2018 IEEE 59th Annual Symposium on Foundations of Computer Science

Spatial Isolation Implies Zero Knowledge Even in a Quantum World

Alessandro Chiesa
UC Berkeley
Berkeley, CA
alexch@berkeley.edu

Michael A. Forbes
University of Illinois at Urbana–Champaign
miforbes@illinois.edu

Tom Gur
UC Berkeley
Berkeley, CA
tom.gur@berkeley.edu

Nicholas Spooner
UC Berkeley
Berkeley, CA
nick.spooner@berkeley.edu

***Abstract*—Zero knowledge plays a central role in cryptography and complexity. The seminal work of Ben-Or et al. (STOC 1988) shows that zero knowledge can be achieved unconditionally for any language in $\mathbf{NEXP}$, as long as one is willing to make a suitable *physical assumption*: if the provers are spatially isolated, then they can be assumed to be playing independent strategies.**

Quantum mechanics, however, tells us that this assumption is unrealistic, because spatially-isolated provers could share a quantum entangled state and realize a non-local correlated strategy. The $\mathbf{MIP}^*$ model captures this setting.

In this work we study the following question: *does spatial isolation still suffice to unconditionally achieve zero knowledge even in the presence of quantum entanglement?*

We answer this question in the affirmative: we prove that every language in $\mathbf{NEXP}$ has a 2-prover *zero knowledge* interactive proof that is sound against entangled provers; that is, $\mathbf{NEXP} \subseteq \mathbf{ZK\text{-}MIP}^*$.

Our proof consists of constructing a zero knowledge interactive PCP with a strong algebraic structure, and then lifting it to the $\mathbf{MIP}^*$ model. This lifting relies on a new framework that builds on recent advances in low-degree testing against entangled strategies, and clearly separates classical and quantum tools.

Our main technical contribution is the development of new algebraic techniques for obtaining unconditional zero knowledge; this includes a zero knowledge variant of the celebrated sumcheck protocol, a key building block in many probabilistic proof systems. A core component of our sumcheck protocol is a new algebraic commitment scheme, whose analysis relies on algebraic complexity theory.

***Keywords*-zero knowledge; multi-prover interactive proofs; quantum entangled strategies; interactive PCPs; sumcheck protocol; algebraic complexity;**

I. Introduction

Zero knowledge, the ability to demonstrate the validity of a claim without revealing any information about it, is a central notion in cryptography and complexity that has received much attention in the last few decades. Introduced in the seminal work of Goldwasser, Micali, and Rackoff [1], zero knowledge was first demonstrated in the model of interactive proofs, in which a resource-unbounded prover interacts with a probabilistic polynomial-time verifier to the end of convincing it of the validity of a statement.

Goldreich, Micali, and Wigderson [2] showed that every language in $\mathbf{NP}$ has a *computational* zero knowledge interactive proof, under the cryptographic assumption that (non-uniform) one-way functions exist. Ostrovsky and Wigderson [3] proved that this assumption is necessary.

Unfortunately, the stronger notion of *statistical* zero knowledge interactive proofs, where both soundness and zero knowledge hold unconditionally, is limited. For example, if $\mathbf{NP}$ had such proofs then the polynomial hierarchy would collapse to its second level [4, 5, 6].

The celebrated work of Ben-Or et al. [7] demonstrated that the situation is markedly different when the verifier interacts with *multiple* provers, in a *classical world* where by spatially isolating the provers we ensure that they are playing independent strategies — this is the model of multi-prover interactive proofs (MIPs). They proved that every language having an MIP (i.e., every language in $\mathbf{NEXP}$ [8]) also has a *perfect* zero knowledge MIP. This result tells us that *spatial isolation implies zero knowledge*.

In light of quantum mechanics, however, we know that spatial isolation *does not* imply independence, because the provers could share an entangled state and realize a strategy that is beyond that of independently acting provers. For example, it is possible for entangled provers to win a game (e.g., the magic square game) with probability 1, whereas independent provers can only win with probability at most $8/9$ [9].

Non-local correlations arising from local measurements on entangled particles play a fundamental role in physics, and their study goes back at least to Bell's work on the Einstein–Podolsky–Rosen paradox [10]. Recent years have seen a surge of interest in MIPs with *entangled provers*, which correspond to the setting in which multiple non-communicating provers share an entangled state and wish to convince a classical verifier of some statement. This notion is captured by $\mathbf{MIP}^*$ protocols, introduced by Cleve et al. [9]. A priori it is unclear whether these systems should be less powerful than standard MIPs, because of the richer class of *malicious* prover strategies, or more powerful, because of the richer class of *honest* prover strategies.

