Multi-Prover Interactive Proof Systems with Leakage

Vahid R. Asadi
National Institute of Informatics, Tokyo, Japan

Atsuya Hasegawa
Graduate School of Mathematics, Nagoya University, Japan

François Le Gall
Graduate School of Mathematics, Nagoya University, Japan

Abstract

It is known that there exist multi-prover interactive protocols (MIP protocols) for the complexity class NEXP, succinct MIP protocols for NP and multi-prover interactive protocols with shared entanglement (MIP^* protocols) for RE. This extraordinary power of multi-prover interactive proof systems comes from the assumption that provers do not communicate with each other during the protocols. If they are allowed to communicate freely, the setting is the same as in the single-prover case, and the computational power of the system becomes significantly weaker.

In this paper, we investigate for the first time the setting where communication (i.e., leakage of information) between provers is allowed but bounded. We introduce two techniques to approach this question and show that multi-prover interactive proof systems are robust against some amount of leakage. Our first technique is based on parallel repetition theorems. We apply it to show that for any polynomial p, we can construct two-prover one-round MIP and MIP^* protocols for NEXP and RE, respectively, that are robust against $p(n)$ bits of leakage. We further derive our second technique to convert any low-soundness PCP construction to a two-prover one-round MIP protocol for NP robust against leakage. We also discuss the relation between robustness against leakage in multi-prover interactive proof systems and the Sliding Scale Conjecture in the PCP literature.

2012 ACM Subject Classification Theory of computation → Interactive proof systems

Keywords and phrases Multi-prover interactive proof systems

Digital Object Identifier 10.4230/LIPIcs.MFCS.2026.60

Related Version *Full Version*: https://arxiv.org/abs/2605.09872

Funding *Vahid R. Asadi*: Supported by a JSPS Postdoctoral Fellowship for Research in Japan. *Atsuya Hasegawa*: Supported by JSPS KAKENHI grant No. 24H00071, 25K24674, 25K24465. *François Le Gall*: Supported by JSPS KAKENHI grant No. 24H00071, 25K24674, 25K24465, MEXT Q-LEAP grant No. JPMXS0120319794, JST ASPIRE grant No. JPMJAP2302 and JST CREST grant No. JPMJCR24I4.

Acknowledgements VRA thanks Shuichi Hirahara and Nobutaka Shimizu for helpful discussions. This research was done when VRA was a visiting researcher at Nagoya University, and a graduate student at the University of Waterloo. AH thanks Srijita Kundu for useful discussions. The authors also thank the anonymous referees for their helpful comments and suggestions

1 Background

Multi-prover interactive proof systems (MIPs) were first introduced by Ben-Or et al. [9] to remove intractability assumptions from zero-knowledge proof systems [21, 20]. In a multi-prover interactive proof system, there are multiple provers who interact with a verifier. The provers cannot communicate with each other during the execution of the protocol, and thus the verifier can "cross-check" their assertions. In 1990, Babai, Fortnow and Lund [4] showed

MIP = NEXP, where NEXP is a set of problems recognized by non-deterministic exponential-time Turing machines. In MIP protocols, the polynomial-time verifier checks exponentially large proofs, and the result of MIP = NEXP led to the discovery of the celebrated PCP theorem [3, 2].

The class MIP*, which is a variant of MIP where provers can share entanglement, was first introduced by Cleve et al. [16] who derived a correspondence with quantum non-local games. Since in non-local games, quantum values are larger than classical values (i.e., $\omega_q(G) \geq \omega_c(G)$) and thus soundness in MIP protocols does not imply one in MIP* protocols, the relation between MIP and MIP* was mysterious. After an extensive effort to investigate the power of MIP* (e.g., [32, 25, 47, 40]), Ji et al. [30] proved MIP* = RE. The RE-complete problem is the Halting Problem (which is undecidable), and shared entanglement between the provers gives extremely strong power to the multi-prover interactive proof systems.

The extraordinary power of multi-prover interactive proof systems comes from the canonical assumption that provers do not communicate with each other during the execution of the protocol[1]. If provers are allowed to communicate freely, the proof system is the same as in the case of the single prover, and the power of the proof system becomes much weaker because IP = PSPACE [48]. If we limit the number of rounds to a constant, it is still the case that there exists a two-prover one-round MIP (resp. MIP*) protocol for NEXP (resp. RE). However, in the single-prover setting (IP), the class is equal to AM [22, 5], which is much weaker than the multi-prover case.

An important line of research in the study of multi-prover interactive proofs is related to the problem of reducing the soundness error to any arbitrary small value. This question motivated the study of the parallel repetition theorem [45, 24] in which the goal is to reduce the soundness error to any arbitrary value by asking the provers to play multiple instances of the same game simultaneously. It was first shown by [45] that the soundness error decays exponentially with respect to the number of copies of the game that are played in parallel.

The parallel repetition theorem also plays an important role in the study of probabilistically checkable proofs (PCPs), where it is used to reduce the soundness error at the cost of increasing the proof length and alphabet size. However, parallel repetition is not the only approach. Gap amplification [18, 6] provides a more efficient way to increase the unsatisfiability gap, thereby reducing the soundness error with smaller overhead. Moreover, similar effects can be achieved through careful composition of PCPs [39, 19].

In this work, we formally show a connection between the no-communication assumption in multi-prover interactive proofs, and the parallel repetition and the low soundness error for PCPs. More precisely, we will use the latter to relax the former and show that if the soundness error is chosen to be a carefully small value, one can allow some bits[2] of communication between the provers.

2 Our contribution

As we have discussed, there are huge gaps in the computational power of multi-prover interactive proof systems with and without communication between provers. Moreover, when we want to realize these systems physically, say for some cryptographic applications, enforcing physical separation may be considered an unrealistic assumption, and in practice, a small amount of information could leak from the devices corresponding to the provers.

[1] Before the protocol starts, they can talk freely and share some strategy.

