Quantum Interactive Proofs with Competing Provers

Gus Gutoski John Watrous

Department of Computer Science
University of Calgary
Calgary, Alberta, Canada

December 13, 2004

Abstract

This paper studies quantum refereed games, which are quantum interactive proof systems with two competing provers: one that tries to convince the verifier to accept and the other that tries to convince the verifier to reject. We prove that every language having an ordinary quantum interactive proof system also has a quantum refereed game in which the verifier exchanges just one round of messages with each prover. A key part of our proof is the fact that there exists a single quantum measurement that reliably distinguishes between mixed states chosen arbitrarily from disjoint convex sets having large minimal trace distance from one another. We also show how to reduce the probability of error for some classes of quantum refereed games.

1 Introduction

A *refereed game* consists of a conversation between a computationally bounded verifier and two computationally unbounded provers regarding some input string x. The two provers use their unbounded computational power to compete with each other: one prover, called the *yes-prover*, attempts to convince the verifier to accept x, while the other prover, called the *no-prover*, attempts to convince the verifier to reject x. At the end of the interaction, the verifier decides whether to accept or reject the input x, effectively deciding which of the provers wins the game. Such games represent games of incomplete information; the messages exchanged between one prover and the verifier are considered to be hidden from the other player.

A language L is said to have a refereed game with error ε if there is a polynomial-time verifier satisfying the the following conditions. For each string $x \in L$, there exists a yes-prover that can always convince the verifier to accept x with probability at least $1-\varepsilon$, regardless of the no-prover's strategy, and for each $x \notin L$, there exists a no-prover that can always convince the verifier to reject x with probability at least $1-\varepsilon$, regardless of the yes-prover's strategy. A *turn* for one of the provers consists of a message from the verifier to that prover, followed by a response from that prover back

to the verifier. One may consider the case where the provers' turns are played sequentially or in parallel.

The refereed games model is based on the interactive proof system model [11, 2, 3, 4], which has a rich history that we will not survey here. The refereed games model, and variations on this model, were considered in the classical case in Refs. [17, 8, 7, 14, 9, 6], among others. Much of what is known about the complexity-theoretic aspects of the classical refereed games model is due to Feige and Kilian [6]. The class of languages having classical refereed games in which the provers may play any polynomial number of turns coincides with EXP (deterministic time $2^{p(n)}$ for some polynomial p). The simulation of EXP by a polynomial-turn refereed game is due to Feige and Kilian [6], and is based on arithmetization technique developed by Lund, Fortnow, Karloff and Nisan [15] and used in proofs of IP = PSPACE [20, 21]. The containment of this class in EXP is due to Koller and Megiddo [14]. On the other hand, the class of languages having games in which the provers play precisely one turn each, with the turns played in parallel, coincides with PSPACE [6]. Apparently little is known about the expressive power of classical refereed games intermediate between these two extremes. For instance, games with a constant number of prover turns may correspond to PSPACE, EXP, or some complexity class between the two.

Similar to the classical case, quantum refereed games are based on the quantum interactive proof system model [22, 13]. Quantum refereed games differ from classical ones in that the provers and the verifier may perform quantum computations and exchange quantum messages. Our two main motives for considering the quantum refereed games model are to better understand the power of quantum interactive proof systems and to examine the effect of quantum information on the complexity of finding strategies for two-player games.

The main result of this paper establishes that any language having a quantum interactive proof system also has a quantum refereed game with exponentially small probability of error wherein each prover plays just one turn (with the yes-prover playing first). An interesting fact about the resulting game from the point of view of understanding quantum interactive proofs is that entanglement between the provers and the verifier does not play any role in this game, and may without loss of generality be assumed not to exist. More specifically, the game we define has the following general form: the yes-prover sends the verifier a mixed quantum state, the verifier processes this state and sends some state to the no-prover, and the no-prover measures the state and sends a classical result to the verifier. The verifier checks the result of the measurement and accepts or rejects.

A key ingredient for our result is an information-theoretic assertion stating that there exists a quantum measurement that can reliably distinguish between states chosen from two disjoint convex sets of quantum states. This assertion generalizes a well-known fact about the relation between the trace distance between two states and their distinguishability, and may be viewed as a quantitative version, from the point of view of quantum information theory, of the fact from convex analysis that disjoint convex sets are separated by some hyperplane.

The remainder of this paper is organized as follows. We begin by defining quantum refereed games in Sect. 2. In Sect. 3 we prove the fact concerning measurements distinguishing convex sets mentioned previously. Using this fact, we then prove in Sect. 4 that a two-turn quantum refereed game exists for any language L having a quantum interactive proof system. In Sect. 5 we describe a

method for error reduction in two-turn quantum refereed games. The paper concludes with Sect. 6, which mentions some open problems about quantum refereed games.

2 Definitions

In this section we define the quantum refereed games model and some complexity classes based on this model. Throughout the paper we assume all strings are over the alphabet $\Sigma = \{0, 1\}$. For $x \in \Sigma^*$, $|x|$ denotes the length of x. We let $poly$ denote the set of polynomial-time computable functions $f : \mathbb{N} \to \mathbb{N} \setminus \{0\}$ for which there exists a polynomial p such that $f(n) \leq p(n)$ for all n. We also let 2^{-poly} denote the set of polynomial-time computable functions ε such that $\varepsilon(n) = 2^{-f(n)}$ for all n for some $f \in poly$.

