Non-Locality and Zero-Knowledge MIPs

Claude Crépeau[1 ⋆] and Nan Yang[2 ⋆⋆]

[1] McGill University, Montréal, Québec, Canada. crepeau@cs.mcgill.ca
[2] Concordia University, Montreal, Quebec, Canada. na_yan@encs.concordia.ca

Abstract. The foundation of zero-knowledge is the *simulator*: a weak machine capable of pretending to be a weak verifier talking with all-powerful provers. To achieve this, simulators need some kind of advantage such as the knowledge of a trapdoor. In existing zero-knowledge multi-prover protocols, this advantage is essentially *signalling*, something that the provers are explicitly forbidden to do. This advantage is stronger than necessary, as it is possible to define a sense in which simulators need much less to simulate. We define a framework in which we can quantify the simulators' *non-local advantage* and exhibit examples of zero-knowledge protocols that are sound against local or entangled provers that are not sound against no-signalling provers precisely because the no-signalling simulation strategy can be adopted by malicious provers.

1 Introduction

An *interactive proof* is a dialog between two parties: a polynomial-time *verifier* and an all-powerful *prover* [1, 2]. They agree ahead of time on some language L and a string x. The prover wishes to convince the verifier that $x \in L$. If this is true, the prover should succeed almost all the time; if not, the prover should fail almost all the time. This is a generalization of the complexity class **NP**, except instead of simply being handed a polynomial-sized witness, the verifier is allowed to quiz the prover. The set of languages that admit an interactive proof is called **IP**.

An interactive proof is *zero-knowledge* if the verifier learns nothing except the truth of "$x \in L$". This is usually defined by saying that a *distinguisher* is unable to tell apart a real conversation between the prover and the verifier, and one which is generated by a lone polynomial-time *simulator*. We will denote sets of zero-knowledge interactive proofs with a **ZK** bold prefix.

The *multi-prover* model was introduced in [3]. This model consists of multiple, non-communicating[⋆⋆⋆] provers talking to a single verifier. We will abbreviate "multi-prover interactive proof" as MIP and the set of languages which can be accepted by MIPs as the boldface **MIP**.

From a *complexity* perspective, the zero-knowledge aspect of interactive proofs is characterized by **IP** = **ZKIP**[†] = **PSPACE** for single-prover IPs ([4–6]), and **MIP** = **ZKMIP** = **NEXP** for multi-prover IPs ([3, 7–12]). The (conjectured) necessity of complexity assumptions for zero-knowledge in the single-prover case was the initial motivation for the multi-prover model.

1.1 A Cryptographic Perspective

The foundation of zero-knowledge is the idea of a *simulator*: a machine, with no more power than the verifier, which can pretend to having interacted with all-powerful provers. Obviously, this simulator cannot accomplish this task without some kind of *advantage* (something independent of knowledge). In single-prover zero-knowledge proofs, this advantage can be in the form of

⋆ Supported in part by FRQNT (INTRIQ) and NSERC (CryptoWorks21 and Discovery grant program).
⋆⋆ Supported in part by Professors Václav Chvátal, Jeremy Clark, Claude Crépeau, and David Ford.
⋆⋆⋆ The precise meaning of these words shall become a lot clearer throughout the rest of this paper.
† This is for computational Zero-Knowledge. For statistical ZK however the corresponding class is **SZK** and is most likely contained in **MA**, contains **BPP** and most likely contains only some part of **NP**.

the ability to *rewind* computation, the ability to discard failed simulations, or knowledge of a trapdoor in a commitment scheme. In multi-prover zero-knowledge proofs, the advantage in existing literature can be summed up as *signalling*: the simulator, acting in the name of several provers, knows secrets which real provers, in a real instance of the protocol, would not because they are unable to communicate.

From a complexity perspective, this simulator advantage can be anything as long as it is truly independent of knowledge – we do not want to exclude anything a priori. But, in practice, zero-knowledge is ultimately applied cryptography and from a cryptographic perspective, *not all advantages are equal*.

1.2 Relativistic Motivation

The need for more nuanced simulators is motivated by relativistic cryptography, an example of which can be found in [13]. Relativistic cryptography exploits the fact that it is impossible to signal faster than light. We can enforce the no-signalling condition of MIPs by spatially separating the provers from each other. In order to enforce the provers' spatial separation during the execution of the protocol, each prover is paired with a verifier of its own, which is located nearby. The verifiers can use the timing of the replies of their respective provers to judge their relative distance.

In practice, this means that we can implement MIPs under relativistic assumptions if the verifier can be "split" into multiple verifiers, each locally interacting with its corresponding prover. An example of relativistic cryptography can be found in [13], where a commitment was sustained for over 24 hours.

Some MIPs have verifiers which, intrinsically, cannot be split. Examples include [3] and [9]. In these examples, the verifier is used to courier an authenticated message between provers. In the relativistic setting, if the verifier has time to pass a message between provers, then the provers just signal between themselves.

Luckily, most MIPs in the literature have verifiers that are *non-adaptive*. These verifiers' questions to one prover are independent of the answers from all the provers. MIPs with non-adaptive verifiers can be rewritten into a format with multiple, split verifiers; this format we will call *locality-explicit*, and will be defined formally in section 4.

As an example of what we mean, consider the following two-prover interactive proof for graph 3-coloring:

Protocol 1. *(Simple MIP, Single-Verifier)*

Two provers P_1, P_2, one verifier V. On input graph G, P_1 and P_2 agree on a 3-coloring.

1. V asks P_1 for the colors of an edge e.
2. V asks P_2 for the colors of one of the nodes of e.

V accepts if and only if the colors of that edge from P_1 are not equal, and P_2 corroborates with P_1's answer by replying with the same color for the same node.

In the above protocol, V's questions to either prover does not depend on answers from any prover. This is what is commonly known as a *non-adaptive* verifier. We can therefore split the above verifier into a two-verifier version:

Protocol 2. *(Simple MIP, Multi-Verifier)*

Two provers P_1, P_2, two verifiers V_1, V_2. On input graph G, P_1 and P_2 agree on a 3-coloring, V_1 and V_2 agree on an edge e.

1. V_1 asks P_1 for the colors of e.
2. V_2 asks P_2 for the colors of one of the nodes of e.

Post execution, V_1 and V_2 confer with each other, and accept if and only if the colors of that edge from P_1 are not equal, and P_2 corroborates with P_1's answer by replying with the same color.

This version of the protocol is naturally suited for relativistic implementation. However, it is not zero-knowledge because even if P_1 and P_2 agreed on a randomly selected 3-coloring each time, a dishonest verifier V_2 may sample a node which is not from e. We can make a zero-knowledge, multi-verifier MIP with the help of the following commitment scheme, which is adapted from [3]:

Protocol 3. *(Multi-Verifier Commitment)*

Two provers P_1, P_2, two verifiers V_1, V_2. The provers share a random string w, and the verifiers share a random string r. Operations are over a finite field. P_1 wishes to commit b.

1. (Commit) V_1 sends P_1 the string r. P_1 replies with $x = w + br$.
2. (Unveil) P_2 sends V_2 the string w.

Post execution, the verifiers confer. They accept if and only if $x + w = r$ or $x + w = 0$.

Combining protocol 3 and the zero-knowledge protocol of [14] gives us a zero-knowledge, multi-verifier MIP.

Protocol 4. *(ZKMIP, Multi-Verifier)*

Two provers P_1, P_2, two verifiers V_1, V_2. On input graph G, P_1 and P_2 agree on a randomly selected 3-coloring and $2|V|$ strings w_i, V_1 and V_2 agree on an edge e and $2|V|$ strings r_i.

1. P_1 commits the colouring of G to V_1 using the $2|V|$ w_i, r_i they pre-agreed.
2. V_2 asks P_2 to unveil the colours of the edge e.

Post execution, V_1 and V_2 confer with each other, and accept if and only if the commitment is valid, and the colors unveiled are not equal.

What makes this protocol zero-knowledge? In the commitment scheme (protocol 3), if P_2 has knowledge of r, then it can break the commitment by unveiling either way (by sending w or $w + r$ as needed). Following the precedents set by existing literature's definition of zero-knowledge, the (*single*) simulator, interacting with both verifiers, learns r. Therefore it can break the commitment and always unveil a color that will be accepted by the verifiers.

1.3 Simulator's Advantage

As mentioned, the (single) simulator's advantage is its ability to interact with both verifiers at once. This is equivalent to having a pair of simulators signaling and, as we will see, is actually a tremendous power. However, it turns out that simulators do not need to signal in order to break the above commitment (section 3); a weaker non-local distribution will do. What we wish is to construct a framework in which this "non-local advantage" of the simulators can be quantified. We do this in section 4.

To see how much overkill signaling is for the simulators, imagine that in the above protocol, the distinguisher were able to eavesdrop on the "conversation" between the (possibly malicious) verifiers and black boxes, inside of which are either real provers, or simulators. This is giving the distinguisher more power than simply reading a transcript; and yet, the (signaling) simulators can succeed not only in generating the transcript, but behave as if they were provers in real-time. If we consider existing zero-knowledge as "transcript-indistinguishable", then we may consider this as "eavesdrop-indistinguishable". We will leave these terms undefined (as intuition) as they are not the focus of this work.

1.4 Our Contributions

In this work, we propose a framework for writing MIPs which is naturally suited for implementation and analysis under relativistic assumptions. We discuss how this framework extends naturally to zero-knowledge protocols and quantifies the non-local advantage which simulators use in many ZKMIPs. We show that **NEXP** can be accepted by MIPs in this form, and discuss the relationship between simulators' non-local advantage and soundness.

We exhibit a MIP for **NP** which, if is zero-knowledge, then cannot be sound; we introduce this as a tool for proving impossibility results of soundness against no-signalling provers but it could be used for for any non-locality class similarly.

2 Previous Work

The early work by Ben-Or, Goldwasser, Kilian and Wigderson asserting that **ZKMIP** = **MIP** from [3] and [9] use multi-round protocols and their (honest) verifiers are inherently signaling. This is precisely why we address the situation in this work. Proving soundness is quite subtle in this case because the provers could use the (signaling) verifier to break binding of the commitments. In particular, soundness will not be valid if the protocol is composed concurrently with other executions of itself or even used as a sub-routine. In recent conversations with Kilian [15], we have learned that controlling the impact of this *signaling* (via the verifier) has been a concern since the early days of MIPs. The protocols as they are might be sound but it is not fully proven anywhere in writing. However, it is also clear that no considerations had been given to the fact that general non-local correlations are possible via the verifier. If soundness rests on the binding property of a commitment scheme (such as those zero-knowledge proofs) and this binding property rests on the inability to achieve a certain non-local correlation then impossibility to achieve this correlation via the verifier must be demonstrated. It is not done or hinted in these papers.

The multi-round issue we address may seem trivial because it is a known fact that multi-round MIPs may be reduced to a single round using techniques of Lapidot-Shamir [16] and Feige-Lovasz [17]. Nevertheless, if interested in *zero-knowledge* MIPs, commitment schemes are generally used to obtain the zero-knowledge property and thus the single-round structure is lost in the process. Although single-round protocols bypass verifier's non-local contamination problems we describe in this work, converting multi-round protocols into single-round ones is highly inefficient and complex. Preserving zero-knowledge while achieving single-round has turned out to be a major challenge. Practically, keeping a multi-round protocol's structure, using only commitments to achieve zero-knowledge is very appealing.