[2] For MIP* protocols, quantum communication (qubits) can always be replaced by classical communication (bits) via teleportation [11] since the provers share entanglement. In this paper, we thus only consider classical communication between provers, even when discussing MIP* protocols.

These facts motivate us to ask the following question:

Assume some bounded amount of communication is allowed between the provers. Then, what is the computational power of the multi-prover interactive proof systems?

Note that it is necessary to assume there is a meaningful bound on the amount of leakage between the provers. In particular, if the amount of leakage is greater than the size of the questions the verifier asks, then trivially the power of the protocol comes down to the power of single-prover interactive proofs.

In this work, we will introduce two techniques to answer this question.

- Our first tool, which uses the power of parallel repetition theorems, can be used to show that it is indeed the case that we can allow the provers to have bounded communication, while preserving the power of multi-prover settings. Of course, there will be a cost to pay, and it shows up in the size of the question and answer sets of the non-local game, but as we will discuss shortly, the blow-up in the question and answer size is moderate.
- Our second technique exploits low-soundness PCP constructions (achieved from gap amplification), and we show how to convert them to leakage-resilient succinct MIP protocols for NP. We achieve better parameters for question, answer, leakage size, and soundness, although this comes at the cost of additional assumptions or restrictions.

Our techniques allow us to show that the extraordinary power of multi-prover interactive proof systems still holds even if we weaken the assumption of no communication.

This, to the best of our knowledge, is the first time one utilizes the power of soundness error reductions towards relaxing the no-communication assumption in the multi-prover interactive proof systems. We will also show connections between the possibility of certain multi-prover interactive proof systems with leakage and the Sliding Scale Conjecture in the PCP literature, with the hope that one can find resolutions for either using techniques from the other.

3 Technical overview

For the rest of this technical overview, we fix the following notation. We denote by $\mathsf{MIP}[q, a, \ell]$ (resp. $\mathsf{MIP}^*[q, a, \ell]$) the class of problems decided by a two-prover one-round MIP (resp. MIP^*) protocol whose question length is q and answer length is a, and the amount of allowed communication between the provers is ℓ (see Section 5.1 for formal definitions). Here, n always denotes the instance size of the problem, and for simplicity, we assume the alphabet is $\{0, 1\}$.

3.1 Techniques based on parallel repetition

Our first contribution is to show how to use parallel repetition theorems for classical and quantum non-local games to remove the no-communication assumption between provers in a multi-prover interactive proof system.

We utilize this technique to first show the following result for MIP.

▶ **Theorem 3.1** (Informal version of Theorem 5.6)**.** *Even if any* $\mathrm{poly}(n)$ *bits of communication is allowed between the provers, there exists a two-prover one-round* MIP *protocol for* NEXP*.*

Note that in the protocol above, the question size is some asymptotically larger polynomial than the leakage size, and the answer size is a sufficiently large linear multiple of the leakage size.

Second, considering the parallel repetition of quantum non-local games, we show the following result for MIP^*.

▶ **Theorem 3.2** (Informal version of Theorem 5.9)**.** *Even if any* $\mathrm{poly}(n)$ *bits of communication is allowed between the provers, there exists a two-prover one-round* MIP^* *protocol for* RE*.*

In the protocol we show above, the question size is a sufficiently large linear multiple of the leakage size, and the answer size is larger than the leakage size and asymptotically the same polynomial up to some $\mathrm{polylog}(n)$ factor.

Our results imply that any $\mathrm{poly}(n)$ amount of leakage between provers in a multi-prover interactive proof can be allowed without changing their computational power.

Proof strategies

Given the correspondence between MIP and MIP^* protocols and non-local games, we can consider an N-times parallel repetition of a two-prover one-round MIP (resp. MIP^*) protocol for NEXP (resp. RE). Roughly speaking, the parallel repetition of a non-local game G is a game where the parties try to win N copies of G simultaneously. The parallel repetition game is denoted by $G^{\otimes N}$. Suppose that the amount of communication allowed between the provers is sublinear in N. Then, intuitively, it is hard for the cheating provers to win all N games. This can be formalized by a direct product theorem in communication complexity, which has a long history of works (e.g., [46, 15, 29, 14]).

For classical communication complexity, it is known that a direct product theorem follows from the parallel repetition theorem for the classical value, first shown by Raz [45]. It is also known that there exists a MIP protocol for NP and NEXP whose answer size is constant from the PCP theorem [3, 2, 18, 39, 6]. Therefore, by considering a sufficiently large $\mathrm{poly}(n)$ times parallel repetition of MIP protocols for NEXP with perfect completeness, we have $\mathsf{MIP}[\mathrm{poly}(n), \mathrm{poly}(n), \mathrm{poly}(n)] \supseteq \mathsf{NEXP}$, which proves Theorem 3.1.

For the entanglement-assisted interactive quantum communication complexity, Jain and Kundu recently [26] showed a direct product theorem for the parallel repetition of non-local games whose question distribution is product (see Lemma 5.7 for a formal description).

In a non-local game, if the distribution of questions is a product distribution, we denote such games by free games. Natarajan and Zhang [42] showed that there exists a succinct quantum free game protocol for the Halting Problem, which meets the conditions we want. Considering $\mathrm{poly}(n)$-times parallel repetition of the protocol in [42], and combining it with the direct product theorem in [26], we have $\mathsf{MIP}^*[\mathrm{poly}(n), \mathrm{poly}(n), \mathrm{poly}(n)] \supseteq \mathsf{RE}$, giving us Theorem 3.2.

We will formalize these proofs in Sections 5.2 and 5.3.

3.2 Techniques based on low-soundness PCPs

As our next contribution, we develop a new technique where we show how to convert any PCP construction with small soundness error into a MIP protocol resilient against leakage. See Theorem 5.11 for a formal statement. For this technique, we need the additional assumption that the leakage between the provers is one-way, meaning that only the first prover can send messages to the second prover. Another way of viewing this scenario is to assume that the verifier sends a question to the first prover, receives the answer, and all the communication between the provers happens at this stage, before the second prover receives its question from the verifier. Therefore, any communication from the second malicious prover to the first is not helpful in winning the game.