The model for quantum computation that provides a basis for quantum refereed games is the quantum circuit model, with which we assume the reader is familiar. As mentioned in Sect. 1, a quantum refereed game has a verifier V and two competing provers Y and N. Each of V, Y, and N is defined by a mapping on input strings $x \in \Sigma^*$ where $V(x)$, $Y(x)$, and $N(x)$ are each sequences of quantum circuits. The circuits in these sequences are assumed to be composed only of gates taken from some universal set of quantum gates. Thus, each of the circuits implements a unitary operation on its input qubits. However, we lose no generality by allowing only unitary operations because arbitrary admissible quantum operations, including measurements, can be simulated by unitary circuits as described in Ref. [1].

For each prover, the qubits upon which that prover's circuits act are partitioned into two sets: one set of qubits is private to that prover and the other is shared with the verifier. These shared qubits act as a quantum channel between the verifier and that prover. No restrictions are placed on the complexity of the provers' circuits, which captures the notion that the provers' computational power is unbounded—each of the provers' circuits can be viewed as an arbitrary unitary operation.

The qubits on which the verifier's circuits act are partitioned into three sets: one set is private to the verifier and two sets are shared with each of the provers. One of the verifier's private qubits is designated as the *output qubit*. At the end of the game, acceptance is dictated by a measurement of the output qubit in the computational basis. We also require that the verifier's sequence of circuits $V(x)$ be generated by a polynomial-time Turing machine on input x. This uniformity constraint captures the notion that the verifier's computational power is limited.

In addition to the verifier and provers, a quantum refereed game consists of a *protocol* that dictates the number and order of turns taken by the provers. The circuits in the verifier's and provers' sequences are applied to the initial state in which each qubit is in state $|0\rangle$ in such a way as to implement the protocol of the game.

The games we study in this paper have the following protocol: a message from the yes-prover to the verifier, a message from the verifier to the no-prover, and a message from the no-prover the the verifier. Quantum refereed games that follow this protocol will be called *short quantum games*. We note that entanglement between the provers and the verifier is immaterial in games of this form—each prover takes only one turn, and thus has no need to remember anything after his turn ends. Thus, when convenient, we may assume that the provers do not have private qubits but instead may perform arbitrary admissible quantum operations (i.e., completely positive trace-

preserving maps) on their message qubits.

We now define the complexity class SQG based on short quantum games of the type just described. For $c, s : \mathbb{N} \rightarrow [0,1]$, the set $\mathsf{SQG}(c,s)$ consists of all languages $L \subseteq \Sigma^*$ for which there exists a verifier V for a short quantum game such that the following conditions hold:

1. There exists a yes-prover Y such that, for all no-provers N and all $x \in L$, $Y(x)$ convinces $V(x)$ to accept x with probability at least $1 - c(|x|)$; and

2. There exists a no-prover N such that, for all yes-provers Y and all $x \notin L$, $N(x)$ convinces $V(x)$ to reject x with probability at least $1 - s(|x|)$.

The functions c and s are called the *completeness error* and *soundness error*, respectively. We define $\mathsf{SQG}\left(2^{-poly}, 2^{-poly}\right)$ to be the set of all languages $L \subseteq \Sigma^*$ such that $L \in \mathsf{SQG}(\varepsilon, \varepsilon)$ for every $\varepsilon \in 2^{-poly}$. Let us also write SQG as shorthand for $\mathsf{SQG}\left(2^{-poly}, 2^{-poly}\right)$.

The class QIP contains all problems having single-prover quantum interactive proof systems as in Ref. [13]. The main complexity-theoretic result of the present paper states that $\mathsf{QIP} \subseteq \mathsf{SQG}$. We prove this result by exhibiting a short quantum game that solves a promise problem called the CLOSE-IMAGES problem, which is known to be complete for QIP [13]. It is convenient for us to use the formulation of this problem based on the one found in Ref. [19].

The promise problem CLOSE-IMAGES is defined for any desired $\varepsilon \in 2^{-poly}$ as follows. Given are descriptions of two mixed state quantum circuits Q_0 and Q_1, which both implement some admissible (i.e., completely positive and trace-preserving) transformation from n qubits to m qubits. The promise is that exactly one of the following conditions holds:

1. There exist n-qubit mixed states ρ_0 and ρ_1 such that $Q_0(\rho_0) = Q_1(\rho_1)$; or

2. For all n-qubit mixed states ρ_0 and ρ_1, the states $Q_0(\rho_0)$ and $Q_1(\rho_1)$ have fidelity squared at most $\varepsilon(n)$.

In other words, the images of Q_0 and Q_1 are either overlapping or are far apart. The goal is to accept when case 1 holds and reject when case 2 holds.

