New Perspectives on Multi-Prover Interactive Proofs

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 existing multi-prover interactive proof framework suffers from incompleteness in terms of soundness and zero-knowledge that is not completely addressed in the literature. The problem is that the existing definitions of what is local, entangled and no-signalling are not rich enough to capture the full generality of multi-prover interaction. In general, existing proofs do not take into account possible *changes in locality* either during a protocol's execution or when protocols are composed together. This is especially problematic for zero-knowledge, as composing commitments is the only known way of achieving zero-knowledge outside of some **NP**-intermediate languages.
In this work, we introduce the *locality hierarchy* for multiparty (multi-round) interaction, and for the first time a complete definition of multi-round multiparty no-signalling distributions and strategies. Within this framework, we define the *locality* of a protocol which involves the provers, verifiers, simulators and distinguishers. We show that an existing protocol for **NEXP** [1] and a zero-knowledge variant we introduce are sound in a local sense, but are zero-knowledge in a sense that is even stronger than usually understood. All prior claims of zero-knowledge proofs in the multi-prover model were actually incorrect. Finally, we present similar constructions for entangled and no-signalling prover sets for **NEXP** and **EXP** based on [2] and [3] using new multi-prover commitment schemes.

1 Introduction

The idea behind multi-prover interactive proofs (MIPs) should have been simple: a weak *verifier* interrogates powerful but isolated *provers* in order to ascertain something outside of its computational powers. The "multi-prover" part of the idea is that the verifier may check the consistency of a prover's answers against another prover, thereby overcoming a fundamental limitation of the verifier's computational weakness in the single-prover model.

The simplicity is lost when we think clearly on what it means for the provers to be "isolated". This is because we would like to prove statements of the form, "If the provers have cheated, then they had talked." This is the basis for soundness in protocols such as [4], [1] and [5]. However, it has been subsequently discovered that the provers can cheat while *technically* not talking to each other; that is, the act of breaking some of these protocols does not reduce to signalling. This problem is compounded by the fact that the verifier may, by design or not, perform a no-signalling task *for* the provers; worse yet, it may even directly courier messages between them. Therefore, not only can we not ignore the locality of the *provers*, we also cannot ignore that of the *verifier*. Analyses of existing multi-prover protocols do not account for any of these problems.

Furthermore, in the time since [4], there have been several works which attempted to augment the multi-prover model by giving the provers non-locality. For instance, [2] looked at entangled provers, while [3] showed that provers with arbitrary no-signalling powers can be relied on to recognize languages up to **EXP**. These claims, in addition to those made in [4], [1] and [5] must now be interpreted carefully, as they never stated explicitly the locality of both the provers *and* the verifiers.

1.1 Our Contribution

In this paper, we define the notion of "locality hierarchy" in an information-theoretical sense, and also the local, entangled and no-signalling classes within that hierarchy. The definition for "no-signalling" is a direct

* Supported in part by FRQNT (via INTRIQ) and NSERC (via CryptoWorks21 and the Discovery grant program).
** Supported in part by Professors V. Chvátal, J. Clark, C. Crépeau, and D. Ford.

generalization of single-round no-signalling distributions; in effect, we now have a working definition of multi-round, multiparty non-local distributions. We define the notion of "zero-knowledge locality" for MIPs, and in doing so we generalize the notion of a simulator in the context of non-local interaction. We combine these two ideas to give a proof that $\mathbf{NEXP} = \mathbf{ZK}^{\neq}\mathbf{MIP}^{\textcircled{\scriptsize ?}}$. That is, there exists a local ZKMIP for **NEXP** which can be simulated if the simulators are given at least no-signalling powers. This is actually the very first complete proof $\mathbf{NEXP} = \mathbf{ZKMIP}$.

2 Preliminaries

The theory of multi-prover interactive proofs (MIPs) originated from the work of Ben-Or, Goldwasser, Kilian and Wigderson [4]. We denote the class of languages with such interactive proofs by **MIP** (and its zero-knowledge counterpart **ZKMIP**). In that paper and subsequent work of Babai, Fortnow and Lund [1], it was claimed that $\mathbf{ZKMIP} = \mathbf{MIP} = \mathbf{NEXP}$.

The proof of security in [4] and many subsequent MIPs reduces the breaking of soundness to signalling. However, in the last decade, two major problems with MIPs/ZKMIPs have emerged. The first is that the provers do not actually need to signal in order to break some MIPs, as demonstrated in the work of Cleve, Høyer, Toner and Watrous [6]; they can perform *no-signalling tasks* which do not allow communication (for example, using shared entanglement). That is, there is a fundamental and yet subtle difference between what is *local* and what is *no-signalling*. The second by Crépeau, Salvail, Simard and Tapp [7] is that while the provers are unable to signal between themselves, the verifier could inadvertently perform a non-local task for them; in the extreme case, the verifier may plainly signal *for* the provers. This can happen while the provers are perfectly isolated and local.

The ZKMIPs as found in [4] and [8] require provers to authenticate a question before the verifier can ask that question to another prover. The provers must also show that the authentication is not an attempt at signalling. This means that a relativistic guarantee of no-signalling (no faster-than-light communication) is insufficient to securely implement these protocols. This is another indication that existing ZKMIPs contain a certain ambiguity about their locality.

For an explicit example of how such ambiguity regarding locality can lead to catastrophic protocol failure, see [9].

The role that the verifier must play in these MIPs was studied in [7]. It was defined and shown that a verifier must be *isolating*, so that it will never (inadvertently or not) perform a *non-local* task (no-signalling or signalling). We have shown in [9] that many existing MIPs do not satisfy isolation, even in a weak sense.

More recently, the model of multi-prover interactive proofs was extended to allow entangled provers and the class of languages accepted under this new setting is called $\mathbf{MIP}^{\star}$ [10]. It was recently shown that $\mathbf{NEXP} \subseteq \mathbf{MIP}^{\star}$ [2] but we do not know whether equality holds. Similarly, the model of multi-prover interactive proofs was extended to allow no-signalling provers and the class of languages accepted under this new setting is called $\mathbf{MIP}^{\mathbf{ns}}$ [10]. We now know that $\mathbf{MIP}^{\mathbf{ns}} = \mathbf{EXP}$ [3]. We use some of these results to illustrate our explanation.

