Delegation for Bounded Space

[Extended Abstract]

Yael Tauman Kalai
Microsoft Research
New England
yael@microsoft.com

Ran Raz[*]
Weizmann Institute of Science
Rehovot, Israel
ran.raz@weizmann.ac.il

Ron D. Rothblum[†]
Weizmann Institute of Science
Rehovot, Israel
ron.rothblum@weizmann.ac.il

ABSTRACT

We construct a 1-round delegation scheme for every language computable in time $t = t(n)$ and space $s = s(n)$, where the running time of the prover is $\mathsf{poly}(t)$ and the running time of the verifier is $\tilde{O}(n + \mathsf{poly}(s))$ (where $\tilde{O}$ hides $\mathsf{polylog}(t)$ factors).

The proof exploits a curious connection between the problem of *computation delegation* and the model of *multi-prover interactive proofs that are sound against no-signaling (cheating) strategies*, a model that was studied in the context of multi-prover interactive proofs with provers that share quantum entanglement, and is motivated by the physical principle that information cannot travel faster than light.

For any language computable in time $t = t(n)$ and space $s = s(n)$, we construct MIPs that are sound against no-signaling strategies, where the running time of the provers is $\mathsf{poly}(t)$, the number of provers is $\tilde{O}(s)$, and the running time of the verifier is $\tilde{O}(s + n)$.

We then show how to use the method suggested by Aiello *et al.* (ICALP, 2000) to convert our MIP into a 1-round delegation scheme, by using a computational private information retrieval (PIR) scheme. Thus, assuming the existence of a sub-exponentially secure PIR scheme, we get our 1-round delegation scheme.

1. INTRODUCTION

The problem of delegating computation considers a setting where one party, the *delegator* (or *verifier*), wishes to delegate the computation of a function f to another party, the *worker* (or *prover*). The challenge is that the delegator may not trust the worker, and thus it is desirable to have the worker "prove" that the computation was done correctly. We require that verifying this proof will be significantly easier than doing the computation itself, that is, the delegator running time will be significantly smaller than the time complexity of f. Moreover, we require that the running time of the worker will not be much larger than the time complexity of f.

The problem of delegating computation became a central problem in cryptography, especially with the increasing popularity of cloud computing, where people (and weak devices) use cloud platforms to run their computations.

We focus on the problem of constructing *one-round* delegation protocols, where the delegator wants to verify a statement of the form $x \in \mathcal{L}$. The delegator sends x to the worker together with some query q; then the worker computes $b = \mathcal{L}(x)$, and based on the query q provides a *non-interactive* proof π for the fact that $b = \mathcal{L}(x)$. The delegator should be able to verify the correctness of the proof π very efficiently. And the worker should run in time polynomial in the time it takes to compute f. Throughout this work (similarly to all previous works that consider the problem of one-round delegation), the security requirement is against *computationally bounded* cheating workers. Namely, we consider the computational setting, where the security (i.e., soundness) of our scheme relies on a cryptographic assumption, and the guarantee is that any cheating worker, who cannot break the underlying assumption, cannot prove the correctness of an incorrect statement.

It is known that (under plausible cryptographic assumptions) any function in LOGSPACE-uniform NC has a one-round delegation scheme [16,23]. More generally, it is known that any function f that can be computed by a LOGSPACE-uniform circuit C of size $t = t(n)$ and depth $d = d(n)$ has a one-round delegation scheme where the communication complexity is $\mathsf{poly}(d, k, \log t)$, the running time of the delegator is $\mathsf{poly}(d, k, \log t) + n \cdot \mathsf{poly}(\log t)$, and the running time of the prover is $\mathsf{poly}(k, t)$, where k is the security parameter. Note however that for circuits with large depth d, this delegation scheme does not satisfy the efficiency criterion.

A fundamental question is: Does there exist an (efficient) 1-round delegation scheme for circuits with large depth? More ambitiously, does there exist an (efficient) 1-round delegation scheme for every function in P? There are several works that (partially) answer this question in the prepro-

[*] Part of this work was done when the author was a member at the Institute for Advanced Study, Princeton. Research supported by an Israel Science Foundation grant, by the I-CORE Program of the Planning and Budgeting Committee and the Israel Science Foundation, and by NSF grants number CCF-0832797, DMS-0835373.

[†] This research was partially supported by the Israel Science Foundation (grant No. 1041/08). Parts of this research were conducted while the author was an intern at Microsoft Research, New England.

STOC'13, June 1âĂŞ4, 2013, Palo Alto, California, USA.
Copyright 2013 ACM 978-1-4503-2029-0/13/06 ...$15.00.

cessing model, or under *non-falsifiable* assumptions.[1] We elaborate on these works in Section 1.4.

In this paper we show a 1-round delegation scheme for any language that can be computed in bounded space. More specifically, we show a delegation scheme for every language computable in time t and space s, where the running time of the verifier depends polynomially on s but only *polylogarithmically* on t. The running time of the prover is polynomial in t.

We note that several widely used algorithms have large depth, and yet only linear space. These include algorithms for linear programming and the perceptron algorithm, to name a few. In general, often algorithms with loops tend to have large depth and only small space.

Our delegation scheme exploits a connection to the seemingly unrelated model of multi-prover interactive proof systems (MIP) in which soundness holds even against *no-signaling* cheating provers. Loosely speaking, no-signaling provers are allowed to use arbitrary strategies (as opposed to local ones, where the reply of each prover is a function only of her own input), as long as their strategies cannot be used for communication between any two disjoint sets of provers.

Our delegation result follows by a new construction of an MIP with soundness against such no-signaling provers, together with a generic transformation of such an MIP into a delegation scheme, using a fully-homomorphic encryption scheme (FHE), or alternatively, a computational private information retrieval (PIR) scheme.

1.1 Multi-Prover Interactive Proofs with No-Signaling Provers

The study of MIPs that are secure against no-signaling provers was motivated by the study of MIPs with provers that share entangled quantum states. Recall that no-signaling provers are allowed to use arbitrary strategies, as long as their strategies cannot be used for communication between any two disjoint sets of provers. By the physical principle that information cannot travel faster than light, a consequence of Einstein's special relativity theory, it follows that all the strategies that can be realized by provers that share entangled quantum states are no-signaling strategies.

Moreover, the principle that information cannot travel faster than light is a central principle in physics, and is likely to remain valid in any future ultimate theory of nature, since its violation means that information could be sent from future to past. Therefore, soundness against no-signaling strategies is likely to ensure soundness against provers that obey a future ultimate theory of physics, and not only the current physical theories that we have, that are known to be incomplete.

The study of MIPs that are secure against no-signaling provers is very appealing also because no-signaling strategies have a simple mathematical characterization.

Loosely speaking, in a no-signaling strategy the answer given by each prover is allowed to depend on the queries to all other provers, as long as for any subset of provers S, and any queries given to the provers in S, the distribution of the answers given by the provers in S is independent of all the other queries. Thus, the answer of each prover can depend on the queries to all other provers as a function, but not as a random variable.

More formally, fix any MIP consisting of n provers, and fix any set of cheating provers $\{P_1^*, \ldots, P_n^*\}$ who may see each other's queries (and thus each answer may depend on the queries sent to all the provers). The provers are said to be *no signaling* if for every subset of provers $\{P_i^*\}_{i \in S}$, and for every two possible query sets $\{q_i\}_{i \in [n]}$ and $\{q_i'\}_{i \in [n]}$ such that $q_i = q_i'$ for every $i \in S$, it holds that the distributions of answers $\{a_i\}_{i \in S}$ and $\{a_i'\}_{i \in S}$ are *identical*, where $\{a_i\}_{i \in S}$ is the answers of the provers in S corresponding to the queries $\{q_i\}_{i \in [n]}$, and $\{a_i'\}_{i \in S}$ is the answers of the provers in S corresponding to the queries $\{q_i'\}_{i \in [n]}$. If we have the slightly weaker guarantee that these two distributions are statistically close, then we say that the provers are *statistically no-signaling*. More specifically, if these two distributions are δ-close, then we say that the provers are δ-no-signaling. We refer the reader to Section 2.3 for details.