3 Distinguishing Convex Sets of Quantum States

We motivate discussion in this section by pointing out that, for any mixed-state quantum circuit Q, the image $\mathcal{A} = \{Q(\rho) : \rho \text{ a mixed state}\}$ of the admissible transformation associated with Q is a compact, convex set of mixed states. If the images of two circuits Q_0 and Q_1 are far apart, then one could reasonably hope that there is a quantum measurement that reliably distinguishes between outputs $Q_0(\rho_0)$ and $Q_1(\rho_1)$ of these transformations, with the measurement depending only on Q_0 and Q_1, and not on the choice of input states ρ_0 and ρ_1. In this section we prove that indeed there always exists such a measurement. More generally, we prove that given any two disjoint convex sets of mixed quantum states, there exists a single measurement that distinguishes states drawn arbitrarily from one set from the other with success probability determined by the minimal trace

distance between the sets. The short quantum game for the CLOSE-IMAGES problem we define in Sect. 4 relies heavily upon the existence of such a measurement.

Let us first begin with some notation. Given a finite dimensional Hilbert space $\mathcal{H}$, let $\mathbf{L}(\mathcal{H})$ denote the set of all linear operators on $\mathcal{H}$, let $\mathbf{H}(\mathcal{H})$ denote the set of all Hermitian operators on $\mathcal{H}$, let $\mathbf{Pos}(\mathcal{H})$ denote the set of all positive semidefinite operators on $\mathcal{H}$, and let $\mathbf{D}(\mathcal{H})$ denote the set of all density operators (i.e., unit trace positive semidefinite operators) on $\mathcal{H}$. For $A, B \in \mathbf{L}(\mathcal{H})$, define $\langle A, B\rangle = \operatorname{tr} A^\dagger B$. This is an inner product on $\mathbf{L}(\mathcal{H})$ that is sometimes called the Hilbert-Schmidt inner product.

For a vector $|\psi\rangle \in \mathcal{H}$, $\||\psi\rangle\|$ denotes the Euclidean norm of $|\psi\rangle$. For an operator $A \in \mathbf{L}(\mathcal{H})$, the operator norm of A, denoted $\|A\|$, is defined by

$$\|A\| = \sup_{|\psi\rangle \in \mathcal{H}\setminus\{0\}} \frac{\|A|\psi\rangle\|}{\||\psi\rangle\|}.$$

The trace norm of A, denoted $\|A\|_{\mathrm{tr}}$, is defined by $\|A\|_{\mathrm{tr}} = \operatorname{tr}\sqrt{A^\dagger A}$. The trace norm and the operator norm are dual to one another with respect to the Hilbert-Schmidt inner product, meaning that the following fact holds.

Fact 1. *For every* $A \in \mathbf{L}(\mathcal{H})$,

$$\begin{aligned} \|A\| &= \max\left\{|\langle B, A\rangle| : B \in \mathbf{L}(\mathcal{H}),\ \|B\|_{\mathrm{tr}} \leq 1\right\}, \\ \|A\|_{\mathrm{tr}} &= \max\left\{|\langle B, A\rangle| : B \in \mathbf{L}(\mathcal{H}),\ \|B\| \leq 1\right\}. \end{aligned}$$

See, for instance, Bhatia [5] for a proof of this fact.

The trace norm characterizes the distinguishability of a given pair of density matrices $\rho_0, \rho_1 \in \mathbf{D}(\mathcal{H})$ in the following sense. There exists a binary-valued quantum measurement such that if $\rho \in \{\rho_0, \rho_1\}$ is chosen uniformly at random, then the measurement correctly determines which of ρ_0 or ρ_1 was given with probability $\frac{1}{2} + \frac{1}{4}\|\rho_0 - \rho_1\|_{\mathrm{tr}}$. Furthermore, such a measurement is optimal in the sense that no other quantum measurement can possibly distinguish between ρ_0 and ρ_1 with a higher success rate. An immediate corollary of this fact is that for a given pair ρ_0 and ρ_1, there exists a measurement that correctly identifies a chosen state $\rho \in \{\rho_0, \rho_1\}$ with probability of correctness at least $\frac{1}{2}\|\rho_0 - \rho_1\|_{\mathrm{tr}}$, even if ρ is chosen by an adversary that knows the measurement.

Consider the following variant of the distinguishability problem: We are given $\rho \in \mathbf{D}(\mathcal{H})$ chosen from one of two disjoint convex sets of density operators $\mathcal{A}_0, \mathcal{A}_1 \subseteq \mathbf{D}(\mathcal{H})$, and we are asked to determine the set from which ρ was chosen. For simplicity we will assume $\mathcal{A}_0$ and $\mathcal{A}_1$ are closed sets. Under this assumption, it is meaningful to define the trace distance $\operatorname{dist}(\mathcal{A}_0, \mathcal{A}_1)$ between $\mathcal{A}_0$ and $\mathcal{A}_1$ as the minimum of the quantity $\|\rho_0 - \rho_1\|_{\mathrm{tr}}$ over all choices of $\rho_0 \in \mathcal{A}_0$ and $\rho_1 \in \mathcal{A}_1$. We prove that there exists a single measurement with the property that if an arbitrary ρ is chosen from $\mathcal{A}_0$ with probability 1/2, and otherwise ρ is chosen from $\mathcal{A}_1$, then the measurement correctly determines which set ρ was chosen from with probability at least $\frac{1}{2} + \frac{1}{4}\operatorname{dist}(\mathcal{A}_0, \mathcal{A}_1)$. This fact therefore generalizes the fact concerning a single pair of quantum states mentioned above, as singleton sets are of course closed and convex. As above, this fact implies that if ρ is chosen from $\mathcal{A}_0 \cup \mathcal{A}_1$ in an arbitrary manner, even depending on the measurement itself, then the measurement

will correctly determine from which of $\mathcal{A}_0$ or $\mathcal{A}_1$ the state ρ was chosen with probability at least $\frac{1}{2}\operatorname{dist}(\mathcal{A}_0, \mathcal{A}_1)$.