2.1 Terminology, Definitions and Previous Work

The terminology we are about to define originates mostly from physics. The acclaimed work from John Bell [11] in the 1960s can be summarized as "It seems that quantum entanglement allows for *non-local* yet *no-signalling* distributions". However, it turns out that quantum physics does not allow *all* such no-signalling distributions. For instance, the CHSH game [12] cannot be achieved perfectly from entanglement alone. The best possible strategy using entanglement can win the game $\cos(\pi/8)^2 \approx 85\%$ of the time, whereas any local strategy can only succeed up to 75% of the time [6]. It is also known that any strategy winning the CHSH game 91% of the time or more can be turned into another strategy winning essentially 100% of the time [13]. In this case, we can see the correspondence between game-winning probabilities and locality clearly: $\leq 75\%$, $\approx 85\%$, and $\geq 91\%$.

The terms "communicating" and "signalling" are used interchangeably throughout this work and should have the obvious meaning of information transfer between two or several parties. Signalling provers is essentially the same as a single prover because we put no restriction whatsoever on their communication (potential

interesting sub-cases arise when we restrict the amount of communication they can actually use, or allow only one of them to signal; we do not address them here). In the context of multiple parties, we consider signalling from subsets of participants to subsets of participants to be, in fact, signalling, even if it does not occur between individual participants.

As soon as we prohibit communication, we need to define what "no-signalling" actually means. The initial intuition [11, 4] was that "non-signalling" = "locality", meaning that the provers are allowed to share arbitrary amount of randomness before being restricted to computations involving only this local randomness. However, it was later understood that certain classes of probability distributions cannot be shared in such a local fashion, such as those which arise from entanglement. The term *no-signalling* was later coined to define "everything but signalling". Of course this includes locality, but also strictly more. Typical examples are the CHSH Game (on inputs a, b output uniform x, y s. t. $x \oplus y = a \times b$) and the Magic Square Game [6]. No-signalling distributions were studied by Khalfin and Tsirelson [14], and Rastall [15]. Popescu and Rohrlich discussed them in the context of non-local games in [16]. MIPs which are secure against no-signalling provers can be found in [17–23] but almost entirely in the single-round model (except for Fehr and Fillinger [24]).

We now define a few interesting sub-cases[‡] and consider MIPs and ZKMIPs in these scenarios.

2.2 Local

Definition (LOCAL (☯)): Let $M_1, ..., M_k$ be k Turing machines. All machines have a read-only input tape, a work tape and a random tape. In addition, $M_1, ..., M_k$ share an infinite read-only random tape. Every M_i has one write-only communication tape on which it writes messages for another machine W_i. Every M_i has one read-only communication tape on which it reads messages from another machine W_i. We call $(M_1, ..., M_k)$ a *local* k-party interactive TM.

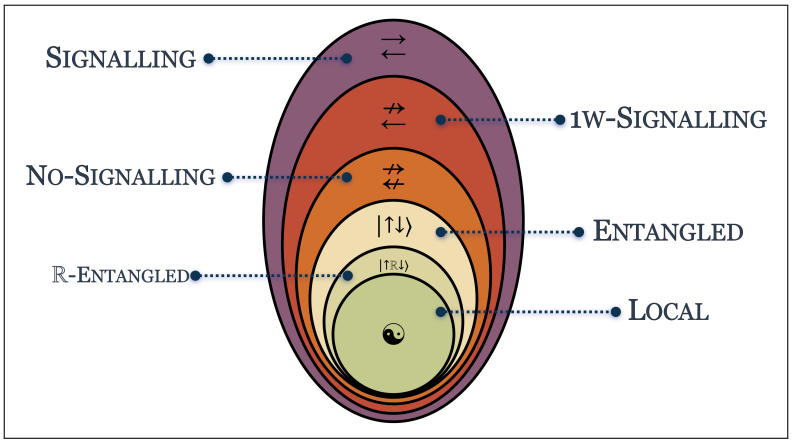

Fig. 2.2.1. locality hierarchy.

Definition (MIP☯): Let $(P_1, ..., P_k)$ be a *local* k-party interactive TM which is computationally unbounded and $(V_1, ..., V_k)$ be another *local* k-party interactive TM which is probabilistic polynomial time bounded. Each P_i can only exchange messages with the corresponding V_i and vice versa. A special party V_0 has read-only access to the k work tapes of $V_1, ..., V_k$ as well as a read-only input tape, a work tape and a random tape. V_0 has two special terminal states *accept* and *reject*. The outcome of a computation involving V_0 is defined as its final state.

[‡] It may seem that the only models which make sense are "local" and "entangled" because they are motivated by physical models of reality. Nevertheless, the no-signalling model turns out to be useful under certain circumstances as explained in [3] section 1.2. Moreover, soundness in relativistic cryptography should be proven in the no-signalling multi-prover model if no-signalling is the sole assumption we want to make.

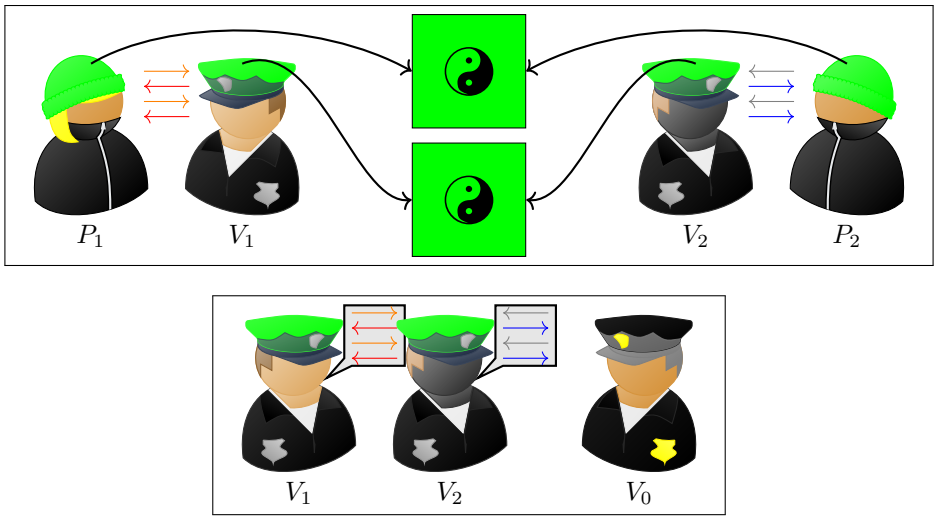

Fig. 2.2.2. Interrogation phase (top) followed by decision phase (bottom).

We call $(P_1, ..., P_k, V_0, V_1, ..., V_k)$ a *local* k-prover interactive protocol. A local multi-prover interactive proof for a language L is a local k-prover interactive protocol $(P_1, ..., P_k, V_0, V_1, ..., V_k)$ satisfying the following extra conditions:

(Completeness) $\forall x \in L, \Pr[(P_1, ..., P_k, V_0, V_1, ..., V_k)(x) = accept] > \frac{2}{3}$,
(Soundness) $\forall (P'_1, ..., P'_k), \forall x \notin L, \Pr[(P'_1, ..., P'_k, V_0, V_1, ..., V_k)(x) = accept] < \frac{1}{3}$.