No-signaling strategies were first studied in physics in the context of Bell inequalities by Khalfin and Tsirelson [25] and Rastall [32], and they gained much attention after they were reintroduced by Popescu and Rohrlich [31]. MIPs that are secure against no-signaling provers were extensively studied in the literature (see for example [3, 4, 19–21, 24, 34]). However, their precise power is still not known. It is known that they contain PSPACE [20] and are contained in EXP.[2] We note that known constructions for MIPs that are secure against no-signaling strategies for PSPACE [21] are *inefficient*, in the sense that the provers in the MIP protocol run in exponential time (even if the computation takes much less time). The blowup in the running time of the provers is particularly undesirable for applications (such as cryptographic applications).

In this work, we show how to construct MIPs that are secure against no-signaling strategies (and more generally, statistically no-signaling strategies) for all of PSPACE, where the provers are *efficient*; i.e., they run in time that is polynomial in the computation time. Specifically, for any language computable in time $t = t(n)$ and space $s = s(n)$, we construct MIPs that are sound against no-signaling strategies, where the running time of the provers is $\mathsf{poly}(t)$, the number of provers is $\tilde{O}(s)$, and the running time of the verifier is $\tilde{O}(s + n)$.

1.1.1 *The Challenges in Proving Soundness Against No-Signaling Strategies*

It is tempting to consider known constructions for MIPs and to try to prove their soundness against no-signaling strategies. However, known constructions for MIPs are usually for NEXP (or the scaled down version for NP). Since MIPs that are secure against no-signaling strategies are contained in EXP, there is no hope to construct such MIPs for NEXP. In particular, all known MIPs for NEXP (or the scaled down version for NP) are not no-signaling.

Indeed, often the trivial strategy, where the provers simply choose random answers that make the verifier accept, is no-signaling. For example, consider the trivial 2-prover interactive proof for graph 3-coloring, where the verifier sends each prover a vertex in the graph, where with probability 1/2 the vertices are the same and with probability 1/2 there

[1] We note that under non-falsifiable assumptions, there are known positive results even for non-deterministic languages. The focus of this work is on deterministic languages.

[2] In a nutshell, one can find the best strategy for the provers by solving an exponential size linear program.

is an edge between these vertices, and the provers reply with the color of these vertices. Suppose the graph is not 3-colorable. We argue that the "random accepting strategy" is a no-signaling strategy that is accepted with probability 1. More specifically, the cheating strategy is the following: If both vertices are the same, choose a random color from the set of three legal colors, and both provers send this color to the verifier. Otherwise, choose two distinct random colors from the set of three legal colors, and each prover sends one of these colors to the verifier. This strategy is clearly accepted with probability 1. Moreover, it is a no-signaling strategy, since the distribution of answers of each prover is uniform, independent of the query to the other party.

This intuitive argument extends to more sophisticated MIPs and demonstrates the difficulty in proving soundness against no-signaling strategies. We note that the foregoing example is based on the work of Dwork *et al.* [11] discussed next.

1.2 From Multi-Prover Interactive Proofs to One-Round Delegation

Aiello *et al.* [1] (based on a heuristic of Biehl *et al.* [5]) suggested a method for converting a 1-round MIP into a 1-round delegation scheme, by using a PIR scheme, (or more generally, by using an FHE scheme).[3] In this work we choose to use the terminology of FHE schemes (as opposed to PIR schemes), because we find this terminology to be simpler. Nevertheless, all our results hold with PIR schemes as well.

In the resulting delegation scheme, the delegator computes all the queries of the MIP verifier, and sends all these queries to the worker, each encrypted under a fresh key, using an FHE scheme. The worker then computes the MIP prover's responses homomorphically over the encrypted answers, that is, underneath the layer of the FHE scheme.

Unfortunately, shortly after this method was introduced, Dwork *et al.* [11] showed that it may, in general, be insecure. In fact, although they used different terminology, [11] essentially show that the 3-coloring MIP, mentioned in Section 1.1.1, is not sound against no-signaling strategies and they argue that it may be possible to implement such strategies under an FHE. We elaborate further on the work of Dwork *et al.* and their connection to no-signaling soundness in Section 1.4.

Motivated by the work of Aiello *et al.*, Kalai and Raz [23] showed that a variant of this method can be used to securely convert any interactive proof into a one-round argument system.[4] The idea is simply to have the verifier send all its (say t) messages in the first round, in the following redundant form: For every $i \in [t]$, all the first i messages are encrypted using a fresh FHE key.[5] The work of [23], together with the interactive delegation scheme of Goldwasser *et al.* [16], gives rise to the 1-round delegation protocol for LOGSPACE-uniform NC, mentioned above.

We show that the method of Aiello *et al.* [1] is secure if the underlying MIP is sound against statistically no-signaling strategies. This result generalizes the work of [23], since any interactive proof can be seen as an MIP where the verifier sends his first i messages to prover i (it is quite easy to verify that the resulting MIP is secure against statistically no-signaling cheating provers). Moreover, this result significantly simplifies the one of [23], which implicitly converts the interactive proof into an MIP scheme and then applies the PIR to the resulting MIP scheme. We believe that due to the lack of the "correct" terminology, the result of [23] was relatively complicated, whereas this current result is significantly simpler and more general.

1.3 Summary of Our Results

We show that when applying the method of Aiello *et al.* [1] to an MIP that is sound against statistically no-signaling cheating provers, then the resulting 1-round delegation protocol is secure (assuming that the underlying PIR is secure against attackers of sub-exponential size).

Informal Theorem 1. *Assume the existence of an* FHE *scheme with sub-exponential security. Then, there exists an efficient way to convert any 1-round* MIP *that is sound against statistically no-signaling cheating provers into a secure 1-round delegation scheme, where the running time of the prover and verifier in the delegation scheme are proportional to the running time of the provers and verifier in the* MIP.

Thus, we reduced the cryptographic problem of constructing secure one-round delegation schemes, to the information theoretical problem of constructing MIP schemes that are secure against statistically no-signaling provers. Such a reduction allows us to "strip off" the cryptography, and to focus on an information theoretic question of constructing an MIP that is secure against statistically no-signaling provers.

As mentioned above, the problem of constructing an MIP that is secure against no-signaling provers, is a problem that is well studied by researchers in the field of quantum complexity. Indeed, it was shown in [20] how to construct such an MIP for all languages in PSPACE. However, the running time of the provers in these MIPs is exponential (even if the computation is significantly more efficient). Therefore, in the resulting 1-round delegation scheme, the prover runs in exponential time, which is more than we can afford in a delegation scheme (in which the prover must run in time polynomial in the time it takes to run the computation).

In this work, we construct an *efficient* MIP that is sound against (statistically) no-signaling strategies.

Informal Theorem 2. *For any language L computable by a Turing machine running in time $t = t(n)$ and space $s = s(n)$, there exists an* MIP *that is secure against statistically no-signaling adversaries. The provers in this* MIP *run in time* $\mathsf{poly}(t)$*, the number of provers is $\tilde{O}(s)$, and the verifier runs in time $\tilde{O}(s+n)$, where $\tilde{O}$ hides* $\mathsf{polylog}(t)$ *factors.*

In particular this theorem gives an MIP for PSPACE with no-signaling soundness *and* efficient provers. We note that our MIP has the additional property that the verifier does not need to know the entire input, but rather only needs to access a few points in the low-degree extension of the input (we refer the reader to full version for the definition of low-

[3] Actually, [1] suggested to use a PCP. However, the work of [11] shows that an MIP is more suitable.

[4] Recall that an argument system is an interactive proof system that only guarantees soundness against computationally bounded adversaries. A delegation scheme is an argument-system in which the focus is on the efficiency of the verifier and the (honest) prover.