The proof of this fact begins with a well-known result from convex analysis, which informally states that there exists a separating hyperplane between any two disjoint convex sets. Typically, the separation result is stated in terms of the vector space $\mathbb{R}^n$, but it translates to $\mathbf{H}(\mathcal{H})$ for a given space $\mathcal{H}$ without complications, as $\mathbf{H}(\mathcal{H})$ may be identified with the vector space $\mathbb{R}^{m^2}$, for $m = \dim(\mathcal{H})$. Here we state a restricted variant of this fact that is most convenient for our purposes—see Rockafellar [18], for instance, for a more general statement.

Fact 2. *Let $\mathcal{A}, \mathcal{B} \subseteq \mathbf{H}(\mathcal{H})$ be disjoint convex sets with $\mathcal{A}$ compact and $\mathcal{B}$ open. Then there exists a Hermitian operator $H \in \mathbf{H}(\mathcal{H})$ and a real number $a \in \mathbb{R}$ such that $\langle H, X\rangle \geq a > \langle H, Y\rangle$ for all $X \in \mathcal{A}$ and $Y \in \mathcal{B}$.*

We are now ready to state and prove the main result of this section.

Theorem 3. *Let $\mathcal{A}_0, \mathcal{A}_1 \subseteq \mathbf{D}(\mathcal{H})$ be closed convex sets of density operators. Then there exist measurement operators $E_0, E_1 \in \mathbf{Pos}(\mathcal{H})$ with $E_0 + E_1 = I$ such that the following holds. For every pair $\rho_0 \in \mathcal{A}_0$ and $\rho_1 \in \mathcal{A}_1$, if ρ is chosen uniformly from $\{\rho_0, \rho_1\}$ and measured via the measurement $\{E_0, E_1\}$, the measurement will correctly determine whether $\rho \in \mathcal{A}_0$ or $\rho \in \mathcal{A}_1$ with probability at least $\frac{1}{2} + \frac{1}{4}\operatorname{dist}(\mathcal{A}_0, \mathcal{A}_1)$.*

Proof. Let $d = \operatorname{dist}(\mathcal{A}_0, \mathcal{A}_1)$. If $d = 0$, the theorem is trivially satisfied by the measurement defined by $E_0 = E_1 = \frac{1}{2}I$ (which is equivalent to a random coin-flip), so assume that $d > 0$. Let

$$\mathcal{A} = \mathcal{A}_0 - \mathcal{A}_1 = \{\rho_0 - \rho_1 : \rho_0 \in \mathcal{A}_0, \rho_1 \in \mathcal{A}_1\}.$$

Then $\mathcal{A}$ is a compact convex set of Hermitian operators and $\|X\|_{\mathrm{tr}} \geq d$ for every $X \in \mathcal{A}$. Let

$$\mathcal{B} = \{Y \in \mathbf{H}(\mathcal{H}) : \|Y\|_{\mathrm{tr}} < d\}$$

denote the open ball of radius d in $\mathbf{H}(\mathcal{H})$ with respect to the trace norm. The sets $\mathcal{A}$ and $\mathcal{B}$ satisfy the conditions of Fact 2, and therefore there exists a Hermitian operator $H \in \mathbf{H}(\mathcal{H})$ and a real number $a \in \mathbb{R}$ such that $\langle H, X\rangle \geq a > \langle H, Y\rangle$ for all $X \in \mathcal{A}$ and $Y \in \mathcal{B}$. Because $Y \in \mathcal{B}$ if and only if $-Y \in \mathcal{B}$ for every Y, it follows that $-a < a$, and therefore $a > 0$.

Let $K = \frac{d}{a}H$. We therefore have that $\langle K, X\rangle \geq d$ for every $X \in \mathcal{A}$ and $\langle K, \frac{1}{d}Y\rangle < 1$ for every $Y \in \mathcal{B}$. As $\frac{1}{d}Y$ ranges over all Hermitian operators with trace norm smaller than 1, this implies $\|K\| \leq 1$ by Fact 1. Now, let $K^+, K^- \in \mathbf{Pos}(\mathcal{H})$ denote the positive and negative parts of K, meaning that they satisfy $K = K^+ - K^-$ and $\langle K^+, K^-\rangle = 0$. As $\|K\| \leq 1$ it follows that $K^+ + K^- \leq I$.

At this point we define $E_0, E_1 \in \mathbf{Pos}(\mathcal{H})$ as follows:

$$E_0 = K^+ + \frac{1}{2}(I - K^+ - K^-) \quad \text{and} \quad E_1 = K^- + \frac{1}{2}(I - K^+ - K^-).$$

The operators E_0 and E_1 are both positive semidefinite and satisfy $E_0 + E_1 = I$, and therefore represent a binary-valued POVM.

