Quantum interactive proofs with short messages

Salman Beigi[*] Peter W. Shor[†] John Watrous[‡]

[*] *Institute for Quantum Information*
California Institute of Technology

[†] *Department of Mathematics*
Massachusetts Institute of Technology

[‡] *Institute for Quantum Computing and School of Computer Science*
University of Waterloo

October 22, 2009

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.

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 roughly 25 years ago [GMR85, GMR89, Bab85, BM88], and as a result much is known about them. (See [AB09] and [Gol08], 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 [JJUW09, LFKN92, Sha92]. However, quantum interactive proof systems may be significantly more efficient than classical interactive proofs in terms of the number of messages that are required by their interactions, as every problem in PSPACE has a quantum interactive proof system requiring just three messages to be exchanged between a prover and verifier [KW00]. This is not possible classically unless AM $=$ PSPACE, which implies the collapse of the polynomial-time hierarchy [BM88, GS89].

In this paper, we consider quantum interactive proof systems in which some of the messages are short, by which we mean that they have logarithmic length. In particular, we consider three variants of quantum interactive proofs with short messages. 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.

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, and quantum interactive proof systems. Sections 3, 4, and 5 then discuss the three of quantum interactive proof systems with short messages described above.

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, general measurements, channels, and so on. 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 concerning 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 of vectors 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 such a space.

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 subset of this space representing the density operators on $\mathcal{Q}$ is denoted $\mathrm{D}(\mathcal{Q})$. A standard inner product on $\mathrm{L}(\mathcal{Q})$ is defined as $\langle X, Y \rangle = \mathrm{Tr}(X^* Y)$ for all $X, Y \in \mathrm{L}(\mathcal{Q})$ (and where X^* denotes the adjoint, or conjugate-transpose, of X). The trace norm of $X \in \mathrm{L}(\mathcal{X})$ 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 $\mathcal{R} = \mathbb{C}(\Sigma^l)$ is a completely positive and trace-preserving linear mapping of the form

$$\Phi : \mathrm{L}(\mathcal{Q}) \rightarrow \mathrm{L}(\mathcal{R}).$$

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 of Φ as

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

In other words, this is the $l + k$ 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 (1), and that Q is in an arbitrary quantum state (and is possibly entangled with other registers other than 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. The fact that this process implements the channel Φ exactly for all density operators ξ implies that also operates correctly in the case that Q is entangled with any 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 $\mathsf{X}_1, \ldots, \mathsf{X}_N$ 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 must exist at least one choice of a set $\{M_a : a \in \Gamma\}$ in $\mathrm{L}(\mathcal{Q})$ 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 $\{M_a : a \in \Gamma\}$ 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. If q is an approximation to the vector of probabilities p, it holds that

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

and 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 not represent a 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:

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

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

This is not an optimal information-complete 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. 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 $a \in \Gamma = \{0,1,2,3\}$.

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 $\{M_x : x \in \Gamma^k\}$ is its corresponding inverse measurement set. 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 $q : \Gamma^k \rightarrow [0,1]$ 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 a bound on the 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. However, for any choice of $\varepsilon > 0$, taking $N \geq 2^{12k}/\varepsilon^3$ (for instance) will guarantee that with probability at least $1 - \varepsilon$, the estimate H satisfies $\|\rho - H\|_1 < \varepsilon$. (This bound, which sacrifices accuracy to give a simple expression, can be established by using Chernoff-type bounds on the closeness of the frequency distribution q to the distribution defined by the probabilities $\langle P_x, \rho \rangle$, which are sampled independently by the measurements.)

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. 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 Quantum interactive proofs

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 polynomial-time generated collection of quantum circuits

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

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

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

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 numbers of qubits sent in each message, as suggested by Figure 1. We do not bother to assign any name to the number of private memory qubits kept by the prover because we have no need to refer to this number anywhere in the paper. While the number of private prover qubits could be arbitrarily large, it is not advantageous for the prover to use more than a polynomial number of them.

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.

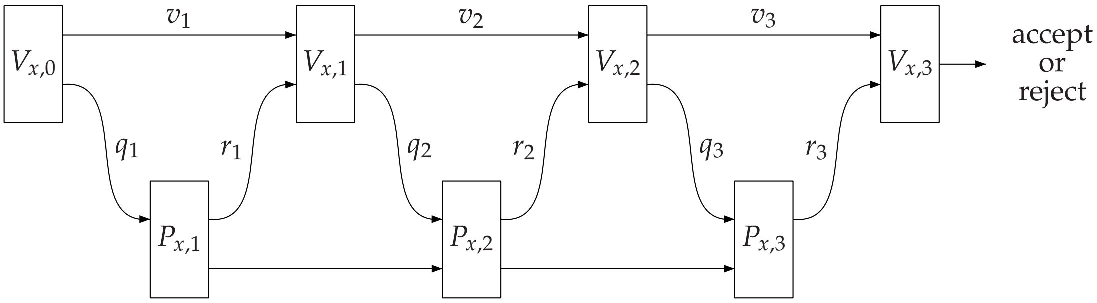

Figure 1: An illustration of an interaction between a prover and verifier in a quantum interactive proof system. In the picture it is assumed that $t = 3$. The labels v_j, q_j and r_j on the arrows refer to the number of qubits represented by each arrow.

