Perfect Parallel Repetition Theorem for Quantum XOR Proof Systems

Richard Cleve
School of Computer Science and IQC
University of Waterloo
Perimeter Institute
Waterloo, Ontario, Canada
cleve@cs.uwaterloo.ca

William Slofstra*
Department of Mathematics
University of California
Berkeley, California, USA
slofstra@math.berkeley.edu

Falk Unger
CWI
Amsterdam, the Netherlands
F.Unger@cwi.nl

Sarvagya Upadhyay
School of Computer Science and IQC
University of Waterloo
Waterloo, Ontario, Canada
supadhya@cs.uwaterloo.ca

Abstract

We consider a class of two-prover interactive proof systems where each prover returns a single bit to the verifier and the verifier's verdict is a function of the XOR of the two bits received. We show that, when the provers are allowed to coordinate their behavior using a shared entangled quantum state, a perfect parallel repetition theorem holds in the following sense. The prover's optimal success probability for simultaneously playing a collection of XOR proof systems is exactly the product of the individual optimal success probabilities. This property is remarkable in view of the fact that, in the classical case (where the provers can only utilize classical information), it does not hold. The theorem is proved by analyzing parities of XOR proof systems using semidefinite programming techniques, which we then relate to parallel repetitions of XOR games via Fourier analysis.

1. Introduction and summary of results

The theory of interactive proof systems has played an important role in the development of computational complexity and cryptography. Also, the impact of quantum information on the theory of interactive proof systems has been shown to have interesting consequences [20, 15]. In [5] a variant of the model of interactive proof system was introduced where there are two provers who have unlimited computational power subject to the condition that they cannot communicate between themselves once the execution of the protocol starts. This model is sufficiently powerful to characterize NEXP [1].

Our present focus is on *XOR interactive proof systems*, which are based on *XOR games*. For a predicate $f : S \times T \rightarrow \{0,1\}$ and a probability distribution π on $S \times T$, define the XOR game $G = (f, \pi)$ operationally as follows.

- The Verifier selects a pair of questions $(s,t) \in S \times T$ according to distribution π.
- The Verifier sends one question to each prover: s to prover Alice and t to prover Bob (who are forbidden from communicating with each other once the game starts).
- Each prover sends a bit back to the Verifier: a from Alice and b from Bob.
- The Verifier accepts if and only if $a \oplus b = f(s,t)$.

A definition that is essentially equivalent to this[1] appears in [8]. In the classical version, the provers have unlimited computing power, but are restricted to possessing classical information; in the quantum version, the provers may possess qubits whose joint state is entangled. In both versions, the communication between the provers and the verifier is classical.

An *XOR interactive proof system* (with soundness probability s and completeness probability $c > s$) for a language L associates an XOR game with every input string x, such that:

*Research carried out while at the University of Waterloo.

[1] Except that *degeneracies* are allowed, where for some (s,t) pairs, the Verifier is allowed to accept or reject independently of the value of $a \oplus b$. All results quoted here apply to nondegenerate games.

- S_x and T_x consist of strings of length polynomial in $|x|$, π_x can be sampled in time polynomial in $|x|$, and f_x can be computed in time polynomial in $|x|$.
- If $x \in L$ then the maximum acceptance probability over prover's strategies is at least c.
- If $x \notin L$ then the maximum acceptance probability over prover's strategies is at most s.

In [8] it is pointed out that results in [4, 13] imply that, in the case of classical provers, these proof systems have sufficient expressive power to recognize every language in NEXP (with soundness probability $s = 11/16 + \epsilon$ and completeness probability $c = 12/16 - \epsilon$, for arbitrarily small $\epsilon > 0$). Thus, although these proof systems appear restrictive, they can recognize any language that an unrestricted multi-prover interactive proof system can. Moreover, in [9, 22] it is shown that any language recognized by a quantum XOR proof system is in EXP. Thus, assuming EXP $\neq$ NEXP, quantum entanglement strictly weakens the expressive power of XOR proof systems.