Now suppose $\rho_0 \in \mathcal{A}_0$ and $\rho_1 \in \mathcal{A}_1$ are chosen arbitrarily, and ρ is chosen uniformly from the set $\{\rho_0, \rho_1\}$. Let C denote the event that the measurement $\{E_0, E_1\}$ correctly determines which of ρ_0 and ρ_1 was selected. We have $\Pr[C] = \frac{1}{2}\langle E_0, \rho_0\rangle + \frac{1}{2}\langle E_1, \rho_1\rangle$, and therefore

$$\Pr[C] - \Pr[\neg C] = \frac{1}{2}\langle E_0 - E_1, \rho_0 - \rho_1\rangle = \frac{1}{2}\langle K, \rho_0 - \rho_1\rangle \geq \frac{d}{2},$$

with the inequality following from the fact that $\rho_0 - \rho_1 \in \mathcal{A}$. Consequently the measurement is correct with probability at least $\frac{1}{2} + \frac{d}{4}$ as required. □

As before, it follows from this theorem that the measurement $\{E_0, E_1\}$ will correctly identify an arbitrary choice of $\rho \in \mathcal{A}_0 \cup \mathcal{A}_1$ with probability at least $\frac{1}{2}\operatorname{dist}(\mathcal{A}_0, \mathcal{A}_1)$.

4 A Short Quantum Game for QIP

In this section, we prove that any language with a quantum interactive proof system also has a short quantum game by solving the QIP-complete problem CLOSE-IMAGES from Sect. 2.

First, let us recall that the fidelity $F(\rho, \xi)$ between two quantum states $\rho, \xi \in \mathbf{D}(\mathcal{H})$ is defined as $F(\rho, \xi) = \left\|\sqrt{\rho}\sqrt{\xi}\right\|_{\mathrm{tr}}$. The following fact, proved by Fuchs and van de Graaf [10], gives one relationship between the fidelity and the trace norm.

Fact 4. *Let $\rho, \xi \in \mathbf{D}(\mathcal{H})$. Then*

$$1 - \frac{1}{2}\|\rho - \xi\|_{\mathrm{tr}} \leq F(\rho, \xi) \leq \sqrt{1 - \frac{1}{4}\|\rho - \xi\|_{\mathrm{tr}}}.$$

We are now ready to state and prove the main result of this section.

Theorem 5. $\mathsf{QIP} \subseteq \mathsf{SQG}\left(1/2, 2^{-poly}\right)$.

Proof. It suffices to show that CLOSE-IMAGES is in $\mathsf{SQG}(1/2, 2^{-poly})$. Suppose the input encodes mixed state quantum circuits Q_0 and Q_1, each mapping n qubits to m qubits. Let $\mathcal{H}$ and $\mathcal{K}$ be Hilbert spaces with dimensions 2^n and 2^m corresponding to the n input qubits and m output qubits respectively. We may view Q_0 and Q_1 as corresponding to admissible transformations $Q_0, Q_1 : \mathbf{D}(\mathcal{H}) \to \mathbf{D}(\mathcal{K})$. Let $\mathcal{A}_i = \{Q_i(\rho) : \rho \in \mathbf{D}(\mathcal{H})\} \subseteq \mathbf{D}(\mathcal{K})$ denote the image of Q_i for $i = 0, 1$. The sets $\mathcal{A}_0$ and $\mathcal{A}_1$ are closed, convex sets of density operators.

Consider the following verifier for a short quantum game:

1. Receive n-qubit registers X_0 and X_1 from the yes-prover.
2. Choose $i \in \{0, 1\}$ uniformly at random and apply Q_i to register X_i. Let the output be contained in an m-qubit register Y, which is then sent to the no-prover.
3. Receive a classical bit b from the no-prover. Accept if $b \neq i$ and reject if $b = i$.

If (Q_0, Q_1) is a "yes" instance of CLOSE-IMAGES then there exist $\rho_0, \rho_1 \in \mathbf{D}(\mathcal{H})$ such that $Q_0(\rho_0) = Q_1(\rho_1)$. The strategy for the yes-prover is to prepare the registers X_0 and X_1 in states ρ_0 and ρ_1, respectively, and to send them to the verifier in step 1 of the verifier's protocol. Because $Q_0(\rho_0) = Q_1(\rho_1)$, the state contained in the register Y is independent of i, so the no-prover can do no better than randomly guessing in step 3. The verifier will therefore accept with probability $1/2$ in this case.

If (Q_0, Q_1) is a "no" instance of CLOSE-IMAGES then for any desired $\varepsilon \in 2^{-poly}$ we are promised that

$$\sqrt{\varepsilon(n)} \geq \max_{\xi_0, \xi_1 \in \mathbf{D}(\mathcal{H})} \{F(Q_0(\xi_0), Q_1(\xi_1))\} \geq 1 - \frac{1}{2} \operatorname{dist}(\mathcal{A}_0, \mathcal{A}_1).$$

It follows that $\operatorname{dist}(\mathcal{A}_0, \mathcal{A}_1) \geq 2 - 2\sqrt{\varepsilon(n)}$.

