quantum channel no matter what the input is, and our proof system for $\mathsf{QSEP\text{-}CHANNEL}_{1,1\text{-LOCC}}$ provides an operational interpretation for the above quantity as well.

5

Conclusion

We have proved that several entanglement or correlation detection problems are complete for BQP, QMA, QIP, QMA(2), and QSZK, as well as giving an entanglement detection promise problem in QIP(2) which is both NP- and QSZK-hard, the first nontrivial example of such. The completeness of these promise problems for a wide range of complexity classes illustrates an important connection between entanglement and quantum computational complexity theory. In hindsight, it is perhaps natural that these problems related to entanglement can capture the expressive power of these classes since entanglement seems to be the most prominent feature which distinguishes classical from quantum computational complexity theory.

It is interesting to note the connection between these problems, and the differences that give rise to problems complete for different quantum interactive proof classes. The differences are sometimes intuitive: a single-prover proof system for QSEP-ISOMETRY would allow unentangled provers to be simulated with a single one, so, under the assumption that QMA is strictly contained in QMA(2), it seems natural that it should not be possible to place QSEP-ISOMETRY in QMA. Some patterns between classes also emerge: it seems as though mixed state separability requires two messages to be added onto a proof system for pure state separability (from BQP to QSZK, and QMA to QIP), to work with the purification of the mixed state (as is the case for both the "state" and "channel" versions of these problems).

Two-message quantum interactive proof systems continue to be somewhat mysterious. Intuitively, $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}$ has the qualities that one would expect for a QIP(2)-complete problem by extrapolating from these results. Despite this, we do not know whether it is QIP(2)-complete or even QMA-hard. However, our work here gives evidence for why $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}$ should not be either QSZK- or QMA-complete—there are other problems very different from it that are complete for these classes (QPROD-STATE and QPROD-ISOMETRY, respectively).

This work can be expanded in a number of directions. A trace-distance version of $\mathsf{QSEP\text{-}CHANNEL}_{1,1\text{-LOCC}}$ may help to understand the relation between QMIP and $\mathsf{QMIP}_{\mathrm{ne}}$, and similarly a trace-distance version of $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}$ may provide further insights. Additionally, it would be worthwhile to characterize the channel version of $\mathsf{QPROD\text{-}STATE}$ in order to map out more of the space of entanglement detection problems. Such an extension may also help to provide a tighter characterization of classes that rely on "unentanglement," such as $\mathsf{QMA}(2)$.

It is satisfying that each of the entanglement detection problems, with the exception of $\mathsf{QSEP\text{-}STATE}_{1,1\text{-LOCC}}$, is complete for a different complexity class. Perhaps by visiting the remaining related problems in terms of the trace distance and mixed product state cases, one may find two different types of entanglement detection problems that are reducible to each other.

A

Approximate k-extendibility

The following proposition applies to l-partite states $\rho_C = \rho_{A_1 \cdots A_l}$ that are approximately k-extendible:

Proposition 26 *Let ρ_C be δ-close to a k-extendible state, in the sense that*

$$\min_{\sigma_C \in \mathcal{E}_k} \|\rho_C - \sigma_C\|_1 \leq \delta, \tag{52}$$

for some $\delta > 0$, where $\mathcal{E}_k$ is the set of k-extendible l-partite states. Then, the following bound holds

$$\|\rho_C - \mathcal{S}\|_{1\text{-}LOCC} \leq \sqrt{\frac{4l^2 \log |C|}{k-l}} + \delta,$$

where the quantity on the left is multipartite 1-LOCC distance (defined in (7)) to the set of fully separable states.

Proof. Let σ'_C be the state that achieves the minimum in (52). Since this state is k-extendible, we have from Theorem 2 of [BaH13] that

$$\min_{\sigma_C \in \mathcal{S}} \|\sigma'_C - \sigma_C\|_{1-\mathrm{LOCC}} \leq \sqrt{\frac{4l^2 \log |C|}{k-l}}. \tag{53}$$

Let σ^*_C be the state achieving the minimum on the left in (53). From the premise of the theorem, it follows that

$$\begin{aligned}
\|\sigma'_C - \sigma^*_C\|_{1-\mathrm{LOCC}} + \delta &> \|\sigma'_C - \sigma^*_C\|_{1-\mathrm{LOCC}} + \|\sigma'_C - \rho_C\|_1 \\
&\geq \|\sigma'_C - \sigma^*_C\|_{1-\mathrm{LOCC}} + \|\sigma'_C - \rho_C\|_{1-\mathrm{LOCC}} \\
&\geq \|\sigma^*_C - \rho_C\|_{1-\mathrm{LOCC}} \\
&\geq \min_{\sigma_C \in \mathcal{S}} \|\sigma_C - \rho_C\|_{1-\mathrm{LOCC}}\,.
\end{aligned}$$