We then use this tool to show a succinct MIP protocol for NP that is sound against constant bits of leakage. Here succinct means that the question and answer size are small, and thus while NP $\subset$ NEXP, this result is not superseded by Theorem 3.1. An informal statement of our result is the following.

▶ **Theorem 3.3** (Informal version of Theorem 3.13 in the full version)**.** *There exists a two-prover one-round* MIP *protocol for* NP *where the question size is* $2\log n + O(\log\log n)$*, the answer size is* $O(1)$*, and the system is sound against* $O(1)$ *bits of one-way leakage.*

Note that in this setting, if we allow the provers to communicate freely (in other words, having a single prover while the questions are $O(\log n)$ in length), the power of the protocol reduces to P, essentially rendering the verification protocol useless. What we show is that we can still maintain the power of verifying NP while the provers are allowed to communicate a constant number of bits, and the overhead to pay is only an additive $O(\log\log n)$ term in the length of the question.

We further show that there exists a MIP protocol for NP that is sound against $O(\log n)$ bits of leakage, but our result comes at the cost of super-constant soundness error.

▶ **Theorem 3.4** (Informal version of Theorem 3.14 in the full version)**.** *For any positive constant* ℓ*, there exists a two-prover one-round* MIP *protocol allowing* $\ell \log n$ *bits of one-way leakage for* NP *where the question size is* $O(\log n)$*, the answer size is* $\log n \cdot \mathrm{poly}(\log\log n)$*, and the system has soundness* $1 - 1/\mathrm{poly}(\log\log n)$*.*

Proof strategies

We start by showing a correspondence between low-soundness PCPs and the construction of MIPs in which the first prover is allowed to communicate a bounded number of bits to the second prover. Roughly speaking, by utilizing a PCP construction in which the number of satisfied constraints in a no instance is very low, we can make sure that even if the provers decide to change their strategy mid-protocol, there are still many constraints that remain unsatisfied even with a new cheating strategy. For simplicity, suppose that the provers are only able to communicate 1 bit. This corresponds to the first prover specifying which of the two strategies the second prover has to respond with (one strategy corresponds to bit 0, and the other one to bit 1). Therefore, if we can amplify the soundness gap to a large enough constant (say, something less than 1/10), then for any two different assignments they agree on before the protocol starts, there are still many constraints that remain unsatisfied by both. Hence, if the verifier is lucky enough to sample one of those, it can catch the cheating provers. We will formalize this intuition in Section 5.4.

It is important to note that it is not clear to us whether Dinur's gap amplification theorem [18] can be used to obtain such a result from our techniques. The main reason is that for our proof to go through, we need an *arbitrarily* small (but still constant) soundness error, which is something that Dinur's protocol cannot achieve, as observed by [12]. The original PCP construction of [2, 3] is also not useful due to the same reason. However, by using the constructions in [39, 6], we can achieve a very succinct MIP protocol for NP that is sound against leakage. This is because these constructions have the property of achieving an arbitrarily small soundness error by paying a moderate cost in the alphabet size and the proof length. To prove Theorem 3.3, we make use of the recent breakthrough of [6]. We can have a similar result with a question length of $2\log n + O(\sqrt{\log n})$ from [39], which is larger than our result of $2\log n + O(\log\log n)$. To prove Theorem 3.4, we will make use of a similar argument, but we need to utilize a PCP construction that achieves $1/\mathrm{poly}(n)$ soundness

error. To this end, we will use the construction of Dinur et al. [19] which has $\text{poly}(\log \log n)$ query complexity, but the soundness error vanishes inverse polynomially in n. This will allow us to afford $O(\log n)$ bits of leakage, without incurring too much overhead in the question size. However, the super-constant query complexity causes the soundness error to not be a constant anymore.

This result also has interesting connections to the Sliding Scale Conjecture in the PCP literature. We will state it below and refer the interested reader to the survey of Moshkovitz [38] for more information.

▶ **Conjecture 3.5** (Sliding Scale Conjecture [8])**.** *For any* $1/\text{poly}(n) \leq s < 1$ *and for all languages in* NP*, there exists a PCP verifier* V *with perfect completeness and soundness* s *such that* V *tosses* $\log n$ *random coins and makes a constant number of queries to a proof over an alphabet of size* $1/\text{poly}(s)$*.*

Assuming that this conjecture is true, we can use it in our technique to improve the soundness error in Theorem 3.4 to a constant, and keep the answer size $O(\log n)$. Hence, it is possible to show that ruling out the existence of such succinct MIP protocols with $O(\log n)$ bits of leakage for NP can give a resolution for Conjecture 3.5. It is an interesting open problem to figure out whether we can achieve constant soundness using other techniques. We will discuss more about open problems in Section 6.

3.3 Comparison of the two approaches

We would like to emphasize that the two approaches that we present in this work are not entirely unrelated. Indeed, the key idea behind both is to reduce the soundness error in a multi-prover proof system to a carefully chosen small constant and utilize this to show that even a bounded amount of leakage does not help the cheating provers to succeed. Note that to obtain a succinct MIP protocol for NP, we can allow leakage directly from our first technique by utilizing a 2-query low-soundness PCPs [39, 6][3] (see Corollary 5.10 for more details). In this case, the parameters are even better than what we have in Theorem 3.3 from our second technique. However, the strength of our second technique is that we can deal with general low-soundness PCPs whose number of queries is more than 2, and in particular, this allows us to achieve interesting parameters such as Theorem 3.4.

Despite the similarity, there are differences in the power of these techniques. First, it is important to note that the second technique is only applicable when we make the additional assumption that the leakage is one-way. Also, it is not clear how to apply the second technique to achieve a result similar to the one presented in Theorem 3.2. It could be the case that one achieves a similar result assuming the quantum PCP conjecture is true, but we will not discuss this possibility here and leave it to future work. Note that although the first technique is more powerful in capturing different scenarios, it inherently suffers from a polynomial blow-up in the game size. Our second technique suggests that this blow-up is not optimal.