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 length) and the second message being normal (at most polynomial length). In particular, let us say that a 1-round verifier V is a [log, poly]-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 $a, b : \mathbb{N} \rightarrow [0, 1]$ we define $\mathrm{QIP}([\log, \mathrm{poly}], a, b)$ to be the class of all promise problems $B = (B_{\text{yes}}, B_{\text{no}})$ for which there exists a [log, 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 that convinces 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 convinces 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 1. *Let $a, b : \mathbb{N} \rightarrow (0, 1)$ be polynomial-time computable functions such that $a(n) - b(n) \geq 1/p(n)$ for some polynomial p. Then* $\mathrm{QIP}([\log, \mathrm{poly}], a, b) = \mathrm{QMA}$.

Proof. It is clear that $\mathrm{QMA} \subseteq \mathrm{QIP}([\log, \mathrm{poly}], a, b)$ for any choice of a and b that satisfies 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 $\mathrm{QIP}([\log, \mathrm{poly}], a, b)$, and let V be a [log, poly] verifier that witnesses this fact. Write $q = q_1$ and $r = r_1$ to denote the number of question qubits the verifier sends and response qubits the verifier V receives, respectively. As V is a [log, 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 $B \in \mathrm{QMA}$, and to do this we will define a verification procedure (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 (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. Receive $N + m$ registers $(\mathsf{R}_1, \mathsf{Q}_1), \ldots, (\mathsf{R}_{N+m}, \mathsf{Q}_{N+m})$ from Merlin, 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 $\delta/2$ 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 protocol dictates.

To specify N and δ, we first set

$$\varepsilon = \frac{1}{4\,p\,2^q}$$

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

$$\delta = \varepsilon^2/4 \qquad \text{and} \qquad N = \frac{2^{12q}}{(\delta/2)^3}.$$

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 (1). Then we may define Merlin so that he prepares each of the pairs $(\mathsf{R}_j, \mathsf{Q}_j)$ independently in the state ρ. Step 3 will reject with probability at most $\delta/2$, and step 4 will accept with probability at least $a/2^q$. Thus, Arthur accepts with probability at least $(1 - \delta/2)a/2^q \geq a/2^q - \delta/2$.

Now let us suppose that $x \in B_{no}$, and consider the state of the registers $(\mathsf{R}_1, \mathsf{Q}_1, \ldots, \mathsf{Q}_{N+1})$ after steps 1 and 2 are performed by Arthur. Let us first assume that these registers are in a state of the form

$$\sigma \otimes \xi^{\otimes N}, \tag{2}$$

where ξ is a q-qubit state and σ is a state of $(\mathsf{R}_1, \mathsf{Q}_1)$ satisfying $\sigma(\mathsf{Q}_1) = \xi$. If it is the case that $\|\xi - \mathbb{1}/2^q\|_1 \geq \delta$, then step 3 results in acceptance with probability at most $\delta/2$. If, on the other hand, $\|\xi - \mathbb{1}/2^q\|_1 < \delta$, then there must exist a state ρ of $(\mathsf{R}_1, \mathsf{Q}_1)$ such that $\rho(\mathsf{Q}_1) = \mathbb{1}/2^q$ and $\|\rho - \sigma\|_1 \leq \varepsilon$. Given that $x \in B_{no}$, the state ρ would cause acceptance in step 4 with probability at most $b/2^q$, and therefore σ causes acceptance with probability at most $b/2^q + \varepsilon$. We therefore have that the probability of acceptance is at most $b/2^q + \varepsilon$ for any state of the form (2).

Now in general, the state of $(\mathsf{R}_1, \mathsf{Q}_1, \ldots, \mathsf{Q}_{N+1})$ after steps 1 and 2 have been performed may not be of the form (2). However, it follows from the finite quantum de Finetti theorem [KR05, CKMR07] that there exists a suitable choice of m, polynomial in δ and 2^q, for which the state of these registers is within trace distance ε of a convex combination of such states. By setting m in this way, it follows that the probability of acceptance is at most $b/2^q + 2\varepsilon$ in the general case.