Thus,

$$\begin{aligned}\min_{\sigma_C \in \mathcal{S}} \|\sigma_C - \rho_C\|_{1-\mathrm{LOCC}} &< \|\sigma'_C - \sigma^*_C\|_{1-\mathrm{LOCC}} + \delta \\ &\leq \sqrt{\frac{4l^2 \log |C|}{k-l}} + \delta,\end{aligned}$$

which concludes the proof. ■

B

Bounding pure state distance

Lemma 27 *Given that*

$$\min_{\rho,\sigma_{A_1\cdots A_l}\in\mathcal{S}} \left\| U\left(\rho_S \otimes |0\rangle\langle 0|\right) U^\dagger - \sigma_{A_1\cdots A_l} \right\|_1 \leq \delta_c,$$

there exists a pure state $|\psi\rangle_s$ *and a product state* $|\phi_1\rangle \otimes \cdots \otimes |\phi_l\rangle$ *such that*

$$\left\| U\left(|\psi\rangle\langle\psi|_S \otimes |0\rangle\langle 0|^{\otimes m}\right) U^\dagger - |\phi_1\rangle\langle\phi_1|_{A_1} \otimes \cdots \otimes |\phi_l\rangle\langle\phi_l|_{A_l} \right\|_1 \leq 4\sqrt{\delta_c}. \tag{54}$$

Proof. Since $\sigma_{A_1\cdots A_l}$ is a separable state, it can be written in the following form:

$$\sigma_{A_1\cdots A_l} = \sum_x p_X(x)\, |\phi_1^x\rangle\langle\phi_1^x|_{A_1} \otimes \cdots \otimes |\phi_l^x\rangle\langle\phi_l^x|_{A_l}.$$

Thus, a particular purification of $\sigma_{A_1\cdots A_l}$ is the following state:

$$|\zeta\rangle_{RA_1\cdots A_l} \equiv \sum_x \sqrt{p_X(x)}\, |x\rangle_R \otimes |\phi_1^x\rangle_{A_1} \otimes \cdots \otimes |\phi_l^x\rangle_{A_l}.$$

By Uhlmann's theorem, we then know that there is a purification $|\psi\rangle_{RS}$ of ρ_S such that the following condition holds

$$\left\| U\left(|\psi\rangle\langle\psi|_{RS} \otimes |0\rangle\langle 0|^{\otimes m}\right) U^\dagger - |\zeta\rangle\langle\zeta|_{RA_1\cdots A_l} \right\|_1 \leq 2\sqrt{\delta_c}. \tag{55}$$

We can then write $|\psi\rangle_{RS}$ as follows:

$$\sum_x \sqrt{q(x)}\, |x\rangle_R\, |\psi^x\rangle_S,$$

for some distribution $q(x)$ and states $\{|\psi^x\rangle\}$ (which are not necessarily orthonormal). Applying a dephasing in the basis $\{|x\rangle\}$ to the R system of both states leading to the following inequality:

$$\left\| \sum_x q(x) |x\rangle\langle x|_R \otimes U \left(|\psi^x\rangle\langle\psi^x|_S \otimes |0\rangle\langle 0|^{\otimes m} \right) U^\dagger \right. \\ \left. - \sum_x p_X(x) |x\rangle\langle x|_R \otimes |\phi_1^x\rangle\langle\phi_1^x|_{A_1} \otimes \cdots \otimes |\phi_l^x\rangle\langle\phi_l^x|_{A_l} \right\|_1 \leq 2\sqrt{\delta_c},$$

which follows by applying monotonicity of trace distance under noisy operations to (55). Tracing over the $A_1, \ldots, A_l$ systems then leads to