In [16], Lapidot-Shamir proposed a parallel ZKMIP for **NEXP**, but they removed the zero-knowledge claim in the journal version [18] of their work without any explanation as of why. Feige and Kilian [10] were the last ones to follow this approach combining techniques drawn from Lapidot-Shamir [16], Feige-Lovasz [17] and Dwork, Feige, Kilian, Naor, and Safra, [11] to achieve a "2-prover 1-round 0-knowledge" proof for **NEXP**. As far as we can tell, this is the only paper in the ZKMIP literature that appears to avoid the multi-round problems and the non-local contamination that we discuss. However, note that the analysis of [10] is partly based of that of [16], and the journal version of Feige-Kilian [12] does not contain their prior claim of zero-knowledge either. All other ZKMIPs for **NEXP** in the literature are multi-round, and thus our analysis applies to them.

Similar issues are possible using more recent results such as Ito-Vidick's proof [19] that $\mathbf{NEXP} \subseteq \mathbf{MIP}^*$, Kalai, Raz and Rothblum's proof [20] that $\mathbf{MIP}^{ns} = \mathbf{EXP}$ and Natarajan-Wright's proof [21] that $\mathbf{NEEXP} \subseteq \mathbf{MIP}^*$. The reason why these multi-round constructions may maintain their soundness despite the potential non-locality contamination (via the verifier) is the *non-adaptive* nature of their verifiers. Non-adaptive verifiers cannot take advantage of information acquired in recent rounds to construct new questions to the provers: all their questions are pre-established before the interaction with the provers start. This is a special simpler case of local verifiers. Nowhere in this large literature can one find a single statement observing the non-adaptiveness of the verifiers and its importance to guarantee soundness of those MIPs. Moreover, their multi-round structure requires that any straightforward extensions to $\mathbf{ZKMIP}^*$ or $\mathbf{ZKMIP}^{ns}$ via commitment schemes be analyzed very carefully and the locality of the resulting verifiers be re-established. This is part of the reasons why the ZK version did not follow easily. Recently, Chiesa, Forbes, Gur, and Spooner [22] discovered a proof that $\mathbf{NEXP} \subseteq \mathbf{ZKMIP}^*$. Their construction is based on refinements of Ito-Vidick's proof and along the lines of Feige-Kilian, building on algebraic structures to bypass the need of commitment schemes. Unfortunately, this work is so complicated that we are unable to assess whether their verifier is actually non-adaptive. And of course, this is not mentioned or proven anywhere nor available from the authors... At the time of writing this paper, we just found out that indeed $\mathbf{ZKMIP}^* = \mathbf{MIP}^*$ as proven by Grilo, Slofstra and Yuen [23].

Bellare, Feige, and Kilian [24] considered a multi-verifier model similar to ours in order to analyze the role of randomness in multi-prover proofs. This is completely unrelated to our goal of analyzing verifier non-local contamination. Finally, the notion of relativistic commitment schemes put forward by Kilian [25] and Kent [26] leads to several results [13, 27, 28] where a similar multi-verifier model is necessary in order to assess spatial separation of the provers. The new (non-local) zero-knowledge definition is 100% fresh from this work. No prior work exists at all.

3 The Standard MIP Model

Multi-prover interactive proofs were introduced in [3]. The intuition for their model was that of a detective interrogating two suspects held in different rooms. This was formalized as follows:

Definition 1. *Let $P_1, \ldots, P_k$ be computationally unbounded Turing machines and let V be a probabilistic polynomial-time TM. All machines have a read-only input tape, a read-only auxiliary-input tape, a private work tape and a random tape. The P_i's share a joint, infinitely long, read-only random tape. Each P_i has a write-only communication tape to V, and vice-versa. We call $(P_1, \ldots, P_k, V)$ a k-*prover IP*, or multi-prover interactive proof (MIP).*

This model is essentially equivalent to that of Bell [29] who introduced his famous Bell's inequality to distinguish *local* parties from *entangled* parties.

Zero-knowledge MIPs were also defined in [3]:

Definition 2. *Let* $(P_1, \ldots, P_k, V)$ *be a k-prover IP for language* L*. Let* $\mathbf{view}(P_1, \ldots, P_k, V, x)$ *denote the verifier's incoming and outgoing messages with the provers, and his coin tosses*[‡]*. We say that* $(P_1, \ldots, P_k, V)$ *is* perfect zero-knowledge *for* L *if there exists an expected polynomial-time machine* M *such that for all* V'*,* $\mathbf{view}(P_1, \ldots, P_k, V', x)$ *and* $M(x)$ *are identically distributed.*

Let us call the above two definitions the *standard MIP model*. There have also been augmentations of the model by giving the provers various non-local resources, such as entanglement [19], or arbitrary no-signaling power [20].

Of specific interest to us are standard MIPs which have verifiers that are non-adaptive.

Definition 3. *A verifier is* non-adaptive *if the verifier's questions depend only on its random coins and the input* x*. A MIP with a non-adaptive verifier is a* non-adaptive MIP*.*

Some zero-knowledge MIPs such as [9] *require* that the verifier courier an authenticated message between the provers in order to obtain soundness while ensuring zero-knowledge. The gist of it goes like this:

1. V asks P_1 some questions.
2. V wants to check one of P_1's answers with P_2 for consistency.
3. In order for zero-knowledge to hold, V *must* ask P_2 a question it has already asked P_1.
4. P_1 authenticates a question with a key that was committed at the beginning of the protocol and sends it to V.
5. V sends the question and the authentication to P_2, who proceeds only if it succeeds.

Steps 4 and 5 consists of V sending a message from P_1 to P_2. This is problematic under relativistic assumptions, as discussed in the introduction. Therefore, the no-signaling assumption of standard MIPs are not immediately compatible with the no-faster-than-light-signaling assumption of relativity.

4 Locality-Explicit MIP

We define a framework for writing MIPs guaranteeing compatibility with relativistic assumptions. This framework uses multiple verifiers, each of which talks to a single prover; in turn, each prover talks to that single verifier. There are no communication tapes between the verifiers, nor are there between provers. There is a special verifier V_0 which *only reads* the outputs of the other verifiers; this is the verifier that will decide to accept or reject membership to L. We call this model "locality-explicit" since the provers and verifiers are explicitly local.

Any correlational resources available are explicitly specified via a supplementary *correlator* named $\widehat{P}$ for the provers and $\widehat{V}$ for the verifiers. Examples of these resources include entanglement, no-signalling distributions, or slower-than-light signalling.

Definition 4. *An* interactive Turning machine *(ITM) is augmented with the following tapes:*

- k_1 *read-only incoming communication tapes.*
- k_2 *write-only outgoing communication tapes.*
- *Private work, auxiliary-input, and random tapes.*

An ITM A *can signal to ITM* B *if* A*'s write-only outgoing tape is* B*'s read-only incoming tape.*

[‡] We ignore auxiliary inputs because we are not going to discuss composition.

Definition 5. *Let* $(\widehat{P}, P_1, \ldots, P_k, \widehat{V}, V_0, V_1, \ldots, V_k)$ *be a tuple of ITMs, where the P's are computationally all-powerful and the V's are polynomial-time. For each* i*, there are two-way communication tapes between* V_i *and* P_i*, and that for all* j*, there is a two-way communication tape between* $\widehat{V}$ *and* V_j *and also between* $\widehat{P}$ *and* P_j*. In addition, for each* ℓ*, there is a read-only tape going from* V_ℓ *to* V_0 *(where* V_0 *reads). Then, this is said to be a* locality-explicit multi-prover interactive proof*.*

We call $\widehat{P}$ *and* $\widehat{V}$ correlators *and say that the provers and verifiers are* $\widehat{P}$-local *and* $\widehat{V}$-local *respectively. We define the class of all MIPs with such correlators* $\mathbf{MIP}^{\widehat{P}}_{\widehat{V}}$*.*

It is perhaps easier to understand our definition with the help of figure 1.

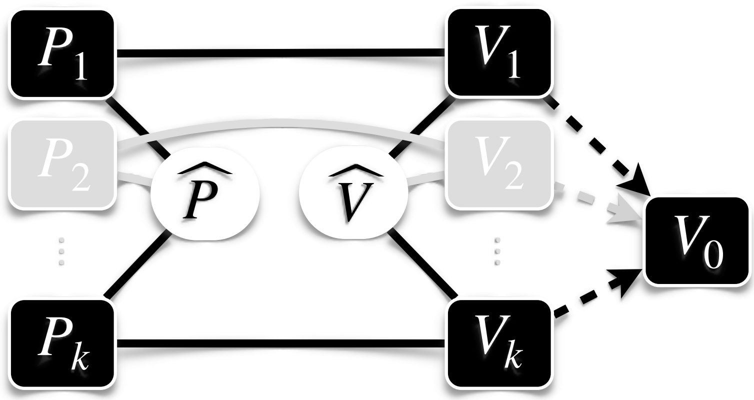

Fig. 1. Locality-Explicit MIP

The solid lines represents two-way communication and the dashed arrows represents one-way communication, with the arrow indicating the direction of information flow.

We can define that an LE-MIP accepts a language L if the usual soundness and completeness conditions hold:

Definition 6. *An LE-MIP* $(\widehat{V}, V_0, V_1, \ldots, V_k, \widehat{P}, P_1, \ldots, P_k)$ *accepts a language* L *if and only if*

- *(completeness)* $\forall x \in L, \mathbf{Pr}[V_0(x, t_1, \ldots, t_k) = \mathsf{accept}] > 2/3$,
- *(soundness)* $\forall x \notin L, \forall P'_1, \ldots, P'_k, \mathbf{Pr}[V_0(x, t_1, \ldots, t_k) = \mathsf{accept}] < 1/3$,

where t_i *is the read-only tape from* V_i *to* V_0 *at the end of* V_i*'s interaction with* P_i *(or* P'_i*) on input* x*.*

Note that we do not quantify over $\widehat{P}$ (nor $\widehat{V}$), as we want to use them not as (possibly malicious) participants to the protocol, but as a description of correlational resources available to the provers and verifiers.

Definition 7. *An LE-MIP is* local *if* $\widehat{V} = \widehat{P} = \varnothing$ *and all of the provers' (resp. verifiers') random tapes are initialized with the same uniformly random string* R *(resp. verifiers with another, independent uniformly random string* S*)*[§]*.*

[§] By $\varnothing$ we mean the empty correlator that provides everyone with nothing at all as output whatever the input is.

MIPs in the standard model (with local provers) are equivalent to LE-MIPs where $\widehat{P} = \varnothing$ and $\widehat{V}$ acts as a bulletin board. That is, a single verifier communicating with multiple provers is equivalent to multiple verifiers individually communicating with a local prover and each among themself.

