

Entanglement-Resistant Two-Prover Interactive Proof Systems and Non-Adaptive Private Information Retrieval Systems

Richard Cleve ^{*} Dmitry Gavinsky

Rahul Jain [†]

Computer Science Department and Institute for Quantum Computing [‡]

Abstract

We show that, for any language in NP, there is an entanglement-resistant constant-bit two-prover interactive proof system with a constant completeness vs. soundness gap. The previously proposed *classical* two-prover constant-bit interactive proof systems are known not to be entanglement-resistant. This is currently the strongest expressive power of any known constant-bit answer multi-prover interactive proof system that achieves a constant gap. Our result is based on an “oracularizing” property of certain private information retrieval systems, which may be of independent interest.

1 Introduction

Properties of interactive proof systems have been shown to change in fundamental ways when the underlying setting changes from classical information to quantum information. The first result along these lines was discovered by Watrous [14], and several subsequent results have occurred.

In the present paper, we are concerned with *multiprover interactive proof systems* (MIPs), first proposed (in the classical setting) by Ben-Or *et al.* [2]. An example of such a system, for 3SAT, is where the first prover is sent a clause of the formula and the second prover is sent a variable from the clause. The first prover must give a partial truth assignment that satisfies the clause and the second prover must give an assignment to the variable that is consistent with the first prover’s (this protocol occurs in several places in the literature, e.g., [8]). Classically, the completeness probability of this system is 1, whereas the soundness probability is at most $1 - \frac{1}{3m}$, where m is the number of clauses. Roughly speaking, the role of the second prover is to “oracularize” the first prover: to make it respond to queries as an oracle would; if the first prover behaves adaptively to queries, this introduces a positive probability that the results between the two provers are not consistent. In the quantum setting, if the provers are allowed to share *a priori* entanglement they can cheat the protocol in the sense that there are *unsatisfiable* 3CNF formulas where the soundness probability of the protocol is 1 (hence the gap is zero) [3]. Thus this particular oracularization technique fails in the setting of quantum information (other examples are also given in [3]). Note that, in the

^{*}Email: cleve@cs.uwaterloo.ca.

[†]Email: rjain@cs.uwaterloo.ca.

[‡]University of Waterloo, 200 University Ave. West, Waterloo, ON N2L 3G1, Canada.

above, quantum information enters the picture by the entanglement between the provers; the verifier and its communication with the provers remains classical. Interesting results have also been obtained for multiprover interactive proof systems where the communication between the verifier and provers is quantum.

A major question is how the expressive power of two-prover interactive proof systems changes when the provers possess entanglement. Without entanglement, this is known to be NEXP [1, 7]. Since entanglement can potentially increase both the completeness and soundness probability, it is not even clear whether the expressive power is a subset or superset of NEXP. In [3] (based on results in [8]), it is shown that a restricted class of MIPs (called \oplus -MIPs or XOR-MIPs) has the property that, classically, their expressive power equals NEXP; whereas, with entangled provers, its expressive power reduces to a subset of EXP (see [15] for a refinement of this result). Thus, for \oplus -MIPs, entanglement strictly reduces their expressive power (unless EXP = NEXP). An \oplus -MIP has the simple form where the verifier makes one polynomial-length query to each prover, and each prover returns a single bit answer to the verifier. The verifier's acceptance condition is a function of the XOR of the two answer bits and the questions.

Our main result is the introduction of a new technique for oracularizing provers, based on properties of certain *private information retrieval systems* (PIRs). A PIR is a system that enables information to be obtained from a database without revealing to the database server(s) what the information is that is being queried. The framework is two (or more) isolated servers who each have a copy of the database, but who cannot communicate with each other. Instead of asking an individual server for the information (which would reveal the query), each server is asked for information and the responses are combined to produce the answer. There are ways of doing this such that no individual server acquires any information about the actual query being made.

Intuitively, this seems like a natural approach to oracularizing provers in a MIP: if the servers have no idea of that is being queried in the first place, how can they make their answers adaptive? Although this sounds intuitively compelling, the non-adaptiveness property is operationally different from the PIR property. This distinction is reminiscent of the distinction between malleable cryptography and non-malleable cryptography [4]. For example, a cryptosystem may be secure in the sense that it is not possible to deduce x from an encryption of x , nevertheless it may be possible from this encryption to construct an encryption of y that is somehow related to x .