Returning to XOR games, quantum physicists have, in a sense, been studying them since the 1960s, when John Bell introduced his celebrated results that are now known as Bell inequality violations [3]. An example is the *CHSH* game, named after the authors of [7]. In this game, $S = T = \{0, 1\}$, π is the unform distribution on $S \times T$, and $f(s, t) = s \wedge t$. It is well known that, for the *CHSH* game, the best possible classical strategy succeeds with probability $3/4$, whereas the best possible quantum strategy succeeds with higher probability of $(1 + 1/\sqrt{2})/2 \approx 0.85$ [7, 18].

Following [8], for an XOR game G, define its *classical value* $\omega_c(G)$ as the maximum possible success probability achievable by a classical strategy. Similarly, define its *quantum value* $\omega_q(G)$ as the maximum possible success probability achievable by a quantum strategy.

1.1 Taking the sum of XOR games

For any two XOR games $G_1 = (f_1, \pi_1)$ and $G_2 = (f_2, \pi_2)$, define their *sum (modulo* 2*)* as the XOR game

$$G_1 \oplus G_2 = (f_1 \oplus f_2, \pi_1 \times \pi_2). \tag{1}$$

In this game, the verifier begins by choosing questions $((s_1, t_1), (s_2, t_2)) \in (S_1 \times T_1) \times (S_2 \times T_2)$ according to the product distribution $\pi_1 \times \pi_2$, sending (s_1, s_2) to Alice and (t_1, t_2) to Bob. Alice and Bob then win if and only if their respective outputs, a and b, satisfy $a \oplus b = f_1(s_1, t_1) \oplus f_2(s_2, t_2)$.

A simple way for Alice and Bob (who may or may not share entanglement) to play $G_1 \oplus G_2$ is to optimally play G_1 and G_2 separately, producing outputs a_1, b_1 for G_1 and a_2, b_2 for G_2, and then to output $a = a_1 \oplus a_2$ and $b = b_1 \oplus b_2$ respectively. It is straightforward to calculate that the above method for playing $G_1 \oplus G_2$ succeeds with probability

$$\omega(G_1)\omega(G_2) + (1 - \omega(G_1))(1 - \omega(G_2)). \tag{2}$$

Is this the optimal way to play $G_1 \oplus G_2$?

The answer is *no* for *classical* strategies. To see why this is so, note that, using this approach for the XOR game *CHSH* $\oplus$ *CHSH*, produces a success probability of $5/8$. A better strategy is for Alice to output $a = s_1 \wedge s_2$ and Bob to output $b = t_1 \wedge t_2$. It is straightforward to verify that this latter strategy succeeds with probability $3/4$.

Our first result is that the answer is *yes* for *quantum* strategies.

Theorem 1 (Additivity). *For any two XOR games G_1 and G_2 an optimal quantum strategy for playing $G_1 \oplus G_2$ is for Alice and Bob to optimally play G_1 and G_2 separately, producing outputs a_1, b_1 for G_1 and a_2, b_2 for G_2, and then to output $a = a_1 \oplus a_2$ and $b = b_1 \oplus b_2$.*

The proof of Theorem 1 uses a known characterization of quantum strategies for individual XOR games as semidefinite programs. Section 2 contains the proof.

1.2 Parallel repetition of XOR games

For any sequence of XOR games $G_1 = (f_1, \pi_1)$, ..., $G_n = (f_n, \pi_n)$, define their *conjunction*, denoted by $\wedge_{j=1}^{n} G_j$, as follows. The verifier chooses questions $((s_1, t_1), \dots, (s_n, t_n)) \in (S_1 \times T_1) \times \cdots \times (S_n \times T_n)$ according to the product distribution $\pi_1 \times \cdots \times \pi_n$, and sends $(s_1, \dots, s_n)$ to Alice and $(t_1, \dots, t_n)$ to Bob. Alice and Bob output bits $a_1, \dots, a_n$ and $b_1, \dots, b_n$, respectively, and win if and only if their outputs simultaneously satisfy these n conditions: $a_1 \oplus b_1 = f_1(s_1, t_1)$, ..., $a_n \oplus b_n = f_n(s_n, t_n)$. (Note that $\wedge_{j=1}^{n} G_j$ is not itself an XOR game for $n > 1$.)