Lemma 1. *If a MIP is non-adaptive, then there exists a local LE-MIP which accepts it.*

This is obvious as a non-adaptive verifier's questions are decided ahead of time, once its random coins are fixed. Therefore, we may split the verifier into one for each prover with a list of predetermined questions.

4.1 Zero-Knowledge LE-MIPs

As discussed in the introduction, zero-knowledge is defined by simulations. The simulator of single-prover IP and standard MIP are equal to the verifier in computational power, but they do have "advantages" – such as the ability to rewind computation.

LE-MIPs makes explicit a new advantage for the simulator: non-local correlations, a very powerful advantage. Using the correct non-local correlations, simulators do not need to rewind, do not need to pretend to be multiple (isolated) provers, and do not need to know any commitment-breaking secrets. In short, they do not need to signal. Multiple, no-signaling simulators can even produce transcripts in "real-time" (example will follow) if the proper correlations are used.

Definition 8. *Let $\mathcal{M} = (\widehat{M}, M_1, \ldots, M_k)$ be a tuple of polynomial-time ITMs. Each machine has a random tape, and every random tape is initialized with the same random bits. For $1 \leq i \leq k$, there is a two-way communication tape between $\widehat{M}$ and M_i. There are no communication tapes between any of the M_i's. Then this is called a tuple of* locality-explicit simulators *and $\widehat{M}$ is the* locality class *of $\mathcal{M}$, which will be abbreviated $\widehat{M}$*-local.

Definition 9. *Let $\mathcal{PV} = (\widehat{P}, P_1, \ldots, P_k, \widehat{V}, V_0, V_1, \ldots, V_k)$ be an LE-MIP for language L. If there exists a tuple of locality-explicit simulators $(\widehat{S}, S_1, \ldots, S_k)$, such that for all verifiers $(\widehat{V}', V_0', V_1', \ldots, V_k')$, such that for all $x \in L$ the transcripts of conversations*

$$(\widehat{P}, P_1, \ldots, P_k, \widehat{V}', V_0', V_1', \ldots, V_k')(x)$$

and those generated by

$$(\{\widehat{S}, \widehat{V}'\}, V_0', S_1^{V_1'}, \ldots, S_k^{V_k'})(x)$$

are identically distributed, then we say that $\mathcal{PV}$ is a $\widehat{S}$-local perfect zero-knowledge LE-MIP *for L. Note that the simulators are responsible for using $\widehat{V}'$, if necessary, to ensure that the verifier oracles*[¶] *receive the necessary inputs.*

We will denote the set of all ZK LE-MIPs where the provers, verifiers, and simulators are $\widehat{P}$-local, $\widehat{V}$-local, and $\widehat{S}$-local by

$$\mathbf{ZK}^{\widehat{S}}\mathbf{MIP}^{\widehat{P}}_{\widehat{V}}.$$

Let $\mathbb{S}, \mathbb{P}, \mathbb{V}$ be sets of correlators. We will denote, by convention,

$$\mathbf{ZK}^{\mathbb{S}}\mathbf{MIP}^{\mathbb{P}}_{\mathbb{V}}$$

as the set of all ZK LE-MIPs where each correlator comes from each of the respective sets.

[¶] Each simulator S_i is restricted to oracle calls to its own corresponding V_i'.

Our motivations for the above definition are twofold.

First, a simulator (or simulators) should not have more power than necessary. If two *local* simulators can output for two *local* verifiers, then it is not necessary to have a single simulator (equivalent to two *signaling* simulators) do the job. In general, finding the minimal $\widehat{S}$ that will allow simulation establishes how little extra is needed to obtain the zero-knowledge property.

Second, the non-locality of simulators is a characterization of the resilience of zero-knowledge. A protocol with local simulators which can withstand arbitrary (malicious) verifiers is more resilient than one in which signaling simulators are needed.

This may be of practical interest, if transcripts are timestamped. For example, under the relativistic assumption that one may not signal faster-than-light, one may be able to distinguish two spatially separated simulators from two spatially separated verifiers, if the simulators need to signal (transmit a commitment-breaking secret) in order to generate a transcript. On the other hand, if two entangled simulators are sufficient to produce the transcript, then they are indistinguishable from real verifiers and provers. Our protocol 8 can be modified as to let entangled simulators do their work, without needing PR-boxes or signaling. Details in section 5

The complexity of LE-MIPs are the same as those of MIP, namely:

Theorem 5. *There exists a LE-MIP which accepts* **NEXP**.

The proof is a line-by-line inspection of the BFL protocol as found in [8], and checking that the verifier is non-adaptive, and therefore can be written as a LE-MIP. We have included a brief summary of the BFL protocol in appendix B.

5 Zero-Knowledge LE-MIP for NEXP

The question which follows naturally is whether there exists a *zero-knowledge*, local LE-MIP for **NEXP** where $\mathbb{S} \subsetneq$ **SIG**. By adapting the protocol from [8], we will exhibit a protocol with the following properties:

1. The provers and verifiers are local: $\widehat{V} = \widehat{P} = \varnothing$.
2. The simulators need only access to instances of **PR**-boxes to work. That is, $\widehat{S}$ simply computes indexed instances of **PR**-boxes. We will abbreviate this as "**PR**-local."

We may succinctly summarize the above as:

Theorem 6. $\mathbf{ZK^{PR}MIP^{\varnothing}_{\varnothing}} = \mathbf{NEXP}$*, where* **PR** *denotes a correlator which simply computes* **PR***-boxes for the simulators.*

We prove the above theorem by constructing an LE-MIP with the right properties: protocol 8. The generic way of turning an interactive proof into a zero-knowledge one is by running it in committed form [3, 9]. With this technique, provers commit their answers instead of directly responding, and use cryptographic techniques to convince the verifier that the answers are correct. As argued previously, this is not possible to enforce from relativistic assumptions alone.

Our solution essentially asks the provers to (strongly-universal-2) hash the selected committed answer with a key that is based on the verifier's question. We force V_2 to behave honestly (to ask a question that V_1 has asked) by making bad questions meaningless. If the verifiers ask the provers the same question, they will receive the same hash of the same answer. Otherwise, they will receive two independent random hash values.

The **PR**-type commitment (protocol 7) is secure in the local setting as previously proved in [26, 30, 13]. It is perfectly concealing and statistically binding. In general, we use the commitment-box notation $\boxed{b}$ as the name of a commitment to bit b in the next two protocols.

Protocol 7. *A statistically binding, perfectly concealing commitment protocol to bit b.*

All parties agree on a security parameter 1^k.
P_1 and P_2 partition their private random tape into two k-bit strings w_1, w_2.

Pre-computation phase:

- V_1 samples two k-bit strings z_1, z_2 independently and uniformly, and provides them to V_2.
- V_1 sends z_1 to P_1 and V_2 sends z_2 to P_2.

Commit phase:

- P_1 commits b to V_1 as $\boxed{b} = (b \times z_1) \oplus w_1$, where $b \times z_1$ is a multiplication in $\mathbb{F}_{2^n}$.
- P_2 sends V_2: $d = (w_1 \times z_2) \oplus w_2$.

Unveiling phase:

- P_1 sends w_1, w_2 to V_1.
- V_1 computes $b = 1$ if $\boxed{b} \oplus w_1 = z_1$, or $b = 0$ if $\boxed{b} = w_1$.
- V_0 **rejects** if $\boxed{b} \oplus w_1$ is anything but z_1 or 0, or if $d \oplus w_2 \neq w_1 \times z_2$ and **accepts** b otherwise.

A note on notation: for a circuit f, we will denote $f(\boxed{x})$ as the gate-by-gate committed circuit evaluated with x as the input. We also use statements such as "P_1 proves to V_1 that $\boxed{\Omega_1}$ was computed correctly". The reader is expected familiarity with zero-knowledge computations on committed circuits as put forward by [31, 32, 5, 9].

Protocol 8. *A local zero-knowledge LE-MIP for oracle-3-SAT*

Let $x = (B, r, s)$, an instance of oracle-3-SAT, be the common input, let $k = |x| = r + 3s + 3$, and let Λ be the verifier's program in protocol B (see appendix).

1. **Pre-computation:**
   (a) V_1 samples two k-bit strings z_1, z_2 independently and uniformly, and provides them to V_2.
   (b) V_1 selects $k + 3$ random bit strings $R_1, ..., R_{k+3}$ (size specified implicitly by Λ) and evaluates the circuit of Λ using the R_i as randomness, resulting in questions $Q_1, ..., Q_{k+3}$, and provides them to V_2
   (c) V_1 randomly chooses i, $1 \leq i \leq k + 3$, the index of an oracle query that will be made to both P_1 and P_2. V_1 provides i to V_2.
   (d) V_1 sends z_1 to P_1 and V_2 sends z_2 to P_2 for future commitments.
   (e) All parties agree on a family of strongly-universal-2 hash functions $\{H_i\}$ indexed by k-bit keys.
   (f) P_1 and P_2 agree on a k-bit index γ to the above family. P_1 commits $\boxed{\gamma}$ to V_1.
2. **Sumcheck with oracle:**
   - Let f be the arithmetization obtained in protocol 12, let z be a string from I^r and $Q_{k+1}, Q_{k+2}, Q_{k+3}$ be strings of I^s as generated in protocol B. V_1 and P_1 execute protocol 12 in committed form. At the end of this phase, P_1 shows that the committed final value is equal to

$$f\left(z, Q_{k+1}, Q_{k+2}, Q_{k+3}, \boxed{A(Q_{k+1})}, \boxed{A(Q_{k+2})}, \boxed{A(Q_{k+3})}\right),$$

     an evaluation in committed form of f using the committed values that were used during the protocol's loop. If this fails, V_1 instructs V_0 to reject.

3. **Multilinearity test:**
   (a) For $1 \leq i \leq k$:
      i. V_1 sends Q_i to P_1,
      ii. P_1 commits his answer as $\boxed{A(Q_i)}$.
   (b) P_1 and V_1 evaluate a circuit description of Λ in committed form with inputs $\boxed{A(Q_1)}, \ldots, \boxed{A(Q_k)}$ to verify proper linearity among them. P_1 unveils the circuit's committed output. If it rejects, V_1 instructs V_0 to reject.
4. **Consistency test:**
   (a) V_1 sends i to P_1.
   (b) P_1 computes $\boxed{\Omega_1} = \boxed{A(Q_i)} \oplus H_{\boxed{\gamma}}(Q_i)$ and sends $\boxed{\Omega_1}$ to V_1.
   (c) P_1 proves to V_1 that $\boxed{\Omega_1}$ was computed correctly, from the existing commitments.
   (d) P_1 unveils $\boxed{\Omega_1}$ for V_1, who gets Ω_1.
   (e) V_2 sends Q_i to P_2 (recall that this was pre-agreed in step 1.(c))
   (f) P_2 responds to V_2 with $\Omega_2 = A(Q_i) \oplus H_\gamma(Q_i)$.
   (g) V_0 accepts if and only if all of the following conditions are met:
      – $\Omega_1 = \Omega_2$
      – All commitments which have been unveiled are valid.
      – V_1 did not reject in the two previous cases.

The proofs of security can be found in appendix A.