The requirement that $(V_1, ..., V_k)$ be local implies that $(P_1, ..., P_k)$ cannot go beyond its own locality via V. Each pair of parties $\langle P_i, V_i \rangle$ is as local as P_i by itself (see Fig. 2.2.2). The reader may choose that restricting $(V_1, ..., V_k)$ to be local is just *a proof technique* to simplify soundness arguments, but we believe that this restriction should be a fundamental part of the definition of multi-prover interaction. Without it, the provers are not necessarily local. Fortunately, most MIPs found in the literature are actually local MIPs: it is simply a matter of "splitting" the honest verifier V as local verifiers $(V_0, V_1, ..., V_k)$ (most of the time simply (V_0, V_1, V_2)). The only non-local part is actually V_0, but it cannot influence the other verifiers since it only reads and cannot communicate.

2.3 Entangled

Definition (Entangled $\left(|\uparrow\downarrow\rangle\right)$): Let $M_1, ..., M_k$ be k Turing machines. All machines have a read-only input tape, a work tape and a random tape. In addition, M_i may access a special read-only *quantum* tape where (entangled) quits are stored with the following properties: each M_i controls the location of its read-head on its own tape and in any location, M_i may measure the qubit to which its head is pointing in one of a small number of possible basis. As a result of this measurement the observed classical bit is stored in the same location and may be read classically from this point on. An initial quantum entangled state $|\Psi\rangle$ is arbitrarily split to the quantum tapes of $M_1, ..., M_k$[***]. Every M_i has one write-only communication tape on which it writes messages for another machine W_i. Every M_i has one read-only communication tape on which it reads messages from another machine W_i. We call $(M_1, ..., M_k)$ a *local* k-party interactive TM.

Definition (MIP$^{|\uparrow\downarrow\rangle}$): Let $(P_1, ..., P_k)$ be an *Entangled* k-party interactive TM which is computationally unbounded and $(V_1, ..., V_k)$ be another *Entangled* k-party interactive TM which is probabilistic polynomial time bounded. Each P_i can only exchange messages with the corresponding V_i and vice versa. A special party V_0 has read-only access to the k work tapes of $V_1, ..., V_k$ as well as a read-only input tape, a work tape

[***] according to a certain literature such a model is equivalent to the circuit model of quantum computers. We leave out all further details.

and a random tape. V_0 has two special terminal states *accept* and *reject*. The outcome of a computation involving V_0 is defined as its final state.

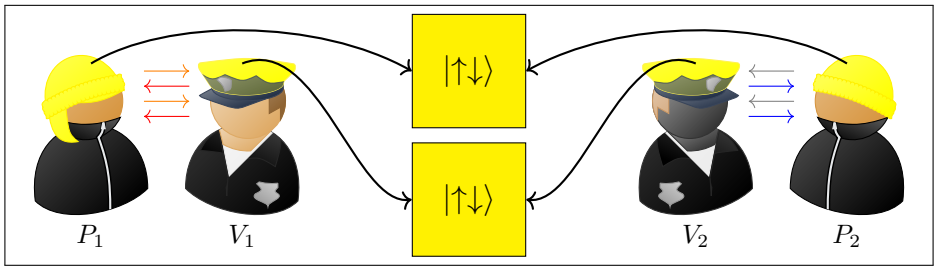

Fig. 2.3.1. Interrogation phase.

We call $(P_1, ..., P_k, V_0, V_1, ..., V_k)$ an *Entangled* k-prover interactive protocol. An *Entangled* multi-prover interactive proof for a language L is an *Entangled* k-prover interactive protocol $(P_1, ..., P_k, V_0, V_1, ..., V_k)$ satisfying the following extra conditions:

$$(Completeness)\forall x \in L, \exists \left|\Psi_P\right\rangle, \left|\Psi_V\right\rangle, \qquad \Pr(P_1, ..., P_k, \left|\Psi_P\right\rangle, V_0, V_1, ..., V_k, \left|\Psi_V\right\rangle)(x) = accept > \tfrac{2}{3},$$

$$(Soundness)\forall_{|\uparrow\downarrow\rangle}(P'_1, ..., P'_k), \forall \left|\Psi_P\right\rangle, \forall x \notin L, \exists \left|\Psi_V\right\rangle,$$
$$\Pr(P'_1, ..., P'_k, \left|\Psi_P\right\rangle, V_0, V_1, ..., V_k, \left|\Psi_V\right\rangle)(x) = accept < \tfrac{1}{3}.$$

2.4 No-Signalling Multi-Party Multi-Round Distributions

In order to study MIPs in the no-signalling setting we must first agree on a definition of *no-signalling multi-party multi-round distributions*. Let $P_1, ..., P_k$ be k parties, let x_j^i be the (IID) input of party $i \leq n$ in round $j \leq k$, and y_j^i be the related output. Let $\mathcal{P} = [n] = \{1, ..., n\}$ be the parties' indices and $\mathcal{R} = [k] = \{1, ..., k\}$ be the rounds' indices.

Definition A multiparty multi-round distribution $\Pr_{Y|X}(y_{\mathcal{R}}^{\mathcal{P}} \left| x_{\mathcal{R}}^{\mathcal{P}} \right.)$ is *no-signalling* if it satisfies the causal structure of a multi-round protocol and if every possible multiparty one-round sub-distribution is itself *no-signalling* according to the generally agreed definition. More technically, $\Pr_{Y|X}(y_{\mathcal{R}}^{\mathcal{P}} \left| x_{\mathcal{R}}^{\mathcal{P}} \right.)$ must satisfy the following three sets of conditions:

Causality Conditions (**C**): Answers from party $i \in \mathcal{I}$ cannot depend on questions he does not know yet.

$$(\mathbf{C})\ \forall \mathcal{I} \subseteq \mathcal{P}, \forall j_1, ..., j_n \leq k, \forall x_{\mathcal{R}}^{\mathcal{P}}, \hat{x}_{\mathcal{R}}^{\mathcal{P}}, \text{ s.t. } x_{[j_{\mathcal{I}}]}^{\mathcal{I}} = \hat{x}_{[j_{\mathcal{I}}]}^{\mathcal{I}}, x_{\mathcal{R}}^{\overline{\mathcal{I}}} = \hat{x}_{\mathcal{R}}^{\overline{\mathcal{I}}}, \qquad \Pr_{Y|X}(y_{j_{\mathcal{I}}}^{\mathcal{I}} \left| x_{\mathcal{R}}^{\mathcal{P}} \right.) = \Pr_{Y|X}(y_{j_{\mathcal{I}}}^{\mathcal{I}} \left| \hat{x}_{\mathcal{R}}^{\mathcal{P}} \right.)$$

Note: causality does not forbid answers to depend on questions of other people; (**NS**) and (**CNS**) will prevent that.