[5] The reason the i'th message is encrypted together with the preceding messages, is since the prover's reply may depend on all these messages.

degree extension). This property, which was also a property of the [16] protocol, is important for applications such as memory delegation [8].

The two theorems above immediately yield the following corollary:

COROLLARY 3. *Assume the existence of an* FHE *scheme with sub-exponential security. Then, there exists a* 1*-round delegation scheme for any function computable by a Turing machine running in time* $t = t(n)$ *and space* $s = s(n)$*. The prover in this delegation scheme runs in time* $\mathsf{poly}(t)$ *and the verifier runs in time* $\tilde{O}(n + \mathsf{poly}(s))$*.*

We note that the bulk of technical contribution of this work is in proving Theorem 2. Indeed, as explained in Section 3, proving this theorem requires overcoming several technical hurdles that do not appear in classical MIP (or PCP). Theorem 1 is mainly a conceptual contribution. Its proof is relatively straightforward, but we find the connection between the seemingly unrelated concepts of delegation and no-signaling soundness to be intriguing.

As a special case, Theorem 2 also gives soundness against provers that share an entangled quantum state, since such provers are no-signaling. This gives a scheme to delegate computation to a group of workers that cannot communicate with each other (where the parameters are as in Theorem 2). The scheme is information theoretically secure even if the workers share an entangled quantum state. Moreover, the scheme remains secure in any future ultimate theory (that may extend quantum theory) as long as the no-signaling principle remains valid. We note that a recent breakthrough by Ito and Vidick constructs MIPs that are secure against provers that share entangled quantum states, for NEXP [22]. However, in their construction the provers are inefficient in the sense that their running time is super-polynomial in the running time of the initial computation.

1.4 Related Work

Our work is greatly inspired by the work of Aiello *et al.* [1], who propose a general methodology of constructing 1-round delegation schemes, by combining an MIP (or a PCP) with a (computational) PIR scheme. Also very relevant to our work is the work of Dwork *et al.* [11], who proved that this method is not sound, by giving an example of a PCP for which the resulting one-round delegation scheme is not sound, no matter which PIR scheme (or FHE scheme) is used.

Moreover, [11] define the notion of a "spooky interaction" which is a behavior of the cheating prover, which on the one hand does not directly contradict the security of the PIR, yet on the other hand is not consistent with answers based on PIR databases. Using our terminology, a spooky behavior is exactly a no-signaling distribution on prover answers that are computed "homomorphically" under the "encrypted" PIR queries.

More importantly, Dwork *et al.* also argue that the soundness of the [1] technique cannot essentially be based on a general MIP (or PCP). In fact, Dwork *et al.* implicitly considered the possibility of applying the [1] technique on MIPs that have no-signaling soundness. However, Dwork *et al.* (and [1]) were focused on constructing 1-round delegation schemes for non-deterministic languages (such as NEXP or the scaled down version of NP). Since they showed that non-deterministic languages cannot have soundness against no-signaling strategies (under very reasonable complexity assumptions), they reasoned that the [1] method is not useful for non-deterministic languages.

We also note that Gentry and Wichs [15] recently showed a negative result, proving that there does not exist a non-interactive delegation scheme for NP with a black-box proof of security under any falsifiable assumption.[6]

However, both of these negative results do not apply to our setting as our delegation scheme is not for all of NP, but rather for a class of languages in P (or, in the scaled up version, in EXP). Thus, by focusing on deterministic classes (as opposed to non-deterministic ones), we manage to show that the [1] method is indeed sound in some cases.

Related work on computation delegation. Beyond the works of [16, 23] which we mentioned earlier, there are many other works on delegating computation that are less relevant to this work. Let us mention a few. In the *interactive* setting, Kilian [26] constructed a 4-message delegation scheme for every function in NEXP. Micali [28] showed that in the so called *random oracle model* this result can be made non-interactive, by relying on the Fiat-Shamir paradigm [12]. There are also several results that construct non-interactive delegation schemes under *non-falsifiable* assumptions (as defined by Naor [29]). These works include [6, 7, 10, 14, 17, 18, 27] and more. Finally, we mentions a series of results that construct non-interactive delegation scheme in the *preprocessing model*, where the verifier is efficient only in the amortized setting. These results include [2, 9, 13, 30]. There are many other results that we do not mention, which consider various different models, or are concerned with practical efficiency.

Organization

In this extended abstract we only include a high-level overview of our proof, see the full version for details. In Section 2 we define MIPs, no-signaling strategies and argument systems. In Section 3 we provide a high-level overview of our techniques. Finally, in Section 4 we formally state our results.

2. PRELIMINARIES

2.1 Notation

For a vector $a = (a_1, \ldots, a_k)$ and a subset $S \subseteq [k]$, we denote by a_S the sequence of elements of a that are indexed by indices in S, that is, $a_S = (a_i)_{i \in S}$. In general, we denote by a_S a sequence of elements indexed by S, and we denote by a_i the i^{th} coordinate of a vector a.

For a distribution $\mathcal{A}$, we denote by $a \in_R \mathcal{A}$ a random variable distributed according to $\mathcal{A}$ (independently of all other random variables).

We will measure the distance between two distributions by their *statistical distance*, defined as half the l_1-distance between the distributions. We will say that two distributions are δ-*close* if their statistical distance is at most δ.

2.2 Multi-Prover Interactive Proofs

Let $\mathcal{L}$ be a language and let x be an input of length n. In a one-round k-prover interactive proof, k computationally unbounded provers, $P_1, \ldots, P_k$, try to convince a (probabilistic) $\mathsf{poly}(n)$-time verifier, V, that $x \in \mathcal{L}$. The input x is known to all parties.

[6] The model of [15] differs from our model in that they allow the (cheating) prover the additional power of choosing the instance x after seeing the first message sent by the verifier.

The proof consists of only one round. Given x and her random string, the verifier generates k queries, $q_1, \ldots, q_k$, one for each prover, and sends them to the k provers. Each prover responds with an answer that depends only on her own individual query. That is, the provers respond with answers $a_1, \ldots, a_k$, where for every i we have $a_i = P_i(q_i)$. Finally, the verifier decides whether to accept or reject based on the answers that she receives (as well as the input x and her random string).

We say that $(V, P_1, \ldots, P_k)$ is a one-round multi-prover interactive proof system (MIP) for $\mathcal{L}$ if the following two properties are satisfied:

1. **Completeness:** For every $x \in \mathcal{L}$, the verifier V accepts with probability 1, after interacting with $P_1, \ldots, P_k$.

2. **Soundness:** For every $x \notin \mathcal{L}$, and any (computationally unbounded, possibly cheating) provers $P_1^*, \ldots, P_k^*$, the verifier V rejects with probability $\geq 1-\epsilon$, after interacting with $P_1^*, \ldots, P_k^*$, where ϵ is a parameter referred to as the *error* or *soundness* of the proof system.

Important parameters of an MIP are the number of provers, the length of queries, the length of answers, and the error.

2.2.1 MIPs with Oracle

We will also consider the model of *one-round k-prover interactive proofs with oracle*, where the verifier V is given access to an oracle that computes some fixed function (that may depend on the language $\mathcal{L}$). We require that all queries, to the oracle and the provers, are done simultaneously.

For every n, let $\phi_n : \{0,1\}^{n'} \rightarrow \{0,1\}^{n''}$ be a function (where n', n'' depend on n). We allow the functions ϕ_n to depend on the language $\mathcal{L}$ (but not on the input x).

We define a *one-round multi-prover interactive proof system for $\mathcal{L}$, relative to the oracle $\{\phi_n\}$*, exactly as before, except that now the verifier V is a (probabilistic, $\mathsf{poly}(n)$-time) oracle machine that on input x of length n has free oracle access to the function ϕ_n. The verifier may base her accept/reject decision on queries to the oracle, but the oracle queries are not adaptive, and we do not allow the queries to the provers to depend on the answers of the oracle or the queries to the oracle to depend on the answers of the provers. In other words, we require that all queries, to the oracle and to the provers, are done simultaneously.

