THEORY OF COMPUTING, Volume 7 (2011), pp. 101–117
www.theoryofcomputing.org

Quantum Interactive Proofs with Short Messages

Salman Beigi*   Peter Shor†   John Watrous‡

Received: March 31, 2010; published: April 5, 2011.

Abstract: This paper considers three variants of quantum interactive proof systems in which short (meaning logarithmic-length) messages are exchanged between the prover and verifier. The first variant is one in which the verifier sends a short message to the prover, and the prover responds with an ordinary, or polynomial-length, message; the second variant is one in which any number of messages can be exchanged, but where the combined length of all the messages is logarithmic; and the third variant is one in which the verifier sends polynomially many random bits to the prover, who responds with a short quantum message.

We prove that in all of these cases the short messages can be eliminated without changing the power of the model, so the first variant has the expressive power of QMA and the second and third variants have the expressive power of BQP. These facts are proved through the use of quantum state tomography, along with the finite quantum de Finetti theorem for the first variant.

ACM Classification: F.1.2, F.1.3

AMS Classification: 68Q12, 68Q10, 81P68

Key words and phrases: quantum interactive proof systems, quantum state tomography, quantum de Finetti theorem, quantum computation

*Research partially supported by NSF under Grant No. PHY-0803371 and by NSA/ARO under Grant No. W911NF-09-1-0442.

†Research supported in part by NSF grant CCF-0829421, and by funds provided by the W. M. Keck Foundation Center for Extreme Quantum Information Theory.

‡Research supported by NSERC, CIFAR, MITACS and QuantumWorks.

DOI: 10.4086/toc.2011.v007a007

1 Introduction

The interactive proof system model extends the notion of efficient proof verification to an interactive setting, where a computationally unrestricted *prover* tries to convince a computationally bounded *verifier* that an input string satisfies a particular fixed property. They have been studied extensively in computational complexity theory since their introduction over 25 years ago [9, 10, 2, 3], and as a result much is known about them. (See [1] and [7], for instance, for further discussions of classical interactive proof systems.)

Quantum interactive proof systems are a natural quantum computational extension of the interactive proof system model, where the prover and verifier can perform quantum computations and exchange quantum information. The expressive power of quantum interactive proofs is no different from classical interactive proofs: it holds that QIP = PSPACE = IP, and therefore any problem having a quantum interactive proof system also has a classical one [14, 19, 22]. However, quantum interactive proof systems may be significantly more efficient than classical interactive proofs in terms of the number of messages they require, as every problem in PSPACE has a quantum interactive proof system requiring just three messages to be exchanged between a prover and verifier [24, 17]. This is not possible classically unless AM = PSPACE, an equality that implies the collapse of the polynomial-time hierarchy [3, 11].

In this paper we consider quantum interactive proof systems in which some of the messages are short, by which we mean that the messages consist of a number of qubits that is logarithmic in the input length. Three particular variants of quantum interactive proofs with short messages are considered. The first variant is one in which the verifier sends a short message to the prover, and the prover responds with an ordinary, or polynomial-length, message. We prove that this model has the expressive power of QMA. The second variant is one in which any number of messages can be exchanged between the prover and verifier, but where the combined length of all the messages is logarithmic. We prove that this model has the expressive power of BQP. The third variant is one in which the verifier sends polynomially many random bits to the prover, who responds with a short quantum message. We prove that this model also has the expressive power of BQP. Thus, in each of these three cases, logarithmic-length messages are effectively worthless and can be removed without changing the power of the model.

It should be noted that classical analogues of these results are immediate. For instance, one can enumerate all logarithmic-length interactions between a classical verifier and prover in polynomial time, and this implies that a classical variant of our second model has the expressive power of BPP. This argument, however, does not work in the quantum case: one would need to iterate over an exponential number of logarithmic-length quantum interactions to be sure that every interaction is closely approximated.

To illustrate the basic idea through which our results are proved, consider the following simplification of our second model: instead of an arbitrary number of messages with logarithmic combined length, there is only one logarithmic-length message, sent by the prover to the verifier. Upon receiving this logarithmic-length message from the prover, the verifier applies a binary measurement $\{P_{\text{acc}}, P_{\text{rej}}\}$ to decide whether to accept or reject. The corresponding complexity class is denoted $\text{QMA}_{\log}$, and was proved to be equal to BQP by Marriott and Watrous [20] using their error reduction method for QMA. Our methods lead to a different proof that $\text{QMA}_{\log} = \text{BQP}$. We first note that the verifier's maximum acceptance probability is the maximum eigenvalue of P_{acc}. Although P_{acc} acts on a logarithmic number of qubits, it is given by a polynomial-size circuit, so one cannot directly compute the matrix representation of P_{acc} in polynomial time. Nevertheless, using quantum process tomography, one can perform the measurement $\{P_{\mathrm{acc}}, P_{\mathrm{rej}}\}$ on polynomially many known states, and then compute an accurate approximation of P_{acc} using the data collected. The matrix representation of P_{acc} has only polynomially many entries, so the approximation may be represented explicitly. One then simply computes the maximum eigenvalue of this approximation. (Instead of considering quantum process tomography on measurements, we actually consider quantum state tomography on the normalized Choi-Jamiołkowski representation of the quantum channel corresponding to the measurement in this paper. The two approaches are equivalent, but the second one unifies the arguments made in the different sections of the paper.)

In addition to quantum state tomography and the Choi-Jamiołkowski representation of quantum channels, the finite quantum de Finetti theorem is a useful tool in this work. Suppose that we are given several copies of an unknown quantum state, and we want to verify that it is close to some given state. Using quantum state tomography on the copies, one can find an approximation of the unknown state and solve the problem. In an adversarial setting represented by a potentially dishonest prover, however, one may not be guaranteed that the states are indeed copies of the same unknown state—and there could even be entanglement among different copies. To overcome these difficulties, we use the finite quantum de Finetti theorem to reduce the problem to the first case.

One possible application of our work is to the design of new quantum algorithms or QMA verification procedures. Although we do not yet have interesting examples, we believe it is possible that an intuition about quantum interactive proof systems with short messages may lead to new problems being shown to be in BQP or QMA, based on characterizations of the sort we prove.

The remainder of this paper has the following organization. Section 2 discusses some of the background information needed for the rest of the paper, including background on the Choi-Jamiołkowski representation of quantum channels, quantum state tomography, the finite quantum de Finetti theorem, and quantum interactive proof systems. Sections 3, 4, and 5 then discuss the three variants of quantum interactive proof systems with short messages described above. A couple of open problems are described in Section 6.

2 Background

We assume the reader is familiar with quantum information and computation, including the basic quantum complexity classes BQP and QMA, simple properties of mixed states, measurements, channels, and so on [16, 21]. The purpose of the present section is to highlight background knowledge on three topics, represented by the three subsections below, that are particularly relevant to this paper. These topics are: the Choi-Jamiołkowski representation of quantum channels, quantum state tomography, and quantum interactive proof systems.