NS Conditions (**NS**): If all the questions to party $i \in \mathcal{I}$ were given together the resulting multiparty one-round distribution would be no-signalling.

$$(\mathbf{NS})\ \forall \mathcal{I} \subseteq \mathcal{P}, \forall x_{\mathcal{R}}^{\mathcal{P}}, \hat{x}_{\mathcal{R}}^{\mathcal{P}}, \text{ s.t. } x_{\mathcal{R}}^{\mathcal{I}} = \hat{x}_{\mathcal{R}}^{\mathcal{I}}, \qquad \Pr_{Y|X}(y_{\mathcal{R}}^{\mathcal{I}} \left| x_{\mathcal{R}}^{\mathcal{P}} \right.) = \Pr_{Y|X}(y_{\mathcal{R}}^{\mathcal{I}} \left| \hat{x}_{\mathcal{R}}^{\mathcal{P}} \right.)$$

Causal NS Conditions (**CNS**): If each party $i \in \mathcal{I}$ is about to answer $Y^i_{j_i}$ to query $x^i_{j_i}$ then $Y^i_{j_i}$ should not depend on any other $x^{i'}_{j_{i'}}$. In other words, given $x^i_{[j_i]}$, $y^i_{[j_i-1]}$, party i replies $Y^i_{j_i}$ independent from any $x^{i'}_{j_{i'}}$. Alternatively, just saying $\Pr_{Y|X}\Big(y^1_{j_1}...y^n_{j_n} \Big| x^1_{[j_1]}...x^n_{[j_n]}, y^1_{[j_1-1]}...y^n_{[j_n-1]}\Big)$ is no-signalling.

$$(\textbf{CNS})\ \forall \mathcal{I} \subseteq \mathcal{P}, \forall j_1, ..., j_n \leq k, \forall x^{\mathcal{P}}_{[j_{\mathcal{P}}]}, y^{\mathcal{P}}_{[j_{\mathcal{P}}-1]}, \hat{x}^{\overline{\mathcal{I}}}_{j_{\overline{\mathcal{I}}}},$$
$$\Pr_{Y|X}\Big(y^{\mathcal{I}}_{j_{\mathcal{I}}} \Big| x^{\mathcal{I}}_{j_{\mathcal{I}}}, x^{\overline{\mathcal{I}}}_{j_{\overline{\mathcal{I}}}}, x^{\mathcal{P}}_{[j_{\mathcal{P}}-1]}, y^{\mathcal{P}}_{[j_{\mathcal{P}}-1]}\Big) = \Pr_{Y|X}\Big(y^{\mathcal{I}}_{j_{\mathcal{I}}} \Big| x^{\mathcal{I}}_{j_{\mathcal{I}}}, \hat{x}^{\overline{\mathcal{I}}}_{j_{\overline{\mathcal{I}}}}, x^{\mathcal{P}}_{[j_{\mathcal{P}}-1]}, y^{\mathcal{P}}_{[j_{\mathcal{P}}-1]}\Big)$$

2.5 No-Signalling Multi-Party Multi-Round Boxes

A no-signalling multiparty multi-round box is a $2k+2$-tape Turing machine, implementing a no-signalling multiparty multi-round distribution. The TM has access to k read-only communication tapes, k write-only communication tapes, one read/write work tape, and a read-only random tape. Each pair of read/write tapes is used to communicate with a specific Turing machine that provides the consecutive inputs and receives the consecutive outputs of a party's no-signalling multi-round Distribution. No-signalling boxes respond immediately when queried by a party because their output distribution is independent from all other inputs.

To check that a $2k+2$-tape Turing machine implements a no-signalling multiparty multi-round distribution the conditions from the previous subsection can be verified one by one. At this point we do not know of an efficient algorithm to check this property, and leave it as an open problem finding one.

2.6 No-Signalling Multi-Prover Interactive Proofs

Definition (NOSIG $\left(\not\equiv\right)$): Let $M_1, ..., M_k$ be k Turing machines. All machines have a read-only input tape, a work tape and a random tape. In addition, $M_1, ..., M_k$ share a (read/write) no-signalling box where an input may be written and a corresponding output be read. Every M_i has one write-only communication tape on which it writes messages for another machine W_i. Every M_i has one read-only communication tape on which it reads messages from another machine W_i. We call $(M_1, ..., M_k)$ a *no-signalling* k-party interactive TM.

Definition (MIP$^{\not\equiv}$): Let $(P_1, ..., P_k)$ be a *no-signalling* k-party interactive TM which is computationally unbounded and $(V_1, ..., V_k)$ be a *no-signalling* k-party interactive Turing machine which is probabilistic polynomial time bounded. Each P_i can only exchange messages with the corresponding V_i and vice versa. A special party V_0 has read-only access to the k work tapes of $V_1, ..., V_k$ as well as a read-only input tape, a work tape and a random tape. V_0 has two special terminal states *accept* and *reject*. The outcome of a computation involving V_0 is defined as its final state.

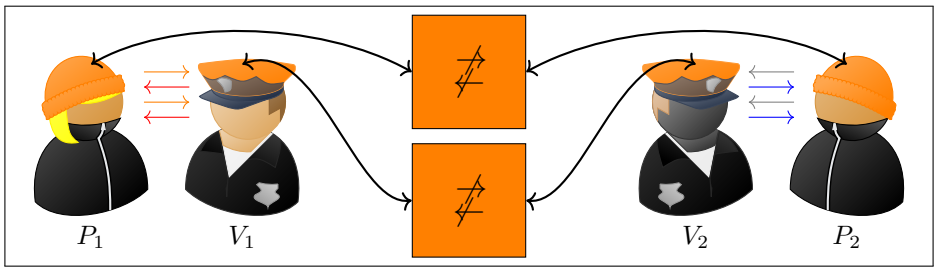

Fig. 2.6.1. Interrogation phase.