We show that certain PIRs are in fact non-adaptive in the sense that, not only do they reveal no information about the items queried to the servers, but they satisfy the additional property that the servers cannot conspire to make their answers satisfy a property that non-trivially depends on the queries made. Remarkably, this property is robust even against *quantum* servers who have the resource of *a priori* entanglement.

Based on this, we show that \oplus -MIP has expressive power at least that of NP for entangled provers. This is the strongest expressive power of any known constant-bit answer MIP that achieves a constant gap.

Related work is [9, 10], where other novel techniques are introduced. In this work, the complexity classes are supersets of NP; however, the gaps between completeness and soundness probability are smaller and the communication from the provers is larger. Hence these results are incomparable with ours.

2 Some notation

We use the following notation in this paper. For $s, t \in \{0, 1\}^m$, let $s \cdot t \in \{0, 1\}$ denote the inner product modulo 2 of s and t , and $s \oplus t \in \{0, 1\}^m$ denote the bitwise exclusive-or of s and t . For $j \in \{1, 2, \dots, m\}$, let $e_j \in \{0, 1\}^m$ denote the characteristic vector of $\{j\}$, which is 1 in component j and 0 in all other components.

3 Our main result

Our main result is as follows.

Theorem 3.1 *For all $\varepsilon > 0$, for any language L in NP, there exists a two-prover protocol \mathcal{P} with entangled provers of the following form. Let $x \in \{0, 1\}^n$ (n large enough depending on ε) be the input received by the provers Alice and Bob and the Verifier V .*

1. V generates messages s and t , each of length polynomial in n , and a private bit δ ; (s, t, δ) chosen from a certain polynomial time samplable joint distribution. V then sends s, t to Alice and Bob respectively.
2. Alice and Bob respond with bits a and b respectively.
3. V accepts x if and only if $a \oplus b = f_x(s \oplus t, \delta)$, where f_x is computable in time polynomial in n .

The protocol satisfies the following soundness/completeness properties:

Completeness: *If $x \in L$ then there exists a strategy for provers Alice and Bob such that V accepts with probability $\geq 1 - \varepsilon$.*

Soundness: *If $x \notin L$ then, for all strategies of provers Alice and Bob, V accepts with probability $\leq \frac{1}{2} + \varepsilon$.*

The above theorem immediately implies $\text{NP} \subseteq \oplus\text{-MIP}^*[2,1]$, where $\oplus\text{-MIP}^*[2,1]$ represents that class of languages acceptable by two entangled prover proof systems with a single round of interaction and in which verifier only uses the xor of the bits answered by the provers to decide.

4 The PCP system

Let $L \in \text{NP}$ and $\varepsilon > 0$. From [8], there exists a probabilistically checkable proof (PCP) system for L of the following form. There is a proof verification procedure V_{PCP} that, for any n -bit string x (n large enough depending on ε), takes an m -bit string w as input (where $m \in n^{O(1)}$) and accepts or rejects w as a certificate of $x \in L$ based on the parity of three bits of w as follows. V_{PCP} probabilistically generates distinct $i, j, k \in \{1, 2, \dots, m\}$ and $\delta \in \{0, 1\}$, from a certain polynomial time (in n) samplable joint distribution, and accepts if and only if $w_i \oplus w_j \oplus w_k = f_x(i, j, k, \delta)$, where f_x is a polynomially computable function. The completeness/soundness properties of the proof system are as follows.

Completeness: For all $x \in L$, there exists a witness string $w \in \{0, 1\}^m$ such that V_{PCP} accepts with probability at least $1 - \varepsilon$.

Soundness: For all $x \notin L$, for all $w \in \{0, 1\}^m$, V_{PCP} accepts with probability at most $\frac{1}{2} + \varepsilon$.

5 Our protocol via the PIR reduction

Using the PCP procedure V_{PCP} we obtain our protocol \mathcal{P} as follows. On receiving input x , V interacts with provers Alice and Bob as follows.