5.1 Minimal Simulator Advantage

What is the minimal simulator advantage needed for achieving zero-knowledge for **NEXP**?

It is clear that signalling simulators can succeed in the above protocol. This is the zero-knowledge simulator of standard MIPs. We can summarize this as

$$\mathbf{ZK^{SIG}MIP}^{\varnothing}_{\varnothing} = \mathbf{NEXP},$$

where **SIG** is a signalling correlator.

Signalling is however unnecessary, as the binding condition of commitment used above (protocol 7) can be broken given **PR**-boxes. This is what the proof of security shows in appendix A. Thus, the simulator's advantage can be lowered to **PR**-boxes, or

$$\mathbf{ZK^{PR}MIP}^{\varnothing}_{\varnothing} = \mathbf{NEXP}.$$

If the verifiers were willing to tolerate approximately 15% of errors in the provers' unveiling string (z_1 or 0), then it is possible to break binding with shared entanglement [33] while maintaining soundness against local provers. Making this slight change in the protocol reduces the simulator advantage further:

$$\mathbf{ZK^{ENT}MIP}^{\varnothing}_{\varnothing} = \mathbf{NEXP},$$

where **ENT** denotes polynomial amount of shared entanglement for the simulators.

Ideally, the simulators would not need any non-local advantage over the verifiers. However, we are unable to find a zero-knowledge MIP where the simulators are *local* which can accept **NEXP**, or prove that it is impossible. We make the following conjecture:

Conjecture 1. $\mathbf{ZK}^{\varnothing}\mathbf{MIP}^{\varnothing}_{\varnothing} = \mathbf{SZK}$, where **SZK** is the set of languages with statistical zero-knowledge interactive proofs without computational assumptions (i.e., graph isomorphism).

5.2 Soundness Against No-Signalling Provers

As a further example of the drastic differences between MIP simulators' non-local advantages and single-prover IP simulators' advantages (e.g., rewinding), consider the following:

Theorem 9. *Suppose that the provers in protocol 8 have access to PR-boxes (thus they are no-signalling, but not local), then the protocol is not sound.*

Proof. The provers adopt the simulators' strategy. Since commitment binding is broken with the aid of PR-boxes, the verifiers will always accept.

This is the sense to which we referred to as "eavesdrop indistinguishable" from "transcript indistinguishable" earlier. A prover having the ability to rewind computations, although enough for simulators in IPs, is not enough to break soundness. We will generalize the above theorem in a future work, on the relationship between zero-knowledge and soundness.

Another Example In appendix E a zero-knowledge protocol for **NP** is extracted from [34]. This protocol is not only sound against local provers but also against entangled provers. It is zero-knowledge in both cases. However, since the ZK simulator (also provided in appendix E) can be implemented as no-signalling simulators, this same protocol cannot be sound against no-signalling provers since they can adopt exactly the simulators' strategy.

6 Conclusions and Future Work

Zero-knowledge simulators need advantages in order to function. In the case of MIPs, it was always implicitly assumed this advantage is necessarily signaling. We have shown that this is not true, and that this aspect of zero-knowledge remains unexplored. LE-MIPs make this explicit, while providing a template for relativistic implementations of the no-signaling assumption.

We close with three open questions.

First, although the provers and verifiers of protocol 8 are local, the simulators are not – they use PR-boxes. We do not know whether it is possible to simulate protocol 8 with *local* simulators. In fact, we conjecture that there does not exist a $\mathbf{ZK}^{\varnothing}\mathbf{MIP}^{\varnothing}_{\varnothing}$ protocol for any language outside **SZK**.

Second, as we have sketched out in section 5.1, by weakening the commitment scheme used, we get $\mathbf{ZK}^{\mathbf{ENT}}\mathbf{MIP}^{\varnothing}_{\varnothing} = \mathbf{NEXP}$. What is a minimal $\widehat{S}$ such that $\mathbf{ZK}^{\widehat{S}}\mathbf{MIP}^{\varnothing}_{\varnothing} = \mathbf{NEXP}$?

Third, what is the relationship between zero-knowledge and soundness in MIPs? As we have shown in section 5.2, some simulators' strategy can be adopted by provers to break soundness, if only the provers had some additional (in this case, non-local) resources. Is there a relationship between the non-local resources needed to achieve zero-knowledge and those that are forbidden in order to achieve soundness?

Acknowledgements

We would like to thank G. Brassard, A. Chailloux, S. Fehr, J. Kilian, S. Laplante, J. Li, A. Leverrier, A. Massenet, S. Ranellucci, L. Salvail, C. Schaffner, and T. Vidick for various discussions about earlier versions of this work. We would also like to thank Jeremy Clark for his insightful comments. Finally, we are grateful to Raphael Phan and Moti Yung for inviting us to publish a lead-up paper to this work as an *Insight Paper* at MyCrypt 2016.

References

1. S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof-systems," *SIAM. J. Computing*, vol. 18, pp. 186–208, Feb. 1989.
2. L. Babai, "Trading group theory for randomness," in *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, pp. 421–429, May 1985.
3. M. Ben-Or, S. Goldwasser, J. Kilian, and A. Wigderson, "Multi-prover interactive proofs: How to remove intractability assumptions," in *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, (New York, NY, USA), pp. 113–131, ACM, 1988.
4. A. Shamir, "IP = PSPACE," *J. ACM*, vol. 39, pp. 869–877, Oct. 1992.
5. R. Impagliazzo and M. Yung, "Direct minimum-knowledge computations," in *Advances in Cryptology: Proceedings of Crypto '87* (C. Pomerance, ed.), vol. 293, pp. 40–51, Springer-Verlag, 1988.
6. M. Ben-Or, O. Goldreich, S. Goldwasser, J. Håstad, J. Kilian, S. Micali, and P. Rogaway, "Everything provable is provable in zero-knowledge," in *Proceedings of the 8th Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '88, (London, UK, UK), pp. 37–56, Springer-Verlag, 1990.
7. L. Fortnow, J. Rompel, and M. Sipser, "On the power of multi-prover interactive protocols," *Theor. Comput. Sci.*, vol. 134, pp. 545–557, Nov. 1994.
8. L. Babai, L. Fortnow, and C. Lund, "Non-deterministic exponential time has two-prover interactive protocols," *Comput. Complex.*, vol. 2, pp. 374–374, Dec. 1992.
9. J. Kilian, *Uses of randomness in algorithms and protocols.* MIT Press, 1990.
10. U. Feige and J. Kilian, "Two prover protocols: low error at affordable rates," in *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada* (F. T. Leighton and M. T. Goodrich, eds.), pp. 172–183, ACM, 1994.
11. C. Dwork, U. Feige, J. Kilian, M. Naor, and S. Safra, "Low communication 2-prover zero-knowledge proofs for NP," in *Advances in Cryptology - CRYPTO '92, 12th Annual International Cryptology Conference, Santa Barbara, California, USA, August 16-20, 1992, Proceedings* (E. F. Brickell, ed.), vol. 740 of *Lecture Notes in Computer Science*, pp. 215–227, Springer, 1992.
12. U. Feige and J. Kilian, "Two-prover protocols - low error at affordable rates," *SIAM J. Comput.*, vol. 30, no. 1, pp. 324–346, 2000.
13. T. Lunghi, J. Kaniewski, F. Bussières, R. Houlmann, M. Tomamichel, S. Wehner, and H. Zbinden, "Practical relativistic bit commitment," *Phys. Rev. Lett.*, vol. 115, p. 030502, Jul 2015.
14. O. Goldreich, S. Micali, and A. Wigderson, "Proofs that yield nothing but their validity or all languages in np have zero-knowledge proof systems," *J. ACM*, vol. 38, pp. 690–728, July 1991.
15. J. Kilian, "Personal e-mail communication," July 2018.
16. D. Lapidot and A. Shamir, "Fully parallelized multi prover protocols for nexp-time (extended abstract)," in *32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991*, pp. 13–18, IEEE Computer Society, 1991.
17. U. Feige and L. Lovász, "Two-prover one-round proof systems: Their power and their problems (extended abstract)," in *Proceedings of the Twenty-fourth Annual ACM Symposium on Theory of Computing*, STOC '92, (New York, NY, USA), pp. 733–744, ACM, 1992.
18. D. Lapidot and A. Shamir, "Fully parallelized multi-prover protocols for nexp-time," *J. Comput. Syst. Sci.*, vol. 54, no. 2, pp. 215–220, 1997.
19. T. Ito and T. Vidick, "A multi-prover interactive proof for nexp sound against entangled provers," in *Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, FOCS '12, (Washington, DC, USA), pp. 243–252, IEEE Computer Society, 2012.
20. Y. T. Kalai, R. Raz, and R. D. Rothblum, "How to delegate computations: The power of no-signaling proofs," in *Proceedings of the Forty-sixth Annual ACM Symposium on Theory of Computing*, STOC '14, (New York, NY, USA), pp. 485–494, ACM, 2014.
21. A. Natarajan and J. Wright, "**NEEXP** in **MIP***," *CoRR*, vol. abs/1904.05870, 2019.
22. A. Chiesa, M. A. Forbes, T. Gur, and N. Spooner, "Spatial isolation implies zero knowledge even in a quantum world," *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 25, p. 44, 2018.
23. A. B. Grilo, W. Slofstra, and H. Yuen, "Perfect zero knowledge for quantum multiprover interactive proofs," *CoRR*, vol. abs/1905.11280, 2019.