We call $(P_1, ..., P_k, V_0, V_1, ..., V_k)$ a *no-signalling* k-prover interactive protocol. A no-signalling multi-prover interactive proof for a language L is a no-signalling k-prover interactive protocol $(P_1, ..., P_k, V_0, V_1, ..., V_k)$ satisfying the completeness and NoSIG-soundness conditions:

(Soundness) $\forall_{\not\equiv}(P'_1,...,P'_k), \forall x \notin L, \Pr[(P'_1,...,P'_k,V_0,V_1,...,V_k)(x) = accept] < \frac{1}{3}$.

The requirement that $(V_1,...,V_k)$ be No-Signalling implies that $(P_1,...,P_k)$ cannot go beyond its own locality via powers provided (willingly or inadvertently) by V. Each pair of parties $\langle P_i, V_i\rangle$ is as no-signalling as P_i by itself (see Fig. 2.6.1). As noted before, most MIPs found in the literature are actually local-verifier MIPs (see Fig. 2.6.2).

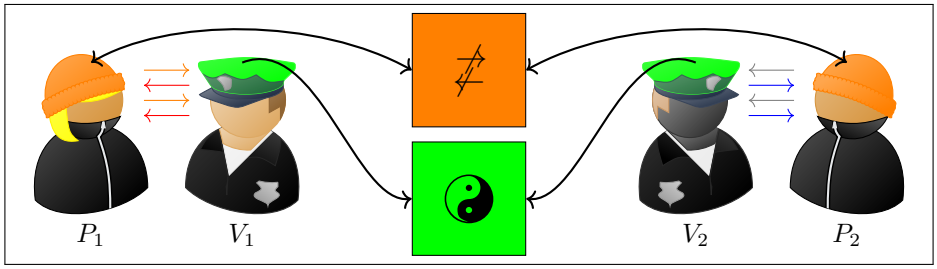

Fig. 2.6.2. Interrogation phase.

The other elements of the locality hierarchy $|\uparrow\mathbb{R}\downarrow\rangle$(Entanglement with real coefficients only), $\overrightarrow{\not\rightleftarrows}$, $\overrightarrow{\rightleftarrows}$ and corresponding complexity classes $\mathrm{MIP}^{|\uparrow\mathbb{R}\downarrow\rangle}$, $\mathrm{MIP}^{\overrightarrow{\not\rightleftarrows}}$, and $\mathrm{MIP}^{\overrightarrow{\rightleftarrows}}$ are defined analogously but are left out of this extended abstract (but fairly unsurprisingly we get $\mathrm{MIP}^{\overrightarrow{\rightleftarrows}}$=IP).

2.7 Locality-Explicit Form

The above definitions do not account for every possible multi-prover interaction setup – they are not meant to. Let us define that a MIP is in *locality-explicit form* if it satisfies the above definitions. By writing a protocol in this form, the signalling powers of every party is explicitly apparent. Thus, we can prove completeness and soundness without having to worry about changes in the parties' localities; one simply needs to rewrite a particular MIP in the appropriate locality-explicit form. If this turns out to be impossible, then we cannot assess the soundness of the resulting protocol, in general. In the case of zero-knowledge MIPs, we need to consider arbitrary (possibly signalling) verifiers, thus we do not limit ourselves to protocols in locality-explicit form. In section 4, we will discuss how it makes sense to have multiple simulators which have their corresponding localities.

2.8 Special Relativity

The work of BCMS [25] and later Kent [26] introduced the idea that locality (or no-signalling) can be enforced by separating the parties involved with sufficient distance. A number of papers along these lines have later explored the power of local and entangled provers in such a scenario and bit commitment was proven possible under these combined assumptions. Similarly, relativistic MIPs with entangled provers have recently been proven secure for languages in NP [27]. The relativistic model is analogous to our definition of MIPs in the sense that the verifier of these proofs must be broken down as separate entities, one verifier per prover, so that each prover can be interrogated locally by its corresponding verifier. At the end of the whole proof, the verifiers get together and then decide whether to accept or reject the membership of a certain string. The locality of the provers and verifiers must be equal.

In actuality, the only guarantee we need from special relativity is that no information travels faster-than-light. The correct cryptographic model we should be using is the no-signalling one. We present in section 5 a protocol for EXP which is zero-knowledge under the no faster-than-light assumption.

We reiterate here that any zero-knowledge MIP which requires the verifier to courier an authenticated question from prover to prover cannot be implemented under the special relativity assumption, for obvious reasons.

3 ZKMIP$^{\textcircled{\scriptsize c}}$ = NEXP

Let us put what we have constructed above together into something tangible. We will describe here a MIP with the following properties (in addition to completeness and soundness):

- Local, as defined in the previous section.
- Zero-knowledge, in the stronger sense discussed in the next section.

The **NEXP**-complete protocol from [1] is local in the sense that the interaction with P_2 involves only a single question, sampled from the pool of questions the verifier has asked P_1. This can be "localized" by a pair of verifiers by asking them to share a random tape before the protocol begins. The phases which involves only P_1 can be made zero-knowledge by executing it in committed form, as in the single-prover case. The phase in which the verifier interacts with P_2 can also be executed in committed form if only the phases were independent from each other. The problem is that during this P_2 phase, the verifier *must* ask a question that it has asked P_1 in order to assure zero-knowledge.

This is addressed in [4] and [8] by asking the first prover to compute the verifier's question in committed form, add a message authentication tag using a secret key that the provers shared and finally ask the verifier to courier the message to P_2, who then checks whether the authentication is valid before executing its part of the protocol. This solution is problematic because it cannot be done by local verifiers. In a word, this protocol is *not local*.

We present a zero-knowledge **MIP** protocol which addresses this problem. Our solution essentially asks the provers to encrypt an answer with a key that is based on the verifier's question. Therefore, if the same question is asked to both provers, the verifiers receive the same answers; if two different questions are asked, then the verifiers receive two answers which, from their points of view, are independently and uniformly random. This way, there is no need for a prover to authenticate questions, nor is there need for questions to be couriered between provers.

We describe the protocol below, and prove its completeness, soundness, zero-knowledge and "localness".

3.1 Protocol

Construction 31 *A statistically binding, perfectly concealing BC protocol. (Entangled Secure)*

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_0 samples two k-bit strings z_1, z_2 independently and uniformly, and provides them to V_1, 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: $d = (w_1 \times z_2) \oplus w_2$.

Unveil 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 $d \oplus w_2 \neq z_1 \times z_2$ and **accepts** b otherwise.

□

Construction 32 *Zero-Knowledge Protocol for* $\mathsf{MIP}^{\textcircled{\scriptsize ?}}$

Let x, an instance of oracle-3-SAT, be the common input, let $|x| = k$, and let Λ be the verifier's program in the BFL protocol.