Before discussing these three topics, it is appropriate to mention a few simple points of notation and terminology. Throughout this paper we let $\Sigma = \{0,1\}$ denote the binary alphabet, and for each $k \in \mathbb{N}$ we write $\mathbb{C}(\Sigma^k)$ to denote the finite-dimensional Hilbert space whose standard basis vectors are indexed by Σ^k (i. e., the Hilbert space associated with a k-qubit quantum register). The Dirac notation is used to describe vectors in spaces of this sort.

For a given space $\mathcal{Q} = \mathbb{C}(\Sigma^k)$, we write $\mathrm{L}(\mathcal{Q})$ to denote the space of all linear mappings from $\mathcal{Q}$ to itself, which is associated with the space of all complex matrices with rows and columns indexed by Σ^k

in the usual way. The subsets of this space representing the positive semidefinite operators and density operators on $\mathcal{Q}$ are denoted $\mathrm{Pos}(\mathcal{Q})$ and $\mathrm{D}(\mathcal{Q})$, respectively. A standard inner product on $\mathrm{L}(\mathcal{Q})$ is defined as $\langle X, Y\rangle = \mathrm{Tr}(X^* Y)$ for all $X, Y \in \mathrm{L}(\mathcal{Q})$ (where X^* denotes the adjoint, or conjugate-transpose, of X). The trace norm of an operator $X \in \mathrm{L}(\mathcal{Q})$ is defined as $\|X\|_1 = \mathrm{Tr}\sqrt{X^* X}$, and the spectral (or operator) norm of X is denoted $\|X\|$.

2.1 Quantum channels and the Choi-Jamiołkowski representation

A *quantum channel* from a k-qubit space $\mathcal{Q} = \mathbb{C}(\Sigma^k)$ to an ℓ-qubit space $\mathcal{R} = \mathbb{C}(\Sigma^\ell)$ is a completely positive and trace-preserving linear mapping of the form $\Phi : \mathrm{L}(\mathcal{Q}) \to \mathrm{L}(\mathcal{R})$. (A mapping Φ is completely positive if $\Phi \otimes \mathbb{1}_{\mathrm{L}(\mathcal{S})}$ is positive for every choice of a space $\mathcal{S} = \mathbb{C}(\Sigma^m)$, meaning that it maps positive semidefinite operators to positive semidefinite operators. To say that Φ is trace-preserving means that $\mathrm{Tr}(\Phi(\rho)) = \mathrm{Tr}(\rho)$ for all operators ρ.) We will write $\mathrm{C}(\mathcal{Q}, \mathcal{R})$ to denote the set of all such quantum channels. For any quantum channel $\Phi \in \mathrm{C}(\mathcal{Q}, \mathcal{R})$ one defines the (normalized) Choi-Jamiołkowski representation [15, 4] of Φ as

$$\rho = \frac{1}{2^k} \sum_{y,z \in \Sigma^k} \Phi(|y\rangle\langle z|) \otimes |y\rangle\langle z| \,. \tag{2.1}$$

In other words, this is the $\ell + k$ qubit state that results from applying Φ to one-half of k pairs of qubits in the $|\phi^+\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$ state.

The action of the mapping Φ can be recovered from its normalized Choi-Jamiołkowski representation in the following way that makes use of *post-selection*. Suppose that Q and Q_0 are k-qubit registers and R is an ℓ-qubit register, that the pair $(\mathsf{R}, \mathsf{Q}_0)$ is initialized to the state ρ as defined by Φ in (2.1), and that Q is in an arbitrary quantum state (and is possibly entangled with additional registers not including Q_0 and R). Consider the following procedure:

1. Measure each qubit of Q together with its corresponding qubit in Q_0 with respect to the Bell basis.
2. If every one of these k measurements results in an outcome corresponding to the Bell state $|\phi^+\rangle$, then output "success," else output "failure."

This procedure gives the outcome "success" with probability 4^{-k}, and conditioned on success the register R is precisely as it would be had it resulted from the channel Φ being applied to Q. (The registers Q and Q_0 can safely be discarded if the procedure succeeds.) To see this, assume first that the joint state of $(\mathsf{R}, \mathsf{Q}_0, \mathsf{Q})$ is $\rho \otimes \xi$ before the measurement takes place. Then the (unnormalized) state of R after the measurements are performed, assuming the end result is "success," is

$$\frac{1}{2^{2k}} \sum_{y,y',z,z' \in \Sigma^k} \Phi(|y\rangle\langle z|)\langle y'|y\rangle\langle z|z'\rangle\langle y'|\xi|z'\rangle = \frac{1}{4^k} \sum_{y,z \in \Sigma^k} \Phi\left(|y\rangle\langle y|\, \xi\, |z\rangle\langle z|\right) = \frac{1}{4^k}\Phi(\xi)\,.$$

The probability of success is therefore 4^{-k}, and conditioned on this outcome the process implements the channel Φ. In our applications, k will be logarithmic in the size of the problem, in which case Φ is implemented with an inverse-polynomial probability of success. The fact that this process implements the channel Φ exactly for all density operators ξ implies that it also operates correctly in the case that Q is entangled with additional registers.

2.2 Quantum state tomography

Quantum state tomography is the process by which an approximate description of an unknown quantum state is obtained by measurements on many independent copies of the unknown state. To be more precise, let $\mathcal{Q} = \mathbb{C}(\Sigma^k)$ denote the space corresponding to a k-qubit register, and suppose that $\mathsf{Q}_1, \ldots, \mathsf{Q}_N$ are k-qubit quantum registers independently prepared in an unknown k-qubit state $\rho \in \mathrm{D}(\mathcal{Q})$. The purpose of quantum state tomography is to obtain an explicit description of a k-qubit state that closely approximates ρ.

One way to perform quantum state tomography is through the use of an *information-complete measurement*. A measurement $\{P_a : a \in \Gamma\}$ on k-qubit registers is information-complete if and only if the set $\{P_a : a \in \Gamma\}$ spans the entire 4^k-dimensional space $\mathrm{L}(\mathcal{Q})$. When such a measurement is performed on a k-qubit state ρ, each measurement outcome is obtained with probability

$$p(a) = \langle P_a, \rho \rangle \, .$$

Based on the assumption that $\{P_a : a \in \Gamma\}$ is information-complete, this vector p of probabilities uniquely determines the state ρ. A close approximation of p, which may be obtained by sufficiently many independent measurements, leads to an approximate description of ρ.

The accuracy of an approximation based on the process just described naturally depends on the choice of an information-complete measurement as well as the specific notion of approximation that is considered. Our interest will be on the trace distance $\|\rho - \sigma\|_1$ between the approximation σ and the true state ρ. To describe the "quality" of an information-complete measurement, it is appropriate to describe the specific process that is used to reconstruct ρ from the vector of probabilities p.