One way for Alice and Bob to play $\wedge_{j=1}^{n} G_j$ is to independently play each game optimally. This succeeds with probability $\prod_{j=1}^{n} \omega(G_j)$. Is this the optimal way to play $\wedge_{j=1}^{n} G_j$?

The answer is *no* for classical strategies. It is shown in [2] that[2] ω_c(*CHSH* $\wedge$ *CHSH*) $= 10/16 > 9/16 = \omega_c$(*CHSH*)$\omega_c$(*CHSH*).

Our second result is that the answer is *yes* for quantum strategies.

Theorem 2 (Parallel Repetition). *For any XOR games $G_1, \dots, G_n$, $\omega_q(\wedge_{j=1}^{n} G_j) = \prod_{j=1}^{n} \omega_q(G_j)$.*

[2] After posing this question about ω_c(*CHSH* $\wedge$ *CHSH*), the answer was first shown to us by S. Aaronson, who independently discovered the classical protocol and then found the prior result in [2].

This is a quantum version of Raz's parallel repetition theorem [17] for the restricted class of XOR games. We call it a *perfect* parallel repetition theorem because the probabilities are multiplicative in the exact sense (as opposed to an asymptotic sense, as in [17]). The proof of Theorem 2 is based on Theorem 1 combined with Fourier analysis techniques for boolean functions. Section 3 contains the proof.

1.3 Comparison with other work

There is no known parallel repetition theorem along the lines of [17] for quantum games (where the players share entanglement). As far as we know, our results represent the first progress in this direction. Recently, Holenstein [14] gave a simplified proof of the parallel repetition theorem that applies to classical and no-signalling strategies. Neither of these cases capture quantum strategies for XOR games (for example, every XOR game has value 1 in the no-signaling model).

For games other than XOR games, the question of parallel repetition remains open. Watrous [21] has shown that, there is a binary game (that is not an XOR game) for which $\omega_q(G) = \omega_q(G \wedge G) = 2/3$, as in the classical case. Thus, a perfect parallel repetition property does not automatically apply to quantum games.

For any XOR game G, the semidefinite programming relaxation of G due to Feige and Lovász [11] has value equal to the quantum value of G. It is not clear whether such a correspondence occurs for games over larger alphabets; this is worth exploring further.

2 Proof of the Additivity Theorem

In this section we prove Theorem 1, which is stated in Section 1.1.

It is convenient to define the quantum *bias* of an XOR game as $\varepsilon_q(G) = 2\omega_q(G) - 1$. Then, due to Eq. 2, to prove Theorem 1, it suffices to show that $\varepsilon_q(G_1 \oplus G_2) = \varepsilon_q(G_1)\varepsilon_q(G_2)$.

Since Alice and Bob can independently play games G_1 and G_2 optimally and then take the parity of their outputs as their outputs for $G_1 \oplus G_2$, we immediately have the following.

Proposition 3. *For two XOR games G_1 and G_2, $\varepsilon_q(G_1 \oplus G_2) \geq \varepsilon_q(G_1)\varepsilon_q(G_2)$.*

The nontrivial part of the proof is the reverse inequality.

A quantum strategy for an XOR game consists of a bipartite quantum state $|\psi\rangle$ shared by Alice and Bob, a set of observables X_s $(s \in S)$ corresponding to Alice's part of the quantum state, and a set of observables Y_t $(t \in T)$ corresponding to Bob's part of the state. The bias achieved by this strategy is given by

$$\sum_{s,t} \pi(s,t)(-1)^{f(s,t)} \langle\psi| X_s \otimes Y_t |\psi\rangle .$$

We make use of a vector characterization of XOR games due to [19] (also pointed out in [8]), which is a consequence of the following.