4 Related works

Some decays of values by parallel repetition (the parallel repetition theorem) for classical values (e.g., [45, 24, 44]) and quantum values (e.g., [17, 33, 28, 49, 7]) are known.

[3] We cannot obtain similar results using PCPs whose soundness error is at least 1/2 [2, 3, 18] for the same reason we discussed earlier.

We usually prohibit communication between the provers during the protocol. One exception, to our knowledge, is a model considered by Ben-Or, Hassidim and Pilpel [10]. They considered a variant of quantum MIP, where the provers do not share entanglement, the communication between the verifier and the provers is quantum, but the provers are allowed to have unlimited classical communication between them. They showed that any language in NEXP can be recognized in the model, which is the same as MIP.

Communication complexity of parallel repetition of non-local games has been considered. Jain and Kundu [26] showed a direct product theorem for the entanglement-assisted communication complexity, and its application for device-independent quantum key distribution [27]. Hasegawa, Le Gall and Modanese [23] considered the parallel repetition of Magic Square games [37, 43, 1], and showed that even if $o(n)$ EPR pairs are allowed to share prior to a communication protocol, $\Omega(n)$ qubits of communication is required to solve all Magic Square games with high probability.

5 Proofs for main results

In the next part, we formally define the notions of multi-prover interactive proofs with leakage, and then proceed to show our results on the power of these complexity classes. Relevant preliminaries are provided in Section 2 of the full version.

5.1 Definitions of MIP and MIP* protocols with leakage

Let us define a $\mathsf{MIP}[q, a, \ell]$ protocol as a two-prover one-round MIP protocol with question size q and answer size a, where the provers conduct ℓ-bit communication between themselves based on their questions received from the verifier. The provers then send some answers back to the verifier based on their communication and questions. We also use $\mathsf{MIP}[q, a, \ell]$ as the set of languages recognized by the $\mathsf{MIP}[q, a, \ell]$ protocols. A formal definition follows.

▶ **Definition 5.1** (MIP with leakage). *We say a language L is in $\mathsf{MIP}_{c,s}[q, a, \ell]$ if and only if, for $z \in \{0,1\}^*$, there exists a family of non-local games $G_z = (\pi, \mathcal{X} \times \mathcal{Y}, \mathcal{A} \times \mathcal{B}, \mathsf{V})$ satisfying the following properties:*

1. *The question size is q and the answer size is a.*
2. *There is an algorithm running in $\mathrm{poly}(|z|)$-time, given z, to sample questions $(x, y) \in \mathcal{X} \times \mathcal{Y}$ from π.*
3. *There is an algorithm running in $\mathrm{poly}(|z|)$-time, given z and $(x, y, a, b) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{A} \times \mathcal{B}$, to output $\mathsf{V}(a, b|x, y) \in \{0, 1\}$.*
4. *If $z \in L$, $\omega_c(G_z) = c$.*
5. *If $z \notin L$, and even if the provers are allowed to communicate at most ℓ bits, we must have $\omega_c(G_z) \leq s$.*

When we do not specify completeness c and soundness s, it should be assumed that $c = 1$ and $s = 1/2$.

Similarly, let us define a $\mathsf{MIP}^*[q, a, \ell]$ protocol and $\mathsf{MIP}^*[q, a, \ell]$ as the set of languages recognized by $\mathsf{MIP}^*[q, a, \ell]$ protocols. In a $\mathsf{MIP}^*[q, a, \ell]$ protocol, the provers can exploit the shared entanglement for their communication.

▶ **Definition 5.2** (MIP* with leakage). *We say a language L is in $\mathsf{MIP}^*[q, a, \ell]$ if and only if, for an input $z \in \{0,1\}^*$, there exists a family of non-local games $G_z = (\pi, \mathcal{X} \times \mathcal{Y}, \mathcal{A} \times \mathcal{B}, \mathsf{V})$ satisfying the following properties:*

1. *The question size is q and the answer size is a.*
2. *There is an algorithm running in* $\mathrm{poly}(|z|)$*-time, given z, to sample questions $(x, y) \in \mathcal{X} \times \mathcal{Y}$ from π.*
3. *There is an algorithm running in* $\mathrm{poly}(|z|)$*-time, given z and $(x, y, a, b) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{A} \times \mathcal{B}$, to output $\mathsf{V}(a, b|x, y) \in \{0, 1\}$.*
4. *If $z \in L$, $\omega_q(G_z) = 1$.*
5. *If $z \notin L$, and even if the provers are allowed to communicate at most ℓ bits, $\omega_q(G_z) \leq \frac{1}{2}$.*

Note that for $\ell \geq q$, $\mathsf{MIP}[q, a, \ell] \subseteq \mathsf{AM}$ and $\mathsf{MIP}^*[q, a, \ell] \subseteq \mathsf{AM}$ because the provers can share their questions by communication and the game value becomes trivial.

We also define a notion of multi-prover interactive proofs, in which the communication between the provers is one-way. In this scenario, the verifier asks a question from the first prover, then the first prover answers the question to the verifier and is allowed to send a bounded number of bits to the second prover. After that, the verifier will ask the second question from the second prover. This time, instead of using the terminology of non-local games, we give the definition in terms of the task of verifying a decision problem. This will make the corresponding results simpler to present. The formal definition follows.

▶ **Definition 5.3.** *A language L is in $\mathsf{MIP}^{\rightarrow}_{c,s}[q, a, \ell]$ if and only if there exists a two-prover one-round interactive proof for L, in which the verifier V runs in time* $\mathrm{poly}(|x|)$ *for every instance $x \in \{0, 1\}^*$ and satisfies the following properties:*

- Completeness: *If $x \in L$, then after exchanging one round of messages first with P_1 and then with P_2, the verifier V accepts with probability at least c.*
- Soundness: *If $x \notin L$, for every set of malicious provers P_1 and P_2, in which P_1 is allowed to send ℓ bits to P_2 after receiving the question from V, the verifier will accept with probability at most s.*

Additionally, the question size is q bits, and the answer size is at most a bits.