For any spanning set $\{P_a : a \in \Gamma\}$ of $\mathrm{L}(\mathcal{Q})$, there exists a set $\{M_a : a \in \Gamma\} \subseteq \mathrm{L}(\mathcal{Q})$ that satisfies

$$\sum_{a \in \Gamma} M_a \langle P_a, X \rangle = X$$

for every $X \in \mathrm{L}(\mathcal{Q})$. (One may find such a set $\{M_a : a \in \Gamma\}$ by solving a system of linear equations.) The set $\{M_a : a \in \Gamma\}$ allows one to reconstruct a given state from the measurement outcome probabilities it generates. This set is uniquely determined when $\{P_a : a \in \Gamma\}$ has exactly 4^k elements (i. e., is a basis), and hereafter we will restrict our attention to this case. Notice that if ρ is a density matrix, the coefficients $p(a) = \langle P_a, \rho \rangle$ form a probability distribution; and if q is a probability vector that represents an approximation to p it holds that

$$\left\| \rho - \sum_{a \in \Gamma} q(a) M_a \right\|_1 = \left\| \sum_{a \in \Gamma} p(a) M_a - \sum_{a \in \Gamma} q(a) M_a \right\|_1 \leq \sum_{a \in \Gamma} |p(a) - q(a)| \, \|M_a\|_1 \leq \|p - q\|_1 \max_{a \in \Gamma} \|M_a\|_1 \, .$$

It is therefore desirable that the maximum trace norm over the set $\{M_a : a \in \Gamma\}$ determined by the measurement $\{P_a : a \in \Gamma\}$ is as small as possible.

There is one additional consideration that is sometimes relevant, which is that the approximation

$$\sum_{a \in \Gamma} q(a) M_a$$

may fail to be positive semidefinite, and therefore fail to represent a valid quantum state. In this situation one can find a quantum state near to the approximation by renormalizing the positive part of the approximation. For the applications of tomography in this paper, however, this issue may safely be disregarded, as non-positive approximations of density operators will still provide valid approximations to the quantities we are interested in.

An example of an information-complete measurement on a single qubit is given by the following matrices:

$$P_0 = \begin{pmatrix} \frac{2+\sqrt{2}}{8} & \frac{1+i}{8} \\ \frac{1-i}{8} & \frac{2-\sqrt{2}}{8} \end{pmatrix}, \qquad P_1 = \begin{pmatrix} \frac{2-\sqrt{2}}{8} & \frac{1-i}{8} \\ \frac{1+i}{8} & \frac{2+\sqrt{2}}{8} \end{pmatrix},$$

$$P_2 = \begin{pmatrix} \frac{2+\sqrt{2}}{8} & \frac{-1-i}{8} \\ \frac{-1+i}{8} & \frac{2-\sqrt{2}}{8} \end{pmatrix}, \qquad P_3 = \begin{pmatrix} \frac{2-\sqrt{2}}{8} & \frac{-1+i}{8} \\ \frac{-1-i}{8} & \frac{2+\sqrt{2}}{8} \end{pmatrix}.$$

The corresponding set $\{M_0, M_1, M_2, M_3\}$ described above is given by

$$M_0 = \begin{pmatrix} \frac{1+\sqrt{2}}{2} & 1+i \\ 1-i & \frac{1-\sqrt{2}}{2} \end{pmatrix}, \qquad M_1 = \begin{pmatrix} \frac{1-\sqrt{2}}{2} & 1-i \\ 1+i & \frac{1+\sqrt{2}}{2} \end{pmatrix},$$

$$M_2 = \begin{pmatrix} \frac{1+\sqrt{2}}{2} & -1-i \\ -1+i & \frac{1-\sqrt{2}}{2} \end{pmatrix}, \quad M_3 = \begin{pmatrix} \frac{1-\sqrt{2}}{2} & -1+i \\ -1-i & \frac{1+\sqrt{2}}{2} \end{pmatrix}.$$

It holds that $\|M_a\|_1 = \sqrt{10} < 4$ for $a \in \Gamma = \{0,1,2,3\}$. The set $\{P_0, P_1, P_2, P_3\}$ is not an optimal information-complete measurement, in the sense that $\max_a \|M_a\|_1$ could be made smaller for an alternate choice of the measurement, but it has the advantage of being simple to describe and can be implemented exactly by a quantum circuit composed of Hadamard, controlled-not, and $\pi/8$-phase gates, and measurement in the standard basis.

An information-complete measurement for k qubits may be obtained by taking tensor products of the above matrices. More specifically, for each $x \in \Gamma^k$, let us define $2^k \times 2^k$ matrices P_x and M_x as

$$P_x = P_{x_1} \otimes \cdots \otimes P_{x_k} \qquad \text{and} \qquad M_x = M_{x_1} \otimes \cdots \otimes M_{x_k}.$$

Then $\{P_x : x \in \Gamma^k\}$ is an information-complete measurement, and its corresponding set is given by $\{M_x : x \in \Gamma^k\}$. By the multiplicativity of the trace norm, it holds that $\|M_x\|_1 = 10^{k/2} < 4^k$ for every k.

Now, let us suppose that ρ is a quantum state on k qubits, and tomography (using the measurements just described) is performed on N copies of ρ. More precisely, the measurement $\{P_x\}$ is performed independently on each of the N copies of ρ, a probability distribution $q : \Gamma^k \to [0,1]$ is taken to be the frequency distribution of the outcomes, and an approximation

$$H = \sum_{x \in \Gamma^k} q(x) M_x$$

to ρ is computed. We require the following bound on the expected accuracy of this approximation. (Of course, nothing can be said in the worst case, as any sequence of measurement outcomes could occur with very small probability in general.)

Lemma 2.1. *For any choice of $\varepsilon > 0$, taking $N \geq 2^{10k}/\varepsilon^3$ will guarantee that with probability at least $1-\varepsilon$, the estimate H satisfies $\|\rho - H\|_1 < \varepsilon$.*

Proof. For any $\delta > 0$, and any fixed choice of $x \in \Gamma^k$, it follows from Hoeffding's inequality that

$$\Pr\left[|q(x) - p(x)| \geq \delta\right] \leq 2\exp\left(-2N\delta^2\right) .$$

By the union bound it follows that

$$\Pr\left[\|q-p\|_1 \geq 4^k\delta\right] \leq \Pr\left[|q(x)-p(x)| \geq \delta \text{ for at least one } x \in \Gamma^k\right] \leq 2^{2k+1}\exp\left(-2N\delta^2\right) .$$

Setting $\delta = \varepsilon/16^k$ and using the inequality $e^{-\alpha} < 1/\alpha$ for all $\alpha > 0$, we have

$$\Pr\left[\|q-p\|_1 \geq \varepsilon/4^k\right] \leq 2^{2k+1}\exp\left(-2^{2k+1}/\varepsilon\right) < \varepsilon .$$

It follows that

$$\Pr[\|\rho - H\|_1 \geq \varepsilon] \leq \Pr[\|q-p\|_1 \geq \varepsilon/4^k] < \varepsilon.$$

□

The notion of *quantum process tomography* has also been considered, where a quantum measurement or channel is approximated through many independent evaluations of an appropriate sort. (See for example [21].) In this paper, however, it is not necessary to consider this sort of tomography as being any different from state tomography. Specifically, we will approximate channels (and measurements, modeled as channels) by evaluating them on maximally entangled states, followed by ordinary quantum state tomography on the normalized Choi-Jamiołkowski representations that result.