24. M. Bellare, U. Feige, and J. Kilian, "On the role of shared randomness in two prover proof systems," in *Third Israel Symposium on Theory of Computing and Systems, ISTCS 1995, Tel Aviv, Israel, January 4-6, 1995, Proceedings*, pp. 199–208, IEEE Computer Society, 1995.
25. J. Kilian, "Strong separation models of multi prover interactive proofs," in *DIMACS Workshop on Cryptography*, 1990.
26. A. Kent, "Unconditionally secure bit commitment," *Phys. Rev. Lett.*, vol. 83, pp. 1447–1450, Aug 1999.
27. E. Adlam and A. Kent, "Deterministic relativistic quantum bit commitment," *CoRR*, vol. abs/1504.00943, 2015.
28. A. Chailloux and A. Leverrier, "Relativistic (or 2-prover 1-round) zero-knowledge protocol for NP secure against quantum adversaries," in *Advances in Cryptology – EUROCRYPT 2017: 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 – May 4, 2017, Proceedings, Part III*, pp. 369–396, Springer International Publishing, 2017.
29. J. S. Bell, "On the Einstein-Podolsky-Rosen paradox," *Physics*, vol. 1, pp. 195–200, 1964.
30. C. Crépeau, L. Salvail, J.-R. Simard, and A. Tapp, "Two provers in isolation," in *Advances in Cryptology – ASIACRYPT 2011: 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*, (Berlin, Heidelberg), pp. 407–430, Springer Berlin Heidelberg, 2011.
31. G. Brassard and C. Crépeau, "Zero-knowledge simulation of boolean circuits (extended abstract)," in *Advances in Cryptology: Proceedings of Crypto '86* (A. M. Odlyzko, ed.), vol. 263, pp. 223–233, Springer-Verlag, 1987.
32. G. Brassard and C. Crépeau, "Non-transitive transfer of confidence: A perfect zero-knowledge interactive protocol for SAT and beyond," in 27^{th} *Symp. of Found. of Computer Sci.*, pp. 188–195, IEEE, 1986.
33. G. Brassard, A. Broadbent, and A. Tapp, "Multi-party pseudo-telepathy," in *Algorithms and Data Structures* (F. Dehne, J.-R. Sack, and M. Smid, eds.), (Berlin, Heidelberg), pp. 1–11, Springer Berlin Heidelberg, 2003.
34. C. Crépeau, A. Y. Massenet-Oshima, L. Salvail, L. S. Stinchcombe, and N. Yang, "Zero-knowledge **MIP**s for **NP** sound against entangled provers using a tiny amount of commitments," in **(submitted to)** *Theory of Cryptography*, Springer International Publishing, 2019.
35. A. Acín, T. Fritz, A. Leverrier, and A. B. Sainz, "A combinatorial approach to nonlocality and contextuality," *Communications in Mathematical Physics*, vol. 334, pp. 533–628, Mar 2015.
36. H. Barnum, C. A. Fuchs, J. M. Renes, and A. Wilce, "Influence-free states on compound quantum systems," *CoRR*, vol. quant-ph/0507108v1, 2005.
37. J. Barrett, N. Linden, S. Massar, S. Pironio, S. Popescu, and D. Roberts, "Nonlocal correlations as an information-theoretic resource," *Phys. Rev. A*, vol. 71, p. 022101, Feb 2005.
38. M. Forster and S. Wolf, "Bipartite units of nonlocality," *Phys. Rev. A*, vol. 84, p. 042112, Oct 2011.
39. T. Ito, H. Kobayashi, D. Preda, X. Sun, and A. C. Yao, "Generalized tsirelson inequalities, commuting-operator provers, and multi-prover interactive proof systems," in *Proceedings of the 23rd Annual IEEE Conference on Computational Complexity, CCC 2008, 23-26 June 2008, College Park, Maryland, USA*, pp. 187–198, IEEE Computer Society, 2008.
40. C. Crépeau and N. Yang, "Multi-prover interactive proofs: Unsound foundations," in *Paradigms in Cryptology – Mycrypt 2016. Malicious and Exploratory Cryptology: Second International Conference, Mycrypt 2016, Kuala Lumpur, Malaysia, December 1-2, 2016, Revised Selected Papers*, pp. 485–493, Springer International Publishing, 2017.

A Proofs of Security for Protocol 8

Locality

Since the protocol is written as an LE-MIP in which $\widehat{P} = \widehat{V} = \varnothing$, the protocol is local by definition 7.

Completeness

Completeness follows from the completeness of the underlying protocol [8], and the fact that the commitment protocol (protocol 7) is well-defined for honest provers (who will never send a commitment that they cannot unveil).

Soundness

Without loss of generality, we may assume that the soundness error in the BFL protocol to be 1/3, through sequential amplification. The probability that our commitment scheme (protocol 7) fails binding is exponentially small in k. Local probabilistic provers are equivalent to local deterministic provers. This is because the success probability α of randomized provers of breaking soundness is an average over the randomized provers' random tapes. Each instance of a random tape represents a deterministic strategy. Therefore there is a deterministic strategy which succeeds with probability at least α, and hence we only need to consider local deterministic provers.

Since P_1 is deterministic, we may unambiguously consider what happens if we were to "rewind" the prover machine. Suppose that at some point P_1 unveils a particular commitment c to 0. We rewind P_1 and let V_1 make different choices before that point. Suppose that, with these alternate choices, P_1 then unveils c to 1 (an attempt to break binding). Because of locality, P_1's behavior is independent of what P_2 receives (namely z_2). Therefore, there is only *one* such z_2 which V_0 will ultimately accept as a valid unveiling of c in both ways (recall that our commitment is statistically binding).

Therefore, in the worst case, for every commitment there exists a sequence of interactions between V_1 and P_1 such that P_1 will attempt to break the binding of that commitment. Each such commitment-breaking corresponds to at most one string z_2 that will actually work.

Let us denote the set of such binding-breaking strings by B. If $z_2 \notin B$, then the provers *will not break binding*, and the soundness error is reduced to that of the underlying protocol (at most 1/3). On the other hand, since $|B| < \mathbf{poly}(k)$, the probability that $z_2 \in B$ is at most $\mathbf{poly}(k)/2^k$.

Therefore, the soundness error of our protocol is at most

$$Pr[z_2 \notin B \text{ and underlying protocol accepts}] + Pr[z_2 \in B] \leq \frac{1}{3} + \frac{\mathbf{poly}(k)}{2^k}.$$

Zero-Knowledge The simulation will be divided in two parts. In the first part, the simulator produces a transcript of the *pre-computation*, *multilinearity test* and *sumcheck with oracle* parts, which involves only interactions with V_1. In the second part, the simulator will fake a valid *consistency test.*

Protocol 10. *(Perfectly Indistinguishable,* **PR***-Local Simulator for Protocol 8, Part 1)*

The setup:
- Let $(\widehat{S}, S_1, S_2)$ be a set of locality-explicit simulators.
- S_1 and S_2 can send $\widehat{S}$ an index along with a bit.
- $\widehat{S}$ completes the indexed **PR** box (protocol 7) for both simulators.

The simulation strategy:
1. The simulators agree on unique indices for every commitment used in the protocol.
2. S_1 interacts with V_1 the way P_1 would. Whenever P_1 should commit, S_1 commits to random bits, just like the single-simulator from section 5.
3. For each commitment, V_2 sends S_2 a string s. S_2 sends to $\widehat{S}$ the index of the commitment and s.

4. $\widehat{S}$ runs the **PR** box (protocol 7) and replies with V_2's half of the output.
5. Whenever S_1 needs to unveil a commitment, it can be unveiled in the way S_1 desires by sending the corresponding index and bit to $\widehat{S}$.
6. $\widehat{S}$ completes the corresponding **PR** box which outputs t. $\widehat{S}$ sends t to S_1.
7. S_1 sends t to V_1.

The second part (the consistency test) can be done by having the simulators ignore the question.

Protocol 11. *(Perfectly Indistinguishable,* **PR***-Local Simulator for Protocol 8, Part 2)*

1. V_1 sends i to S_1.
2. S_1 computes $\boxed{\Omega_1} = H_{\boxed{\gamma}}(Q_i)$.
3. Using $\widehat{S}$ to break binding, S_1 convinces V_1 that $\boxed{\Omega_1}$ is actually $\boxed{A(Q_i)} \oplus H_{\boxed{\gamma}}(Q_i)$.
4. S_1 unveils $\boxed{\Omega_1}$ for V_1, who gets $\Omega_1 = H_\gamma(Q_i)$.
5. V_2 sends Q'_i to S_2.
6. S_2 responds with $\Omega_2 = H_\gamma(Q'_i)$.

By the properties of the strongly-universal-2 hash H, if $Q_i = Q'_i$ then $\Omega_1 = \Omega_2$. Otherwise $\Omega_1 \neq \Omega_2$ with probability exponentially close to one. This produces the result as desired. The simulators then feed the transcripts to V_0, and terminates simulation.

B Babai, Fortnow and Lund's MIP for Languages in NEXP

This section describes a variant of the multi-prover protocol for oracle-3-SAT found in [8]. We refer to this as the BFL protocol, or BFL classic.

Definition 10. *Let $r, s > 0$ be integers. Let z, b_1, b_2, b_3 be strings of variables, where $|z| = r$ and $|b_i| = s$. Let $B(z, b_1, b_2, b_3, t_1, t_2, t_3)$ be a Boolean formula in $r + 3s + 3$ variables. A Boolean function $A : \{0,1\}^s \to \{0,1\}$ is a 3-satisfying oracle for B if*

$$B(z, b_1, b_2, b_3, A(b_1), A(b_2), A(b_3)) = 1$$

for every string z, b_1, b_2, b_3.

B is oracle-3-satisfiable if such a function A exists.

The Oracle-3-SAT problem (B, r, s) asks whether a Boolean formula B is oracle-3-satisfiable, where r and s denote the lengths of z and b_i, as above.

Lemma 2. *Oracle-3-SAT is* **NEXP***-complete.*

Definition 11. *Let $\mathbb{F}$ be an arbitrary field. Let $\phi : \{0,1\}^m \to \{0,1\}$ be a Boolean function. An* arithmetization *of ϕ is a polynomial $f(x_1, \ldots, x_m) \in \mathbb{F}[X_1, \ldots, X_m]$ such that for all $z \in \{0,1\}^m$, $\phi(z) = 0 \Leftrightarrow f(z) = 0$. A specific one is given in [8], proposition 3.1 .*

Equivalently, the $\phi(z) = 0 \Leftrightarrow f(z) = 0$ condition can be replaced with $\phi(z) = 1 \Leftrightarrow f(z) = 0$.

Protocol 12. *(Sumcheck Protocol)*

Let $\phi(x_1, \ldots, x_m)$ be the 3-CNF formula which the prover P is trying to show to be a tautology to a verifier V. Let $\mathbb{F}$ be a field of sufficient size (of order at least $(3c+1)m$ will suffice where c is the number of clauses of ϕ).

1. V takes ϕ and computes its arithmetization f according to [8] Proposition 3.1 and sends it to P.
2. V and P agree on a set $I \subset \mathbb{F}$ of size at least $2dm$ where d is the degree of f.
3. V assigns $b_0 = 0$, which is supposed to be equal to the sum

$$\sum_{x_1=0}^{1} \cdots \sum_{x_m=0}^{1} f(x_1, \ldots, x_m)^2 = 0$$

4. $i \leftarrow 1$.
5. P sends the coefficients of the univariate polynomial in x,

$$g_i(x) = h(r_1, \ldots, r_{i-1}, x) = \sum_{x_{i+1}=0}^{1} \cdots \sum_{x_m=0}^{1} f(r_1, \ldots, r_{i-1}, x, x_{i+1}, \ldots, x_m)^2$$

6. V checks whether $b_{i-1} = g_i(0) + g_i(1)$. If not, abort.
7. V chooses a random $r_i \in I$, computes $b_i = g_i(r_i)$ and sends r_i to P.
8. If $i \leq m$ then $i \leftarrow i + 1$ and go to step 4.
9. V checks whether $b_m = f(r_1, \ldots, r_m)^2$.

Protocol 13. *(Babai, Fortnow and Lund's MIP for Oracle-3-SAT)*

Given (B, r, s) as common input.

1. (sumcheck with oracle) V and P_1 execute protocol 12. Let $(Q_{k+1}, Q_{k+2}, Q_{k+3}) = (r_{r+1}...r_{r+s}, r_{r+s+1}...r_{r+2s}, r_{r+2s+1}...r_{r+3s}) \in (I^s)^3$ be V's questions during this phase.
2. (multilinearity test) V asks P_1 to simulate an oracle storing the function A. V queries P_1 with random, linearly related values in I^s. If any response does not satisfy linearity, abort protocol. Let $Q_1, \ldots, Q_k \in I^s$ be V's questions during this phase.
3. (non-adaptiveness test) V chooses uniformly at random an i such that $1 \leq i \leq k+3$ and asks Q_i to P_2. If P_2's answer differs from that of P_1, reject. Otherwise accept.