We require the same completeness and soundness properties as before.

2.3 No-Signaling MIPs

We will consider a variant of the MIP model, where the cheating provers are more powerful. In the MIP model, each prover answers her own query locally, without knowing the queries that were sent to the other provers. The no-signaling model allows each answer to depend on all the queries, as long as for any subset $S \subset [k]$, and any queries q_S for the provers in S, the distribution of the answers a_S, conditioned on the queries q_S, is independent of all the other queries.

Intuitively, this means that the answers a_S do not give the provers in S information about the queries of the provers outside S, except for information that they already have by seeing the queries q_S.

Formally, denote by D the alphabet of the queries and denote by Σ the alphabet of the answers. For every $q = (q_1, \ldots, q_k) \in D^k$, let $\mathcal{A}_q$ be a distribution over Σ^k. We think of $\mathcal{A}_q$ as the distribution of the answers for queries q.

We say that the family of distributions $\{\mathcal{A}_q\}_{q \in D^k}$ is *no-signaling* if for every subset $S \subset [k]$ and every two sequences of queries $q, q' \in D^k$, such that $q_S = q'_S$, the following two random variables are identically distributed:

- a_S, where $a \in_R \mathcal{A}_q$
- a'_S where $a' \in_R \mathcal{A}_{q'}$

If the two distributions are δ-close, rather than identical, we say that the family of distributions $\{A_q\}_{q \in D^k}$ is *δ-no-signaling.*

An MIP, $(V, P_1, \ldots, P_k)$ for a language $\mathcal{L}$ (possibly, relative to an oracle $\{\phi_n\}$) is said to have soundness ϵ against no-signaling strategies (or provers) if the following (more general) soundness property is satisfied:

2. **Soundness:** For every $x \notin \mathcal{L}$, and any no-signaling family of distributions $\{\mathcal{A}_q\}_{q \in D^k}$, the verifier V rejects with probability $\geq 1-\epsilon$, where on queries $q = (q_1, \ldots, q_k)$ the answers are given by $(a_1, \ldots, a_k) \in_R \mathcal{A}_q$, and ϵ is the error parameter.

If the property is satisfied for any δ-no-signaling family of distributions $\{A_q\}_{q \in D^k}$, we say that the MIP has soundness ϵ against δ-no-signaling strategies (or provers).

2.4 Interactive Arguments Systems

An interactive argument for a language L consists of an efficient verifier that wishes to verify a statement of the form $x \in L$ and a prover that helps the verifier to decide. The goal is to have the verifier use less computational resources than the straightforward computation of $x \in L$ (that does not involve the prover). In addition to the answer, the verifier wants assurance that the prover did not cheat and so the prover provides a short certificate proving its claim. In an argument system (in contrast to a proof system) we are only concerned with cheating provers that are *computationally bounded.* That is, we only require that it is *infeasible* to convince the verifier to accept some $x \notin L$.

We focus on 1-round argument systems. That is, the verifier sends a single message to the prover, gets a response, and decides whether to accept or reject. We say that (V, P) is a one-round argument-system for $\mathcal{L}$ if the following two properties are satisfied:

1. **Completeness:** For every $x \in \mathcal{L}$ and every security parameter $\tau > 0$, the verifier $V(1^\tau, x)$ accepts with probability 1, after interacting with $P(1^\tau, x)$.

2. **Soundness:** For every $x \notin \mathcal{L}$, and any family of circuits $\{P_\tau^*\}_\tau$ of size $2^{O(\tau)}$, the verifier $V(1^\tau, x)$ rejects with probability $\geq 1-\epsilon$, after interacting with $P_t^*(1^\tau, x)$, where ϵ is a parameter referred to as the *error* or *soundness* of the argument system.

We note that usually argument systems are defined with respect to polynomial-time adversaries. Since (using subexponential assumptions) we achieve soundness against subexponential cheating provers, in this paper, for simplicity, we define argument-systems as having soundness against subexponential size cheating provers.

3. OUR TECHNIQUES

Our techniques can be separated into two parts. The main part is the construction of a statistically no-signaling MIP for any function computable in time t and space s, where the number of provers grows linearly with s (independent of t), each prover runs in time at most $\mathsf{poly}(t)$, and the verifier runs in time that depends only on the input size, the security parameter and the space s but is independent of t.[7] This part, described in Section 3.1, is information theoretic, and does not rely on any cryptographic assumptions.

Then, in Section 3.2 we show how to convert a statistically no-signaling MIP into a 1-round delegation scheme. The soundness of the resulting delegation scheme assumes the existence of a fully homomorphic encryption (FHE) scheme with sub-exponential security.

3.1 Statistically No-Signaling MIP

We start by giving an overview of our MIP, and then give the high-level idea for why soundness holds against statistically no-signaling cheating provers. Our MIP is a variant of "standard" constructions. Our main contribution is in proving soundness against (statistically) no-signaling provers. This requires a different approach than the ones taken to prove classical soundness. Indeed, all known MIPs for NEXP (or the scale down version of NP) are not sound against no-signaling adversaries (see discussion in Section 1.1.1).

The main difference between a classical MIP and a no-signaling MIP is that in a classical MIP once a prover fixes its random tape (if at all he uses randomness), then his answer is a deterministic function of his query. This is not the case in the no-signaling setting, since a prover's answer can depend on the other queries. It is required that the answer of the prover is independent of the other queries *as a random variable* , but it may certainly depend on the other queries as a function. This makes the soundness proof significantly more challenging.

Before presenting the high level ideas of this proof, we first give a high level overview of our MIP. As a first step in the construction of our MIP, we would like to assume for simplicity that any set of (possibly malicious) provers behave *symmetrically*; namely, any two (possibly malicious) provers, who are asked the same questions, answer similarly. Of course, we cannot ensure such a thing, since cheating provers may behave arbitrarily. Thus, we ensure this by defining a new model of no-signaling PCP, as opposed to no-signaling MIP.

Intuitively, a no-signaling PCP is defined like a classical PCP, but where soundness is required to hold also against a *no-signaling* prover, who may be adaptive. Loosely speaking, a no-signaling prover, upon receiving any set of queries Q, may reply with a set of answers, where each answer may depend on all the queries in Q *as a function*, but not as a random variable. Namely, for any set of queries Q and for any subset $Q' \subseteq Q$, the *distribution* of the answers corresponding to the queries Q', should be independent of queries in $Q \setminus Q'$.

Formally, a no-signaling PCP consists of a family of distributions $\{A_Q\}_Q$, where there is one distribution for every "possible" set of queries Q, and the requirement is that for every subset of queries $Q' \subseteq Q$, the distribution of $(A_Q)|_{Q'}$ (which is the distribution of answers A_Q restricted to queries in Q') is independent of queries in $Q \setminus Q'$. More generally, a δ-no-signaling PCP has the property that for every possible set of queries Q_1 and Q_2 such that $Q' \subseteq Q_1$ and $Q' \subseteq Q_2$, the distributions $(A_{Q_1})|_{Q'}$ and $(A_{Q_2})|_{Q'}$ are δ-close. We emphasize that in a δ-no-signaling PCP we think of a set of queries Q as an *unordered* set, thus achieving the desired *symmetry*; i.e., the answers do not depend on the order of the queries.

We note, however, that the definition of δ-no-signaling PCP given above, is not complete. One needs to define what is a "possible set of queries". We define it to be all the query sets with at most $k_{\max}$ queries. k_{max} is an important parameter. The larger k_{max} is, the more limited the cheating provers are. We denote such a PCP by (k_{max}, δ) no-signaling PCP, and define it formally in the full version. We devote most of the technical sections to constructing a (k_{max}, δ)-no-signaling PCP and proving its soundness.