5.2 The power of MIP with leakage

We say a communication protocol is a (randomized) public coin communication protocol if the parties share some randomness before they receive inputs. A public coin communication protocol can be regarded as a convex mixture of deterministic protocols over the probabilistic distribution of public coins.

The following statement follows from the parallel repetition theorem. A similar statement is known in previous work (see, e.g., Corollary 7.3 in [26] and Theorem 8 in [23]).

▶ **Lemma 5.4.** *Let G be a non-local game such that $\omega_c(G) = 1 - \varepsilon$ for a constant $\varepsilon > 0$ and the answer size is constant. Let $\mathcal{P}$ be a public-coin 2-way classical communication protocol for $G^{\otimes N}$ and suppose that the communication amount of $\mathcal{P}$ is cN for a sufficiently small constant $c > 0$. Then, the success probability of $\mathcal{P}$ is at most $2^{-\Omega(N)}$.*

Proof. Let $c > 0$ be a sufficiently small constant and $\mathcal{P}$ be a public-coin classical 2-way communication protocol for $G^{\otimes N}$ whose communication amount is cN. Let p be the success probability of $\mathcal{P}$.

Let us consider a communication protocol $\mathcal{P}'$ in which the parties share uniformly random cN bits in addition to the public coins, but have no communication. In the beginning of the protocol, Alice and Bob check that their random bits are equal to the cN bits of the communication protocol, and if not, they abort. Let $\not\perp_A$ be the event that Alice does not abort during the protocol $\mathcal{P}$, and let $\not\perp_B$ be the event that Bob does not abort during the protocol $\mathcal{P}$.

Define $W_{\mathcal{P}'}$ to be the event that $\mathcal{P}'$ outputs wins all the games (hence, without abortions). Then, from the definition of $\mathcal{P}'$, we have

$$\Pr[W_{\mathcal{P}'} | \not\perp_A, \not\perp_B] = p \ .$$

Since randomness is sampled uniformly and independently from the public coins of the communication protocols, we also have

$$\Pr[\not\perp_A, \not\perp_B] = 2^{-cN} \ ,$$

which implies

$$\Pr[W_{\mathcal{P}'}] = 2^{-cN} p \ .$$

Considering abortions, the success probability of $\mathcal{P}'$ is at least $2^{-cN}p$. From the parallel repetition theorem with exponential decay for classical values (see Lemma 2.1 in the full version for a formal statement), and since shared randomness does not change the value of the games, the success probability of $\mathcal{P}'$ is also at most $2^{-c'N}$ for a constant c'. Taking $c > 0$ such that $c < c'$ and combining the upper and lower bound of the success probability, $p \leq 2^{-(c'-c)N}$. ◀

▶ **Theorem 5.5.** *For any $\ell \in \mathbb{N}$ such that $\ell \leq \mathrm{poly}(n)$,*

$$\mathsf{MIP}[O(\ell \log n), O(\ell), \ell] \supseteq \mathsf{NP} \ .$$

Proof. Let G be a non-local game corresponding to the MIP protocol for NP from the PCP theorem. The question size of G is $O(\log n)$ and the answer size is $O(1)$ and perfect completeness holds for G (see Fact 2.16 and Fact 2.17 in the full version for a formal statement). Then, let $k \in \mathbb{N}$ be a sufficiently large constant and consider a MIP protocol corresponding to $G^{\otimes k\ell}$ ($(k\ell)$-wise parallel repetition of G). Since perfect completeness holds for G, it also holds in the parallel repetition game. Applying Lemma 5.4 will prove the soundness also holds. The question and answer size of $G^{\otimes k\ell}$, given that k is chosen to be some large constant, are $O(\ell \log n)$ and $O(\ell)$ respectively. ◀

▶ **Theorem 5.6.** *For any $\ell \in \mathbb{N}$ such that $\ell \leq \mathrm{poly}(n)$,*

$$\mathsf{MIP}[O(\ell) \cdot \mathrm{poly}(n), O(\ell), \ell] \supseteq \mathsf{NEXP} \ .$$

Proof. Let G be a non-local game corresponding to the MIP protocol for NEXP from the scaled-up version of the PCP theorem. The question size of G is $O(\mathrm{poly}(n))$ and the answer size is $O(1)$ and perfect completeness holds for G (see Fact 2.18 in the full version for a formal statement). Then, let $k \in \mathbb{N}$ be a sufficiently large constant and consider a MIP protocol corresponding to $G^{\otimes k\ell}$. Since perfect completeness holds for G, it also holds by parallel repetition. Applying Lemma 5.4 will show that the soundness also holds. The question and answer size of $G^{\otimes k\ell}$, given that k is chosen to be some large constant, are $O(\ell) \cdot \mathrm{poly}(n)$ and $O(\ell)$ respectively. ◀

5.3 The power of MIP* with leakage

Next, we consider MIP* protocols with leakage. Jain and Kundu [26] showed the following direct product theorem for the entanglement-assisted quantum communication complexity.

▶ **Lemma 5.7** (Part of Theorem 1.1 in [26]). *For any predicate* V *on* $(\mathcal{A}\times\mathcal{B})\times(\mathcal{X}\times\mathcal{Y})$ *and any product probability distribution* π *on the question set* $\mathcal{X}\times\mathcal{Y}$*, let* $\mathcal{P}$ *be an interactive entanglement-assisted 2-party quantum communication protocol for* $G^{\otimes N}$*. If* $\mathcal{P}$ *has total communication* cN *for* $c<1$*, the success probability of* $\mathcal{P}$ *is at most*

$$\left(1-\frac{\nu}{2}+2\sqrt{c}\right)^{\Omega\left(\frac{\nu^2 N}{\log(|\mathcal{A}|\cdot|\mathcal{B}|)}\right)}$$

where $\nu=1-\omega_q(G)$.

The above statement can also be regarded as a parallel repetition theorem for the quantum value with leakage.