$$\|q - p_X\|_1 \leq 2\sqrt{\delta_c}.$$

We then find that

$$\begin{aligned} &\left\| \sum_x p_X(x) |x\rangle\langle x|_R \otimes \left[U \left(|\psi^x\rangle\langle\psi^x|_S \otimes |0\rangle\langle 0|^{\otimes m} \right) U^\dagger \right. \right. \\ &\qquad \left. \left. - |\phi_1^x\rangle\langle\phi_1^x|_{A_1} \otimes \cdots \otimes |\phi_l^x\rangle\langle\phi_l^x|_{A_l} \right] \right\|_1 \\ &\leq \left\| \sum_x q(x) |x\rangle\langle x|_R \otimes U \left(|\psi^x\rangle\langle\psi^x|_S \otimes |0\rangle\langle 0|^{\otimes m} \right) U^\dagger \right. \\ &\quad \left. - \sum_x p_X(x) |x\rangle\langle x|_R \otimes U \left(|\psi^x\rangle\langle\psi^x|_S \otimes |0\rangle\langle 0|^{\otimes m} \right) U^\dagger \right\|_1 \\ &\quad + \left\| \sum_x q(x) |x\rangle\langle x|_R \otimes U \left(|\psi^x\rangle\langle\psi^x|_S \otimes |0\rangle\langle 0|^{\otimes m} \right) U^\dagger \right. \\ &\quad \left. - \sum_x p_X(x) |x\rangle\langle x|_R \otimes |\phi_1^x\rangle\langle\phi_1^x|_{A_1} \otimes \cdots \otimes |\phi_l^x\rangle\langle\phi_l^x|_{A_l} \right\|_1 \\ &\leq \|q - p_X\|_1 + 2\sqrt{\delta_c} \\ &\leq 4\sqrt{\delta_c}. \end{aligned}$$

This implies that

$$\sum_x p_X(x) \left\| U\left(|\psi^x\rangle\langle\psi^x|_S \otimes |0\rangle\langle 0|^{\otimes m}\right) U^\dagger - |\phi_1^x\rangle\langle\phi_1^x|_{A_1} \otimes \cdots \otimes |\phi_l^x\rangle\langle\phi_l^x|_{A_l} \right\|_1 \leq 4\sqrt{\delta_c}.$$

from which we can conclude that the inequality is satisfied for at least one choice of $|\psi^x\rangle$, $|\phi_1^x\rangle \ldots |\phi_l^x\rangle$, implying (54). ■

Bibliography

[Aar13] Scott Aaronson. *Quantum Computing since Democritus*. Cambridge University Press, March 2013.

[ABD+09] Scott Aaronson, Salman Beigi, Andrew Drucker, Bill Fefferman, and Peter Shor. The power of unentanglement. *Theory of Computing*, 5(1):1–42, 2009. arXiv:0804.0802.

[ADHW09] Anura Abeyesinghe, Igor Devetak, Patrick Hayden, and Andreas Winter. The mother of all protocols: Restructuring quantum information's family tree. *Proceedings of the Royal Society A*, 465(2108):2537–2563, August 2009. arXiv:quant-ph/0606225.

[AKN98] Dorit Aharonov, Alexei Kitaev, and Noam Nisan. Quantum circuits with mixed states. In *Proceedings of the thirtieth annual ACM Symposium on Theory of Computing*, STOC '98, pages 20–30, New York, NY, USA, May 1998. ACM. arXiv:quant-ph/9806029.

[Bab85] László Babai. Trading group theory for randomness. In *Proceedings of the Seventeenth Annual ACM Symposium on the Theory of Computing*, pages 421–429, 1985.

[BaC12] Fernando G. S. L. Brandão and Matthias Christandl. Detection of multiparticle entanglement: Quantifying the search for symmetric extensions. *Physical Review Letters*, 109:160502, October 2012. arXiv:1105.5720.

[BACS07] Dominic W. Berry, Graeme Ahokas, Richard Cleve, and Barry C. Sanders. Efficient quantum algorithms for simulating sparse hamiltonians. *Communications in Mathematical Physics*, 270(2):359–371, 2007.

[BaH13] Fernando G. S. L. Brandão and Aram W. Harrow. Quantum de Finetti theorems under local measurements with applications. In *Proceedings of the 45th annual ACM Symposium on the Theory of Computing*, pages 861–870. ACM, 2013. arXiv:1210.6367.

[BB84] Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, volume 175. New York, 1984.

[BBC+93] Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K. Wootters. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters*, 70(13):1895–1899, March 1993.

[BBD+97] Adriano Barenco, Andre Berthiaume, David Deutsch, Artur Ekert, Richard Jozsa, and Chiara Macchiavello. Stabilization of quantum computations by symmetrization. *SIAM Journal on Computing*, 26(5):1541–1557, October 1997. arXiv:quant-ph/9604028.

[BCY11a] Fernando G. S. L. Brandão, Matthias Christandl, and Jon Yard. Faithful squashed entanglement. *Communications in Mathematical Physics*, 306:805–830, September 2011. arXiv:1010.1750.