C Non-Locality – an introduction

In this section we solely focus on the two-party single-round games and strategies that are sufficient to discuss and analyze most of the MIPs. Definitions and proofs for complete generalizations to multi-party multi-round games and strategies will appear in a forthcoming paper with co-author Adel Magra.

Games: Let V be a predicate on $A \times B \times X \times Y$ (for some finite sets A, B, X, and Y) and let π be a probability distribution on $A \times B$. Then V and π define a (single-round) game G as follows: A pair of questions (a, b) is randomly chosen according to distribution π, and $a \in A$ is sent to Alice and $b \in B$ is sent to Bob. Alice must respond with an answer $x \in X$ and Bob with an answer $y \in Y$. Alice and Bob win if V evaluates to 1 on (a, b, x, y) and lose otherwise.

Strategies: Two-Party Channels A strategy for Alice and Bob is simply a probability distribution $P_{(x,y|a,b)}$ describing exactly how they will answer (x, y) on every pair of questions (a, b). We now breakdown the set of all possible strategies for Alice and Bob according to their *non-locality.*

Deterministic and Local Strategies: A strategy $P_{(x,y|a,b)}$ is *deterministic* if there exists functions $f_A : A \to X, f_B : B \to Y$ such that

$$P_{(x,y|a,b)} = \begin{cases} 1 & \text{if } x = f_A(a) \text{ and } y = f_B(b) \\ 0 & \text{otherwise} \end{cases} .$$

A deterministic strategy corresponds to the situation where Alice and Bob agree on their individual actions before any knowledge of the values a, b is provided to them. In this case they use only their own input to determine their individual output.

A strategy $P_{(x,y|a,b)}$ is *local* if there exists a finite set R and functions $f_A : A \times R \to X, f_B : B \times R \to Y$ such that

$$P_{(x,y|a,b)} = \frac{|\{r \in R : x = f_A(a,r) \text{ and } y = f_B(b,r)|}{|R|}.$$

A local strategy corresponds to the situation where Alice and Bob agree on a deterministic strategy selected uniformly among $|R|$ such possibilities. The choice r of Alice and Bob's strategy, and the choice of inputs (a, b) provided to Alice and Bob are generally agreed to be statistically independent random variables.

C.1 Local Reducibility

We now turn to the notion of locally reducing a strategy to another, that is how Alice and Bob limited to local strategies but equipped with a particular (not necessarily local) strategy U' are able to achieve another particular (not necessarily local) strategy U. For this purpose we introduce a notion of distance between strategies in order to analyze strategies that are approaching each other asymptotically.

Distances between Strategies: Several distances could be selected here as long as their meaning as it approaches zero are the same. In the definitions below, U, U' are strategies and $\mathcal{U}'$ is a finite set of strategies.

Definition 12. $|U, U'| = \sum_{a,b,x,y} |P_U(x, y|a, b) - P_{U'}(x, y|a, b)|$

Definition 13. $|U, \mathcal{U}'| = \min_{U' \in \mathcal{U}'} |U, U'|$

Local extensions of Strategies: For natural integer n, we define the set $\mathrm{LOC}^n(U)$ of strategies that are local extensions (of order n) of U to be all the strategies Alice and Bob can achieve using local strategies where strategy U may be used up to n times as sub-routine calls[‖]. If we restrict all the functions used to be polynomial-time computable we analogously define $\underset{\text{poly}}{\mathrm{LOC}^n}(U)$.

Definition 14. *U' Locally (poly-)Reduces to U ($U' \leq_{\underset{(poly)}{\mathrm{LOC}}} U$) iff* $\lim_{n\to\infty} |U', \underset{(poly)}{\mathrm{LOC}^n}(U)| = 0.$

Definition 15. *U' is Locally (poly-)Equivalent to U ($U' =_{\underset{(poly)}{\mathrm{LOC}}} U$) iff $U' \leq_{\underset{(poly)}{\mathrm{LOC}}} U \leq_{\underset{(poly)}{\mathrm{LOC}}} U'$.*

[‖] Done by selecting functions $f^0_A : A \times R \to A,\ f^1_A : A \times X \times R \to A, ...,\ f^{n-1}_A : A \times X^{n-1} \times R \to A,$ $f^n_A : A \times X^n \times R \to X$ to determine the input of each sub-routine from input a and previous outputs.

Non-Adaptive extensions of Strategies: For natural integer n, we define the set $\mathrm{NAD}^n(U)$ of strategies that are Non-Adaptive extensions (of order n) of U to be all the strategies Alice and Bob can achieve using Non-Adaptive strategies where strategy U may be used up to n times as sub-routine calls**. If we restrict the functions used to be poly-time computable we get $\underset{\text{poly}}{\mathrm{NAD}^n}(U)$.

Definition 16. *U' Non-Adaptively (poly-)Reduces to U ($U' \leq_{\substack{\mathrm{NAD}\\ (poly)}} U$) iff* $\lim_{n\to\infty} |U', \underset{(poly)}{\mathrm{NAD}^n}(U)| = 0$.

Definition 17. *U' is Non-Adaptively (poly-)Equivalent to U ($U' =_{\substack{\mathrm{NAD}\\ (poly)}} U$) iff $U' \leq_{\substack{\mathrm{NAD}\\ (poly)}} U \leq_{\substack{\mathrm{NAD}\\ (poly)}} U'$.*

In general, Non-Adaptive reducibility is a weaker notion than local reducibility. However, for certain distributions $\mathbf{U}$ it may result that $\{D|D \leq_{\substack{LOC\\ \text{(poly)}}} \mathbf{U}\} = \{D|D \leq_{\substack{NAD\\ \text{(poly)}}} \mathbf{U}\}$ as follows.

C.2 Locality

We now define the lowest of the non-locality classes $\mathbb{LOC}$. We could define it directly from the notion of local strategies as defined above, but for analogy with the other classes we later define, $\mathbb{LOC}$ is defined as all those strategies locally reducible to a *complete* strategy we call **ID** (see **Fig. 2**). Of course, any strategy is complete for this class.

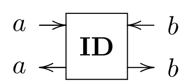

Fig. 2. an **ID**-box

Definition 18. $\mathbb{LOC} = \{U|U \leq_{\mathrm{LOC}} \mathbf{ID}\}$ *and* $\underset{poly}{\mathbb{LOC}} = \{U|U \leq_{\substack{\mathrm{LOC}\\ poly}} \mathbf{ID}\}$

Note: $\mathbb{LOC}$ is the class of strategies that John Bell [29] considered as classical hidden-variable theories that he compared to entanglement. It is also the class of strategies that BenOr, Goldwasser, Kilian and Wigderson [3] chose to define classical Provers in Multi-Provers Interactive Proof Systems. $\mathbb{LOC}$ is also those strategies Non-Adaptively reducible to **ID**

Definition 19. *Alternatively,* $\mathbb{LOC} = \{U|U \leq_{\mathrm{NAD}} \mathbf{ID}\}$ *and* $\underset{poly}{\mathbb{LOC}} = \{U|U \leq_{\substack{\mathrm{NAD}\\ poly}} \mathbf{ID}\}$

Alternatively, we can also define $\mathbb{LOC}$ from an empty box as used in the core of this paper

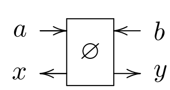

Fig. 3. an $\varnothing$-box where $x \in X$ and $y \in Y$ are uniform and independent of everything else

Definition 20. *Alternatively,* $\mathbb{LOC} = \{U|U \leq_{\mathrm{NAD}} \varnothing\} = \{U|U \leq_{\mathrm{LOC}} \varnothing\}$

** Done by selecting functions $f_A^0 : A \times R \to A,\ f_A^1 : A \times R \to A, ...,\ f_A^{n-1} : A \times R \to A,$ $f_A^n : A \times X^n \times R \to X$ to determine the input of each sub-routine from input a only.

C.3 One-Way Signalling

We now turn to One-Way Signalling which allows communication from one side to the other. We name the directions arbitrarily Left and Right. We define **R**-$\mathbb{SIG}$ (resp. **L**-$\mathbb{SIG}$) as all those strategies locally reducible to a *complete* strategy we call **R-SIG** (see **Fig. 4**) (resp. **L-SIG** (see **Fig. 5**)). These classes are useful to define what it means for a strategy to *signal* as well as the notion of *No-Signalling* strategies.

Fig. 4. an **R-SIG**-box

Definition 21. $\mathbf{R}\text{-}\mathbb{SIG} = \{U | U \leq_{\mathrm{LOC}} \mathbf{R\text{-}SIG}\}$ *and* $\mathbf{R}\text{-}\underset{poly}{\mathbb{SIG}} = \{U | U \underset{poly}{\leq_{\mathrm{LOC}}} \mathbf{R\text{-}SIG}\}$

Definition 22. *We say that U Right Signals (is* **R**-$\mathbb{SIG}$*-verbose*[††]*) iff* $\mathbf{R\text{-}SIG} \leq_{\mathrm{LOC}} U$*.*

Fig. 5. an **L-SIG**-box

Definition 23. $\mathbf{L}\text{-}\mathbb{SIG} = \{U | U \leq_{\mathrm{LOC}} \mathbf{L\text{-}SIG}\}$ *and* $\mathbf{L}\text{-}\underset{poly}{\mathbb{SIG}} = \{U | U \underset{poly}{\leq_{\mathrm{LOC}}} \mathbf{L\text{-}SIG}\}$

Definition 24. *We say that U Left Signals (is* **L**-$\mathbb{SIG}$*-verbose) iff* $\mathbf{L\text{-}SIG} \leq_{\mathrm{LOC}} U$*.*

Definition 25. *We say that U Signals iff U Right Signals or Left Signals.*

We prove a first result that is intuitively obvious. We show that the complete strategy **R-SIG** cannot be approximated in **L-**$\mathbb{SIG}$ and the other way around.

Theorem 14. $\mathbf{R\text{-}SIG} \notin \mathbf{L}\text{-}\mathbb{SIG}$ *and* $\mathbf{L\text{-}SIG} \notin \mathbf{R}\text{-}\mathbb{SIG}$*.*

Proof. Follows from a simple capacity argument. For all n, all the channels in $\mathrm{LOC}^n(\mathbf{R\text{-}SIG})$ have zero left-capacity, while **L-SIG** has non-zero left-capacity. And vice-versa.

C.4 Signalling

We are now ready to define the largest of the non-locality classes $\mathbb{SIG}$. Indeed every possible strategy is in $\mathbb{SIG}$.

Definition 26. $\mathbb{SIG} = \{U | U \leq_{\mathrm{LOC}} \mathbf{SIG}\}$ *and* $\underset{poly}{\mathbb{SIG}} = \{U | U \underset{poly}{\leq_{\mathrm{LOC}}} \mathbf{SIG}\}$

Definition 27. *We say that U Fully Signals (is* $\mathbb{SIG}$*-verbose) iff U Right Signals and Left Signals.*

[††] We define the notion of $\mathbb{L}$-verbose in analogy to $\mathbb{NP}$-hard: it means "as verbose as any distribution in non-locality class $\mathbb{L}$". In consequence, a distribution U is $\mathbb{L}$-complete if $U \in \mathbb{L}$ and U is $\mathbb{L}$-verbose.