Regardless of the state of the registers X_0 and X_1 sent to the verifier by the yes-prover, we must have that the reduced state of the register Y sent to the no-prover is given by some state $\xi \in \mathcal{A}_0 \cup \mathcal{A}_1$, and moreover that $\Pr[\xi \in \mathcal{A}_0] = \Pr[\xi \in \mathcal{A}_1] = 1/2$. By Theorem 3 there exists a quantum measurement $\{E_0, E_1\}$ that correctly determines whether $\rho \in \mathcal{A}_0$ or $\rho \in \mathcal{A}_1$ with probability at least

$$\frac{1}{2} + \frac{1}{4} \operatorname{dist}(\mathcal{A}_0, \mathcal{A}_1) \geq 1 - \frac{\sqrt{\varepsilon(n)}}{2}.$$

The strategy for the no-prover is to perform the quantum measurement $\{E_0, E_1\}$ and send the result to the verifier in step 3. This causes the verifier to reject with probability at least $1 - \sqrt{\varepsilon(n)}/2$. As this argument holds for every $\varepsilon \in 2^{-poly}$, we have that the soundness error is 2^{-poly} as required. □

5 Error Reduction

Suppose that both the completeness and soundness error c and s of a refereed game are bounded below $1/2$ by an inverse polynomial. Then it follows from Chernoff bounds that these error probabilities can be made exponentially close to zero by repeating the game a polynomial number of times in succession and taking a majority vote. Of course, sequential repetition necessarily increases the number of turns in the game and so it is natural to ask if error reduction can be achieved without affecting the turn complexity of the game.

A natural approach to this task is to run many copies of the refereed game in parallel and to accept or reject based on the outcomes of the repetitions. This technique is purely classical and has been successfully applied to classical single- and multi-prover interactive proof systems (see for example Ref. [16] and the references therein). A potential problem with this technique is that the provers need not treat each repetition independently—they might try to correlate the parallel repetitions (or entangle them in the quantum case) in some devious way such that the completeness and/or soundness error does not decrease as desired.

In the quantum setting, the general case of this problem has not been completely solved. But for three-message single-prover quantum interactive proof systems with zero completeness error, Ref. [13] proves that parallel repetition followed by a unanimous vote does indeed achieve the

exponential reduction in soundness error that one might expect, regardless of any possible entanglement by the prover among the parallel copies.

In this section, we prove that parallel repetition followed by a unanimous vote can be used to improve the error bounds for short quantum games by reducing the problem to error reduction for single-prover quantum interactive proof systems with three or fewer messages. The reduction is achieved by fixing a yes- or no-prover P that is guaranteed to win with a certain probability. By viewing the verifier-prover pair (V, P) as a new composite verifier, we are left with what is now effectively a one- or two-message quantum interactive proof system in which the opposing prover is the lone prover. We define a verifier-prover pair (V', P') that runs many copies of (V, P) in parallel and accepts based on a unanimous vote. We can then employ the error reduction result of Ref. [13] to prove that the error of the new game decreases exponentially in the number of repetitions.

We formalize this argument shortly, but first we require additional notation. Given finite-dimensional Hilbert spaces $\mathcal{H}$ and $\mathcal{K}$, let $\mathbf{L}(\mathcal{H}, \mathcal{K})$ denote the set of all linear operators mapping $\mathcal{H}$ to $\mathcal{K}$ and let $\mathbf{T}(\mathcal{H}, \mathcal{K})$ denote the set of all linear operators mapping the vector space $\mathbf{L}(\mathcal{H})$ to $\mathbf{L}(\mathcal{K})$. The trace norm can be extended to $\mathbf{T}(\mathcal{H}, \mathcal{K})$ as follows. For $T \in \mathbf{T}(\mathcal{H}, \mathcal{K})$,

$$\|T\|_{\mathrm{tr}} = \sup_{X \in \mathbf{L}(\mathcal{H}) \setminus \{0\}} \frac{\|T(X)\|_{\mathrm{tr}}}{\|X\|_{\mathrm{tr}}}.$$

Let $\mathcal{L}$ be a Hilbert space with $\dim(\mathcal{L}) = \dim(\mathcal{H})$ and let $I_{\mathbf{L}(\mathcal{L})}$ denote the identity transformation on $\mathbf{L}(\mathcal{L})$. Then for $T \in \mathbf{T}(\mathcal{H}, \mathcal{K})$, the *diamond norm* $\|T\|_\diamond$ of T is given by $\|T\|_\diamond = \left\|T \otimes I_{\mathbf{L}(\mathcal{L})}\right\|_{\mathrm{tr}}$. Further information on the diamond norm may be found in Kitaev, Shen, and Vyalyi [12]. The diamond norm satisfies several nice properties that the trace norm (extended to $\mathbf{T}(\mathcal{H}, \mathcal{K})$) does not. For example, the diamond norm is multiplicative with respect to tensor products: $\|T_1 \otimes T_2\|_\diamond = \|T_1\|_\diamond \|T_2\|_\diamond$ for any choice of transformations T_1 and T_2.