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

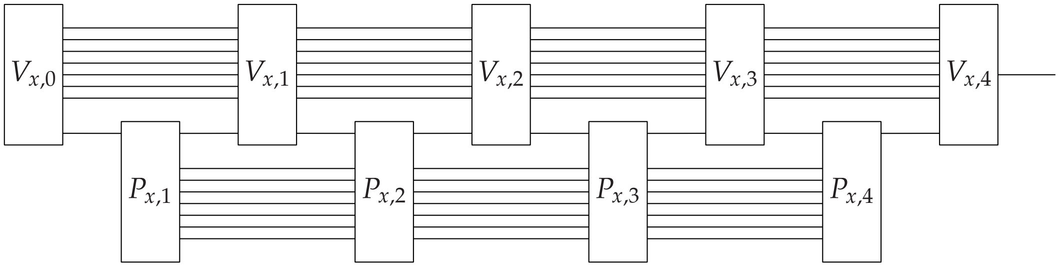

Figure 2: Illustration of a quantum interactive proof in which the messages are single bits.

4 Quantum interactive proofs with only short messages

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 $\text{QMA}_{\log} = \text{BQP}$ proved in [MW05]. 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 [GW07].

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 consider a t-round quantum interactive proof system in which each message consisting of a single qubit (i.e., $q_1 = r_1 = \cdots = q_t = r_t = 1$). We will write $\text{QIP}_{\log}(a, b)$ 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 2. *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 $\text{QIP}_{\log}(a, b) = \text{BQP}$.*

Proof. It is clear that $\text{BQP} \subseteq \text{QIP}_{\log}(a, b)$, 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 $\text{QIP}_{\log}(a, b)$, 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 Hilbert spaces $\mathbb{C}(\Sigma)$ 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. Thus, Φ maps states on the space $\mathcal{R}_1 \otimes \cdots \otimes \mathcal{R}_t$ to states on the space $\mathcal{Q}_1 \otimes \cdots \otimes \mathcal{Q}_t \otimes \mathcal{A}$, where $\mathcal{A}$ denotes the single-qubit space

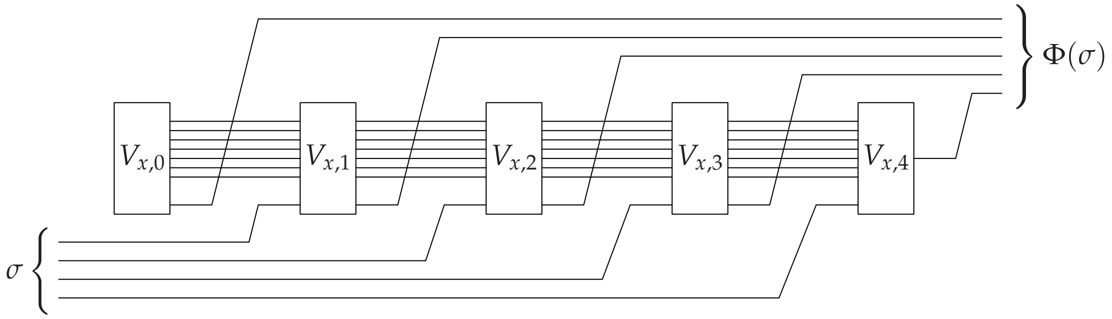

Figure 3: The channel Φ associated with the quantum interactive proof from Figure 2.

associated with the acceptance qubit. Figure 3 illustrates this channel for the protocol pictured in Figure 2.

Next, let

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

This is the normalized Choi–Jamiołkowski representation of Φ (with the input and output spaces in reverse order from the presentation in Section 2.1). The state ρ is obviously efficiently preparable given a description of V. By independently preparing $N = 2^{12t}/\varepsilon^3$ copies of ρ, for $\varepsilon > 0$ to be specified later, and performing quantum state tomography, one obtains a Hermitian operator H on $\mathcal{R}_1 \otimes \cdots \otimes \mathcal{R}_t \otimes \mathcal{Q}_1 \otimes \cdots \otimes \mathcal{Q}_t \otimes \mathcal{A}$ that satisfies $\|H - \rho\|_1 < \varepsilon$ with probability at least $1 - \varepsilon$. Let us also define

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

to denote the portion of these operators that correspond to the acceptance qubit $\mathcal{A}$ giving result 1 (accept).