Natarajan and Zhang [42] showed that there exists a quantum free game protocol for RE. Denote a two-prover one-round MIP^* protocol as an $\mathsf{AM}^*(2)$ protocol if the question distribution is uniform. Note that uniform distribution is a special case of product distributions.

▶ **Lemma 5.8** ([42], see also Theorem 2.2 in [36]). *There exists an* $\mathsf{AM}^*(2)$ *protocol for the Halting Problem with constant length questions and* $\mathrm{poly}\log(n)$ *length answers, perfect completeness and soundness* $\frac{1}{2}$.

▶ **Theorem 5.9.** *There exist constants* c *and* c' *such that for any* $\ell\le\mathrm{poly}(n)$*,*

$$\mathsf{MIP}^*[O(\max\{\ell,\log^c n\}),O(\max\{\ell,\log^c n\}\cdot\log^{c'} n),\ell]\supseteq\mathsf{RE}\ .$$

Proof. Let G be a non-local game corresponding to the MIP^* protocol for RE in Lemma 5.8. Then, let c and k be sufficiently large constants and consider a MIP^* protocol corresponding to $G^{\otimes k\cdot\max\{\ell,\log^c n\}}$. Since perfect completeness holds in Lemma 5.8, it also holds in the parallel repetition game. Applying Lemma 5.7 will prove that soundness[4] also holds. The question size of $G^{\otimes k\cdot\max\{\ell,\log^c n\}}$ is $O(\max\{\ell,\log^c n\})$ and the answer size is $O(\max\{\ell,\log^c n\}\cdot\log^{c'} n)$ for some $c'>0$ from Lemma 5.8. ◀

5.4 Low-soundness PCPs and one-way leakage

We have seen from the proof of Lemma 5.4 that we can allow ℓ bits of leakage as long as we have a soundness parameter $s\le 2^{-\ell}$. Therefore, from the game version of the PCP theorem [6] (see Fact 2.16 in the full version for a formal statement), we have the following corollary.

▶ **Corollary 5.10.** $\mathsf{MIP}_{1,s\cdot 2^{\ell}}[\log n+O(\log\log n),O(1),\ell]\supseteq\mathsf{NP}$.

However, this method only works for 2-query PCPs. In this section, we present a more general method for converting a PCP into a MIP protocol that is resilient to leakage. First, we prove the following theorem for arbitrary CSPs.

▶ **Theorem 5.11.** *Let* ℓ *be a positive constant, and* Ψ *be a CSP instance of size* $\mathrm{poly}(n)$ *on* n *variables, such that every constraint is of arity* k *over an alphabet* Σ*. Suppose* Ψ *is promised to be either fully satisfiable or any assignment to the variables will satisfy at most* $\frac{1}{2^{\ell+1}}$*-fraction of its constraints. Then,*

$$\Psi\in\mathsf{MIP}^{\rightarrow}_{1,1-\frac{1}{2k}}[O(\log n),k\log|\Sigma|,\ell]\ .$$

[4] To be precise, we here consider soundness against ℓ-bit communication which can be simulated using l-qubit communication.

Proof. The verifier for the satisfiability of the CSP instance Ψ works as follows.

Algorithm 1 MIP protocol for verifying the satisfiability of Ψ.

The verifier V having access to Ψ

1. Sample a constraint $e = (u_1, ..., u_k)$ of Ψ uniformly at random and send it P_1.
2. After receiving the assignments $a(u_1), ..., a(u_k)$ sample a random element $i \in [k]$.
3. Ask P_2 for the assignment to u_i.
4. Accept if and only if the answers of the two provers are consistent and if the assignment satisfies the constraint.

In the protocol above, the question size is $O(\log n)$ because the size of CSP is $\text{poly}(n)$, and thus the verifier can specify a constraint and a vertex with $O(\log n)$ bits.

It is clear that if Ψ is satisfiable, then the verifier accepts after interacting with honest provers. Now suppose that Ψ is unsatisfiable, and in particular, by the assumption of the theorem, we know that at any assignment will satisfy at most $2^{-(\ell+1)}$-fraction of the constraints. Therefore, for *every set of* 2^ℓ *assignments*, $1/2$ of the constraints are not satisfied by all of them.

We can assume the provers are deterministic. Therefore, P_2 upon receiving the leakage ℓ fixes its strategy[5], which we denote by P_2^ℓ. Note that for each value of ℓ, the strategy P_2^ℓ corresponds to an assignment to the variables. Hence, regardless of the set of assignments the provers agree on before the protocol starts, a $\frac{1}{2}$-fraction of constraints will remain unsatisfied. Therefore, in the first step, with probability $1/2$, the verifier samples such a constraint and forces P_1 to lie. After P_2 receives the leakage, there is at least $1/k$ chance that the verifier samples a variable that P_2 will answer inconsistently with respect to P_1's answer. Hence, the soundness is at most $1 - 1/2k$, which concludes the proof. ◀

In Section 3.5 of the full version, we provide concrete examples of these theorems.

6 Concluding remarks and open problems

In this paper, we investigated the problem of finding the complexity of multi-prover interactive proof systems where the provers are allowed to communicate a bounded amount of messages to each other. We proposed and developed two techniques, one based on the parallel repetition results and the other based on gap amplification and low-soundness PCPs, to address the problem. In summary, using these techniques, we showed that we can still maintain the power of multi-prover interactive proof systems by increasing the question and answer sizes moderately. These results show that in particular, for real-world applications, we may not need to enforce spatial separation between provers to cross-examine them. Instead, we can still keep the verification power as long as we have a predetermined bound on the amount of communication they are allowed to have. We also showed an interesting connection between this scenario and the Sliding Scale Conjecture in the PCP literature, hoping that this might be a novel approach for making progress on that front. In addition, there are three ways one can interpret our results:

[5] Here, we are slightly abusing the notation and using ℓ for the leakage content as well as its length.