1. V simulates V_{PCP} in the generation of $i, j, k \in \{1, 2, \dots, m\}$ and $\delta \in \{0, 1\}$.
2. V chooses $s \in \{0, 1\}^m$, uniformly distributed and independently of i, j, k, δ , and set $t = s \oplus e_i \oplus e_j \oplus e_k$.
3. V sends s to Alice and t to Bob, receiving one-bit answers a and b from them respectively.
4. V accepts if and only if $a \oplus b = f_x(i, j, k, \delta)$.

5.1 Completeness

If $x \in L$ then we know from the PCP procedure that there exists a PCP-witness $w \in \{0, 1\}^m$. Consider the strategy in which Alice, on input s outputs $a = w \cdot s$, and Bob, on input t , outputs $b = w \cdot t$. Now,

$$\begin{aligned} a \oplus b &= w \cdot (s \oplus t) \\ &= w \cdot (e_i \oplus e_j \oplus e_k) \\ &= w_i \oplus w_j \oplus w_k. \end{aligned}$$

Therefore V accepts whenever V_{PCP} accepts the PCP string w , and hence the probability of acceptance of V is at least $1 - \varepsilon$.

5.2 Soundness

The proof of soundness employs a result about certain XOR games that are similar to those analyzed by Linden *et al.* [13]. Lets define a *transversal XOR game* as an interactive protocol between a verifier and two entangled provers, Alice and Bob, that is specified by a function $g : \{0, 1\}^m \times \{0, 1\}^l \rightarrow \{0, 1\}$ and a distribution π on $\{0, 1\}^m \times \{0, 1\}^l$. The operation of the game is as follows.

1. The verifier generates $(z, r) \in \{0, 1\}^m \times \{0, 1\}^l$ according to distribution π . Then the verifier produces two shares of z , s and t , by generating $s \in \{0, 1\}^n$ uniformly and independently of z and setting $t = s \oplus z$. The verifier sends s to Alice and t to Bob.
2. Alice and Bob produce bits a and b , respectively, and send them to the verifier.
3. The verifier accepts if and only if $a \oplus b = g(s \oplus t, r)$.

The following is a slight generalization of a result in [13] and its proof appears in the Appendix A.

Theorem 5.1 *Let G be a transversal XOR game specified by $g : \{0, 1\}^m \times \{0, 1\}^l \rightarrow \{0, 1\}$ and the distribution π . Then the optimal strategy that maximizes $\Pr[a \oplus b = g(s \oplus t, r)]$ does not use any entanglement and is of the following form. For some $u \in \{0, 1\}^m$ and $\gamma \in \{0, 1\}$ (that depend on g and π), Alice responds with $a = (u \cdot s) \oplus \gamma$ and Bob responds with $b = u \cdot t$.*

Now in our protocol \mathcal{P} , the verifier on receiving x could be thought of as playing a transversal XOR game with the provers Alice and Bob by letting $z \triangleq e_i \oplus e_j \oplus e_k$, $r \triangleq \delta$ and $g : \{0, 1\}^m \times \{0, 1\} \rightarrow \{0, 1\}$ be such that $g(e_i \oplus e_j \oplus e_k, \delta) \triangleq f_x(i, j, k, \delta)$.

Let $x \notin L$. Now from Theorem 5.1 the optimal strategy for the provers in which they are trying to maximize the acceptance probability of the verifier V would be as follows. Alice and Bob ignore the entanglement and for some $u \in \{0, 1\}^m$ and $\gamma \in \{0, 1\}$, Alice outputs $a = (u \cdot s) \oplus \gamma$ and Bob outputs $a = u \cdot t$.

Now it can be easily shown that $\Pr[V \text{ accepts } x] = \Pr[a \oplus b = g(x, s \oplus t)] \leq 1/2 + \varepsilon$. Consider the following PCP witness w . For all $j \in \{1, 2, \dots, m\}$, set $w_j \triangleq u_j \oplus \gamma$. Note that this witness satisfies

$$\begin{aligned} w_i \oplus w_j \oplus w_k &= u_i \oplus u_j \oplus u_k \oplus \gamma \\ &= u \cdot (e_i \oplus e_j \oplus e_k) \oplus \gamma \\ &= u \cdot (s \oplus t) \oplus \gamma \\ &= a \oplus b. \end{aligned}$$

Combining this with the fact that $f_x(i, j, k, \delta) = g(s \oplus t, \delta)$ enables us to conclude that

$$\Pr[a \oplus b = g(x, s \oplus t)] = \Pr[w_i \oplus w_j \oplus w_k = f_x(i, j, k, \delta)] \leq \frac{1}{2} + \varepsilon.$$

The last inequality comes from the soundness property of the PCP procedure V_{PCP} . Thus $\Pr[V \text{ accepts } x]$ in the protocol \mathcal{P} is at most $\frac{1}{2} + \varepsilon$ and hence the soundness property is satisfied.