Theorem 4 ([19, 8]). *Let S and T be finite sets, and let $|\psi\rangle$ be a pure quantum state with support on a bipartite Hilbert space $\mathcal{H} = \mathcal{A} \otimes \mathcal{B}$ such that dim($\mathcal{A}$) = dim($\mathcal{B}$) = n. For each $s \in S$ and $t \in T$, let X_s and Y_t be observables on $\mathcal{A}$ and $\mathcal{B}$ with eigenvalues ± 1 respectively. Then there exists real unit vectors x_s and y_t in $\mathbb{R}^{2n^2}$ such that*

$$\langle\psi| X_s \otimes Y_t |\psi\rangle = x_s \cdot y_t,$$

for all $s \in S$ and $t \in T$.
Conversely, suppose that S and T are finite sets, and x_s and y_t are unit vectors in $\mathbb{R}^N$ for each $s \in S$ and $t \in T$. Let $\mathcal{A}$ and $\mathcal{B}$ be Hilbert space of dimension $2^{\lceil N/2 \rceil}$, $\mathcal{H} = \mathcal{A} \otimes \mathcal{B}$ and $|\psi\rangle$ be a maximally entangled state on $\mathcal{H}$. Then there exists observables X_s and Y_t with eigenvalues ± 1, on $\mathcal{A}$ and $\mathcal{B}$ respectively, such that

$$\langle\psi| X_s \otimes Y_t |\psi\rangle = x_s \cdot y_t,$$

for all $s \in S$ and $t \in T$.

Using Theorem 4, we can characterize Alice and Bob's quantum strategies by a choice of unit vectors $\{x_s\}_{s \in S}$ and $\{y_t\}_{t \in T}$. Using this characterization, the bias becomes

$$\varepsilon_q(G) = \max_{\{x_s\},\{y_t\}} \sum_{s,t} \pi(s,t)(-1)^{f(s,t)}\, x_s \cdot y_t. \qquad (3)$$

The *cost matrix* for the game is defined as the matrix A with entries $A_{s,t} = \pi(s,t)(-1)^{f(s,t)}$.

Note that any matrix A, with the provision that the absolute values of the entries sum to 1, is the cost matrix of an XOR game. If G_1 and G_2 are XOR games with cost matrices A_1 and A_2 respectively, then the cost matrix of $G_1 \oplus G_2$ is $A_1 \otimes A_2$. Also, for $0 \leq \lambda \leq 1$, define the convex combination $\lambda G_1 + (1-\lambda) G_2$ to be the XOR game with cost matrix

$$\begin{pmatrix} 0 & \lambda A_1 \\ (1-\lambda) A_2 & 0 \end{pmatrix}.$$

This convex combination can be interpreted as the game where, with probability λ, game G_1 is played and, with probability $1-\lambda$, game G_2 is played (and Alice and Bob are informed about which game is occurring). Also, for a game G with cost matrix A, define G^T to be the game with cost matrix A^T. In other words, Alice and Bob switch places to play G^T. The next proposition summarizes some simple facts.

Proposition 5.
1. $\varepsilon_q(G_1 \oplus G_2) = \varepsilon_q(G_2 \oplus G_1)$ *and* $\varepsilon_q(G) = \varepsilon_q(G^T)$.
2. *For all* $0 \leq \lambda \leq 1$,
$\varepsilon_q\left(\lambda G_1 + (1-\lambda)\, G_2\right) = \lambda \varepsilon_q(G_1) + (1-\lambda)\, \varepsilon_q(G_2)$
and
$G_1 \oplus (\lambda G_2 + (1-\lambda) G_3) = \lambda(G_1 \oplus G_2) + (1-\lambda)(G_1 \oplus G_3)$.

The bias of a quantum XOR game may be stated as a semidefinite programming problem (SDP). We refer to Boyd and Vandenberghe [6] for a detailed introduction to semidefinite programming. For cost matrix A, the bias is equivalent to the objective value of problem

$$\max\ \mathrm{Tr}\left(A^T U_1^T U_2\right) : \mathrm{diag}\left(U_1^T U_1\right) = \mathrm{diag}\left(U_2^T U_2\right) = \bar{e},$$

where $\{x_s\}$ and $\{y_t\}$ appear as the columns of U_1 and U_2 respectively. Here diag(M) denotes the column vector of diagonal entries of the matrix M, and $\bar{e}$ is the column vector $(1, \ldots, 1)^T$. We begin by considering the game $\frac{1}{2}G + \frac{1}{2}G^T$, whose cost matrix