[BCY11b] Fernando G. S. L. Brandão, Matthias Christandl, and Jon Yard. A quasipolynomial-time algorithm for the quantum separability problem. *Proceedings of ACM Symposium on Theory of Computation*, pages 343–351, June 2011. arXiv:1011.2751.

[Bei10] Salman Beigi. NP vs QMA log(2). *Quantum Information and Computation*, 10(1):141–151, January 2010. arXiv:0810.5109.

[Bel64] John Stewart Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1964.

[BM88] László Babai and Shlomo Moran. Arthur-Merlin games: a randomized proof system, and a hierarchy of complexity classes. *Journal of Computer and System Sciences*, 36(2):254–276, 1988.

[Boo12] Adam D. Bookatz. QMA-complete problems. *arXiv preprint arXiv:1212.6312*, 2012.

[BSST99] Charles H. Bennett, Peter W. Shor, John A. Smolin, and Ashish V. Thapliyal. Entanglement-assisted classical capacity of noisy quantum channels. *Physical Review Letters*, 83(15):3081–3084, October 1999. arXiv:quant-ph/9904023.

[BSST02] Charles H. Bennett, Peter W. Shor, John A. Smolin, and Ashish V. Thapliyal. Entanglement-assisted capacity of a quantum channel and the reverse Shannon theorem. *IEEE Transactions on Information Theory*, 48:2637–2655, October 2002. arXiv:quant-ph/0106052.

[BW92] Charles H. Bennett and Stephen J. Wiesner. Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states. *Physical Review Letters*, 69(20):2881–2884, November 1992.

[CAH13] Lin Chen, Martin Aulbach, and Michal Hajdusek. A comparison of old and new definitions of the geometric measure of entanglement. August 2013. arXiv:1308.0806.

[CB97] Richard Cleve and Harry Buhrman. Substituting quantum entanglement for communication. *Physical Review A*, 56(2):1201–1204, 1997.

[CLMW10] Toby S. Cubitt, Debbie Leung, William Matthews, and Andreas Winter. Improving zero-error classical communication with entanglement. *Physical Review Letters*, 104:230503, June 2010. arXiv:0911.5300.

[CN10] Stephen Cook and Phuong Nguyen. *Logical Foundations of Proof Complexity*. Cambridge University Press, 2010.

[Coo71] Stephen Cook. The complexity of theorem proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing*, pages 151–158, 1971.

[CS12] André Chailloux and Or Sattath. The complexity of the separable Hamiltonian problem. In *Proceedings of the 2012 IEEE Conference on Computational Complexity*, pages 32–41, Porto, Portugal, June 2012. arXiv:1111.5247.

[CyTW04] Richard Cleve, Peter Høyer, Ben Toner, and John Watrous. Consequences and limits of nonlocal strategies. *Proceedings of the 19th Annual IEEE Conference on Computational Complexity*, pages 236–249, June 2004. arXiv:quant-ph/0404076.

[Dev05] Igor Devetak. The private classical capacity and quantum capacity of a quantum channel. *IEEE Transactions on Information Theory*, 51:44–55, January 2005. arXiv:quant-ph/0304127.

[DN06] Christopher M. Dawson and Michael A. Nielsen. The Solovay-Kitaev algorithm. *Quantum Information and Computation*, 6(1):81–95, January 2006. arXiv:quant-ph/0505030.

[DPS02] Andrew C. Doherty, Pablo A. Parrilo, and Federico M. Spedalieri. Distinguishing separable and entangled states. *Physical Review Letters*, 88:187904, April 2002. arXiv:quant-ph/0112007.

[DPS04] Andrew C. Doherty, Pablo A. Parrilo, and Federico M. Spedalieri. Complete family of separability criteria. *Physical Review A*, 69:022308, February 2004. arXiv:quant-ph/0308032.

[DPS05] Andrew C. Doherty, Pablo A. Parrilo, and Federico M. Spedalieri. Detecting multipartite entanglement. *Physical Review A*, 71:032333, Mar 2005.

[Eke91] Artur K. Ekert. Quantum cryptography based on Bell's theorem. *Physical Review Letters*, 67(6):661–663, August 1991.

[Fuc96] Christopher A. Fuchs. Distinguishability and accessible information in quantum theory. *arXiv preprint quant-ph/9601020*, 1996.

[FvdG99] Christopher A. Fuchs and Jeroen van de Graaf. Cryptographic distinguishability measures for quantum-mechanical states. *IEEE Transactions on Information Theory*, 45:1216, May 1999. arXiv:quant-ph/9712042.