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] 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.
- [3] 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.
- [4] D. Dolev, C. Dwork and M. Naor. Non-malleable cryptography. *SIAM Journal on Computing* 30(2):391–437, 2000.
- [5] 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.
- [6] 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.

- [7] L. Fortnow, J. Rompel, and M. Sipser. On the power of multi-prover interactive protocols. *Theoretical Computer Science*, 134:545–557, 1994.
- [8] J. Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001.
- [9] J. Kempe and T. Vidick. On the power of entangled quantum provers. arXiv:quant-ph/0612063, 2006.
- [10] J. Kempe, H. Kobayashi, K. Matsumoto, B. Toner, and T. Vidick. On the power of entangled provers: immunizing games against entanglement. Manuscript, 2007.
- [11] 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.
- [12] 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.
- [13] N. Linden, S. Popescu, A. J. Short, and A. Winter. No quantum advantage for nonlocal computation. arXiv:quant-ph/0610097, 2006.
- [14] 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.
- [15] S. Wehner. Entanglement in interactive proof systems with binary answers. In *Proceedings of STACS 2006*, pages 162–171, 2006.

A Proof of Theorem 5.1

Our proof follows very much in the lines of the proof of Linden et al. [13]. Let Alice and Bob share a pure quantum state $|\phi\rangle$ between them. Let Z, R be a pair of random variables jointly distributed according to π . Let $S, T \in \{0, 1\}^n$ represent the random variables corresponding to the questions of verifier V to Alice and Bob respectively. Let A, B represent the random variables corresponding to the answers by Alice and Bob respectively. Note that in our case $Z = S \oplus T$ and S is uniformly distributed and is independent of (Z, R) . It is well known that $\Pr[g(z, r) = A \oplus B \mid (S, Z, R) = (s, z, r)]$ can be expressed as follows,

$$\Pr[g(z, r) = A \oplus B \mid (S, Z, R) = (s, z, r)] = \frac{1}{2}(1 + (-1)^{g(z, r)} \langle \phi | A_s \otimes B_{s \oplus z} | \phi \rangle)$$

, where $A_s, B_{s \oplus z}$ are Hermitian operators with eigenvalues in $\{-1, 1\}$, (which are also sometimes referred to as *observables*). Therefore we have,

$$\begin{aligned} \Pr[V \text{ accepts}] &= \sum_{s, z, r} \Pr[(S, Z, R) = (s, z, r)] \Pr[g(z, r) = A \oplus B \mid (S, Z, R) = (s, z, r)] \\ &= \sum_{s, z, r} \Pr[(S, Z, R) = (s, z, r)] \cdot \frac{1}{2}(1 + (-1)^{g(z, r)} \langle \phi | A_s \otimes B_{s \oplus z} | \phi \rangle) \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{2} + \sum_{s,z,r} \Pr[(S, Z, R) = (s, z, r)] \cdot \frac{1}{2} (-1)^{g(z,r)} \langle \phi | A_s \otimes B_{s \oplus z} | \phi \rangle \\
&= \frac{1}{2} + \sum_{s,z} \frac{1}{2^n} \left(\sum_r \Pr[(R, Z) = (r, z)] \right) \cdot \frac{1}{2} (-1)^{g(z,r)} \langle \phi | A_s \otimes B_{s \oplus z} | \phi \rangle
\end{aligned}$$

Now let,

$$\begin{aligned}
\theta_z &\triangleq \sum_r \Pr[(R, Z) = (r, z)] \cdot \frac{1}{2} (-1)^{g(z,r)} \\
|\alpha\rangle &\triangleq \frac{1}{\sqrt{2^n}} \sum_s (A_s \otimes I) (|\phi\rangle \otimes |s\rangle) \\
|\beta\rangle &\triangleq \frac{1}{\sqrt{2^n}} \sum_t (I \otimes B_t) (|\phi\rangle \otimes |t\rangle) \\
\Phi &\triangleq \sum_{s,z} \theta_z |s\rangle \langle s \oplus z|
\end{aligned}$$