Now, using the methods of [GW07], it can be shown that the maximum probability with which V can be made to accept is given by the semidefinite program

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

where $\mathcal{S}_t \subset \operatorname{Pos}(\mathcal{R}_1 \otimes \cdots \otimes \mathcal{R}_t \otimes \mathcal{Q}_1 \otimes \cdots \otimes \mathcal{Q}_t)$ is defined recursively as $\mathcal{S}_0 = 1$ and

$$\mathcal{S}_t = \{X \geq 0 \,:\, \operatorname{Tr}_{\mathcal{R}_t}(X) = Y \otimes \mathbb{1}_{\mathcal{Q}_t},\, Y \in \mathcal{S}_{t-1}\}.$$

(In other words, the set of operators S_t represent valid *strategies* for the prover.) It is clear that $\operatorname{Tr}(X) = 2^t$ for every $X \in \mathcal{S}_t$, and therefore

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

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

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

for instance, and substituting H_1 for ρ_1, one may therefore distinguish the cases $x \in A_{\text{yes}}$ and $x \in A_{\text{no}}$ with probability $1 - \varepsilon$ by performing tomography and solving the semidefinite program described above. □

We note that precisely the same argument allows one to conclude that *quantum refereed games*, as defined in [GW07], allowing for at most a logarithmic number of qubits of communication offer no computational power beyond BQP. In other words, $\text{QRG}_{\log} = \text{BQP}$, for $\text{QRG}_{\log}$ 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 a difficult to show: the graph non-isomorphism problem, which is not known to be in QMA, has a simple and well-known classical protocol [GMW91] 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 [Wat02].)

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. (An analogous result for classical interactive proof systems is obvious.)

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 $\text{QAM}([\text{poly}, \log], a, b)$ 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 3. *Let $a, b : \mathbb{N} \to (0, 1)$ be polynomial-time computable functions such that $a - b \geq 1/p$ for some polynomial-bounded function p. Then* $\text{QAM}([\text{poly}, \log], a, b) = \text{BQP}$.

Sketch of proof. Assume B is a promise problem in $\text{QAM}([\text{poly}, \log], a, b)$, 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 send by Merlin with respect to some binary-valued measurement $\{P_0^x, P_1^x\}$ that depends on x. This measurement may of course be described as a quantum channel

$$\Phi_x(\sigma) = \langle P_0^x, \sigma \rangle \, |0\rangle\langle 0| + \langle P_1^x, \sigma \rangle \, |1\rangle\langle 1| \, ,$$

which is easily implemented given a description of Arthur.

Now, to prove $B \in \text{BQP}$, we consider a quantum algorithm that operates as follows:

1. Randomly choose y uniformly (just as Arthur does).
2. Let
$$\varepsilon = \frac{1}{2^{r+2}p} \qquad \text{and} \qquad N = \frac{2^{12r}}{\varepsilon^3}.$$
Prepare N copies of the state ρ, defined to be the normalized Choi–Jamiołkowski representation of Φ_x, and perform quantum state tomography of ρ. Let H denote the result.
3. Compute the value
$$\alpha = 2^r \, \| (\langle 1| \otimes \mathbb{1}) H (|1\rangle \otimes \mathbb{1}) \| \, .$$
If $\alpha \geq 1$ then *accept*. Otherwise, *accept* with probability α and *reject* otherwise.

This procedure has acceptance probability within $1/(4p)$ of the maximum acceptance probability of Arthur, and therefore $A \in \mathrm{BQP}$. □

References

[AB09] S. Arora and B. Barak. *Complexity Theory: A Modern Approach*. Cambridge University Press, 2009.

[Bab85] L. Babai. Trading group theory for randomness. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing*, pages 421–429, 1985.

[BM88] L. Babai and S. 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.

[CKMR07] M. Christandl, R. König, G. Mitchison, and R. Renner. One-and-a-half quantum de Finetti theorems. *Communications in Mathematical Physics*, 273(2):473–498, 2007.

[GMR85] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing*, pages 291–304, 1985.

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

[GMW91] O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM*, 38(1):691–729, 1991.

[Gol08] O. Goldreich. *Computational Complexity – A Conceptual Perspective*. Cambridge University Press, 2008.

[GS89] S. Goldwasser and M. Sipser. Private coins versus public coins in interactive proof systems. In S. Micali, editor, *Randomness and Computation*, volume 5 of *Advances in Computing Research*, pages 73–90. JAI Press, 1989.

[GW07] G. Gutoski and J. Watrous. Toward a general theory of quantum games. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 565–574, 2007.

[JJUW09] R. Jain, Z. Ji, S. Upadhyay, and J. Watrous. QIP = PSPACE. Manuscript, 2009. Available as arXiv.org e-Print 0907.4737.

[KR05] R. König and R. Renner. A de Finetti representation for finite symmetric quantum states. *Journal of Mathematical Physics*, 46:122108, 2005.

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

[LFKN92] C. Lund, L. Fortnow, H. Karloff, and N. Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM*, 39(4):859–868, 1992.

[MW05] C. Marriott and J. Watrous. Quantum Arthur-Merlin games. *Computational Complexity*, 14(2):122–152, 2005.

[Sha92] A. Shamir. IP = PSPACE. *Journal of the ACM*, 39(4):869–877, 1992.

[Wat02] J. Watrous. Limits on the power of quantum statistical zero-knowledge. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 459–468, 2002.