We are now prepared to give the main result of this section, whose proof is based on the proof of Theorem 6 of Ref. [13].

Theorem 6. $\mathsf{SQG}(c, s) \subseteq \mathsf{SQG}(kc, s^k) \cap \mathsf{SQG}(c^k, ks)$ *for any choice of* $c, s : \mathbb{N} \to [0, 1]$ *and* $k \in poly$.

Proof. We first prove that $\mathsf{SQG}(c, s) \subseteq \mathsf{SQG}(kc, s^k)$. Let $L \in \mathsf{SQG}(c, s)$ and let $V(x) = (V(x)_1, V(x)_2)$ be a verifier witnessing this fact. For the remainder of this proof, we assume that the input $x \in \Sigma^*$ is fixed. For brevity we drop the argument and write $V = (V_1, V_2)$ and use similar notation for the provers.

Let $V' = (V_1^{\otimes k}, V_2^{\otimes k})$ be a verifier that runs k copies of the protocol of V in parallel and accepts if and only if every one of the k copies accepts. We must show that V' has completeness error at most kc and soundness error at most s^k.

First consider the case $x \in L$. Let $Y = (Y_1)$ be a yes-prover that always convinces V to accept with probability at least $1 - c$. Let $Y' = (Y_1^{\otimes k})$ be a yes-prover that runs k independent copies of the protocol of Y in parallel. Then no no-prover can win any one of the k copies with probability greater than c and so by the union bound we know that the completeness error of the repeated game is at most kc.

Next consider the case $x \notin L$. Let $N = (N_1)$ be a no-prover that always convinces V to reject with probability at least $1 - s$. Let $N' = (N_1^{\otimes k})$ be a no-prover that runs k independent copies of the protocol of N in parallel. We now show that no yes-prover can win against N' using verifier V' with probability greater than s^k.

Let Π_{init} denote the projection of the entire system onto the all-$|0\rangle$ initial state. Then the projection $\Pi'_{init} = \Pi_{\text{init}}^{\otimes k}$ corresponds to the initial state of the repeated game. Let Π_{acc} denote the projection onto the states for which the output qubit belonging to V is 1. Then the projection $\Pi'_{\text{acc}} = \Pi_{\text{acc}}^{\otimes k}$ corresponds to the accepting state of V'. Let $\mathcal{V}_N$ denote the Hilbert space corresponding to the private qubits of V and the private and message qubits of N and let $\mathcal{M}_Y$ denote the Hilbert space corresponding to the yes-prover's message qubits. Define $T_N \in \mathbf{T}(\mathcal{V}_N \otimes \mathcal{M}_Y, \mathcal{M}_Y)$ as

$$T_N(X) = \text{tr}_{\mathcal{V}_N}(\Pi_{\text{init}})X(\Pi_{\text{acc}}V_2N_1V_1).$$

As mentioned earlier, we may view (V, N) as a new composite verifier and the yes-prover as the lone prover for some one-message quantum interactive proof system (i.e., a message from the prover to (V, N)). In this context, Lemma 7 of Ref. [13] asserts that the maximum probability with which any prover could convince the verifier (V, N) to accept x is precisely $\|T_N\|_\diamond^2$. Because (V, N) has soundness error at most s, we have $\|T_N\|_\diamond^2 \leq s$.

Define a similar transformation $T'_N \in \mathbf{T}((\mathcal{V}_N \otimes \mathcal{M}_Y)^{\otimes k}, \mathcal{M}_Y^{\otimes k})$ using V', N', Π'_{init}, and Π'_{acc}. It follows that $T'_N = T_N^{\otimes k}$. From the multiplicativity of the diamond norm, it follows that the maximum probability with which any prover could convince (V', N') to accept x is

$$\|T'_N\|_\diamond^2 = \left\|T_N^{\otimes k}\right\|_\diamond^2 = \|T_N\|_\diamond^{2k} \leq s^k,$$

which establishes the desired result.

Due to the symmetric nature of quantum refereed games, we can modify the above proof to show that $\mathsf{SQG}(c, s) \subseteq \mathsf{SQG}(c^k, ks)$. In particular, define the verifier V'' so that he rejects if and only if all k copies reject. For the case $x \notin L$, the proof that V'' has soundness error ks is completely symmetric to the proof that V' has completeness error kc.

For the case $x \in L$, we let Y and Y' be yes-players as above. Define the Hilbert spaces $\mathcal{V}_Y$ and $\mathcal{M}_N$ and the projections Π_{rej} and Π'_{rej} in the appropriate symmetric manner as per the above proof. The transformation $T_Y \in \mathbf{T}(\mathcal{V}_Y \otimes \mathcal{M}_N, \mathcal{M}_N)$ is defined as

$$T_Y(X) = \text{tr}_{\mathcal{V}_Y}(V_1Y_1\Pi_{\text{init}})X(\Pi_{\text{rej}}V_2).$$

As before, we may view (V, Y) as a new composite verifier and the no-prover as the lone prover for some quantum interactive proof system. The differences here are that the quantum interactive proof is now a two-message proof instead of a one-message proof (i.e., a message from (V, Y) to the prover followed by the prover's reply to (V, Y)) and that the prover's goal is now to convince the verifier (V, Y) to reject x instead of to accept x.