Investigating proof systems with entangled adversaries

2575-8454/18/$31.00 ©2018 IEEE
DOI 10.1109/FOCS.2018.00077

not only sharpens our understanding of entanglement as a computational resource, but also contributes insights to hardness of approximation and cryptography in a post-quantum world. However, while the last three decades saw the development of powerful ideas and tools for designing and analyzing proof systems with classical adversaries, despite much effort, there are only a handful of tools available for dealing with quantum entangled adversaries, and many fundamental questions remain open.

MIP* protocols were studied in a long line of work, culminating in a breakthrough result of Ito and Vidick [11], who in a technical tour-de-force showed that $\mathbf{NEXP} \subseteq \mathbf{MIP}^*$;[1] this result was further improved in [12, 13]. However, it is unknown whether these MIP protocols can achieve zero knowledge, which is the original motivation behind the classical MIP model. In sum, in this paper we pose the following question:

To what extent does spatial isolation imply unconditional zero knowledge in a quantum world?

A. Our results

Our main result is a strong positive answer to the foregoing question, namely, we show that the $\mathbf{NEXP} \subseteq \mathbf{MIP}^*$ result of Ito and Vidick [11] continues to hold even when we require zero knowledge.

Theorem I.1. *Every language in* $\mathbf{NEXP}$ *has a perfect zero knowledge* 2*-prover MIP**. *In more detail,*

$$\mathbf{NEXP} \subseteq \mathbf{PZK\text{-}MIP}^* \left[\begin{array}{rl} \textit{number of provers:} & 2 \\ \textit{round complexity:} & \mathrm{poly}(n) \\ \textit{communication complexity:} & \mathrm{poly}(n) \\ \textit{soundness error:} & 1/2 \end{array} \right].$$

We stress that the MIP* protocols of Theorem I.1 enjoy both unconditional soundness against entangled provers as well as unconditional (perfect) zero knowledge against *any* (possibly malicious) verifier.

B. Other notions of quantum zero knowledge

To the best of our knowledge, this work is the first to study the notion of zero knowledge with entangled provers, as captured by the MIP* model. Nevertheless, zero knowledge has been studied in other settings in the quantum information and computation literature; we now briefly recall these.

Watrous [14] introduced *honest-verifier* zero knowledge for quantum interactive proofs (interactive proofs in which the prover and verifier are quantum machines), and studied the resulting complexity class $\mathbf{QSZK_{HV}}$. Kobayashi [15] studied a non-interactive variant of this notion. Damgård, Fehr, and Salvail [16] achieve zero knowledge for $\mathbf{NP}$ against malicious quantum verifiers, but only via *arguments* (i.e., computationally sound proofs) in the common reference string model. Subsequently, Watrous [17] constructed quantum interactive proofs that remain zero knowledge against malicious quantum verifiers.

Zero knowledge for quantum interactive proofs has since then remained an active area of research, and several aspects and variants of it were studied in recent works, including the power of public-coin interaction [18], quantum proofs of knowledge [19], zero knowledge in the quantum random oracle model [20], zero knowledge proof systems for $\mathbf{QMA}$ [21], and oracle separations for quantum statistical zero knowledge [22].

All the above works consider protocols between a *single* quantum prover and a quantum verifier. In particular, they do not study entanglement as a shared resource between two (or more) provers.

In contrast, the MIP* protocols that we study differ from the protocols above in two main aspects: (1) our proof systems have multiple spatially-isolated provers that share an entangled state, and (2) it suffices that the honest verifier is a *classical* machine. Indeed, we show that, analogously to the classical setting, MIP* protocols can achieve *unconditional* zero knowledge for a much larger complexity class (namely, $\mathbf{NEXP}$) than possible for QSZK protocols (since $\mathbf{QSZK} \subseteq \mathbf{QIP} = \mathbf{PSPACE}$).

II. Techniques

We begin by discussing the challenge that arises when trying to prove that $\mathbf{NEXP} \subseteq \mathbf{PZK\text{-}MIP}^*$, by outlining a natural approach to obtaining zero knowledge MIP* protocols, and considering why it fails.