Before we give an overview of this construction, let us mention how we convert a δ-no-signaling PCP into a δ-no-signaling MIP. This is relatively straightforward, and is done formally in the full version. The basic idea is that the MIP verifier emulates the PCP verifier, and sends each query to a random prover (that was not yet asked any query). Each prover answers by simulating the (honest) PCP.

In what follows, we give a high level overview of our (k_{max}, δ) no-signaling PCP. The PCP we use is very similar to known PCPs (see, e.g., [33]). The main point of distinction of our PCP is that we repeat each test k times, where k is the security parameter. We mention that (for simplicity) the high-level description of our PCP slightly differ from our actual PCP, presented formally in the full version.

Overview of the underlying PCP. Suppose the provers need to prove that $x \in \mathcal{L}$, where x is an n-bit string and $\mathcal{L}$ is a language computable by a (deterministic) Turing machine running in time $t(n)$ and space $s(n)$. The underlying PCP consists of several low-degree multi-variate polynomials. The first polynomial is the low-degree extension of the entire computation. More specifically, let $\mathcal{C}_n$ be a circuit of size $N = O(t(n)s(n))$ that computes $\mathcal{L}$ on inputs of length n. It is known that this circuit $\mathcal{C}_n$ can be made layered, with $O(s(n))$ gates in each layer, and such that there exists a space $O(\log N)$ Turing machine that on input n outputs the description of the circuit $\mathcal{C}_n$.

Assume that the gates of the circuit are indexed by the numbers $1, \ldots, N$, in an order that agrees with the layers of the circuit. In particular, for every gate, the index of the gate is larger than the indexes of its children. We assume that $1, \ldots, n$ are the indexes of the n input variables and that N is the index of the output gate. Let $x_1, \ldots, x_N$ be the values of the N wires of the circuit $\mathcal{C}_n$ when computed with input $x = (x_1, \ldots, x_n)$.

The entire computation $x_1, \ldots, x_N$ appears in the PCP encoded using an error correcting code (specifically, using the low-degree extension encoding), so that if a single bit in the computation is incorrect it causes a global effect on the encoding. Let $H = \{0, 1, \ldots, \log N - 1\}$ and let $m =$

[7] As a matter of fact, even though the number of queries is linear in s, most of these queries contain no information (and are fixed to some arbitrary value). Indeed, the verifier ignores the responses to these "dummy" queries and they exist solely for our soundness proof to go through. Thus, arguably, the verifier's running time actually depends only on n and the security parameter.

$\frac{\log N}{\log\log N}$.[8] Note that $|H|^m = N$. Thus we can identify $[N]$ with H^m (say, by the lexicographic order on H^m), and can view $x_1, \ldots, x_N$ as a function $X : H^m \to \{0,1\}$. Let $\mathbb{F}$ be a finite field of size $\mathsf{polylog}(N)$, containing the set H. Let $\hat{X} : \mathbb{F}^m \to \mathbb{F}$ be the *low-degree extension* of X (see the full version for the definition). Namely, $\hat{X}$ is the (unique) multivariate polynomial of degree $|H| - 1$ in each variable, that agrees with X on each element in H^m. To be consistent with the body of the paper, we abuse notation and denote the low-degree extension by $X : \mathbb{F}^m \to \mathbb{F}$. This low-degree multi-variate polynomial X is part of the PCP. Note that its truth table is of size $|F|^m = \mathsf{poly}(N)$.

In addition, the PCP contains the low-degree multi-variate polynomials $P_0, P_1, \ldots, P_\ell$, which are defined as follows. Let φ be a 3-CNF boolean formula (which depends on x), where $\varphi(w_1, \ldots, w_N)$ checks that its input, $w_1, \ldots, w_N$, is the correct computation of $\mathcal{C}_n$ on input $x_1, \ldots, x_n$, and that the computation is accepting. Namely, $\varphi(w_1, \ldots, w_N)$ checks that the computation of every gate in the circuit is performed correctly, that $w_i = x_i$ for every $i \in [n]$, and that $w_N = 1$. Let $\phi : (H^m)^3 \times \{0,1\}^3 \to \{0,1\}$ be the function where $\phi(i_1, i_2, i_3, b_1, b_2, b_3) = 1$ if and only if the clause $(w_{i_1} = b_1) \vee (w_{i_2} = b_2) \vee (w_{i_3} = b_3)$ appears in φ. Extend ϕ to be a function $\phi : H^{3m+3} \to \{0,1\}$ by setting it to be 0 for inputs outside of $(H^m)^3 \times \{0,1\}^3$. We denote $\ell \triangleq 3m + 3$. Let $\hat{\phi} : \mathbb{F}^\ell \to \mathbb{F}$ be the low-degree extension of ϕ. Thus, $\hat{\phi}$ is a multi-variate polynomial of degree $|H| - 1$ in each variable that agrees with ϕ on H^ℓ. Note that $x \in \mathcal{L}$ if and only if $\varphi(x_1, \ldots, x_N) = 1$. Therefore, if $x \in \mathcal{L}$ then for every $z = (i_1, i_2, i_3, b_1, b_2, b_3) \in H^\ell$, we have

$$\hat{\phi}(z) \cdot (X(i_1) - b_1) \cdot (X(i_2) - b_2) \cdot (X(i_3) - b_3) = 0. \qquad (1)$$

For $z = (i_1, i_2, i_3, b_1, b_2, b_3) \in \mathbb{F}^\ell$ we define

$$P_0(z) \triangleq \hat{\phi}(z) \cdot (X(i_1) - b_1) \cdot (X(i_2) - b_2) \cdot (X(i_3) - b_3).$$

Equation (1) implies that $P_0|_{H^\ell} \equiv 0$. Moreover, the fact that X and $\hat{\phi}$ have degree $< |H|$ in each variable implies that P_0 has degree $< 2|H|$ in each variable, and hence has total degree $< 2|H|\ell$. Next we define $P_1 : \mathbb{F}^\ell \to \mathbb{F}$. For every $z = (z_1, \ldots, z_\ell) \in \mathbb{F}^\ell$, let

$$P_1(z) = \sum_{h \in H} P_0(h, z_2, \ldots, z_\ell) {z_1}^h.$$

Note that if $x \in \mathcal{L}$ then $P_1|_{\mathbb{F} \times H^{\ell-1}} \equiv 0$. More generally, we define by induction $P_1, \ldots, P_\ell : \mathbb{F}^\ell \to \mathbb{F}$ where for every $z = (z_1, \ldots, z_\ell) \in \mathbb{F}^\ell$,

$$P_i(z) = \sum_{h \in H} P_{i-1}(z_1, \ldots, z_{i-1}, h, z_{i+1}, \ldots, z_\ell) {z_i}^h.$$

Note that $P_1, \ldots, P_{\ell-1}$ have degree $< 2|H|$ in each variable, and hence total degree $< 2|H|\ell$. Note also that if $x \in \mathcal{L}$ then $P_i|_{\mathbb{F}^i \times H^{\ell-i}} \equiv 0$, and in particular $P_\ell \equiv 0$.

The PCP proof for $x \in \mathcal{L}$ consists of the polynomial $X : \mathbb{F}^m \to \mathbb{F}$ and the $\ell + 1$ polynomials $P_i : \mathbb{F}^\ell \to \mathbb{F}$, for $i = 0, \ldots, \ell$.[9] The verifier sends the following queries to the PCP:

1. *low-degree test.* The verifier chooses k random lines $\ell_1, \ldots, \ell_k : \mathbb{F} \to \mathbb{F}^m$. It queries X on all the points $\{L_j(t)\}_{j \in [k], t \in \mathbb{F}}$, and checks that for every $j \in [k]$, the univariate polynomial $X \circ L_j : \mathbb{F} \to \mathbb{F}$ is of degree $< m|H|$.