Note that as defined above $|\alpha\rangle, |\beta\rangle$ are unit vectors. Also Φ is Hermitian. Therefore from above we have,

$$\begin{aligned}
\Pr[V \text{ accepts}] &= \frac{1}{2} + \sum_{s,z} \frac{1}{2^n} \theta_z \langle \phi | A_s \otimes B_{s \oplus z} | \phi \rangle \\
&= \frac{1}{2} + \langle \alpha | (I \otimes \Phi) | \beta \rangle \\
&\leq \frac{1}{2} + \|\langle \alpha | \|_2 \| (I \otimes \Phi) \|_\infty \| |\beta\rangle \|_2 \\
&= \frac{1}{2} + \| (I \otimes \Phi) \|_\infty \\
&= \frac{1}{2} + \|\Phi\|_\infty
\end{aligned}$$

Above $\|\Phi\|_\infty$ represents the highest singular value of Φ and since Φ is Hermitian it means the highest modulus eigenvalue.

Now we show below that the eigenvectors of Φ are precisely the *Hadamard* vectors $|u\rangle \triangleq \sum_{v \in \{0,1\}^n} (-1)^{u \cdot v} |v\rangle$ (for $u \in \{0,1\}^n$) with eigenvalues $\lambda_u = \sum_z (-1)^{u \cdot z} \theta_z$. Consider,

$$\begin{aligned}
\Phi |u\rangle &= \left(\sum_{s,z} \theta_z |s\rangle \langle s \oplus z| \right) \left(\sum_{v \in \{0,1\}^n} (-1)^{u \cdot v} |v\rangle \right) \\
&= \sum_{s,z} (-1)^{u \cdot (s \oplus z)} \theta_z |s\rangle \\
&= \left(\sum_z (-1)^{u \cdot z} \theta_z \right) \sum_s (-1)^{u \cdot s} |s\rangle \\
&= \lambda_u |u\rangle
\end{aligned}$$

Next we show that there exists a classical strategy by Alice and Bob such that $\Pr[V \text{ accepts}] = \frac{1}{2} + \|\Phi\|_\infty$. Let $|w\rangle$ be the eigenvector of Φ corresponding to the highest modulus eigenvalue.

Let $\gamma = 0$ if $\lambda_w \geq 0$ and 1 otherwise. Now let Alice answer with $(w \cdot s) \oplus \gamma$ to question s and let Bob answer with $(w \cdot t)$ to question t . Then we see that,

$$\begin{aligned}
\Pr[V \text{ accepts}] &= \sum_{s,z,r} \Pr[(S, Z, R) = (s, z, r)] \Pr[g(z, r) = (w \cdot s) \oplus \gamma \oplus (w \cdot (s \oplus z))] \\
&= \sum_{s,z,r} \Pr[(S, Z, R) = (s, z, r)] \Pr[g(z, r) = (w \cdot z) \oplus \gamma] \\
&= \sum_{s,z,r} \Pr[(S, Z, R) = (s, z, r)] \left(\frac{1}{2} + \frac{1}{2} (-1)^{g(z,r) + (w \cdot z) + \gamma} \right) \\
&= \frac{1}{2} + \sum_{s,z,r} \Pr[(S, Z, R) = (s, z, r)] \cdot \frac{1}{2} \cdot (-1)^{g(z,r) + (w \cdot z) + \gamma} \\
&= \frac{1}{2} + \sum_{z,r} \Pr[(Z, R) = (z, r)] \cdot \frac{1}{2} \cdot (-1)^{g(z,r) + (w \cdot z) + \gamma} \\
&= \frac{1}{2} + \sum_z \left(\sum_r \frac{1}{2} \cdot \Pr[(Z, R) = (z, r)] \cdot (-1)^{g(z,r)} \right) (-1)^{(w \cdot z) + \gamma} \\
&= \frac{1}{2} + \sum_z \theta_z \cdot (-1)^{(w \cdot z) + \gamma} \\
&= \frac{1}{2} + |\lambda_w| = \frac{1}{2} + \|\Phi\|_\infty
\end{aligned}$$

■