A. The challenge

We know that every language in $\mathbf{NEXP}$ has a (perfect) zero knowledge MIP protocol, namely, that $\mathbf{NEXP} \subseteq \mathbf{PZK\text{-}MIP}$ [7]. We also know that every language in $\mathbf{NEXP}$ has an MIP* protocol, namely, that $\mathbf{NEXP} \subseteq \mathbf{MIP}^*$ [11]. Is it then not possible to simply combine these two facts and deduce that every language in $\mathbf{NEXP}$ has a (perfect) zero knowledge MIP*?

The challenge is that the standard techniques used to construct zero knowledge MIP protocols do not seem compatible with those used to construct MIP* protocols for large classes.[2] In fact, the former are precisely the type of techniques that prove to be very limited for obtaining soundness against entangled provers.

In more detail, while constructions of MIP (and PCP) protocols typically capitalize on an *algebraic* structure, known constructions of *zero knowledge* MIPs are of a *combinatorial* nature. For example, the zero knowledge

[1] While this is the popular statement of the result, [11] show a stronger result, namely, that $\mathbf{NEXP}$ is exactly the class of languages decided by MIPs *sound against entangled provers*. Their honest provers are classical, and soundness holds also against entangled provers. This is also the case in our protocols. It remains unknown whether entanglement grants provers additional power: there is *no* known reasonable upper bound on $\mathbf{MIP}^*$.

[2] For example, Crépeau et al. [23] showed that the commitment scheme in [7] is *not* sound against entangled adversaries.

MIP in [7] is based on a multi-prover information-theoretic commitment scheme, which can be thought of as a CHSH-like game. The zero knowledge MIP in [24] is obtained via the standard transformation from zero knowledge PCPs, which is a form of consistency game. Unfortunately, these types of constructions do not appear resistant to entangled provers, nor is it clear how one can modify them to obtain this resistance without leveraging some algebraic structure.

Indeed, initial attempts to show that $\mathbf{NEXP} \subseteq \mathbf{MIP}^*$ (e.g., [25, 26, 27]) tried to apply some black box transformation to an arbitrarily structured (classical) MIP protocol to force the provers to behave as if they are not entangled, and then appeal to standard MIP soundness. These works were only able to obtain limited protocols (e.g., with very large soundness error).

In their breakthrough paper, Ito and Vidick [11] overcame this hurdle and showed that $\mathbf{NEXP} \subseteq \mathbf{MIP}^*$ by taking a different route: rather than a black box transformation, they modified and reanalyzed a particular proof system, namely the MIP protocol for $\mathbf{NEXP}$ in [8], while leveraging and crucially using its algebraic structure. (Subsequent works [12, 13] improved this result by reducing the number of provers and rounds to a minimum, showing MIP* protocols for $\mathbf{NEXP}$ with two provers and one round.)

In sum, the challenge lies in the apparent incompatibility between techniques used for zero knowledge and those used for soundness against entangled provers.

B. High-level overview

Our strategy for proving our main result is to bridge the aforementioned gap by isolating the role of algebra in granting soundness against entangled provers, and developing new algebraic techniques for zero knowledge. Our proof of Theorem I.1 thus consists of two parts.

1) **Lifting lemma:** a black box transformation from algebraically-structured classical protocols into corresponding MIP* protocols, which preserves zero knowledge.
2) **Algebraic zero knowledge:** a new construction of zero knowledge algebraically-structured protocols for any language in $\mathbf{NEXP}$.

The first part is primarily a conceptual contribution, and it deals with quantum aspects of proof systems. The second part is our main technical contribution, and it deals with classical protocols (it does not require any background in quantum information). We briefly discuss each of the parts, and then provide an overview of the first part in Section II-C and of the second part in Section II-D.

In the first part of the proof, we build on recent advances in low-degree testing against entangled provers, and provide an abstraction of techniques in [11, 12, 13]. We prove a *lifting lemma* (see full version [28] for a precise statement) that transforms a class of algebraically-structured classical protocols into MIP* protocols, while preserving zero knowledge. This provides a generic framework for constructing MIP* protocols, while decoupling the mechanisms responsible for soundness against entangled provers from other classical components.