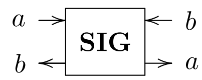

Fig. 6. a **SIG**-box

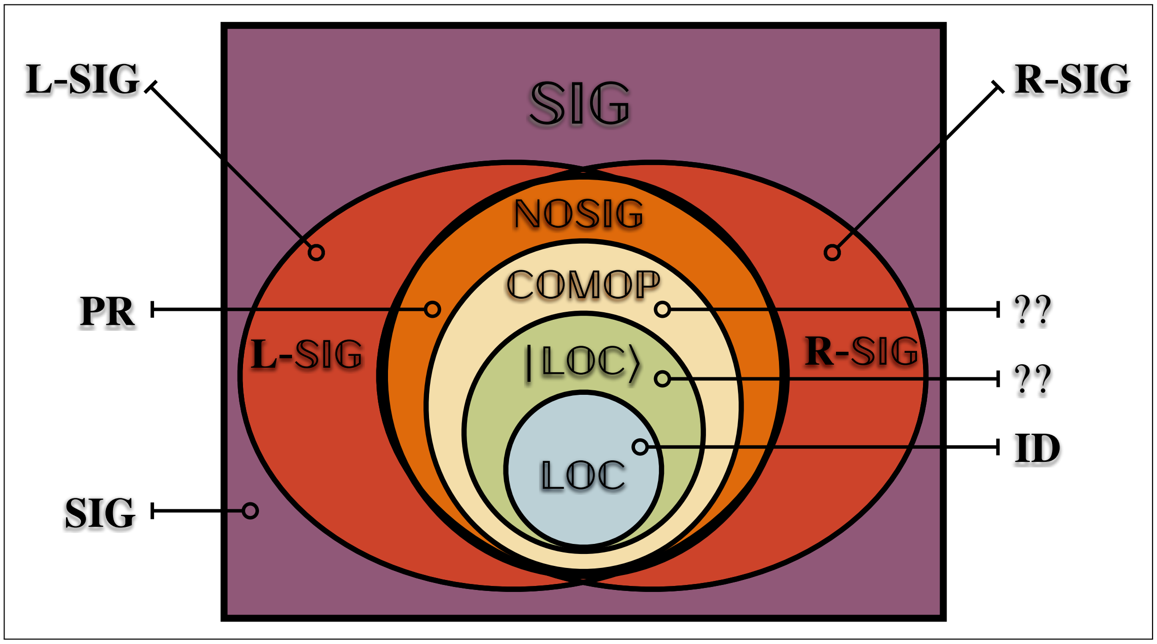

Fig. 7. Non-locality Hierarchy and complete (two-party) distributions in each class.

C.5 No-Signalling

We finally define the less intuitive non-locality class $\mathbb{NOSIG}$ in relation to classes defined above.

Definition 28. $\mathbb{NOSIG} = \mathbf{R}\text{-}\mathbb{SIG} \bigcap \mathbf{L}\text{-}\mathbb{SIG}$ *and* $\underset{poly}{\mathbb{NOSIG}} = \mathbf{R}\text{-}\underset{poly}{\mathbb{SIG}} \bigcap \mathbf{L}\text{-}\underset{poly}{\mathbb{SIG}}$.

A similar characterization may be found in [35] Section 3 and [36] Corollary 3.5.

Theorem 15. . *The above definition of* $\mathbb{NOSIG}$ *exactly coincides with the* traditional *notion of No-Signalling [37].*

Intuitively, a distribution $P(x, y|a, b)$ is No-Signalling as long as for every a the $x|b$ and for every b the $y|a$ channels have zero capacity.

Note: Forster and Wolf [38] have proved that **PR** (see **Fig. 8**) is complete for $\mathbb{NOSIG}$ distributions under an asymptotic definition similar to ours.

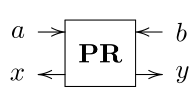

Fig. 8. a **PR**-box satisfying the CHSH condition,that $a \wedge b = x \oplus y$, uniformly among solutions

Fig. 7 shows the relation of these classes as well as the case obtained via quantum entanglement ($|\mathbb{LOC}\rangle$) as considered by Bell [29] and via commuting-operators ($\mathbb{COMOP}$) as defined by

Ito, Kobayashi, Preda, Sun, and Yao [39]. We include those for completeness but will not discuss these particular classes any further in this work.

Definition 29. *We say that U does not Signal iff U does not Right Signal nor Left Signal iff $U \in \mathbb{NOSIG}$.*

D Visual description of the new model

D.1 Local Multi-Prover Interactive Proofs

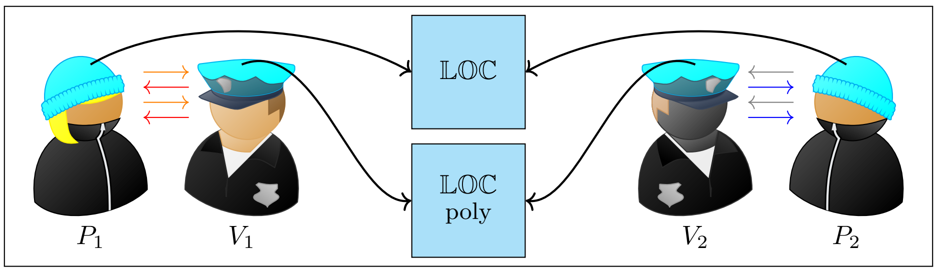

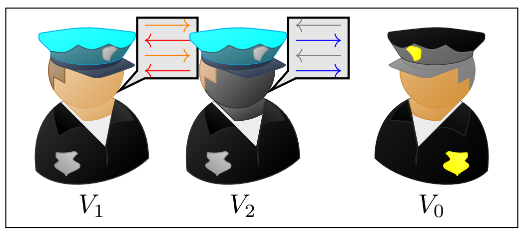

Fig. 9. Interrogation phase (top) followed by decision phase (bottom).

In the Interrogation phase (see **Fig. 9**) $V_1, ..., V_k$ (equipped with an arbitrary local correlator) individually interrogate $P_1, ..., P_k$ (equipped with an arbitrary local correlator). At the end of the interactive part, all the $V_1, ..., V_k$ report to V_0 who takes the final decision. The corresponding complexity class is $\mathbf{MIP} = \mathbf{MIP}^{\mathbb{LOC}}_{\mathbb{LOC}_{\text{poly}}} = \mathbf{NEXP}$.

D.2 Entangled Multi-Prover Interactive Proofs

In the Interrogation phase (see **Fig. 10**) $V_1, ..., V_k$ (equipped with an arbitrary entangled correlator) individually interrogate $P_1, ..., P_k$ (equipped with an arbitrary entangled correlator). At the end of the interactive part, all the $V_1, ..., V_k$ report to V_0 who takes the final decision. The corresponding complexity class is $\mathbf{MIP}^* = \mathbf{MIP}^{|\mathbb{LOC}\rangle}_{|\mathbb{LOC}\rangle_{\text{poly}}} \supseteq \mathbf{NEXP}$.

D.3 No-Signalling Multi-Prover Interactive Proofs

In the Interrogation phase (see **Fig. 11**) $V_1, ..., V_k$ (equipped with an arbitrary No-Signalling correlator) individually interrogate $P_1, ..., P_k$ (equipped with an arbitrary No-Signalling correlator). At the end of the interactive part, all the $V_1, ..., V_k$ report to V_0 who takes the final decision. The corresponding complexity class is $\mathbf{MIP}^{ns} = \mathbf{MIP}^{\mathbb{NOSIG}}_{\mathbb{NOSIG}_{\text{poly}}} = \mathbf{EXP}$.

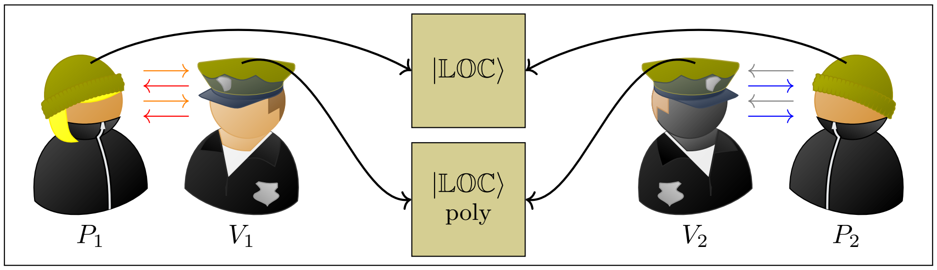

Fig. 10. Interrogation phase.

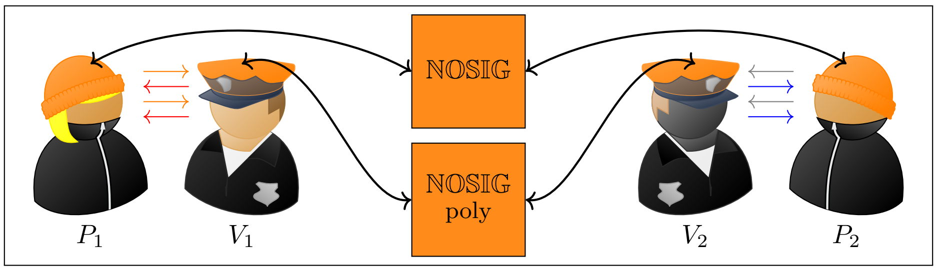

Fig. 11. Interrogation phase.

As noted before, most MIPs found in the literature are actually (non-adaptive) local-verifier MIPs (see **Fig. 12**) yielding for instance $\mathbf{MIP}^{ns} = \mathbf{MIP}^{\text{NOSIG}}_{\substack{\text{LOC} \\ \text{poly}}}$.

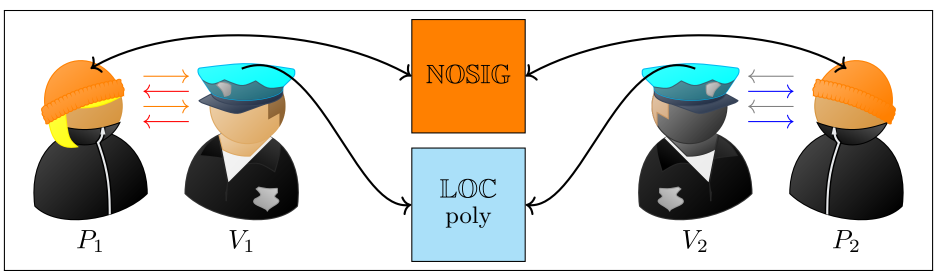

Fig. 12. Interrogation phase.