2.3 The finite quantum de Finetti theorem

Suppose that $\mathsf{Q}_1, \ldots, \mathsf{Q}_N$ are all k-qubit registers. A state on $(\mathsf{Q}_1, \ldots, \mathsf{Q}_N)$ is said to be *symmetric* if it is invariant under any permutation of its registers. For instance, any product state of the form $\rho^{\otimes N}$ is symmetric, and any convex combination of such states is symmetric as well. Note, however, that there are symmetric states that cannot be written as convex combinations of symmetric product states as above; the pure state

$$|\psi\rangle = \frac{1}{\sqrt{2^k}} \sum_{x \in \Sigma^k} |x\rangle \otimes \cdots \otimes |x\rangle$$

provides one such example. Nevertheless, by tracing out any subsystem of $|\psi\rangle$, the resulting reduced density matrix is in the convex hull of the set of symmetric product states. The following theorem [18, 5] generalizes this observation.

Theorem 2.2 (The finite quantum de Finetti theorem)**.** *Let ρ_{N+m} be a symmetric state over registers $(\mathsf{Q}_1, \ldots, \mathsf{Q}_{N+m})$, and let $\rho_N = \operatorname{Tr}_{\mathsf{Q}_{N+1}\cdots\mathsf{Q}_{N+m}}(\rho_{N+m})$. There exist states $\{\xi_j\}$ and associated probabilities $\{p_j\}$ such that*

$$\left\|\rho_N - \sum_j p_j \xi_j^{\otimes N}\right\|_1 \leq \frac{2^{k+1}N}{N+m} .$$

2.4 Quantum interactive proof systems

Quantum interactive proof systems are a natural quantum analogue of ordinary, classical interactive proof systems, where the prover and verifier may process and exchange quantum information. We will only consider quantum interactive proof systems having an even number of messages in this paper, so for simplicity we will restrict our discussion to this case.

For t being a function of the form $t : \mathbb{N} \to \mathbb{N}$, we define a t-round (or $(2t)$-message) quantum verifier V to be a collection of quantum circuits

$$V = \left\{ V_{x,j} : x \in \Sigma^*,\ 0 \leq j \leq t(|x|) \right\}$$

that can be generated in polynomial-time given x and j. We will generally write t rather than $t(|x|)$ hereafter in this paper, keeping in mind that t might vary with the input length. We assume that the verifier's circuits are composed of standard unitary quantum gates (controlled-not, Hadamard, and $\pi/8$-phase gates, let us say), as well as ancillary and erasure gates. Included in the description of these circuits is a specification of which input and output qubits are to be considered *private memory qubits* and which are considered *message* qubits. The message qubits refer to qubits that are sent to or received from a prover (to be described shortly). The following properties are required of the circuits describing a verifier:

1. For each x, the circuit $V_{x,0}$ takes no input qubits, and the circuit $V_{x,t}$ produces a single output qubit (called the *acceptance qubit*).
2. There exist functions $v_1, v_2, \ldots$ such that $V_{x,j-1}$ outputs $v_j(|x|)$ private memory qubits and $V_{x,j}$ inputs $v_j(|x|)$ private memory qubits for $1 \leq j \leq t$.
3. There exist functions $q_1, q_2, \ldots$ and $r_1, r_2, \ldots$ that specify the number of message qubits the verifier sends to or receives from the prover on each round, for a given input length. More precisely, each circuit $V_{x,j-1}$ outputs $q_j(|x|)$ message qubits and each circuit $V_{x,j}$ inputs $r_j(|x|)$ message qubits, for $1 \leq j \leq t$.

Similar to the function t, we will often omit the argument $|x|$ from the functions v_j, q_j, and r_j for the sake of readability. When it is convenient, we will refer to the message qubits sent from the verifier to the prover as *question qubits* and qubits sent from the prover to the verifier as *response qubits*.

A t-round (or $(2t)$-message) prover is defined in a similar way to a t-round verifier, but no computational restrictions are made. Specifically, a t-round prover is a collection of quantum channels

$$P = \left\{ P_{x,j} : x \in \Sigma^*,\ 1 \leq j \leq t(|x|) \right\}.$$

Again, the input and output qubits of these channels are specified as private memory qubits or message qubits. When a particular prover P is considered to interact with a given verifier V, one naturally assumes that they agree on the number of messages and the number of qubits sent in each message, as suggested by Figure 1. There is no restriction on p_j, the number of private memory qubits used by the prover at the j-th round—but although each p_j could in principle be unbounded, it is not difficult to show that for any choice of a prover, there is an equivalent one giving rise to precisely the same behavior as the first that uses at most a polynomial number of private memory qubits [13].

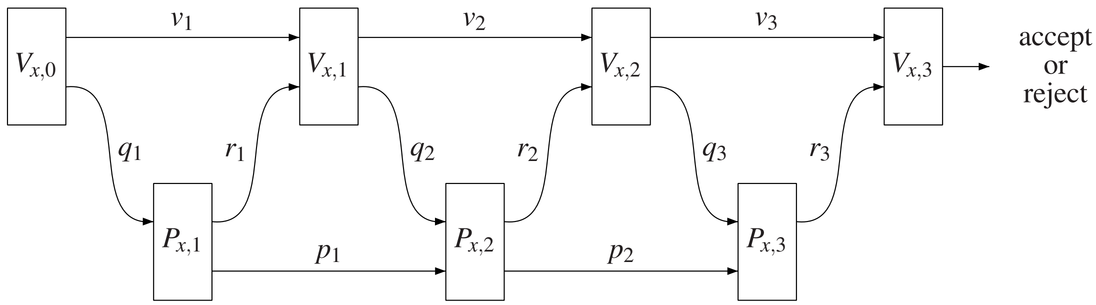

Figure 1: An illustration of an interaction between a prover and verifier in a quantum interactive proof system. In the picture it is assumed that $t = 3$. The labels v_j, p_j, q_j and r_j on the arrows refer to the number of qubits represented by each arrow.

Now, on a given input string x, the prover P and verifier V have an interaction by composing their circuits/channels as described in Figure 1. The *maximum acceptance probability* for a given verifier V on an input x refers to the maximum probability for the circuit $V_{x,t}$ to output 1, assuming it is measured in the standard basis, over all choices of a compatible prover P. It is always the case that a maximal probability is achieved by some prover.

Classes of promise problems may be defined by quantum interactive proof systems in a variety of ways. We will delay the definitions of the classes we consider to the individual sections in which they are discussed.

3 Two-message quantum interactive proofs with short questions

The first specific variant of quantum interactive proof systems we consider are those in which just a single round of communication takes place, with the first message being short (at most logarithmic in length) and the second message being normal (at most polynomial in length). In particular, let us say that a 1-round verifier V is a $[\log, \text{poly}]$ quantum verifier if the number $q = q_1$ of question qubits it sends during the first and only round of communication satisfies $q(n) = O(\log n)$. For functions of the form $a, b : \mathbb{N} \to [0,1]$ we define $\text{QIP}([\log, \text{poly}], a, b)$ to be the class of all promise problems $B = (B_{\text{yes}}, B_{\text{no}})$ for which there exists a $[\log, \text{poly}]$ quantum verifier V with completeness and soundness probability bounds a and b, respectively. In other words, V satisfies the following properties:

1. For every string $x \in B_{\text{yes}}$, there exists a prover P compatible with V that causes V to accept x with probability at least $a(|x|)$.
2. For every string $x \in B_{\text{no}}$, and every prover P compatible with V, it holds that P causes V to accept x with probability at most $b(|x|)$.