In the second part of the proof, we construct an algebraically-structured zero knowledge classical protocol, which we refer to as a *low-degree interactive PCP*, to which we apply the lifting lemma, completing the proof. At the heart of our techniques is a strong zero knowledge variant of the sumcheck protocol [29] (a fundamental subroutine in many probabilistic proof systems), which we deem of independent interest. In turn, a key component in our zero knowledge sumcheck is a new algebraic commitment scheme, whose hiding property is guaranteed by algebraic query complexity lower bounds [30, 31]. These shed more light on the connection of zero knowledge to algebraic complexity theory.

C. Part I: lifting classical proof systems to MIP*

The first step towards obtaining a generic framework for transforming classical protocols into corresponding MIP* protocols is making a simple, yet crucial, observation. Namely, while the result in [11] is stated as a white box modification of the MIP protocol in [8], we observe that the techniques used there can in fact be applied more generally. That is, we observe that *any* "low-degree interactive PCP", a type of algebraically structured proof system that underlies (implicitly and explicitly) many constructions in the probabilistic proof systems literature, can be transformed into a corresponding MIP* protocol.

The first part of the proof of Theorem I.1 formalizes this idea, identifying sufficient conditions to apply the techniques of [11, 12], and showing a lifting lemma that transforms protocols satisfying these conditions into MIP* protocols. We relate features of the original protocol to those of the resulting MIP* protocols, such as round complexity and, crucially, zero knowledge.

To make this discussion more accurate, we next define and discuss low-degree interactive PCPs.

1) Low-degree interactive PCPs: An *Interactive PCP* (IPCP), a proof system whose systematic study was initiated by Kalai and Raz [32], naturally extends the notions of a probabilistically checkable proof (PCP) and an interactive proof (IP). An r-round IPCP is a two-phase protocol in which a computationally unbounded *prover* P tries to convince a polynomial-time *verifier* V that an input x, given to both parties, is in a language $\mathscr{L}$. First, the prover sends to the verifier a PCP oracle (a purported proof that $x \in \mathscr{L}$), which the verifier can query at any time. Second, the prover and verifier engage in an r-round IP, at the end of which the verifier either accepts or rejects.[3] Completeness and soundness are defined in the usual way.

[3]Alternatively, an IPCP can be viewed as a PCP that is verified *interactively* (by an IP, instead of a randomized algorithm).

In this work we consider a type of algebraically-structured IPCP, which we call a *low-degree IPCPs*. This notion implicitly (and semi-explicitly) underlies many probabilistic proof systems in the literature. Informally, a low-degree IPCP is an IPCP satisfying the following: (1) *low-degree completeness*, which states that the PCP oracle sent by the (honest) prover is a polynomial of low (individual) degree; (2) *low-degree soundness*, which relaxes soundness to hold only against provers that send PCP oracles that are low-degree polynomials.

Low-degree completeness and soundness can be viewed as a promise that the PCP oracle is a low-degree polynomial. Indeed, these conditions are designed to capture "compatibility" with low-degree testing: only protocols with low-degree completeness will pass a low-degree test with probability 1; moreover, adding a low-degree test to an IPCP with low-degree soundness results (roughly) in an IPCP with standard soundness.

2) From low-degree IPCP to MIP:* We show that any low-degree IPCP can be transformed into a corresponding MIP* protocol, in a way that preserves zero knowledge (for a sufficiently strong notion of zero knowledge IPCP). To this end, we use an entanglement-resistant low degree test, which allows us to essentially restrict the provers usage of the entangled state to strategies that can be approximately implemented via randomness shared among the provers. Informally, the idea is that by carefully invoking such a test, we can let one prover take on the role of the PCP oracle, and the other to take the role of the IPCP prover, and then emulate the entire IPCP protocol.

In more detail, we show a zero-knowledge-preserving transformation of low-degree IPCPs to MIP* protocols, which is captured by the following lifting lemma.

Lemma II.1 (informally stated, see full version [28])**.** *There exists a transformation T that takes an* r*-round low-degree IPCP (P', V') for a language $\mathscr{L}$, and outputs a* 2*-prover* (r^*+2)*-round MIP* $(P_1, P_2, V) := T(P', V')$ for $\mathscr{L}$, where* $\mathsf{r}^* = \max\{\mathsf{r}, 1\}$*. Moreover, this transformation preserves zero knowledge.*[4]

We stress that the simplicity of the lifting lemma is a key feature since, as we describe below, it requires us to only make small structural changes to the IPCP protocol. This facilitates the preservation of various complexity measures and properties, such as zero knowledge.