Adaptiveness in MIPs

It is a well-known fact that the power of multi-prover interactive proofs, and in particular, one-round MIPs lies in the fact that the provers' answers are non-adaptive. Essentially, this follows from the fact that each prover is blind to the question the other prover receives. At the same time, if we allow complete adaptiveness between the messages, we are back to the power of single-prover interactive proofs, which is incomparable to the former (AM vs. NEXP). However, one could ask, *how much adaptiveness* is required for the cheating provers to be able to convince the verifier to accept an incorrect statement? Could it be the case that a single bit of adaptiveness is sufficient? Our results address this problem negatively. In a way, we show that if the amount of adaptiveness is bounded, then the power of multi-prover interactive proofs remains intact. However, there is more to be done, as it is still unclear where this transition in complexity appears.

Another application of PCPs

PCPs undoubtedly have many interesting applications in complexity theory and cryptography, and our result can also be interpreted as a new one. We showed that one of the applications of reducing soundness error in PCPs, whether through parallel repetition or through novel constructions, can be viewed as mildly relaxing the no-communication assumption in the MIP view. Whether this can be useful for showing no-go results in PCP constructions or whether this idea can be used in real-world cryptographic constructions is a possibility that we leave for future research. In particular, it would be very appealing to know whether our approach can bring us any closer to finding a resolution for the Sliding Scale Conjecture in the PCP literature.

Proof of quantumness

Recently, several protocols for proving quantumness have been proposed based on compiled non-local games [31, 41, 34]. The core mechanism involves simulating a two-prover game using a single (supposedly quantum) prover by homomorphically encrypting the first question [13, 35]. Ideally, an honest quantum prover can respond within the encrypted domain, whereas a dishonest classical prover cannot extract information from the first question to gain an advantage on the second. However, the security of these protocols relies on cryptographic assumptions. Our work provides an alternative approach. In the compiled setting, leakage between provers effectively requires the adversary to store information in memory for use upon receiving the subsequent question. This allows us to frame the compiled non-local game as sound against *bounded-space* adversaries. While our current bounds may not yet yield a practical real-world protocol, given the low cost of classical memory, they demonstrate a theoretical path toward replacing cryptographic assumptions with complexity-theoretic ones, specifically space-complexity bounds.

Open problems

Our work raises several open problems, which we leave for future work. We conclude this paper by naming a few of them here.

- An immediate open problem is to find the optimal tradeoff for allowing constant bits (or qubits) of leakage.
- Our technique based on low-soundness PCPs requires the additional assumption that the communication is one-way. We believe that this requirement is not necessary, and it is a limitation of the technique. It would be interesting to generalize these tools so that they do not rely on this assumption anymore.

- Another open problem is to find a succinct MIP protocol for NP that has constant soundness against $O(\log n)$ bits of leakage. As mentioned in Section 3, we can achieve this if there is a positive answer to the Sliding Scale Conjecture. However, we do not know whether the other direction holds. Is it possible to show such a result without relying on Conjecture 3.5?

References

1. Padmanabhan K Aravind. A simple demonstration of Bell's theorem involving two observers and no probabilities or inequalities. *arXiv preprint quant-ph/0206070*, 2002.
2. Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM (JACM)*, 45(3):501–555, 1998. doi:10.1145/278298.278306.
3. Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM (JACM)*, 45(1):70–122, 1998. doi:10.1145/273865.273901.
4. László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational complexity*, 1:3–40, 1991. doi:10.1007/BF01200056.
5. László Babai and Shlomo 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. doi:10.1016/0022-0000(88)90028-1.
6. Mitali Bafna, Dor Minzer, Nikhil Vyas, and Zhiwei Yun. Quasi-linear size PCPs with small soundness from HDX. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC 2025)*, pages 45–53, 2025. doi:10.1145/3717823.3718197.
7. Mohammad Bavarian, Thomas Vidick, and Henry Yuen. Anchored parallel repetition for nonlocal games. *SIAM Journal on Computing*, 51(2):214–253, 2022. doi:10.1137/21M1405927.
8. Mihir Bellare, Shafi Goldwasser, Carsten Lund, and Alexander Russell. Efficient Probabilistically Checkable Proofs and Applications to Approximations. In *Proceedings of the 25th annual ACM symposium on Theory of computing (STOC 1993)*, pages 294–304, 1993. doi:10.1145/167088.167174.
9. Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: how to remove intractability assumptions. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC 1988)*, pages 113–131, 1988. doi:10.1145/62212.62223.
10. Michael Ben-Or, Avinatan Hassidim, and Haran Pilpel. Quantum Multiprover Interactive Proofs with Communicating Provers. *SIAM Journal on Computing*, 43(3):987–1011, 2014. doi:10.1137/090777724.
11. Charles H Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K Wootters. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters*, 70(13):1895, 1993.
12. Andrej Bogdanov. Gap amplification fails below 1/2. *Comment on ECCC TR05-046, can be found at http://eccc. uni-trier. de/eccc-reports/2005/TR05-046/commt01. pdf*, 2005.
13. Zvika Brakerski. Quantum FHE (almost) as secure as classical. In *Proceedings of the Annual International Cryptology Conference (CRYPTO 2018)*, pages 67–95. Springer, 2018. doi:10.1007/978-3-319-96878-0_3.
14. Mark Braverman, Anup Rao, Omri Weinstein, and Amir Yehudayoff. Direct products in communication complexity. In *Proceedings of 2013 IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS 2013)*, pages 746–755, 2013. doi:10.1109/FOCS.2013.85.
15. Amit Chakrabarti, Yaoyun Shi, Anthony Wirth, and Andrew Yao. Informational complexity and the direct sum problem for simultaneous message complexity. In *Proceedings of 42nd IEEE Symposium on Foundations of Computer Science (FOCS 2001)*, pages 270–278, 2001.

16 Richard Cleve, Peter Hoyer, Benjamin Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *Proceedings of 19th IEEE Annual Conference on Computational Complexity (CCC 2004)*, pages 236–249, 2004.

17 Richard Cleve, William Slofstra, Falk Unger, and Sarvagya Upadhyay. Perfect parallel repetition theorem for quantum XOR proof systems. *Computational Complexity*, 17:282–299, 2008. doi:10.1007/S00037-008-0250-4.