   Similarly for each $P_i \in \{P_0, P_1, \ldots, P_\ell\}$, the verifier chooses k random lines $\ell_1, \ldots, \ell_k : \mathbb{F} \to \mathbb{F}^\ell$. It queries P_i on all the points $\{L_j(t)\}_{j \in [k], t \in \mathbb{F}}$, and checks that for every $j \in [k]$, the univariate polynomial $P_i \circ L_j : \mathbb{F} \to \mathbb{F}$ is of degree $< \ell|H|$.

2. *Consistency check.* For every $i \in [\ell]$, the verifier chooses k random points in $\mathbb{F}^\ell$. For each of these points $z = (z_1, \ldots, z_\ell) \in \mathbb{F}^\ell$, it queries P_i and P_{i-1} on all the points $\{(z_1, \ldots, z_{i-1}, t, z_{i+1}, \ldots, z_\ell)\}_{t \in \mathbb{F}}$, and checks that for every $t \in \mathbb{F}$,

   $$P_i(z_1, \ldots, z_{i-1}, t, z_{i+1}, \ldots, z_\ell) = \sum_{h \in H} P_{i-1}(z_1, \ldots, z_{i-1}, h, z_{i+1}, \ldots, z_\ell) t^h.$$

   The verifier also chooses k random points in $\mathbb{F}^\ell$. For each of these points

   $$z = (i_1, i_2, i_3, b_1, b_2, b_3) \in (\mathbb{F}^m)^3 \times \mathbb{F}^3 = \mathbb{F}^\ell,$$

   it queries P_0 on the point z and it queries X on the points i_1, i_2, i_3, and checks that

   $$P_0(z) = \hat{\phi}(z) \cdot (X(i_1) - b_1) \cdot (X(i_2) - b_2) \cdot (X(i_3) - b_3).$$

   Note that $\hat{\phi}(z)$ is not necessarily efficiently computable, and thus it is not clear how the verifier can carry out this check. For this overview, suppose for simplicity that the verifier has oracle access to $\hat{\phi}$. We note that we make this simplifying assumption also in our formal construction. Thus, when converting our δ-no-signaling PCP to a δ-no-signaling MIP, we get an MIP which assumes that the verifier has oracle access to $\hat{\phi}$.

   The basic idea for removing the oracle is that the provers will compute $\hat{\phi}$ for the verifier, and will prove the correctness of the computation via a statistically no-signaling MIP. At first, it may seems that this brings us back to square one, since our goal in this work is to construct a no-signaling MIP. However, it turns out that we already have a no-signaling MIP for $\hat{\phi}$, for the following reason: $\hat{\phi}$ can be decomposed into two parts: $\hat{\phi} = \hat{\phi}_x + \hat{\phi}_{\mathcal{C}}$, where $\hat{\phi}_x$ depends only on the input x and $\hat{\phi}_{\mathcal{C}}$ depends only on the circuit $\mathcal{C}$. The function $\hat{\phi}_x$ can be computed in time $\tilde{O}(n)$, and thus the verifier can compute it on his own. Regarding the function $\hat{\phi}_{\mathcal{C}}$, as we argue in Section 3.1, it can be computed in $O(\log N)$ space. Moreover, in the full version we show that any space s computation has an (inefficient) statistically no-signaling MIP, where the provers run in time $2^{O(s)}$, and the verifier runs in time $\mathsf{poly}(s, n)$. Thus for $s = O(\log N)$, we get a statistically no-signaling MIP where the provers run in time $\mathsf{poly}(N)$ and the verifier runs in time $\mathsf{poly}(n, \log N)$. We refer the reader to the technical sections for details.

This concludes the description of the underlying PCP.

Soundness against δ-no-signaling cheating provers. In what follows we give a proof overview that slightly differs from the formal proof. However, it conveys the main

[8] Suppose for simplicity that these are integers.

[9] Note that since $P_\ell \equiv 0$ it does not need to be part of the PCP. We include P_ℓ in the PCP only for the sake of simplicity.

ideas in the proof. The proof is by contradiction. Suppose that there exists a δ-no-signaling strategy $\{A_Q\}_Q$ that proves (with non-negligible probability) that $x \in \mathcal{L}$, even though this is not the case.

As a first step, we amplify the soundness, and claim that the δ-no-signaling strategy succeeds in convincing the verifier that $x \in \mathcal{L}$ with probability close to 1. The soundness amplification, which is a crucial step in our proof, is achieved by having the verifier V repeat each test k times, and accept if and only if all tests accept. We define a "relaxed" verifier V' that makes the exact same queries as V, but accepts if and only if for each (repeated) test, at least r of the k repetitions are accepting, where r is a parameter. Loosely speaking, we prove that if the verifier V accepts with probability ϵ then the relaxed verifier accepts with probability $1 - \frac{\tilde{O}(2^{-r})}{\epsilon}$, where $\tilde{O}$ hides $\mathsf{polylog}(N)$ factors.

To prove this we argue that if V and V' choose their queries independently then the probability that V accepts and V' rejects is very small. This is true because for each group of k tests we can first choose the $2k$ tests for both V and V', and only then decide which tests go to V and which ones go to V'. Consider the answers for these $2k$ tests. (It is important here that k_{max} is greater than the total number of queries in these $2k$ tests, so that all these queries can be asked simultaneously.) If among the $2k$ tests many are rejected then V rejects with high probability. On the other hand, if among the $2k$ tests only few are rejected (say, less than r) then V' always accepts. We refer the reader to the full version for details. Throughout this overview, we assume for simplicity that the verifier V accepts with probability close to 1.

We then argue that the low-degree test implies that *for every* $z \in \mathbb{F}^\ell$, when choosing k random lines $\ell_1, \ldots, \ell_k$ through z, such that $\ell_i(0) = z$, and querying the PCP at all the points $\{X \circ \ell_i(t)\}_{i \in [k], t \in \mathbb{F}}$, there exists a value v such that for most lines ℓ_i, when computing $X \circ \ell_i$ at point 0 via interpolation, we get v.[10] The same can be proved for $P_0, P_1 \ldots, P_\ell$. Throughout this overview we denote by $X'(z) = v$, the fact that when $X(z)$ is computed via interpolation as above, most interpolations give the value v. We emphasize that this notation may be misleading, since X' is *not* a function. The value of $X'(z)$ depends on the chosen lines $\ell_1, \ldots, \ell_k$. In other words, the value v above may be different for different lines. This is indeed what makes our analysis significantly harder than the analysis in the classical setting. We use a similar (misleading) notation for $P'_0, P'_1, \ldots, P'_\ell$.

Next, we use this notion to argue that *for every* $z \in H^\ell$, it holds that $P'_0(z) = 0$. To this end, we prove the more general statement that for every $i \in \{0, 1, \ldots, \ell\}$ and for every $z \in \mathbb{F}^i \times H^{\ell - i}$, it holds that $P'_i(z) = 0$. We prove this by backward induction, starting with the base case that $P_\ell \equiv 0$. The induction step follows from the consistency test, together with the fact that $P'_0, P'_1, \ldots, P'_\ell$ are well defined. More specifically, recall that in the consistency test the verifier checks that for every $i \in [\ell]$ and for a random $z = (z_1, \ldots, z_\ell) \in \mathbb{F}^\ell$ it holds that for every $t \in \mathbb{F}$,

$$P_i(z_1, \ldots, z_{i-1}, t, z_{i+1}, \ldots, z_\ell) = \sum_{h \in H} P_{i-1}(z_1, \ldots, z_{i-1}, h, z_{i+1}, \ldots, z_\ell) t^h.$$

One can argue that the low-degree test implies that for every $z \in \mathbb{F}^\ell$, and thus in particular for every $z = (z_1, \ldots, z_i, h_{i+1}, \ldots, h_\ell) \in \mathbb{F}^i \times H^{\ell - i}$, it holds that

$$P'_i(z_1, \ldots, z_{i-1}, t, h_{i+1}, \ldots, h_\ell) = \sum_{h \in H} P'_{i-1}(z_1, \ldots, z_{i-1}, h, h_{i+1}, \ldots, h_\ell) t^h.$$

Thus, if $P'_i|_{\mathbb{F}^i \times H^{\ell - i}} \equiv 0$ then $P'_{i-1}|_{\mathbb{F}^{i-1} \times H^{\ell - i + 1}} \equiv 0$. We refer the reader to the full version for details.