To prove this lemma, a key tool that we use is a new low-degree test by Natarajan and Vidick [13],[5] which adapts the celebrated plane-vs-point test of Raz and Safra [33] to the MIP* model. A *low-degree test* is a procedure used to determine if a given function $f\colon \mathbb{F}^m \to \mathbb{F}$ is close to a low-degree polynomial or if, instead, it is far from all low-degree polynomials, by examining f at very few locations. In the plane-vs-point test, the verifier specifies a random 2-dimensional plane in $\mathbb{F}^m$ to one prover and a random point on this plane to the other prover; each prover replies with the purported value of f on the received plane or point; then the verifier checks that these values are consistent.

Informally, the analysis in [13] asserts that every entangled strategy that passes this test with high probably must satisfy an algebraic structure; more specifically, to pass this test the provers can only use their shared entangled state to (approximately) agree on a low-degree polynomial according to which they answer. We use the following soundness analysis of the this protocol. (See full version [28] for the standard quantum notation used in the theorem below.)

Theorem II.2 ([13, Theorem 2], informally stated)**.** *There exists an absolute constant $c \in (0,1)$ such that, for every soundness parameter $\varepsilon > 0$, number of variables $m \in \mathbb{N}$, degree $d \in \mathbb{N}$, and finite field $\mathbb{F}$, there exists a low-degree test T for which the following holds. For every* symmetric *entangled prover strategy and measurements $\{A_\alpha^z\}_{z\in\mathbb{F},\alpha\in\mathbb{F}^m}$ that are accepted by T with probability at least $1-\varepsilon$, there exists a measurement $\{L^Q\}_Q$, where Q is an m-variate polynomial of degree d, such that:*

1) Approximate consistency with $\{A_\alpha^z\}$:

$$\mathbb{E}_{\alpha\in\mathbb{F}^m} \sum_Q \sum_{z\neq Q(\alpha)} \langle\Psi| A_\alpha^z \otimes L^Q |\Psi\rangle \leq \varepsilon^c$$

.

2) Self-consistency of $\{L^Q\}$:

$$\sum_Q \langle\Psi| L^Q \otimes (\mathrm{Id} - L^Q) |\Psi\rangle \leq \varepsilon^c$$

.

In fact, we actually use a more refined version, which tests a polynomial's *individual* degree rather than its *total* degree. In the classical setting, such a test is implicit in [34] via a reduction from individual-degree to total-degree testing. Informally, this reduction first invokes the test for low total degree, then performs univariate low-degree testing with respect to a random axis-parallel line in each axis. We extend this reduction and its analysis to the setting of MIP*. (See full version [28] for details.) The analysis of the low individual degree test was communicated to us by Thomas Vidick, to whom we are grateful for allowing us to include it here.

With the foregoing low-degree test at our disposal, we are ready to outline the simple transformation from low-degree IPCPs to MIP* protocols. We begin with a preprocessing step. Note that the low individual degree test provides us with means to assert that the provers can (approximately) only use their entangled state to choose a low-degree polynomial Q,

[4]More accurately, we require the given IPCP to be zero knowledge with query bound that is roughly quadratic in the degree of the PCP oracle. See full version [28] for details.

[5]If we do not aim to obtain the optimal number of provers in our MIP* protocols, then it it suffices to use (an adaptation of) the low-degree test in [12].

and answer the verifier with the evaluation of Q on a *single*, uniformly distributed point (or plane). Thus, it is important that the IPCP verifier (which we start from) only makes a single uniform query to its oracle. By adapting techniques from [32], we can leverage the algebraic structure of the low-degree IPCP and capitalize on the interaction to ensure the IPCP verifier has this property, at essentially the cost of increasing the round complexity by 1.[6]

Thus we have a low-degree IPCP, with prover P and verifier V, in which the verification takes place as follows. Both P and V receive an explicit input x that is allegedly in the language $\mathscr{L}$. In addition, V is granted oracle access to a purported low-degree polynomial R, whose full description is known to P. The parties engage in an r-round interaction, at the end of which V is allowed to make a single uniform query to R and decide whether $x \in \mathscr{L}$ (with high probability).