$$B = \begin{pmatrix} 0 & \frac{1}{2}A \\ \frac{1}{2}A^T & 0 \end{pmatrix}$$

has useful structural properties, one of them being that it is symmetric. Proposition 5 implies that $\varepsilon_q(\frac{1}{2}G + \frac{1}{2}G^T) = \varepsilon_q(G)$. This enables us to express the value of game G in terms of the SDP (P_B) defined by

$$\max\ \mathrm{Tr} BX \quad : \quad \mathrm{diag}(X) = \bar{e}, \quad X \succeq 0.$$

The notation $X \succeq Y$ means that the matrix $X - Y$ lies in the cone of positive semidefinite matrices. That (P_B) is equivalent to problem (2) follows from the fact that a semidefinite matrix X can be written as $(U_1, U_2)^T (U_1, U_2)$ for some matrices U_1 and U_2.

To show that an optimal solution for (P_B) exists, we can examine the Lagrange-Slater dual of (P_B). The dual, denoted by (D_B), is defined as

$$\min\ (x, y)\bar{e} \quad : \quad \Delta(x, y) \succeq B,$$

where $\Delta(x, y)$ denotes the diagonal matrix with entries given by the (row) vectors x, y. Both (P_B) and (D_B) have Slater points—that is, feasible points in the interior of the semidefinite cone. Explicitly, the identity matrix is a Slater point for (P_B), and $\bar{e}$ is a Slater point for (D_B). Therefore, by the strong duality theorem, the optimal values of (P_B) and (D_B) are the same and both problems have optimal solutions attaining this value.

The next lemma establishes the upper bound for the game $(\frac{1}{2}G_1 + \frac{1}{2}G_1^T) \oplus (\frac{1}{2}G_2 + \frac{1}{2}G_2^T)$ (which we will show afterwards has the same bias as $G_1 \oplus G_2$).

Lemma 6. *If* G_1 *and* G_2 *are XOR games, then* $\varepsilon_q((\frac{1}{2}G_1 + \frac{1}{2}G_1^T) \oplus (\frac{1}{2}G_2 + \frac{1}{2}G_2^T)) \leq \varepsilon_q(G_1)\varepsilon_q(G_2)$.

Proof. Let G_1 and G_2 be two games with cost matrices A_1 and A_2, respectively, and let

$$B_1 = \begin{pmatrix} 0 & \frac{1}{2}A_1 \\ \frac{1}{2}A_1^T & 0 \end{pmatrix} \text{ and } B_2 = \begin{pmatrix} 0 & \frac{1}{2}A_2 \\ \frac{1}{2}A_2^T & 0 \end{pmatrix}. \qquad (4)$$

Let (x_1, y_1) and (x_2, y_2) be optimal solutions to (D_{B_1}) and (D_{B_2}), respectively, which implies $\Delta(x_i, y_i) - B_i \succeq 0$ and $\varepsilon_q(G_i) = (x_i, y_i)\bar{e}$, for $i = 1, 2$. It suffices to show that $(x_1, y_1) \otimes (x_2, y_2)$ is a solution to $(\mathrm{D}_{B_1 \otimes B_2})$, since $B_1 \otimes B_2$ is the cost matrix of $(\frac{1}{2}G_1 + \frac{1}{2}G_1^T) \oplus (\frac{1}{2}G_2 + \frac{1}{2}G_2^T)$. Note that, for *arbitrary* B_1 and B_2, $\Delta(x_1, y_1) \succeq B_1$ and $\Delta(x_2, y_2) \succeq B_2$ does *not* imply that $\Delta(x_1, y_1) \otimes \Delta(x_2, y_2) \succeq B_1 \otimes B_2$ (a simple counterexample is when $\Delta(x_1, y_1) = \Delta(x_2, y_2) = 0$ and $B_1 = B_2 = -I$). We make use of the structure of B_1 and B_2 arising from Eq. 4. For each i, $\Delta(x_i, y_i) - B_i \succeq 0$ implies that, for all (row) vectors u, v,