D.4 A New, Stronger Flavour of Zero-Knowledge

Traditionally zero-knowledge is defined as a property of the honest provers for all (polynomial-time) verifiers

$$\forall_{\text{poly}} V' \; \exists_{\text{poly}} S \; \forall x \in L \; \forall w \;\; \mathbf{VIEW}_{V'}[P_1, ..., P_k, V'](w, x) = S(w, x).$$

However, in the present context, the fact that the simulation of V''s view via a single centralized simulator S, achieving zero-knowledge is rather easy because such an S can cheat the binding property of the commitments at will. The intuition behind the original definition is that the verifier is unable to convince a third party (a Judge J_0) because the **VIEW** he reports (see **Fig. 13**) could have been equally created (with the same distribution) by a simulator. Nevertheless, a stronger flavour of zero-knowledge is achieved if the simulator is not invoking its full signalling power whenever the verifier does not use such power.

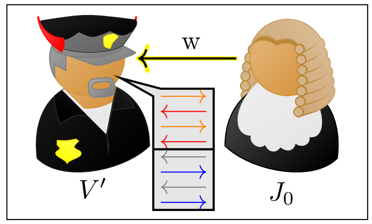

Fig. 13. (Interac/Simula)tion-Distinction phase.

For all non-locality levels starting with $\widehat{\mathbb{S}}$ and up, the simulators S_i do not need more non-local power than the verifiers V_i'. The ultimate (strongest) notion of "$\mathbb{LOC}$-local ZK" being $\mathbf{ZK}^{\overset{\text{poly}}{\mathbb{LOC}}}$ poly because at all levels V' is simulated by a simulator with no extra non-local power, whereas at the opposite end of the spectrum $\mathbf{ZK}^{\overset{\text{poly}}{\mathbb{SIG}}}$ is what is generally considered zero-knowledge with a single simulator or a group of signalling simulators.

This stronger notion of zero-knowledge is particularly interesting in the relativistic bit-commitment scenario where a pair of judges may provide separate auxiliary-inputs to spatially separated verifiers pretending to be speaking to powerful provers. If the verifiers can report their conversation fast enough to the judges (but not interact with the judges however), they must be able to do so without invoking signalling because of the distance separating them. If a pair of simulators can produce the same distribution of views in the same context, we obtain a stronger flavour of zero-knowledge (See **Fig. 14**).

The results of this paper, depending on the specific bit commitment used, may be achieved under a stronger flavour of zero-knowledge if a member of the non-locality class $\widehat{\mathbb{S}}$ is enough to break the binding property of the commitments. For instance, the result of section 5 is really $\mathbf{ZK}^{\overset{\text{poly}}{\mathbb{NOSIG}}}\mathbf{MIP}^{\mathbb{LOC}}_{\underset{\text{poly}}{\mathbb{LOC}}} = \mathbf{NEXP}$ although existing proofs usually mean $\mathbf{ZK}^{\overset{\text{poly}}{\mathbb{SIG}}}\mathbf{MIP}^{\mathbb{LOC}}_{\underset{\text{poly}}{\mathbb{LOC}}} = \mathbf{NEXP}$. Using the bit commitment scheme based on the magic square game of [40] we can also obtain $\mathbf{ZK}^{\overset{\text{poly}}{|\mathbb{LOC}\rangle}}\mathbf{MIP}^{\mathbb{LOC}}_{\underset{\text{poly}}{\mathbb{LOC}}} = \mathbf{NEXP}$.

Some interesting questions resulting from this definition is whether any higher class such as $\mathbf{ZK}^{\overset{\text{poly}}{\mathbb{LOC}}}\mathbf{MIP}^{\mathbb{LOC}}_{\underset{\text{poly}}{\mathbb{LOC}}}$ or $\mathbf{ZK}^{\overset{\text{poly}}{\mathbb{NOSIG}}}\mathbf{MIP}^{\mathbb{NOSIG}}_{\underset{\text{poly}}{\mathbb{NOSIG}}}$ contains more than the natural examples such as GRAPH ISO or CODE EQUIV already found in the most natural class $\mathbf{ZK}^{\overset{\text{poly}}{\mathbb{SIG}}}\mathbf{MIP}^{\mathbb{SIG}}_{\underset{\text{poly}}{\mathbb{SIG}}} = \mathbf{ZKIP}$.

D.5 A note on notation

$$\mathbf{ZK}^{\mathbb{S}}\mathbf{MIP}^{\mathbb{P}}_{\mathbb{V}}$$

is the complexity class of Zero-Knowledge Multi-provers Interactive Proofs where (honest and dishonest) provers are restricted to non-locality class $\mathbb{P}$ (important for soundness), where the

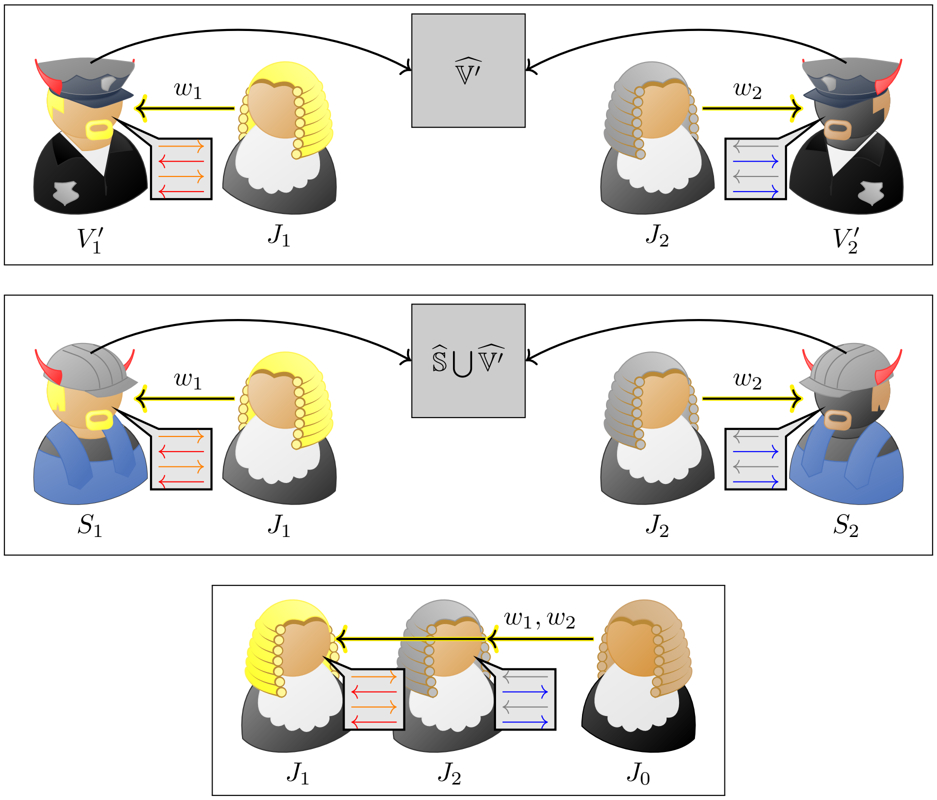

Fig. 14. Interrogation or Simulation phase (top) followed by Distinction phase (bottom).

honest verifier is from non-locality class $\mathbb{V}$ (also important for soundness), and where the Zero-Knowledge simulators are from non-locality class $\mathbb{S}$ unless $\widehat{V'}$ is outside of $\mathbb{S}$ in which case they are from the class of $\widehat{V'}$.

E CMOSSY 3-COL Honest-Verifer Zero-Knowledge Interactive Proof

---

Protocol 16. *Two-out-of-Three-Prover, 3-COL.*

The verifiers V_1, V_2, V_3 pre-agree on random edges (n_0, n_1) and (n_2, n_3), random strings $r_0, r_1, r_2, r_3 \neq 0$ and the provers P_1, P_2, P_3 pre-agree on random values $b_{n_i} : n_i \in V$ and a random 3-colouring of G: $\{c_{n_i} \in \{0, 1, 2\} : n_i \in V\}$ such that $(n_i, n_j) \in E \implies c_{n_j} \neq c_{n_i}$. They also pre-compute an array $W[n_i, r] := b_{n_i} \cdot r + c_{n_i} : n_i \in V, r \in \{1, 2\}$. The values $(n_0, n_1, r_0, r_1), (n_2, n_3, r_2, r_3)$ are selected under one of three constraints: either

$$(n_0, n_1) = (n_2, n_3), r_0 \neq r_2, r_1 \neq r_3 \text{ or}$$

$$\exists i, j \in \{0, 1\} \times \{2, 3\} : n_i = n_j, r_i = r_j \text{ or}$$

$$(n_0, n_1) = (n_2, n_3), (r_0, r_1) = (r_2, r_3).$$

The verifiers V_1, V_2, V_3 pre-select P_A, P_B at random from P_1, P_2, P_3.
Commit phase:

– P_A receives nodes n_0, n_1, strings r_0, r_1 from V_A and if $(n_0, n_1) \in E$, replies $W[n_0, r_0], W[n_1, r_1]$.
– P_B receives nodes n_2, n_3, strings r_2, r_3 from V_B and if $(n_2, n_3) \in E$, replies $W[n_2, r_2], W[n_3, r_3]$.

Check phase:

Consistency Test:

– if $(n_0, n_1) = (n_2, n_3), (r_0, r_1) = (r_2, r_3)$ then V_A, V_B accept iff

$$(W[n_0, r_0], W[n_1, r_1]) = (W[n_2, r_2], W[n_3, r_3]).$$

Edge-Verification Test:

– if $(n_0, n_1) = (n_2, n_3), r_0 \neq r_2, r_1 \neq r_3$ then V_A, V_B accept iff

$$W[n_0, r_0] + W[n_2, r_2] \neq W[n_1, r_1] + W[n_3, r_3].$$

Well-Definition Test:

– if $\exists i, j \in \{0, 1\} \times \{2, 3\} : n_i = n_j, r_i = r_j$ then V_A, V_B accept iff $W[n_i, r_i] = W[n_j, r_j]$.

Protocol 17. *HV Two-prover simulation.*

Commit phase:

– Let π be a uniform permutation of $\{0, 1, 2\}$ and let $coco := 0$.
– $\forall n \in V, r \in \{0, 1, 2\}$, let $mark[n, r] := false$, $count[n] := 0$, $colour[r] := \pi(r)$.
– S runs V_1, V_2, V_3 until it receives $(n_{2A-2}, n_{2A-1}, r_{2A-2}, r_{2A-1})$, $(n_{2B-2}, n_{2B-1}, r_{2B-2}, r_{2B-1})$ from V_A, V_B.
– Whenever $(n_{2i-2}, n_{2i-1}) \in E$ is provided by V_i, S replies (w_{2i-2}, w_{2i-1}), both computed as follows for $k \in \{2i-2, 2i-1\}$:
  - If $\neg mark[n_k, r_k]$ then
    - ∗ If $count[n_k] = 0$ then pick $W[n_k, r_k]$ uniformly in $\{0, 1, 2\}$.
    - ∗ If $count[n_k] = 1$ then
      - · Let $W[n_k, r_k] := -colour[coco] - W[n_k, -r_k]$
      - · Let $coco := coco + 1$.
    - ∗ Let $mark[n_k, r_k] := true$, $count[n_k] := count[n_k] + 1$.
  - Let $w_k := W[n_k, r_k]$.