1. (pre-computation)
   (a) V_0 samples two k-bit strings z_1, z_2 independently and uniformly, and provides them to V_1, V_2.
   (b) V_0 selects k random K-bit strings $r_1, ..., r_k$ and evaluates the circuit of Λ on input r_i, resulting in questions $Q_1, ..., Q_k$, and provides them to V_1, V_2
   (c) V_0 randomly chooses $1 \leq i \leq k+3$, the index of an oracle query that will be made to both P_1 and P_2. V_0 provides i to V_1, V_2.
   (d) V_1 sends z_1 to P_1 and V_2 sends z_2 to P_2.
   (e) P_1 (via P_2) commits $\boxed{\gamma}$ to V_1 (via V_2).
2. (multilinearity test) Let k be the number of oracle queries in this phase.
   For $1 \leq i \leq k$:
   (a) V_1 sends Q_i to P_1, the question that would be asked by Λ with coins r_i.
   (b) P_1 commits his answer as $\boxed{A(Q_i)}$.
   (c) After committing all of his answers, P_1 and V_1 evaluate a circuit description of Λ in committed form with inputs $\boxed{A(Q_1)}, \ldots, \boxed{A(Q_k)}$. P_1 unveils the circuit's output. If it rejects, V_1 rejects.
3. (sumcheck with oracle):
   – Let $g(z, Q_{k+1}, Q_{k+2}, Q_{k+3}, A(Q_{k+1}), A(Q_{k+2}), A(Q_{k+3}))$
   be the arithmetization obtained by section 3.2 of [1], let z be a string of length r and $Q_{k+1}, Q_{k+2}, Q_{k+3}$ be strings of length s, as appropriate. V_1 and P_1 execute the protocol of section 3.3 of [1] in committed form, using coins selected by P_1 whenever randomness is required. At the end of this phase, P_1 shows that the committed final value is equal to

$$g\Big(z, Q_{k+1}, Q_{k+2}, Q_{k+3}, \boxed{A(Q_{k+1})}, \boxed{A(Q_{k+2})}, \boxed{A(Q_{k+3})}\Big),$$

   an evaluation in committed form of g using the committed random bits that were used during the protocol's loop. If this fails, V_1 rejects.
4. (non-adaptiveness test):
   (a) V_1 sends i to P_1 who responds with $\Omega_1 = A(Q_i) \oplus H_\gamma(Q_i)$, and proves to V_1 that $\Omega_1 = \boxed{A(Q_i)} \oplus H_{\boxed{\gamma}}(Q_i)$ was computed correctly.
   (b) V_2 sends Q_i to P_2.
   (c) P_2 responds to V_2 with $\Omega_2 = A(Q_i) \oplus H_\gamma(Q_i)$.
   (d) V_0 accepts if and only if all of the following conditions are met:
      – $\Omega_1 = \Omega_2$
      – All commitments which have been opened are valid.
      – V_1 did not reject in the two previous cases

3.2 Proof Sketches of Security

Localness

We have defined "locality classes" and "locality" in section 2 to mean the different classes of no-signalling correlations and interactions. Therefore, we will use the word "localness" to denote that a protocol has the property of being local.

In both the commitment protocol (construction 31) and the ZKMIP protocol (construction 32), the provers do not interact with each other, making them trivially local as per definition 2.2; V_0 does not

interact with the remaining verifiers *once the protocol begins*; V_1 and V_2 do not interact with the provers *once the protocol ends*; therefore, since neither the provers nor the verifiers share any additional resources, the verifiers are local.

Completeness

Completeness follows from the completeness of the underlying protocol [1], and the fact that the commitment protocol (construction 31) is well-defined for honest provers (who will never send a commitment that they cannot open).

Soundness

The scaffolding for our soundness proof will be the following. Suppose that there is a trusted third party handling commitments that are perfectly binding, and that this third party sends the receiver a unique identifying receipt for every commitment. If we were to use this commitment to execute a protocol in committed form (as we do in construction 32), then to each committed transcript of an execution there would be a unique real transcript (of the BFL protocol in our case). In this case, a committed transcript shows that it ends in *accept* if and only if the associated real transcript, which one would obtain by running the base protocol, also ends in *accept*. Thus, if the base protocol has a soundness error of, say, 1/3, then so does this committed version.

We can remove the scaffolding by replacing the trusted third-party with construction 31. This does not change the argument, except with exponentially small change to the soundness error. To begin with, local probabilistic provers are equivalent to local deterministic provers. This is because the success probability α of randomized provers 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. Therefore, we only need to consider local deterministic provers.

Let us analyze our commitment protocol. Since P_1 is deterministic, we may unambiguously consider what happens if we were to "rewind" it. Suppose that at some point P_1 opens a particular commitment c to 0, and that by rewinding P_1 and have V_1 make different choices, P_1 then opens the *same* commitment c to 1 (an attempt to break binding). Because of localness, P_1's behaviour 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 opening of c in both ways (recall that our commitment is statistically binding).

In the worst case, for every commitment there exists a sequence of interaction 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

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

Zero-Knowledge

Our use of the multi-verifier setup as a guarantee of localness leads us to an unforeseen problem. The participants of our protocol are two pairs of prover-verifiers and a "decision" verifier. In attempting to prove zero-knowledge in this new picture of MIPs, we must decide on the purpose of the simulator.

However, a *single* simulator makes zero-knowledge trivial. One might think that, since a zero-knowledge protocol must deal with arbitrarily malicious verifiers, we should give the simulator as much power as possible in order to "simulate" the said behavior. But that would be going the wrong way: Instead, the question should be, "How *little* extra power does the simulator need in order to do its job?"

We begin with the simple observation that *if the provers can signal, then they can break the commitment of construction 31.* This is the basis of our simulation strategy, because if the provers can break the binding condition of commitments, then committed circuit evaluations become meaningless; put another way, if the simulator knows everything that both verifiers tell their respective provers, then it can make the outcome of any committed circuit to be whatever it wants. In particular, since our protocol is perfectly complete, it will simply simulate the protocol as being "accepted" by V_0, regardless of whether $x \in L$ or $x \notin L$. We leave the details to the reader.

Now consider a modified simulation setup. Suppose that instead of a single simulator, we have two simulators S_1 and S_2, interacting with V_1 and V_2 respectively. If the simulators can signal, it is plain that the above argument applies and that they can simulate anything, given that they can break commitment. However, *the simulators do not need to signal in order to break the commitment.* This is a well-known result [7] regarding non-local games: CHSH boxes can break our commitment scheme without giving the parties the ability to signal. Therefore, our simulation strategy of "breaking the commitment" can work with a pair of simulators augmented with a polynomial number of these CHSH boxes.