It remains to argue that the fact that $P_0|_{H^\ell} \equiv 0$ implies that the computation X is correct. At first this seems to follow immediately from the definition of P_0. Recall that for any $z = (i_1, i_2, i_3, b_1, b_2, b_3) \in (H^m)^3 \times H^3 = H^\ell$,

$$P_0(z) \triangleq \hat{\phi}(z) \cdot (X(i_1) - b_1) \cdot (X(i_2) - b_2) \cdot (X(i_3) - b_3).$$

Thus, it seems that by definition, X satisfies every clause of ϕ. Indeed, proving that this is the case is relatively straightforward (though is quite tedious to do formally). We note, however, that this does not immediately imply that the computation X is correct. It implies local consistency, that each gate is consistent, but does not imply global consistency, as needed for correctness. Namely, this implies that when querying X at three wires corresponding to a gate in $\mathcal{C}_n$ (two input wires and one output wire) then indeed the answers are going to satisfy the gate. However, it turns out that proving that this local consistency implies global consistency, or that X is correct, is quite tricky. Indeed it is for this part that we need to assume that $k_{max} > s$ (which corresponds to the number of provers in the MIP being larger than s, the space of the computation).

At first one might try the following approach for proving the correctness of X. Recall that we assume that the verifier accepts the δ-no-signaling strategy with probability close to 1 (actually it is important that the probability is at least $1 - \frac{1}{\mathsf{poly}(N)}$). This implies that with high probability, the input layer of the circuit $\mathcal{C}_n$ is correct; i.e., $(X(1), \ldots, X(n)) = (x_1, \ldots, x_n)$. Next, we can argue that the first level of the circuit is correct. This follows from the local consistency. Then, we would like to continue to the next level by induction, and so forth. The problem with this approach is that it incurs an exponential blowup in the error, since the output of a gate is correct if *both* its children are correct.

In order to avoid this exponential blowup in error, we take a different route. We assume that k_{max} (and thus the number of provers in the resulting MIP) is twice as large as the size of the longest layer of $\mathcal{C}_n$. Namely, we take $k_{max} = \Omega(s)$. We emphasize that, as described above, the honest verifier only sends $\mathsf{poly}(k, \log N)$ queries (and thus in the resulting MIP only queries $\mathsf{poly}(k, \log N)$ provers). We need $k_{max} = \Omega(s)$ only to argue no-signaling soundness. The verifier can now query the value of x corresponding to an *entire* layer of the circuit $\mathcal{C}_n$. The local consistency implies that the input level is correct with high probability. Then, we can argue that if level i is correct, then level $i+1$ is correct. This is the case, since otherwise, the verifier will ask the provers for the values of the entire layer i and $i+1$,

[10] Typically, in the PCP literature, it is proved that the low-degree test implies that the PCP is close to a low-degree polynomial. In contrast, in this work we only prove local consistency, which seems to be significantly easier.

and if there is an inconsistency, it will contradict the local consistency. This incurs only a linear (in the number of layers) blowup in the error. We refer the reader to the full version for details.

3.2 Converting a δ-No-Signaling MIP into a 2-message Argument

In this section we show that the method of Aiello *et al.* [1], of using a fully homomorphic (FHE) scheme to convert a 1-round MIP into a 1-round delegation scheme, is *sound* if the underlying MIP is secure against δ-no-signaling provers, where the value of δ affects the security requirement of the FHE scheme.[11]

Let us start by recalling their method. Aiello *et al.* proposed to take any MIP and convert it into the following 1-round delegation scheme: The delegator computes all the queries that the MIP verifier would send to the MIP provers, and sends all of these queries to the worker (prover), each encrypted under a fresh and independent key, using an FHE scheme. The worker then answers on behalf of each MIP prover, where each answer is computed *homomorphically* on the corresponding encrypted query.

As mentioned in the introduction, shortly after this method was introduced, Dwork *et al.* [11] showed that it may, in general, be insecure. In this work, we show that this method in fact is *secure* if the underlying MIP is sound against δ-no-signaling provers.

In a nutshell, our result is obtained by proving that if there exists a cheating prover P^* that breaks the soundness of the 1-round argument, then this prover can be used to construct a δ-no-signaling prover P^{NS} that breaks the soundness of the MIP scheme.

The prover P^{NS} uses P^* in the obvious way: Given a set of queries $(q_1, \ldots, q_\ell)$ it encrypts these queries using fresh and independent keys, and sends the encrypted queries to P^*; upon receiving encrypted answers, it decrypts these answers and sends the decrypted answers $(a_1, \ldots, a_\ell)$ to the MIP verifier.

Clearly this strategy breaks the soundness of the MIP verifier, but we need to argue that it is δ-no-signaling. Indeed, we argue that if P^{NS} is *not* δ-no-signaling then the prover P^* can be used to break the underlying FHE scheme. Loosely speaking, by the definition of δ-no-signaling (see Section 2.3), if P^{NS} is *not* δ-no-signaling then there is a subset $S \subset [\ell]$ such that the distribution of the answers $(a_i)_{i \in S}$, conditioned on the corresponding queries $(q_i)_{i \in S}$, depends on the other queries $(q_i)_{i \notin S}$. In other words, these answers give information on the other queries. If this is the case, then indeed one can use P^* to break the FHE scheme.

We note that the above break may take time exponential in the communication complexity of the underlying MIP scheme, since the information obtained from the answers $(a_i)_{i \in S}$, is not necessarily efficiently computable. Therefore, we need the security parameter of the underlying FHE scheme to be bigger than the communication complexity of the MIP scheme, and we need to rely on the sub-exponential hardness of the FHE scheme.

4. OUR RESULTS

We show a general result on MIP proof systems that are secure against no-signaling strategies and use the latter to construct a new delegation scheme (a.k.a, a 1-round argument-system).

THEOREM 4. *Suppose that $\mathcal{L}$ can be computed in time $t = t(n)$ and space $s = s(n)$ with respect to inputs of length n. Then, for any integer $k \geq (\log t)^c$, where c is some (large enough) universal constant, there exists an MIP for $\mathcal{L}$ with $s \cdot k$ provers and with soundness error ϵ against δ-no-signaling strategies where $\epsilon \leq 2^{-k}$ and $\delta \geq 2^{-\frac{k}{\mathsf{polylog}(t)}}$.*

The resulting MIP verifier runs in time $s \cdot k + n \cdot \mathsf{polylog}(t)$ and the resulting provers run in time $\mathsf{poly}(t)$. Each query and answer is of length at most $\mathsf{polylog}(t)$.

At first it may seem a bit confusing that δ is *lower* bounded by some value but indeed, the *larger* δ is, the more powerful δ-no-signaling strategies become and therefore the soundness guarantee is stronger.