$$\begin{aligned} 0 \quad \leq \quad & \begin{pmatrix} u & v \end{pmatrix} \begin{pmatrix} \Delta(x_i) & -\frac{1}{2}A_i \\ -\frac{1}{2}A_i^T & \Delta(y_i) \end{pmatrix} \begin{pmatrix} u^T \\ v^T \end{pmatrix} \\ = \quad & \begin{pmatrix} u & -v \end{pmatrix} \begin{pmatrix} \Delta(x_i) & +\frac{1}{2}A_i \\ +\frac{1}{2}A_i^T & \Delta(y_i) \end{pmatrix} \begin{pmatrix} u^T \\ -v^T \end{pmatrix}, \end{aligned}$$

which in turn implies that $\Delta\left(x_i, y_i\right) + B_i \succeq 0$ also holds. Therefore,

$$\begin{aligned} (\Delta(x_1, y_1) - B_1) \otimes (\Delta(x_2, y_2) + B_2) \quad & \succeq \quad 0 \quad \text{and} \\ (\Delta(x_1, y_1) + B_1) \otimes (\Delta(x_2, y_2) - B_2) \quad & \succeq \quad 0, \end{aligned}$$

which, by averaging, yields

$$\Delta(x_1, y_1) \otimes \Delta(x_2, y_2) - B_1 \otimes B_2 \succeq 0.$$

Therefore, $(x_1, y_1) \otimes (x_2, y_2)$ is a feasible point in the dual $(\mathrm{D}_{B_1 \otimes B_2})$, which obtains the objective value $\varepsilon_q(G_1)\varepsilon_q(G_2)$, which implies the Lemma. □

Now we may complete the proof of Theorem 1. Using Proposition 3 for line (5), Lemma 6 for line (5) and Proposition 5 and some easy algebra for the rest we can derive the following

$$\begin{aligned} & \varepsilon_q\left(G_1 \oplus G_2\right) \\ & \quad \geq \quad \varepsilon_q\left(G_1\right) \varepsilon_q\left(G_2\right) \\ & \quad \geq \quad \varepsilon_q((\tfrac{1}{2}G_1 + \tfrac{1}{2}G_1^T) \oplus (\tfrac{1}{2}G_2 + \tfrac{1}{2}G_2^T)) \\ & \quad = \quad \varepsilon_q\left(\tfrac{1}{4}(G_1 \oplus G_2) + \tfrac{1}{4}(G_1 \oplus G_2^T)\right. \\ & \qquad\qquad \left. + \tfrac{1}{4}(G_1^T \oplus G_2) + \tfrac{1}{4}(G_1^T \oplus G_2^T)\right) \\ & \quad = \quad \varepsilon_q\Big(\tfrac{1}{2}\left[\tfrac{1}{2}(G_1 \oplus G_2) + \tfrac{1}{2}(G_1 \oplus G_2^T)\right] \\ & \qquad\qquad + \tfrac{1}{2}\left[\tfrac{1}{2}(G_1 \oplus G_2) + \tfrac{1}{2}(G_1 \oplus G_2^T)\right]^T\Big) \\ & \quad = \quad \tfrac{1}{2}\varepsilon_q\left(G_1 \oplus G_2\right) + \tfrac{1}{2}\varepsilon_q\left(G_1 \oplus G_2^T\right). \end{aligned}$$

Therefore $\varepsilon_q(G_1 \oplus G_2) \geq \varepsilon_q(G_1 \oplus G_2^T)$. By symmetry, $\varepsilon_q(G_1 \oplus G_2^T) \geq \varepsilon_q(G_1 \oplus G_2)$, as well, which means that all of the above inequalities must be equalities. This completes the proof of Theorem 1.

3 Parallel repetition theorem

In this section we prove Theorem 2, which is stated in Section 1.2.

We begin with the following simple probabilistic lemma.