We would like to define a new property (or attribute) concerning ZKMIPs. We define the *zero-knowledge locality* of a ZKMIP to be the lowest locality hierarchy class in which there exists simulators for that ZKMIP. If the exact locality is not known, then we can say that it is upper-bounded some best-known locality class.

The zero-knowledge locality of construction 32 is upper-bounded by the *no-signalling* locality class. If we were to define more fine-grained locality classes, then we would say it is exactly the locality of local provers augmented by the CHSH task; however, we will stay with the few broad classes we have defined previously for now. Whether it is possible to simulate construction 32 with *local* simulators is an open question. We explore this new notion of zero-knowledge in the next section.

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

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

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 Figure 4.0.1) could have been equally created with the same distribution by a simulator. Nevertheless, a stronger flavour of zero-knowledge is achievable if the simulator is not invoking its full signalling power whenever the verifier does not use such power. We define

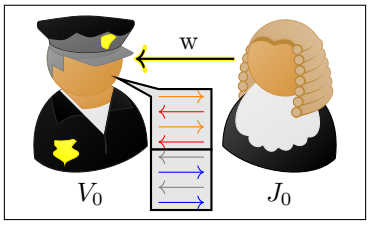

Fig. 4.0.1. (Interac/Simula)tion-Distinction phase.

H-zero-knowledge (ZK^H) for the locality hierarchy member $H \subseteq SIG$ by the extra requirement that for all locality level H' with $H \subseteq H'$ we have $V' \in H' \rightarrow S \in H'$ as well. In words, it means that for all locality levels starting with H, the simulators do not need more signalling power than the verifiers V'. The ultimate (strongest) notion of "local ZK" being $\text{ZK}^{\textcircled{\scriptsize 0}}$ because at all levels V' is simulated by a simulator with no extra signalling power, whereas at the opposite end of the spectrum $\text{ZK}^{\rightleftarrows}$ 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, 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. 4.0.2).

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 locality class H is enough to break the binding property of the commitments. For instance, the result of section 3.1 is really $\mathrm{ZK}^{\not\rightleftarrows}\mathrm{MIP}^{\text{◕}} = \mathrm{NEXP}$, whereas the same protocol using the bit commitment of construction 51 proves only $\mathrm{ZK}^{\rightleftarrows}\mathrm{MIP}^{\text{◕}} = \mathrm{NEXP}$. Using the bit commitment scheme based on the magic square game of [9] we can also obtain $\mathrm{ZK}^{|\uparrow\downarrow\rangle}\mathrm{MIP}^{\text{◕}} = \mathrm{NEXP}$.

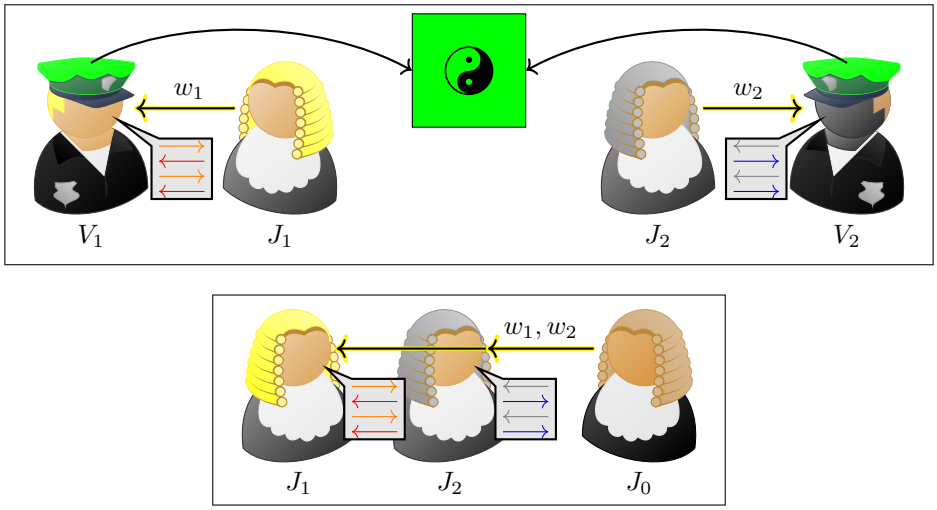

Fig. 4.0.2. Interrogation/Simulation phase (top) followed by decision phase (bottom).

Some interesting questions resulting from this definition is whether any higher class such as $\mathrm{ZK}^{\text{◕}}\mathrm{MIP}^{\text{◕}}$ or $\mathrm{ZK}^{\not\rightleftarrows}\mathrm{MIP}^{\not\rightleftarrows}$ contains more than the natural examples such as GRAPH ISOMORPHISM or CODE EQUIVALENCE already found in the most natural class $\mathrm{ZK}^{\rightleftarrows}\mathrm{MIP}^{\rightleftarrows} = \mathrm{ZKIP}$.

5 Protocols for NEXP $\subseteq$ ZKMIP$^{|\uparrow\downarrow\rangle}$ & ZKMIP$^{\not\rightleftarrows}$ = EXP.

We believe the tools we have developed in this paper, together with a few extra constructions allow us to extend our results of Section 3.1 to the entangled and no-signalling scenarios: our solutions involve the $\mathrm{MIP}^{|\uparrow\downarrow\rangle}$ protocol for NEXP of [2], the $\mathrm{MIP}^{\not\rightleftarrows}$ protocol for EXP of [3], a new three-prover commitment scheme resistant to no-signalling attacks inspired from that of [24], and constructions similar to protocol 32.

5.1 Proposed Protocols

The two protocols have essentially the same structure as construction 32 above: one prover runs the protocol of [2] or [3] by himself acting as all the provers, proving in zero-knowledge that all the answers are computed correctly; the verifier can get the other provers to corroborate what the main prover answers for them. This uses a bit-commitment scheme from the main prover to the verifier that is secure against entangled or no-signalling provers. The CHSH bit commitment (construction 31) used in the local scenario is also secure in the entangled setting as previously proved in [30]. In the no-signalling case the following new bit-commitment protocol (based on the extended CHSH 3-player game as in [24]) is used instead (see next page).

The other differences between the local construction and the entangled/no-signalling cases is the number of provers: four provers in the protocols of Ito-Vidick [2], and polynomially many provers in the protocol of Kalai-Raz-Rothblum [3]. To obtain zero-knowledgeness of the resulting protocols (via the one-time-pad encryption of the oracle queries), the hash function used in the entangled version of the protocol must be strongly-universal-5 and strongly-universal-poly in the no-signalling version. But these are minor changes. These constructions automatically inherit localness, completeness and zero-knowledgeness of the local protocol. However proving their soundness remains a major challenge that we leave as open questions.