18 Irit Dinur. The PCP theorem by gap amplification. *Journal of the ACM (JACM)*, 54(3):12–es, 2007.

19 Irit Dinur, Prahladh Harsha, and Guy Kindler. Polynomially Low Error PCPs with polyloglog n Queries via Modular Composition. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC 2015)*, pages 267–276, 2015. doi:10.1145/2746539.2746630.

20 Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM (JACM)*, 38(3):690–728, 1991. doi:10.1145/116825.116852.

21 Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The Knowledge Complexity of Interactive Proof Systems. *SIAM Journal on Computing*, 18(1):186–208, 1989. doi:10.1137/0218012.

22 Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In *Proceedings of the 18th annual ACM symposium on Theory of computing (STOC 1986)*, pages 59–68, 1986. doi:10.1145/12130.12137.

23 Atsuya Hasegawa, François Le Gall, and Augusto Modanese. Maximum Separation of Quantum Communication Complexity With and Without Shared Entanglement. *arXiv preprint*, 2025. doi:10.48550/arXiv.2505.16457.

24 Thomas Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing*, 5(8):141–172, 2009. doi:10.4086/TOC.2009.V005A008.

25 Tsuyoshi Ito and Thomas Vidick. A multi-prover interactive proof for NEXP sound against entangled provers. In *Proceedings of 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS 2012)*, pages 243–252, 2012. doi:10.1109/FOCS.2012.11.

26 Rahul Jain and Srijita Kundu. A direct product theorem for quantum communication complexity with applications to device-independent cryptography. *SIAM Journal on Computing*, 54(4):964–1020, 2025. doi:10.1137/23M1549353.

27 Rahul Jain, Carl A. Miller, and Yaoyun Shi. Parallel Device-Independent Quantum Key Distribution. *IEEE transactions on information theory*, 66(9):5567–5584, 2020. doi:10.1109/TIT.2020.2986740.

28 Rahul Jain, Attila Pereszlényi, and Penghui Yao. A parallel repetition theorem for entangled two-player one-round games under product distributions. In *Proceedings of 2014 IEEE 29th Conference on Computational Complexity (CCC 2014)*, pages 209–216, 2014. doi:10.1109/CCC.2014.29.

29 Rahul Jain, Jaikumar Radhakrishnan, and Pranab Sen. Prior entanglement, message compression and privacy in quantum communication. In *Proceedings of 20th Annual IEEE Conference on Computational Complexity (CCC 2005)*, pages 285–296, 2005. doi:10.1109/CCC.2005.24.

30 Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP*= RE. *arXiv preprint*, 2020. arXiv:2001.04383.

31 Yael Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Lisa Yang. Quantum advantage from any non-local game. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC 2023)*, pages 1617–1628, 2023. doi:10.1145/3564246.3585164.

32 Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, Ben Toner, and Thomas Vidick. Entangled games are hard to approximate. *SIAM Journal on Computing*, 40(3):848–877, 2011. doi:10.1137/090751293.

33 Julia Kempe and Thomas Vidick. Parallel repetition of entangled games. In *Proceedings of the 43rd annual ACM symposium on Theory of computing (STOC 2011)*, pages 353–362, 2011. doi:10.1145/1993636.1993684.

34 Alexander Kulpe, Giulio Malavolta, Connor Paddock, Simon Schmidt, and Michael Walter. A bound on the quantum value of all compiled nonlocal games. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC 2025)*, pages 222–233, 2025. doi:10.1145/3717823.3718237.

35 Urmila Mahadev. Classical homomorphic encryption for quantum circuits. *SIAM Journal on Computing*, 52(6):FOCS18–189, 2020.

36 Kieran Mastel and William Slofstra. Two prover perfect zero knowledge for MIP*. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC 2024)*, pages 991–1002, 2024. doi:10.1145/3618260.3649702.

37 N David Mermin. Simple unified form for the major no-hidden-variables theorems. *Physical review letters*, 65(27):3373, 1990.

38 Dana Moshkovitz. Sliding Scale Conjectures in PCP. *SIGACT News*, 50(3):25–33, 2019. doi:10.1145/3364626.3364632.

39 Dana Moshkovitz and Ran Raz. Two-query PCP with subconstant error. *Journal of the ACM (JACM)*, 57(5):1–29, 2008.

40 Anand Natarajan and John Wright. NEEXP is contained in MIP*. In *Proceedings of 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS 2019)*, pages 510–518, 2019. doi:10.1109/FOCS.2019.00039.

41 Anand Natarajan and Tina Zhang. Bounding the quantum value of compiled nonlocal games: from CHSH to BQP verification. In *Proceedings of 2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS 2023)*, pages 1342–1348, 2023. doi:10.1109/FOCS57990.2023.00081.

42 Anand Natarajan and Tina Zhang. Quantum free games. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC 2023)*, pages 1603–1616, 2023. doi:10.1145/3564246.3585208.

43 Asher Peres. Incompatible results of quantum measurements. *Physics Letters A*, 151(3-4):107–108, 1990.

44 Anup Rao. Parallel repetition in projection games and a concentration bound. *SIAM Journal on Computing*, 40(6):1871–1891, 2011. doi:10.1137/080734042.

45 Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, 1998. doi:10.1137/S0097539795280895.

46 Alexander A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385–390, 1992. doi:10.1016/0304-3975(92)90260-M.

47 Ben W Reichardt, Falk Unger, and Umesh Vazirani. Classical command of quantum systems. *Nature*, 496(7446):456–460, 2013. doi:10.1038/NATURE12035.

48 Adi Shamir. IP = PSPACE. *Journal of the ACM (JACM)*, 39(4):869–877, 1992. doi:10.1145/146585.146609.

49 Henry Yuen. A Parallel Repetition Theorem for All Entangled Games. In *Proceedings of 43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, volume 55 of *LIPIcs*, pages 77:1–77:13, 2016. doi:10.4230/LIPIcs.ICALP.2016.77.