Lemma 7. *For any sequence of binary random variables $X_1, X_2, \ldots, X_n$,*

$$\frac{1}{2^n}\sum_{M\subseteq[n]} \mathrm{E}\left[(-1)^{\oplus_{j\in M}X_j}\right] = \Pr[X_1\ldots X_n = 0\ldots 0].$$

Proof. By the linearity of expectation,

$$\begin{aligned} &\frac{1}{2^n}\sum_{M\subseteq[n]} \mathrm{E}\left[(-1)^{\oplus_{j\in M}X_j}\right] \\ &= \mathrm{E}\left[\frac{1}{2^n}\sum_{M\subseteq[n]}(-1)^{\oplus_{j\in M}X_j}\right] \\ &= \mathrm{E}\left[\prod_{j=1}^{n}\left(\frac{1+(-1)^{X_j}}{2}\right)\right] \\ &= \Pr\left[X_1\ldots X_n = 0\ldots 0\right], \end{aligned}$$

where the last equality follows from the fact that

$$\prod_{j=1}^{n}(1+(-1)^{X_j}) \neq 0$$

only if $X_1\ldots X_n = 0\ldots 0$. □

We introduce the following terminology. For any strategy $\mathcal{S}$ (classical or quantum) for any game G, define $\omega(\mathcal{S}, G)$ as the success probability of strategy $\mathcal{S}$ on game G. Similarly, define the corresponding bias as $\varepsilon(\mathcal{S}, G) = 2\omega(\mathcal{S}, G) - 1$.

Now let $\mathcal{S}$ be any protocol for the game $\wedge_{j=1}^{n} G_j$. For each $M \subseteq [n]$, define the protocol $\mathcal{S}_M$ (for the game $\oplus_{j\in M} G_j$) as follows.

1. Run protocol $\mathcal{S}$, yielding $a_1, \ldots, a_n$ for Alice and $b_1, \ldots, b_n$ for Bob.
2. Alice outputs $\oplus_{j\in M} a_j$ and Bob outputs $\oplus_{j\in M} b_j$.

Lemma 8.

$$\frac{1}{2^n}\sum_{M\subseteq[n]} \varepsilon(\mathcal{S}_M, \oplus_{j\in M}G_j) = \omega(\mathcal{S}, \wedge_{j=1}^{n}G_j).$$

Proof. For all $j \in [n]$, define $X_j = a_j \oplus b_j \oplus f_j(s_j, t_j)$. Then, for all $M \subseteq [n]$, we have $\mathrm{E}[(-1)^{\oplus_{j\in M}X_j}] = \varepsilon(\mathcal{S}_M, \oplus_{j\in M}G_j)$, and $\Pr[X_1\ldots X_n = 0\ldots 0] = \omega(\mathcal{S}, \wedge_{j=1}^{n}G_j)$. The result now follows from Lemma 7. □

Corollary 9.

$$\omega_c(\wedge_{j=1}^{n}G_j) \leq \frac{1}{2^n}\sum_{M\subseteq[n]} \varepsilon_c(\oplus_{j\in M}G_j) \tag{5}$$

and

$$\omega_q(\wedge_{j=1}^{n}G_j) \leq \frac{1}{2^n}\sum_{M\subseteq[n]} \varepsilon_q(\oplus_{j\in M}G_j). \tag{6}$$

Now, to complete the proof of Theorem 2, using Theorem 1, we have

$$\begin{aligned} \frac{1}{2^n}\sum_{M\subseteq[n]} \varepsilon_q(\oplus_{j\in M}G_j) &= \frac{1}{2^n}\sum_{M\subseteq[n]}\prod_{j\in M}\varepsilon_q(G_j) \\ &= \prod_{j=1}^{n}\left(\frac{1+\varepsilon_q(G_j)}{2}\right) \\ &= \prod_{j=1}^{n}\omega_q(G_j). \end{aligned} \tag{7}$$

Combining this with Eq. 6, we deduce $\omega_q(\wedge_{j=1}^{n}G_j) = \prod_{j=1}^{n}\omega_q(G_j)$, which completes the proof of Theorem 2.