For a wide range of choices of a and b, these classes coincide with QMA as the following theorem states.

Theorem 3.1. *Let $a, b : \mathbb{N} \to (0,1)$ be polynomial-time computable functions such that $a(n) - b(n) \geq 1/p(n)$ for some polynomial p. Then* $\text{QIP}([\log, \text{poly}], a, b) = \text{QMA}$.

Proof. It is clear that QMA $\subseteq$ QIP($[\log,\text{poly}], a, b$) for any choice of a and b that satisfy the conditions of the theorem, so our goal is to prove the reverse containment.

Let $B = (B_{\text{yes}}, B_{\text{no}})$ be a promise problem in QIP($[\log,\text{poly}], a, b$), and let V be a $[\log,\text{poly}]$ verifier that witnesses this fact. We write q (as above) to denote the number of question qubits the verifier V sends, and write r to denote the number of response qubits V receives. As V is a $[\log,\text{poly}]$ verifier it holds that $q(n) = O(\log n)$. For a fixed input x, we will write $\mathcal{Q} = \mathbb{C}(\Sigma^q)$ to denote the *question space* and $\mathcal{R} = \mathbb{C}(\Sigma^r)$ to denote the *response space* for V, corresponding to the question and response qubits in the obvious way.

Our goal is to prove that $B \in$ QMA, and to do this we will define a verification procedure (to be referred to as *Arthur*) that demonstrates this fact. Suppose P is a prover that interacts with V. For a fixed input string x, the action of P may be identified with a quantum channel $\Phi \in \mathrm{C}(\mathcal{Q}, \mathcal{R})$, and any such channel defines a quantum state $\rho \in \mathrm{D}(\mathcal{R} \otimes \mathcal{Q})$ according to its normalized Choi-Jamiołkowski representation (2.1). We will define Arthur so that he expects to receive many independent copies of this state. He will check its validity using quantum state tomography, and will use the state to apply the mapping Φ himself through post-selection.

More specifically, we define Arthur so that he performs the following actions:

1. Input $N+m$ registers $(\mathsf{R}_1, \mathsf{Q}_1), \ldots, (\mathsf{R}_{N+m}, \mathsf{Q}_{N+m})$, where N and m are polynomials in the input length n to be specified below.
2. Randomly permute the pairs $(\mathsf{R}_1, \mathsf{Q}_1), \ldots, (\mathsf{R}_{N+m}, \mathsf{Q}_{N+m})$, according to a uniformly chosen permutation $\pi \in S_{N+m}$, and discard all but the first $N+1$ pairs.
3. Perform quantum state tomography on the registers $(\mathsf{Q}_2, \ldots, \mathsf{Q}_{N+1})$, and *reject* if the resulting approximation is not within trace-distance $\delta/2$ of the completely mixed state $\mathbb{1}/2^q$, for δ to be specified below.
4. Simulate the original protocol (P, V) by post-selection using the register pair $(\mathsf{R}_1, \mathsf{Q}_1)$. *Reject* if the post-selection fails, and otherwise *accept* or *reject* as the outcome of the proof system dictates.

To specify N, m and δ, we first set

$$\varepsilon = \frac{1}{p4^{q+1}}$$

for p being the polynomial whose reciprocal separates the completeness and soundness probability bounds a and b. Now set

$$\delta = \frac{\varepsilon^2}{4}, \qquad N = \frac{2^{10q}}{(\delta/2)^3} \qquad \text{and} \qquad m = \frac{2N4^q}{\varepsilon}.$$

Given that q is logarithmic, it holds that N, m, $1/\varepsilon$ and $1/\delta$ are polynomially bounded.

Suppose first that $x \in B_{\text{yes}}$, which implies that there exists a prover P that causes V to accept x with probability at least a. Let Φ denote the quantum channel that describes the behavior of P, and let ρ be the normalized Choi-Jamiołkowski representation of Φ as described in (2.1). Then for each of the register pairs $(\mathsf{R}_j, \mathsf{Q}_j)$ being prepared independently in the state ρ, it holds that Arthur rejects in step 3 with probability at most $\delta/2$ (by Lemma 2.1), and accepts in step 4 with probability at least $a/4^q$ (conditioned on not having rejected in step 3). Arthur therefore accepts with probability at least

$$\left(1 - \frac{\delta}{2}\right)\frac{a}{4^q} > \frac{a}{4^q} - \varepsilon\,.$$

Now let us suppose that $x \in B_{\text{no}}$. We first consider the situation in which the state of the registers $(\mathsf{Q}_1, \ldots, \mathsf{Q}_{N+1})$ at the beginning of step 3 has the form

$$\xi^{\otimes(N+1)}$$

for some density operator $\xi \in \mathrm{D}(\mathcal{Q})$. There are two cases to consider: one is that $\|\xi - \mathbb{1}/2^q\|_1 < \delta$ and the other is that $\|\xi - \mathbb{1}/2^q\|_1 \geq \delta$.

If it is the case that $\|\xi - \mathbb{1}/2^q\|_1 < \delta$, then by the Fuchs-van de Graaf inequalities [6] there must exist a state $\rho \in \mathrm{D}(\mathcal{R} \otimes \mathcal{Q})$ satisfying $\mathrm{Tr}_{\mathcal{R}}(\rho) = \mathbb{1}/2^q$ that is within trace distance ε of the state of $(\mathsf{R}_1, \mathsf{Q}_1)$. To be more precise, consider a fixed purification $|\psi\rangle$ of the state of $(\mathsf{R}_1, \mathsf{Q}_1)$ with an auxiliary register E. As $\xi = \mathrm{Tr}_{\mathcal{R}_1 \otimes \mathcal{E}}(|\psi\rangle\langle\psi|)$ has a high fidelity with $\mathbb{1}/2^q$, and due to the characterization of fidelity in terms of purifications, there exists a purification of $\mathbb{1}/2^q$ over $(\mathsf{R}, \mathsf{Q}, \mathsf{E})$ that has a large inner product with $|\psi\rangle$. Then ρ can be chosen as the reduced density matrix of this pure state over (R, Q). Given that $x \in B_{\text{no}}$, the state ρ would cause acceptance in step 4 with probability at most $b/4^q$, and therefore acceptance may occur in the case at hand with probability at most $b/4^q + \varepsilon$.

If, on the other hand, it holds that $\|\xi - \mathbb{1}/2^q\|_1 \geq \delta$, then rejection must occur in step 3 with probability at least $1 - \delta/2$, so Arthur accepts with probability at most $\delta/2$ (which of course is smaller than $b/4^q + \varepsilon$).

Thus, in both cases, acceptance occurs with probability at most $b/4^q + \varepsilon$. It follows that if the registers $(\mathsf{Q}_1, \ldots, \mathsf{Q}_{N+1})$ are, at the beginning of step 3, in any state of the form

$$\sum_j p_j \xi_j^{\otimes(N+1)} \tag{3.1}$$

(i. e., a convex combination of states of the form just discussed), acceptance may occur with probability at most $b/4^q + \varepsilon$.