Fortunately, it is still straightforward to apply Lemma 7 of Ref. [13] to this quantum interactive proof system and so we may claim that the maximum probability with which any prover could convince the verifier (V, Y) to reject x is precisely $\|T_Y\|_\diamond^2$. That V'' has completeness error c^k follows as before. □

The proof of Theorem 6 can be extended to allow for the slightly more general protocol wherein the verifier sends a message to the yes-prover (via some circuit V_{init}) before the short quantum game commences. This extension follows from the fact that we can apply Lemma 7 of Ref. [13] to the augmented transformations

$$\begin{aligned} T_N(X) &= \text{tr}_{\mathcal{V}_N}(V_{\text{init}}\Pi_{\text{init}})X(\Pi_{\text{acc}}V_2N_1V_1), \\ T_Y(X) &= \text{tr}_{\mathcal{V}_Y}(V_1Y_1V_{\text{init}}\Pi_{\text{init}})X(\Pi_{\text{rej}}V_2). \end{aligned}$$

Combining Theorems 5 and 6 we obtain the following corollary, which is the main result of this paper.

Corollary 7. $\mathsf{QIP} \subseteq \mathsf{SQG}$.

Proof. Given a desired error bound 2^{-p} where $p \in poly$, choose $\varepsilon \in 2^{-poly}$ so that $p\varepsilon \leq 2^{-p}$. We have $\mathsf{QIP} \subseteq \mathsf{SQG}\left(1/2, \varepsilon\right) \subseteq \mathsf{SQG}\left(2^{-p}, 2^{-p}\right)$. □

6 Conclusion

We introduced in this paper the quantum refereed game model of computation and gave a short quantum game with exponentially small error for languages with single-prover quantum interactive proof systems. However, we have only scratched the surface of the quantum games model, and many questions about it remain unanswered. Some examples follow.

- The two-turn game presented in this paper has an asymmetric protocol. Is there also a two-turn quantum refereed game for QIP in which the no-prover sends the first message, or in which the provers play one turn in parallel?
- It is known that $\mathsf{QIP} \subseteq \mathsf{EXP}$. How does SQG relate to EXP?
- We mentioned in Sect. 1 that classical refereed games characterize EXP [6], which implies that many-turn quantum refereed games are at least as powerful as EXP. What upper bounds can be proved on the power of refereed quantum games?
- We demonstrated that parallel repetition followed by a unanimous vote can reduce error for short quantum games. Is there a way to reduce the error in *any* quantum refereed game without affecting the number of turns in the game?

Acknowledgments

This research was supported by Canada's NSERC, the Canada Research Chairs program, the Canadian Institute for Advanced Research (CIAR), and a graduate student scholarship from the Province of Alberta.

References

[1] D. Aharonov, A. Kitaev, and N. Nisan. Quantum circuits with mixed states. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, pages 20–30, 1998.

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

[3] 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.

[4] M. Ben-Or, S. Goldwasser, J. Kilian, and A. Wigderson. Multi-prover interactive proofs: how to remove intractability assumptions. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, pages 113–131, 1988.

[5] R. Bhatia. *Matrix Analysis*. Springer, 1997.

[6] U. Feige and J. Kilian. Making games short. In *Proceedings of the Twenty-Ninth annual ACM Symposium on Theory of Computing*, pages 506 – 516, 1997.

[7] U. Feige and A. Shamir. Multi-oracle interactive protocols with constant space verifiers. *Journal of Computer and System Sciences*, 44:259–271, 1992.

[8] U. Feige, A. Shamir, and M. Tennenholtz. The noisy oracle problem. In *Advances in Cryptology – Proceedings of Crypto'88*, volume 403 of *Lecture Notes in Computer Science*, pages 284 – 296. Springer–Verlag, 1990.

[9] J. Feigenbaum, D. Koller, and P. Shor. A game-theoretic classification of interactive complexity classes. In *Proceedings of the 10th Conference on Structure in Complexity Theory*, pages 227–237, 1995.

[10] C. Fuchs and J. van de Graaf. Cryptographic distinguishability measures for quantum-mechanical states. *IEEE Transactions on Information Theory*, 45(4):1216–1227, 1999.

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

[12] A. Kitaev, A. Shen, and M. Vyalyi. *Classical and Quantum Computation*, volume 47 of *Graduate Studies in Mathematics*. American Mathematical Society, 2002.

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

[14] D. Koller and N. Megiddo. The complexity of two-person zero-sum games in extensive form. *Games and Economic Behavior*, 4:528–552, 1992.

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

[16] R. Raz. A parallel repetition theorem. *SIAM Journal of Computing*, 27(3):763–803, 1998.

[17] J. Reif. The complexity of two-player games of incomplete information. *Journal of Computer and System Sciences*, 29:274–301, 1984.

[18] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.

[19] B. Rosgen and J. Watrous. On the hardness of distinguishing mixed-state quantum computations. arXiv.org e-Print cs.CC/0407056, 2004.

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

[21] A. Shen. IP = PSPACE: simplified proof. *Journal of the ACM*, 39(4):878–880, 1992.

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