Construction 51 *A statistically binding, perfectly concealing BC protocol (no-signalling secure)*

All parties agree on a security parameter 1^k.
P_1, P_2 and P_3 partition their private random tape into three k-bit strings w_1, w_2, w_3.

Pre-computation phase:

- V_0 samples three k-bit strings z_1, z_2, z_3 independently and uniformly, and provides them to V_1, V_2, V_3.
- V_1 sends z_1 to P_1, V_2 sends z_2 to P_2 and V_3 sends z_3 to P_3.

Commit phase:

- P_1 commits b to V_1 as $\boxed{b} = b \times z_1 \oplus w_1$.
- P_2 sends V_2: $d_2 = w_1 \times z_2 \oplus w_2$.
- P_3 sends V_3: $d_3 = w_1 \times z_3 \oplus w_3$.

Unveil phase:

- P_1 sends w_1, w_2, w_3 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 $\boxed{b} \oplus w_1 \notin \{0, z_1\}$ or if $d_2 \neq w_1 \times z_2 \oplus w_2$, or if $d_3 \neq w_1 \times z_3 \oplus w_3$
  but **accepts** b otherwise.

□

6 Discussion and Open Questions

Considering the many issues described in the previous sections, we believe that there is a need to rethink MIPs/ZKIPs with respect to locality, particularly this new notion of a protocol's zero-knowledge locality. Non-local interaction in general is still poorly understood, and a good direction of research would be to find an operational description of non-locality (in the same sense as the postulates of quantum mechanics), as opposed to the information-theoretical description which we have here.

Acknowledgements

We would like to thank Gilles Brassard, André Chailloux, Serge Fehr, Max Fillinger, Sophie Laplante, Justin Li, Anthony Leverrier, Samuel Ranellucci, Louis Salvail, and Christian Schaffner for various fruitful discussions about early versons of this work. Finally, we are grateful to Raphael C.-W. Phan and Moti Yung for inviting us to publish a preliminary version of this work as an *Insight Paper* at MyCrypt 2016.

References

1. 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.

2. 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.
3. 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.
4. 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.
5. 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.
6. R. Cleve, P. Hoyer, B. Toner, and J. Watrous, "Consequences and limits of nonlocal strategies," in *Proceedings of the 19th IEEE Annual Conference on Computational Complexity*, CCC '04, (Washington, DC, USA), pp. 236–249, IEEE Computer Society, 2004.
7. C. Crépeau, L. Salvail, J.-R. Simard, and A. Tapp, *Two Provers in Isolation*, pp. 407–430. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011.
8. J. Kilian, *Uses of randomness in algorithms and protocols*. MIT Press, 1990.
9. C. Crépeau and N. Yang, *Multi-prover Interactive Proofs: Unsound Foundations*, pp. 485–493. Cham: Springer International Publishing, 2017.
10. T. Ito, "Polynomial-space approximation of no-signaling provers," in *Proceedings of the 37th International Colloquium Conference on Automata, Languages and Programming*, ICALP'10, (Berlin, Heidelberg), pp. 140–151, Springer-Verlag, 2010.
11. J. S. Bell, "On the Einstein-Podolsky-Rosen paradox," *Physics*, vol. 1, pp. 195–200, 1964.
12. J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, "Proposed experiment to test local hidden-variable theories," *Phys. Rev. Lett.*, vol. 23, pp. 880–884, Oct 1969.
13. G. Brassard, H. Buhrman, N. Linden, A. A. Méthot, A. Tapp, and F. Unger, "Limit on nonlocality in any world in which communication complexity is not trivial," *Phys. Rev. Lett.*, vol. 96, p. 250401, Jun 2006.
14. L. A. Khalfin and B. S. Tsirelson, "Quantum and quasi-classical analogs of bell inequalities," in *Symposium on the Foundations of Modern Physics*, pp. 441–460, 1985.
15. P. Rastall, "Locality, bell's theorem, and quantum mechanics," in *Found. of Physics*, vol. 15, pp. 963–972, 1985.
16. S. Popescu and D. Rohrlich, "Quantum nonlocality as an axiom," *Foundations of Physics*, vol. 24, no. 3, pp. 379–385, 1994.
17. B. Toner, "Monogamy of non-local quantum correlations," *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 465, no. 2101, pp. 59–69, 2009.
18. 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.
19. 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*, vol. 39, no. 36, p. 11283, 2006.
20. J. Kempe, H. Kobayashi, K. Matsumoto, B. Toner, and T. Vidick, "Entangled games are hard to approximate," in *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '08, (Washington, DC, USA), pp. 447–456, IEEE Computer Society, 2008.
21. T. Ito, H. Kobayashi, and K. Matsumoto, "Oracularization and two-prover one-round interactive proofs against nonlocal strategies," in *2009 24th Annual IEEE Conference on Comp. Complexity*, pp. 217–228, July 2009.
22. T. Holenstein, "Parallel repetition: Simplification and the no-signaling case," *Theory of Computing*, vol. 5, no. 8, pp. 141–172, 2009.
23. T. Ito, "Polynomial-space approximation of no-signaling provers," in *Proceedings of the 37th International Colloquium Conference on Automata, Languages and Programming*, ICALP'10, (Berlin, Heidelberg), pp. 140–151, Springer-Verlag, 2010.
24. S. Fehr and M. Fillinger, *Multi-prover Commitments Against Non-signaling Attacks*, pp. 403–421. Berlin, Heidelberg: Springer Berlin Heidelberg, 2015.
25. G. Brassard, C. Crépeau, D. Mayers, and L. Salvail, "Defeating classical bit commitments with a quantum computer." arXiv:quant-ph/9806031, June 1998.
26. A. Kent, "Unconditionally secure bit commitment," *Phys. Rev. Lett.*, vol. 83, pp. 1447–1450, Aug 1999.
27. A. Chailloux and A. Leverrier, *Relativistic (or 2-Prover 1-Round) Zero-Knowledge Protocol for NP Secure Against Quantum Adversaries*, pp. 369–396. Cham: Springer International Publishing, 2017.
28. C. Crépeau, J. v. d. Graaf, and A. Tapp, "Committed oblivious transfer and private multi-party computation," in *Proceedings of the 15th Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '95, (London, UK, UK), pp. 110–123, Springer-Verlag, 1995.
29. N. Yang, "Zero-knowledge multi-prover interactive proofs," Master's thesis, Concordia University, 2013.
30. 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.