Finally, by the finite quantum de Finetti theorem (Theorem 2.2), the state of $(\mathsf{Q}_1, \ldots, \mathsf{Q}_{N+1})$ after step 2 is within trace-distance ε of a state of the form (3.1), and therefore the probability of acceptance is at most $b/4^q + 2\varepsilon$ in the general case.

Given that $a/4^q - \varepsilon$ and $b/4^q + 2\varepsilon$ are efficiently computable and separated by the reciprocal of a polynomial, it holds that B is in QMA as claimed. □

4 Quantum interactive proofs with short interactions

Next we consider quantum interactive proof systems restricted so that the total number of qubits exchanged by the prover and verifier is logarithmic. We prove that any problem having such a quantum interactive proof system is contained in BQP. This fact represents a significant generalization of the equality $\mathrm{QMA}_{\log} = \mathrm{BQP}$ proved in [20]. Like the result of the previous section, our proof of this fact is based on quantum state tomography. In addition we will make use of the *quantum games* framework of [13].

It is clear that any quantum interactive proof system allowing at most a logarithmic number of qubits to be exchanged can be simulated by one in which a logarithmic number of single-qubit messages are permitted, because any number of these messages could consist of meaningless "dummy" qubits that are interspersed with the qubits sent by the other party. To be more precise, let $t(n) = O(\log n)$ and

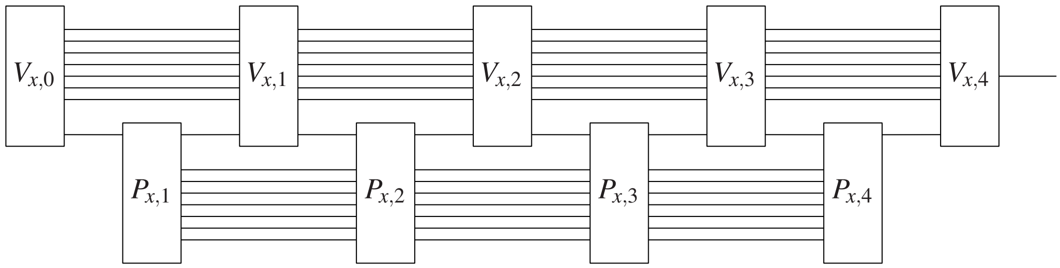

Figure 2: Illustration of a quantum interactive proof in which the messages are single bits.

consider a t-round quantum interactive proof system in which each message consists of a single qubit (i. e., $q_1 = r_1 = \cdots = q_t = r_t = 1$). We will write $\mathrm{QIP}_{\log}(a,b)$ to denote the class of problems having quantum interactive proof systems of this sort having completeness and soundness probability bounds a and b, respectively. As the following theorem states, this model offers no computational advantage over BQP.

Theorem 4.1. *Let $a,b : \mathbb{N} \to (0,1)$ be polynomial-time computable functions such that $a(n) - b(n) \geq 1/p(n)$ for some polynomial p. Then* $\mathrm{QIP}_{\log}(a,b) = \mathrm{BQP}$.

Proof. It is clear that $\mathrm{BQP} \subseteq \mathrm{QIP}_{\log}(a,b)$, and so it remains to prove the reverse containment. To this end let $B = (B_{\text{yes}}, B_{\text{no}})$ be a promise problem in $\mathrm{QIP}_{\log}(a,b)$, and let V be a verifier that witnesses this fact. As above, let $t(n) = O(\log n)$ denote the number of rounds of communication this verifier exchanges with any compatible prover. For a fixed input string x, we will write $\mathcal{Q}_1, \ldots, \mathcal{Q}_t$ to denote copies of the space $\mathbb{C}(\Sigma)$ associated with the t single-qubit messages that V sends to a given prover P, and we will write $\mathcal{R}_1, \ldots, \mathcal{R}_t$ to denote copies of the same space $\mathbb{C}(\Sigma)$ corresponding to the response qubits of P.

The action of V, on a given input string x, is determined by $t+1$ quantum circuits $V_{x,0}, \ldots, V_{x,t}$ as defined in Section 2. Figure 2 illustrates an interaction between V and a prover P for the case that $t = 4$. Now consider the channel Φ obtained from the circuits $V_{x,0}, \ldots, V_{x,t}$ by setting all of the response qubits the verifier receives from the prover as input qubits and setting all of the question qubits sent by the verifier to the prover as output qubits. More precisely, Φ maps states on the space $\mathcal{R}_1 \otimes \cdots \otimes \mathcal{R}_t$ to states on the space $\mathcal{A} \otimes \mathcal{Q}_1 \otimes \cdots \otimes \mathcal{Q}_t$, where $\mathcal{A}$ denotes the single-qubit space associated with the acceptance qubit. Figure 3 illustrates this channel for the protocol pictured in Figure 2.

Next, let

$$\rho = \frac{1}{2^t} \sum_{y,z \in \Sigma^t} \Phi(|y\rangle\langle z|) \otimes |y\rangle\langle z|$$

be the normalized Choi-Jamiołkowski representation of Φ. The state ρ is obviously efficiently preparable given a description of V. By independently preparing $N = 2^{10(2t+1)}/\varepsilon^3$ copies of ρ, for $\varepsilon > 0$ to be specified later, and performing quantum state tomography, one obtains a Hermitian operator H on $\mathcal{A} \otimes \mathcal{R}_1 \otimes \cdots \otimes \mathcal{R}_t \otimes \mathcal{Q}_1 \otimes \cdots \otimes \mathcal{Q}_t$ that satisfies $\|H - \rho\|_1 < \varepsilon$ with probability at least $1 - \varepsilon$. Let us also define

$$\rho_1 = (\langle 1| \otimes \mathbb{1})\, \rho\, (|1\rangle \otimes \mathbb{1}) \qquad \text{and} \qquad H_1 = (\langle 1| \otimes \mathbb{1})\, H\, (|1\rangle \otimes \mathbb{1})$$

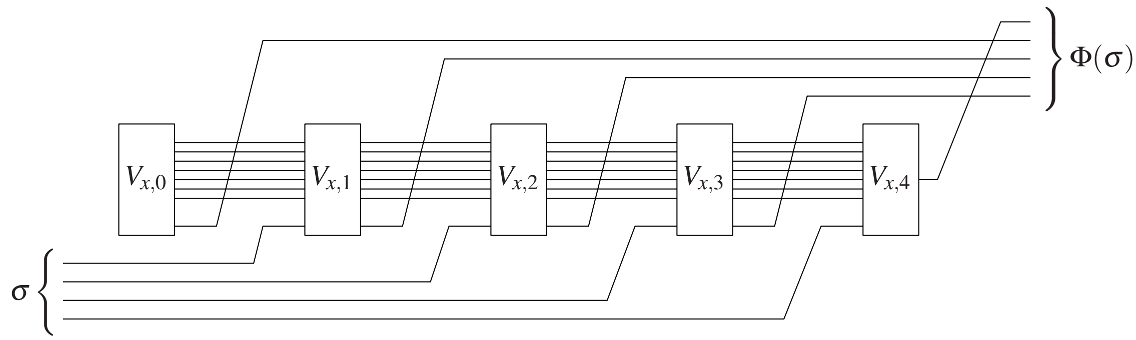

Figure 3: The channel Φ associated with the quantum interactive proof from Figure 2.

to denote the projection of these operators on the subspace in which the qubit $\mathcal{A}$ is $|1\rangle$ (corresponding to accept).