[Gha10] Sevag Gharibian. Strong NP-hardness of the quantum separability problem. *Quantum Information and Computation*, 10(3):343–360, March 2010. arXiv:0810.4507.

[GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. In *Proceedings of the Seventeenth Annual ACM Symposium on the Theory of Computing*, pages 291–304, 1985.

[GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on computing*, 18(1):186–208, 1989.

[Gur03] Leonid Gurvits. Classical deterministic complexity of Edmonds' problem and quantum entanglement. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, pages 10–19, June 2003. arXiv:quant-ph/0303055.

[Hel69] Carl W. Helstrom. Quantum detection and estimation theory. *Journal of Statistical Physics*, 1:231–252, June 1969.

[Hel76] Carl W. Helstrom. *Quantum Detection and Estimation Theory*. Academic, New York, 1976.

[HHHH09] Ryszard Horodecki, Paweł Horodecki, Michał Horodecki, and Karol Horodecki. Quantum entanglement. *Reviews of Modern Physics*, 81:865–942, June 2009. arXiv:quant-ph/0702225.

[HM10] Aram Harrow and Ashley Montanaro. An efficient test for product states with applications to quantum Merlin-Arthur games. In *51st Annual IEEE Symposium on the Foundations of Computer Science (FOCS)*, pages 633–642, 2010. arXiv:1001.0017.

[HM13] Aram W. Harrow and Ashley Montanaro. Testing product states, quantum Merlin-Arthur games and tensor optimization. *Journal of the ACM*, 60(1):Article No. 3, February 2013. arXiv:1001.0017.

[HMW13] Patrick Hayden, Kevin Milner, and Mark M. Wilde. Two-message quantum interactive proofs and the quantum separability problem. In *Proceedings of the 18th Annual IEEE Conference on Computational Complexity*, pages 156–167, 2013. arXiv:1211.6120.

[Hol72] Alexander S. Holevo. An analog of the theory of statistical decisions in noncommutative theory of probability. *Trudy Moscov Mat. Obsc.*, 26:133–149, 1972. English translation: Trans. Moscow Math Soc. 26, 133–149 (1972).

[Ioa07] Lawrence M. Ioannou. Computational complexity of the quantum separability problem. *Quantum Information and Computation*, 7(4):335–370, May 2007.

[JJUW10] Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. QIP = PSPACE. *Communications of the ACM*, 53(12):102–109, December 2010. arXiv:0905.1300.

[JUW09] Rahul Jain, Sarvagya Upadhyay, and John Watrous. Two-message quantum interactive proofs are in PSPACE. *50th Annual IEEE Symposium on Foundations of Computer Science*, pages 534–543, October 2009. arXiv:0905.1300.

[Kit95] Alexei Kitaev. Quantum measurements and the abelian stabilizer problem. *arXiv preprint*, November 1995. arXiv:quant-ph/9511026.

[Kit99] Alexei Kitaev. Quantum NP. Talk at the 2nd Workshop on Algorithms in Quantum Information Processing, DePaul University, Chicago, January 1999.

[KKR06] Julia Kempe, Alexei Kitaev, and Oded Regev. The complexity of the local hamiltonian problem. *SIAM Journal on Computing*, 35(5):1070–1097, 2006.

[KMY01] Hirotada Kobayashi, Keiji Matsumoto, and Tomoyuki Yamakami. Quantum certificate verification: Single versus multiple quantum certificates. 2001. arXiv:quant-ph/0110006.

[Kni95] Emmanuel Knill. Approximation by quantum circuits. August 1995. arXiv:quant-ph/9508006.

[KNY08] Masaru Kada, Harumichi Nishimura, and Tomoyuki Yamakami. The efficiency of quantum identity testing of multiple states. *Journal of Physics A: Mathematical and Theoretical*, 41(39):395309, October 2008. arXiv:0809.2037.

[KSV02] Alexei Kitaev, Alexander Shen, and Mikhail N. Vyalyi. *Classical and Quantum Computation*. Number 47. American Mathematical Society, 2002.

[KW00] Alexei Kitaev and John Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof systems. In *In Proceedings of the 32nd ACM Symposium on Theory of Computing*, pages 608–617, 2000.

[LW12] Cécilia Lancien and Andreas Winter. Distinguishing multi-partite states by local measurements. June 2012. arXiv:1206.2884.