By setting the parameters $t = \mathsf{poly}(n)$, $k = O(\mathsf{polylog}(n))$ and $s = O(n)$ we obtain the following corollary:

COROLLARY 5. *Suppose that the language $\mathcal{L}$ can be computed in polynomial time and linear space. Then, there exists an MIP for $\mathcal{L}$ with $\tilde{O}(n)$ provers, and with soundness error $2^{-\mathsf{polylog} n}$ against $2^{-\mathsf{polylog} n}$-no-signaling strategies. The verifier runs in time $\tilde{O}(n)$ and the prover runs in time $\mathsf{poly}(n)$.*

Alternatively, setting $s = k = \mathsf{poly}(n)$ we obtain the following result:

COROLLARY 6. *If $\mathcal{L} \in \mathsf{PSPACE}$ and $\mathcal{L}$ can be computed in time $t = t(n)$ then there exists an MIP for $\mathcal{L}$ with $\mathsf{poly}(n)$ provers and with soundness error $2^{-\mathsf{poly}(n)}$ against $2^{-\mathsf{poly}(n)}$-no-signaling strategies. The verifier runs in time $\mathsf{poly}(n)$ and the prover runs in time $\mathsf{poly}(t)$.*

Using Theorem 4 we are also able to construct an argument-system as follows:

THEOREM 7. *If there exists a sub-exponentially secure FHE then there exists a polynomial T_{FHE} (that depends only on the encryption scheme) such that the following holds. Suppose that $\mathcal{L}$ can be computed in time $t = t(n)$ and space $s = s(n)$. Then, for any integer $k \geq (\log t)^c$, where c is some (large enough) universal constant, there exists a 1-round argument-system for $\mathcal{L}$ with soundness error 2^{-k} against provers of size $2^{O(k)}$.*

The verifier runs in time $(n + T_{\mathsf{FHE}}(s, k)) \cdot \mathsf{polylog}(t)$ and the prover runs in time $\mathsf{poly}(t, k)$.

We stress that the running time of the verifier in Theorem 7 only depends *poly-logarithmically* on the time that it takes to compute $\mathcal{L}$.

By setting $t = \mathsf{poly}(n)$, $s = n^\epsilon$ where $1/\epsilon$ is the degree of s in T_{FHE}, we obtain the following:

COROLLARY 8. *If there exists a sub-exponentially secure FHE then there exists a polynomial T_{FHE} and a constant $\epsilon > 0$ (that depend only on the encryption scheme) such that the following holds. Suppose that $\mathcal{L}$ can be computed in time $\mathsf{poly}(n)$ and space n^ϵ. Then, there exists a 1-round argument-system for $\mathcal{L}$ with soundness error 2^{-k} against provers of size $2^{O(k)}$ where $k \geq \mathsf{polylog}(n)$. The verifier runs in time $\tilde{O}(n) \cdot T_{\mathsf{FHE}}(k)$ and the prover runs in time $\mathsf{poly}(n)$.*

[11] Aiello *et al.* originally suggested to use a PCP together with a private information retrieval (PIR) scheme to construct a 1-round delegation scheme. We use FHE instead of PIR only for the sake of simplicity of notation. Our results hold also when replacing FHE by PIR (see the full version).

5. REFERENCES

[1] W. Aiello, S. Bhatt, R. Ostrovsky, and S. R. Rajagopalan. Fast verification of any remote procedure call: Short witness-indistinguishable one-round proofs for NP. In *ICALP: Annual International Colloquium on Automata, Languages and Programming*, 2000.

[2] B. Applebaum, Y. Ishai, and E. Kushilevitz. From secrecy to soundness: Efficient verification via secure computation. In *ICALP (1)*, pages 152–163, 2010.

[3] D. Avis, H. Imai, and T. Ito. On the relationship between convex bodies related to correlation experiments with dichotomic observables. *Journal of Physics A: Mathematical and General, 39(36)*, 39(36):11283, 2006.

[4] J. Barrett, N. Linden, S. Massar, S. Pironio, S. Popescu, and D. Roberts. Nonlocal correlations as an information-theoretic resource. *Physical Review A, 71(022101)*, 71(2):022101, 2005.

[5] I. Biehl, B. Meyer, and S. Wetzel. Ensuring the integrity of agent-based computations by short proofs. In *Proceedings of the Second International Workshop on Mobile Agents*, MA '98, pages 183–194, London, UK, UK, 1999. Springer-Verlag.

[6] N. Bitansky, R. Canetti, A. Chiesa, and E. Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In *ITCS*, pages 326–349, 2012.

[7] N. Bitansky, R. Canetti, A. Chiesa, and E. Tromer. Recursive composition and bootstrapping for snarks and proof-carrying data. *IACR Cryptology ePrint Archive*, 2012:95, 2012.

[8] K.-M. Chung, Y. T. Kalai, F.-H. Liu, and R. Raz. Memory delegation. In *CRYPTO*, pages 151–168, 2011.

[9] K.-M. Chung, Y. T. Kalai, and S. P. Vadhan. Improved delegation of computation using fully homomorphic encryption. In *CRYPTO*, pages 483–501, 2010.

[10] I. Damgård, S. Faust, and C. Hazay. Secure two-party computation with low communication. In *TCC*, pages 54–74, 2012.

[11] C. Dwork, M. Langberg, M. Naor, K. Nissim, and O. Reingold. Succinct proofs for NP and spooky interactions. Unpublished manuscript, available at `http://www.cs.bgu.ac.il/~kobbi/papers/spooky_sub_crypto.pdf`, 2004.

[12] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, pages 186–194, 1986.

[13] R. Gennaro, C. Gentry, and B. Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *CRYPTO*, pages 465–482, 2010.

[14] R. Gennaro, C. Gentry, B. Parno, and M. Raykova. Quadratic span programs and succinct nizks without pcps. *IACR Cryptology ePrint Archive*, 2012:215, 2012.

[15] C. Gentry and D. Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *STOC*, pages 99–108, 2011.

[16] S. Goldwasser, Y. T. Kalai, and G. N. Rothblum. Delegating computation: interactive proofs for muggles. In *STOC*, pages 113–122, 2008.

[17] S. Goldwasser, H. Lin, and A. Rubinstein. Delegation of computation without rejection problem from designated verifier cs-proofs. *IACR Cryptology ePrint Archive*, 2011:456, 2011.

[18] J. Groth. Short pairing-based non-interactive zero-knowledge arguments. In *ASIACRYPT*, pages 321–340, 2010.

[19] T. Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing*, 5(1):141–172, 2009.

[20] T. Ito. Polynomial-space approximation of no-signaling provers. In *ICALP (1)*, pages 140–151, 2010.

[21] T. Ito, H. Kobayashi, and K. Matsumoto. Oracularization and two-prover one-round interactive proofs against nonlocal strategies. In *IEEE Conference on Computational Complexity*, pages 217–228, 2009.

[22] T. Ito and T. Vidick. A multi-prover interactive proof for NEXP sound against entangled provers. *CoRR*, abs/1207.0550, 2012.

[23] Y. T. Kalai and R. Raz. Probabilistically checkable arguments. In *CRYPTO*, pages 143–159, 2009.

[24] J. Kempe, H. Kobayashi, K. Matsumoto, B. Toner, and T. Vidick. Entangled games are hard to approximate. In *FOCS*, pages 447–456, 2008.

[25] L. A. Khalfin and B. S. Tsirelson. Quantum and quasi-classical analogs of Bell inequalities. In *In Symposium on the Foundations of Modern Physics*, pages 441–460, 1985.

[26] J. Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *STOC*, pages 723–732, 1992.

[27] H. Lipmaa. Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In *TCC*, pages 169–189, 2012.

[28] S. Micali. Cs proofs (extended abstracts). In *FOCS*, pages 436–453, 1994.

[29] M. Naor. On cryptographic assumptions and challenges. In *CRYPTO*, pages 96–109, 2003.

[30] B. Parno, M. Raykova, and V. Vaikuntanathan. How to delegate and verify in public: Verifiable computation from attribute-based encryption. In *TCC*, pages 422–439, 2012.

[31] S. Popescu and D. Rohrlich. Quantum nonlocality as an axiom. *Foundations of Physics*, 24(3):379–385, 1994.

[32] P. Rastall. Locality, Bell's theorem, and quantum mechanics. *Foundations of Physics*, 15(9):963–972, 1985.

[33] M. Sudan. Probabilistically checkable proofs - lecture notes, 2000. Available at `http://people.csail.mit.edu/madhu/pcp/pcp.ps`.

[34] B. Toner. Monogamy of non-local quantum correlations. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 465(2101):59–69, 2009.