Using the terminology of [13], ρ_1 is a *co-strategy* that describes the verifier's action, and the prover optimizes the acceptance probability corresponding to ρ_1 over all *strategies*. (Strategies are defined in a similar way to co-strategies as above, but with respect to the prover's action.) Given any strategy X of the prover, the acceptance probability is proportional to the inner product of ρ_1 and X. More precisely, the maximum acceptance probability is the optimal value of the following optimization problem:

$$\begin{aligned} \text{maximize:} \quad & 2^t \langle \rho_1, X \rangle \\ \text{subject to:} \quad & X \in \mathcal{S}_t \end{aligned}$$

where $\mathcal{S}_t \subset \mathrm{Pos}(\mathcal{R}_1 \otimes \cdots \otimes \mathcal{R}_t \otimes \mathcal{Q}_1 \otimes \cdots \otimes \mathcal{Q}_t)$ is the space of all strategies. It is shown in [13] that $\mathcal{S}_t$ is characterizes as $\mathcal{S}_0 = 1$ and

$$\mathcal{S}_j = \left\{ X \geq 0 : \mathrm{Tr}_{\mathcal{R}_j}(X) = Y \otimes \mathbb{1}_{\mathcal{Q}_j}, Y \in \mathcal{S}_{j-1} \right\}$$

for $j \geq 1$. This characterization of $\mathcal{S}_t$ turns the above optimization problem to a semidefinite program—at which point we simply replace ρ_1 with its approximation H_1.

It is clear that $\mathrm{Tr}(X) = 2^t$ for every $X \in \mathcal{S}_t$, and therefore

$$\left| 2^t \langle \rho_1, X \rangle - 2^t \langle H_1, X \rangle \right| \leq 2^t \|X\| \, \|\rho_1 - H_1\|_1 \leq 4^t \|\rho_1 - H_1\|_1 \leq 4^t \|\rho - H\|_1$$

for every $X \in \mathcal{S}_t$. By taking

$$\varepsilon = \frac{1}{4^{t+1} p}$$

for instance, one may therefore distinguish the cases $x \in B_{\text{yes}}$ and $x \in B_{\text{no}}$ with probability $1 - \varepsilon$, by approximately solving the semidefinite program described above (using the ellipsoid method [12] for the sake of concreteness). □

We note that precisely the same argument allows one to conclude that *quantum refereed games*, as defined in [13], allowing for at most a logarithmic number of qubits of communication, offer no computational power beyond BQP. In other words, $\mathrm{QRG}_{\log} = \mathrm{BQP}$, for $\mathrm{QRG}_{\log}$ defined appropriately. The details are left to the reader.

5 Two-message quantum interactive proofs with short answers

In light of the results of Section 3, one may ask if two-message quantum interactive proofs with short *answers* (as opposed to short *questions*) have the power of QMA or even BQP. If this is true it is likely to be difficult to show: the graph non-isomorphism problem, which is not known to be in QMA, has a simple and well-known classical protocol [8] requiring polynomial-length questions and constant-length answers. (Indeed, every problem in QSZK has a two-message quantum interactive proof system with a constant-length message from the prover to the verifier, for any choice of constant completeness and soundness errors [23].)

We can show, however, that *public-coin* quantum interactive proofs in which the verifier sends polynomially many random bits to the prover, followed by a logarithmic-length quantum message response from the prover, have only the power of BQP.

Following a similar terminology to the classical case, we refer to a quantum interactive proof system in which the verifier's messages to the prover consist of uniformly-generated random bits as *quantum Arthur–Merlin games*. Let us write $\mathrm{QAM}([\mathrm{poly},\log],a,b)$ to denote the class of promise problems having two-message quantum Arthur–Merlin games with completeness and soundness probability bounds a and b, in which Merlin's response to Arthur has logarithmic length.

Theorem 5.1. *Let $a,b:\mathbb{N}\to(0,1)$ be polynomial-time computable functions such that $a(n)-b(n)\geq 1/p(n)$ for some polynomial-bounded function p. Then* $\mathrm{QAM}([\mathrm{poly},\log],a,b)=\mathrm{BQP}$.

Proof. Assume that B is a promise problem in $\mathrm{QAM}([\mathrm{poly},\log],a,b)$, and consider a choice of Arthur that witnesses this fact. For $r(n)=O(\log n)$, and for any choice of an input string x, Arthur chooses a random string y with length polynomial in $|x|$, and then measures $r=r(|x|)$ qubits sent by Merlin with respect to some binary-valued measurement $\{P_0^{x,y},P_1^{x,y}\}$ that depends on x and y. Thus, assuming that the randomly chosen string is y, the maximum acceptance probability of Arthur is equal to the spectral norm of $P_1^{x,y}$. So, to find $P_1^{x,y}$ (and its norm), we perform quantum state tomography on the normalized Choi-Jamiołkowski representation of the channel

$$\Phi_{x,y}(\sigma)=\left\langle P_0^{x,y},\sigma\right\rangle|0\rangle\langle 0|+\left\langle P_1^{x,y},\sigma\right\rangle|1\rangle\langle 1|\,,$$

which describes Arthur's measurement.

The following algorithm shows that $B\in\mathrm{BQP}$.

1. Choose y uniformly at random (just as Arthur does).
2. Let
$$\varepsilon=\frac{1}{2^{r+3}p}\qquad\text{and}\qquad N=\frac{2^{10(r+1)}}{\varepsilon^3}\,.$$
Prepare N copies of the state ρ, defined to be the normalized Choi-Jamiołkowski representation of $\Phi_{x,y}$, and perform quantum state tomography on ρ. Let H denote the result. Then, by Lemma 2.1, with probability at least $1-\varepsilon$ it holds that $\|\rho-H\|_1\leq\varepsilon$.
3. Compute the value
$$\alpha_y=2^r\,\|(\langle 1|\otimes\mathbb{1})H(|1\rangle\otimes\mathbb{1})\|\,.$$
(It can easily be shown that α_y is an approximation of $\left\|P_1^{x,y}\right\|$.) If $\alpha_y\geq 1$ then *accept*. Otherwise, *accept* with probability α_y and *reject* otherwise.

Because the maximum acceptance probability of Arthur is equal to the expectation value of $\left\|P_1^{x,y}\right\|$ over the random choice of y, and with probability $1-\varepsilon$ the approximation α_y is within distance $2^r\varepsilon$ of $\left\|P_1^{x,y}\right\|$, the above procedure has acceptance probability within $1/(4p)$ of the maximum acceptance probability of Arthur. It follows that $B \in \mathrm{BQP}$. □

6 Open problems

The following two questions have been raised by the anonymous referees. The first one is to find the expressive power of the following model: the verifier and prover send messages to one other consisting of a logarithmic number of qubits in total, and at the end the prover sends a polynomial-length message to the verifier. This model contains QMA, and it might be possible to prove the reverse containment by combining the ideas in Sections 3 and 4. The second question is a similar one involving an interactive protocol in which the verifier always sends polynomial-length public-coin messages to the prover, and the prover replies with quantum messages whose combined length is logarithmic.