[MCKB05] Florian Mintert, André R.R. Carvalho, Marek Kuś, and Andreas Buchleitner. Measures and dynamics of entangled states. *Physics Reports*, 415(4):207–259, 2005. arXiv:quant-ph/0505162.

[MGHW13] Kevin Milner, Gus Gutoski, Patrick Hayden, and Mark M. Wilde. Quantum interactive proofs and the complexity of entanglement detection. *arXiv preprint arXiv:1308.5788*, 2013.

[MW05] Chris Marriott and John Watrous. Quantum Arthur-Merlin games. *Computational Complexity*, 14(2):122–152, 2005. arXiv:cs/0506068.

[MWW09] William Matthews, Stephanie Wehner, and Andreas Winter. Distinguishability of quantum states under restricted families of measurements with an application to quantum data hiding. *Communications in Mathematical Physics*, 291(3):813–843, November 2009. arXiv:0810.2327.

[NC00] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[NWZ09] Daniel Nagaj, Pawel Wocjan, and Yong Zhang. Fast amplification of QMA. *Quantum Information & Computation*, 9(11):1053–1068, 2009. arXiv:0904.1549.

[Pap94] Christos H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994. Theorem 16.3.

[Ros09] Bill Rosgen. *Computational Distinguishability of Quantum Channels*. PhD thesis, University of Waterloo, September, 2009. arXiv:0909.3930.

[Ros11] Bill Rosgen. Testing non-isometry is QMA-complete. In *Theory of Quantum Computation, Communication, and Cryptography*, pages 63–76. Springer, 2011.

[RW05] Bill Rosgen and John Watrous. On the hardness of distinguishing mixed-state quantum computations. *Proceedings of the 20th IEEE Conference on Computational Complexity*, pages 344–354, June 2005. arXiv:cs/0407056.

[Sip96] Michael Sipser. *Introduction to the Theory of Computation*. International Thomson Publishing, 1996.

[Sti55] W. F. Stinespring. Positive functions on C*-algebras. *Proceedings of the American Mathematical Society*, 6:211–216, 1955.

[Sto76] Larry J. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):1–22, 1976.

[SV97] Amit Sahai and Salil Vadhan. A complete promise problem for statistical zero-knowledge. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, FOCS '97, pages 448–457, Washington, DC, USA, October 1997. IEEE Computer Society.

[Uhl76] Armin Uhlmann. The "transition probability" in the state space of a *-algebra. *Reports on Mathematical Physics*, 9(2):273–279, 1976.

[VV12] Umesh Vazirani and Thomas Vidick. Fully device independent quantum key distribution. October 2012. arXiv:1210.1810.

[Wat00] John Watrous. Succinct quantum proofs for properties of finite groups. *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pages 537–546, 2000. arXiv:cs/0009002.

[Wat02] John Watrous. Limits on the power of quantum statistical zero-knowledge. *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 459–468, November 2002. arXiv:quant-ph/0202111.

[Wat03] John Watrous. PSPACE has constant-round quantum interactive proof systems. *Theoretical Computer Science*, 292(3):575–588, January 2003.

[Wat04] John Watrous. Lecture 17: LOCC distinguishability of sets of states. Theory of Quantum Information (course lecture notes), 2004.

[Wat09a] John Watrous. Quantum computational complexity. *Encyclopedia of Complexity and System Science*, 2009. arXiv:0804.3401.

[Wat09b] John Watrous. Zero-knowledge against quantum attacks. *SIAM Journal on Computing*, 39(1):25–58, 2009. arXiv:quant-ph/0511020.

[Wer89a] Reinhard F. Werner. An application of Bell's inequalities to a quantum state extension problem. *Letters in Mathematical Physics*, 17:359–363, 1989.

[Wer89b] Reinhard F. Werner. Quantum states with Einstein-Podolsky-Rosen correlations admitting a hidden-variable model. *Physical Review A*, 40:4277–4281, October 1989.

[WG03] Tzu-Chieh Wei and Paul M. Goldbart. Geometric measure of entanglement and applications to bipartite and multipartite quantum states. *Physical Review A*, 68:042307, October 2003. arXiv:quant-ph/0307219.

[Wil11] Mark M. Wilde. *From Classical to Quantum Shannon Theory.* 2011. arXiv:1106.1445.

[Wil13] Mark M. Wilde. *Quantum Information Theory.* Cambridge University Press, 2013.

[ZZF00] Paolo Zanardi, Christof Zalka, and Lara Faoro. Entangling power of quantum evolutions. *Physical Review A*, 62:030301, August 2000. arXiv:quant-ph/0005031.