Comments: Although Eq. 6 is used to prove a tight upper bound on $\omega_q(\wedge_{j=1}^{n}G_j)$, Eq. 5 cannot be used to obtain a tight upper bound on $\omega_c(\wedge_{j=1}^{n}G_j)$ for general XOR games. This is because $\varepsilon_c(\mathit{CHSH}) = \varepsilon_c(\mathit{CHSH} \oplus \mathit{CHSH}) = 1/2$ and it can be shown that $\varepsilon_c(\mathit{CHSH} \oplus \mathit{CHSH} \oplus \mathit{CHSH}) = 5/16$. Therefore, for $G_1 = G_2 = G_3 = \mathit{CHSH}$, the right side of Eq. 5 is $\frac{1}{8}\sum_{M\subseteq[3]}\varepsilon_c(\oplus_{j\in M}G_j) = 34.5/64$, whereas $\omega_c(\wedge_{j=1}^{3}G_j)$ must be expressible as an integer divided by 64 (in fact[3], $\omega_c(\wedge_{j=1}^{3}G_j) = 31/64$).

Acknowledgments

We would like to thank Scott Aaronson, Ben Toner, John Watrous, and Ronald de Wolf for helpful discussions. RC and SU acknowledge support from Canada's NSERC, CIAR and MITACS, and the U.S. ARO/DTO. WS acknowledges support from Canada's NSERC. FU acknowledges support from the EU project QAP (IST-2005-15848).

References

[1] L. Babai, L. Fortnow, and C. Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1(1):3–40, 1991.

[2] J. Barrett, D. Collins, L. Hardy, A. Kent, and S. Popescu. Quantum nonlocality, Bell inequalities and the memory loophole. *Physical Review A* 66:042111, 2002.

[3] This was independently calculated by S. Aaronson and B. Toner, by searching over a finite number of deterministic classical strategies.

[3] J. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1(3):195–200, 1964.

[4] M. Bellare, O. Goldreich, and M. Sudan. Free bits, PCPs, and non-approximability — towards tight results. *SIAM Journal on Computing*, 27(3):804–915, 1998.

[5] 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.

[6] S. Boyd and L. Vandenberghe. Semidefinite programming. *SIAM Review*, 38(1):49–95, 1996.

[7] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt. Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23(15):880–884, 1969.

[8] R. Cleve, P. Høyer, B. Toner, J. Watrous. Consequences and limits of nonlocal strategies. In *Proceedings of the 19th IEEE Conference on Computational Complexity*, pages 236–249, 2004.

[9] R. Cleve, P. Høyer, B. Toner, J. Watrous. Consequences and limits of nonlocal strategies. Presentation given at *Ninteenth IEEE Conference on Computational Complexity*, June 2004.

[10] U. Feige. On the success probability of two provers in one-round proof systems. In *Proceedings of the Sixth Annual Conference on Structure in Complexity Theory*, pages 116–123, 1991.

[11] U. Feige and L. Lovász. Two-prover one-round proof systems: their power and their problems. In *Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing*, pages 733–744, 1992.

[12] L. Fortnow, J. Rompel, and M. Sipser. On the power of multi-prover interactive protocols. *Theoretical Computer Science*, 134:545–557, 1994.

[13] J. Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001.

[14] T. Holenstein. Parallel repetition: simplifications and the no-signaling case. In *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*, 2007.

[15] A. Kitaev and J. Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof systems. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, pages 608–617, 2000.

[16] H. Kobayashi and K. Matsumoto. Quantum multi-prover interactive proof systems with limited prior entanglement. *Journal of Computer and System Sciences*, 66(3):429–450, 2003.

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

[18] B. S. (Tsirelson) Cirel'son. Quantum generalizations of Bell's inequality. *Letters in Mathematical Physics*, 4(2):93–100, 1980.

[19] B. S. (Tsirelson) Tsirel'son. Quantum analogues of the Bell inequalities: The case of two spatially separated domains. *Journal of Soviet Mathematics*, 36:557–570, 1987.

[20] J. Watrous. PSPACE has constant-round quantum interactive proof systems. in *Proceedings of the Fourtieth Annual Symposium on Foundations of Computer Science*, pages 112–119, 1999.

[21] J. Watrous. Personal communication, 2004.

[22] S. Wehner. Entanglement in interactive proof systems with binary answers. In *Proceedings of STACS 2006*, pages 162–171, 2006.