References

[1] SANJEEV ARORA AND BOAZ BARAK: *Complexity Theory: A Modern Approach*. Cambridge University Press, 2009. 102

[2] LÁSZLÓ BABAI: Trading group theory for randomness. In *Proc. 17th STOC*, pp. 421–429. ACM Press, 1985. [doi:10.1145/22145.22192] 102

[3] LÁSZLÓ BABAI AND SHLOMO MORAN: Arthur-Merlin games: A randomized proof system, and a hierarchy of complexity classes. *J. Comput. System Sci.*, 36(2):254–276, 1988. [doi:10.1016/0022-0000(88)90028-1] 102

[4] MAN-DUEN CHOI: Completely positive linear maps on complex matrices. *Linear Algebra Appl.*, 10(3):285–290, 1975. [doi:10.1016/0024-3795(75)90075-0] 104

[5] MATTHIAS CHRISTANDL, ROBERT KÖNIG, GRAEME MITCHISON, AND RENATO RENNER: One-and-a-half quantum de Finetti theorems. *Comm. Math. Phys.*, 273(2):473–498, 2007. [doi:10.1007/s00220-007-0189-3] 107

[6] CHRISTOPHER A. FUCHS AND JEROEN VAN DE GRAAF: Cryptographic distinguishability measures for quantum-mechanical states. *IEEE Trans. Inform. Theory*, 45(4):1216–1227, 1999. [doi:10.1109/18.761271] 111

[7] ODED GOLDREICH: *Computational Complexity – A Conceptual Perspective*. Cambridge University Press, 2008. 102

[8] ODED GOLDREICH, SILVIO MICALI, AND AVI WIGDERSON: Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *J. ACM*, 38(1):691–729, 1991. [doi:10.1145/116825.116852] 114

[9] SHAFI GOLDWASSER, SILVIO MICALI, AND CHARLES RACKOFF: The knowledge complexity of interactive proof systems. In *Proc. 17th STOC*, pp. 291–304. ACM Press, 1985. [doi:10.1145/22145.22178] 102

[10] SHAFI GOLDWASSER, SILVIO MICALI, AND CHARLES RACKOFF: The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989. [doi:10.1137/0218012] 102

[11] SHAFI GOLDWASSER AND MICHAEL SIPSER: Private coins versus public coins in interactive proof systems. In S. MICALI, editor, *Randomness and Computation*, volume 5 of *Advances in Computing Research*, pp. 73–90. JAI Press, 1989. 102

[12] MARTIN GRÖTSCHEL, LÁSZLÓ LOVÁSZ, AND ALEXANDER SCHRIJVER: *Geometric Algorithms and Combinatorial Optimization*. Springer–Verlag, second corrected edition, 1993. 113

[13] GUS GUTOSKI AND JOHN WATROUS: Toward a general theory of quantum games. In *Proc. 39th STOC*, pp. 565–574. ACM Press, 2007. [doi:10.1145/1250790.1250873] 108, 111, 113

[14] RAHUL JAIN, ZHENGFENG JI, SARVAGYA UPADHYAY, AND JOHN WATROUS: QIP = PSPACE. In *Proc. 42nd STOC*, pp. 573–581. ACM Press, 2010. [doi:10.1145/1806689.1806768] 102

[15] A. JAMIOŁKOWSKI: Linear transformations which preserve trace and positive semidefiniteness of operators. *Rep. Math. Phys.*, 3(4):275–278, 1972. [doi:10.1016/0034-4877(72)90011-0] 104

[16] A. YU. KITAEV, A. H. SHEN, AND M. N. VYALYI: *Classical and Quantum Computation*. Volume 47 of *Graduate Studies in Mathematics*. American Mathematical Society, 2002. 103

[17] ALEXEI KITAEV AND JOHN WATROUS: Parallelization, amplification, and exponential time simulation of quantum interactive proof system. In *Proc. 32nd STOC*, pp. 608–617. ACM Press, 2000. [doi:10.1145/335305.335387] 102

[18] ROBERT KÖNIG AND RENATO RENNER: A de Finetti representation for finite symmetric quantum states. *J. Math. Phys.*, 46:122108, 2005. [doi:10.1063/1.2146188] 107

[19] CARSTEN LUND, LANCE FORTNOW, HOWARD KARLOFF, AND NOAM NISAN: Algebraic methods for interactive proof systems. *J. ACM*, 39(4):859–868, 1992. [doi:10.1145/146585.146605] 102

[20] CHRIS MARRIOTT AND JOHN WATROUS: Quantum Arthur-Merlin games. *Comput. Complexity*, 14(2):122–152, 2005. [doi:10.1007/s00037-005-0194-x] 102, 111

[21] MICHAEL A. NIELSEN AND ISAAC L. CHUANG: *Quantum Computation and Quantum Information*. Cambridge University Press, 2000. 103, 107

[22] ADI SHAMIR: IP = PSPACE. *J. ACM*, 39(4):869–877, 1992. [doi:10.1145/146585.146609] 102

[23] JOHN WATROUS: Limits on the power of quantum statistical zero-knowledge. In *Proc. 43rd FOCS*, pp. 459–468. IEEE Comp. Soc. Press, 2002. [doi:10.1109/SFCS.2002.1181970] 114

[24] JOHN WATROUS: PSPACE has constant-round quantum interactive proof systems. *Theoret. Comput. Sci.*, 292(3):575–588, 2003. Preliminary version in *Proc. 40th FOCS*, 1999, pp. 112–119. [doi:10.1016/S0304-3975(01)00375-9] 102

AUTHORS

Salman Beigi
School of Mathematics
Institute for Research in Fundamental Sciences (IPM)
Tehran, Iran
salman.beigi@gmail.com

Peter Shor
Department of Mathematics
Massachusetts Institute of Technology
Cambridge, Massachusetts, USA
shor@math.mit.edu
http://www-math.mit.edu/~shor/

John Watrous
Institute for Quantum Computing and School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada
watrous@cs.uwaterloo.ca
http://www.cs.uwaterloo.ca/~watrous/

ABOUT THE AUTHORS

SALMAN BEIGI received his B. S. at Sharif University of Technology, Tehran in 2004. He finished his Ph. D. at the MIT Math Department in 2009 under the direction of Peter Shor, and continued his research as a postdoc at the Institute for Quantum Information at Caltech. He has recently started a new job at the Institute for Research in Fundamental Sciences. His interests include quantum complexity theory, quantum coding theory, photography, and playing *daf,* a traditional Persian musical instrument.

PETER SHOR is a professor at MIT. He is known for his factoring algorithm.

JOHN WATROUS is an associate professor in the School of Computer Science at the University of Waterloo, and is a member of the University of Waterloo's Institute for Quantum Computing. He received his Ph. D. in Computer Science from the University of Wisconsin in 1998, under the supervision of Eric Bach.