We transform this IPCP into a 2-prover MIP* by considering the following protocol. First, the verifier chooses uniformly at random whether to (1) invoke a low-degree test, in which it asks one prover to evaluate R on a random plane or axis-parallel line and the other prover to evaluate R on a random point on this plane or line, or (2) emulate the IPCP protocol, in which one prover plays the role of the IPCP prover and the other acts as lookup for R.

We use the approximate consistency condition of Theorem II.2 to assert that the lookup prover approximately answers according to a low-degree polynomial, and use the self-consistency condition to ensure that both provers are consistently answering according to the *same* low-degree polynomial.[7]

We remark that preserving zero knowledge introduces some subtle technicalities (which we resolve), the main of which is that because the analysis of the entanglement-resistant low individual degree test requires that the provers employ *symmetric* strategies, we need to perform a non-standard symmetrization (since standard symmetrization turns out to break zero knowledge in our case). See full version [28] for details.

3) Towards zero knowledge MIP* for nondeterministic exponential time: Equipped with the lifting lemma, we are left with the task of constructing classical zero knowledge low-degree IPCPs for all languages in $\mathbf{NEXP}$. We first explain why current constructions do *not* suffice for this purpose.

The first thing to observe is that the classical protocol for the $\mathbf{NEXP}$-complete language *Oracle 3SAT* by Babai, Fortnow, and Lund [8] (neglecting the multilinearity test) can be viewed as low-degree IPCP. Indeed, in [8] the protocol is stated as an "oracle protocol", which is equivalent to an IPCP. The oracle is encoded as a low-degree polynomial, and so low-degree completeness is satisfied. Alas, the foregoing protocol is *not* zero knowledge. We remark that since the MIP* protocol in [11] relies on the protocol in [8], the former inherits the lack of zero knowledge from the latter.

Proceeding to consider classical *zero knowledge* proof systems, for example the protocols in [24, 35, 36], we observe that while some of these proof systems can be viewed as IPCPs, they are not *low-degree* IPCPs. This is because they achieve zero knowledge via combinatorial techniques that do *not* admit the algebraic structure that we require. We stress that the natural way of endowing an IPCP with algebraic structure by taking the low-degree extension of the PCP oracle does *not* necessarily preserve zero knowledge.[8] Correspondingly, the MIP* protocols in [12, 13], which rely on applying the low-degree extension code to a PCP, do not preserve zero knowledge for this reason.

Finally, we observe that recent advances in algebraic zero knowledge [37] (building on techniques from [38]) already provide us with a classical proof system that is compatible with our framework, and can thus be used to derive a zero knowledge MIP* protocol, albeit only for languages in $\#\mathbf{P}$.

To strengthen the aforementioned result and show that $\mathbf{NEXP} \subseteq \mathbf{PZK\text{-}MIP}^*$ (matching the $\mathbf{NEXP} \subseteq \mathbf{MIP}^*$ containment, and showing that zero knowledge can, in a sense, be obtained for "free" in the setting of MIP* protocols), we need to construct a *much stronger* zero knowledge low-degree IPCP. The second part of Theorem I.1, which is our main technical contribution, provides exactly that. We proceed to provide an overview of the techniques that we use to construct such protocols.

D. Part II: new algebraic techniques for zero knowledge

The techniques discussed thus far tell us that, if we wish to obtain a zero knowledge MIP* for $\mathbf{NEXP}$, it suffices to obtain a zero knowledge *low-degree* IPCP for $\mathbf{NEXP}$ (an IPCP wherein the oracle is a low-degree polynomial). Doing so is the second part of our proof of Theorem I.1, and for this we develop new algebraic techniques for obtaining zero knowledge protocols. Our techniques, which build on recent developments [38, 37], stand in stark contrast to other known constructions of zero knowledge PCPs and interactive PCPs (such as [24, 35, 36]). We remind the reader that this part of our work only deals with classical protocols, and does not require any knowledge of quantum information.

1) A zero knowledge low-degree IPCP for $\mathbf{NEXP}$: Our starting point is the protocol of Babai, Fortnow, and Lund [8] (the "BFL protocol"). We first recall how the BFL protocol works, in order to explain its sources of information leakage and how one could prevent them via algebraic techniques.

[6] Indeed, if the original IPCP verifier makes a single uniform query to its oracle, then we can save a round in Lemma II.1; that is, we obtain an MIP* with round complexity $\mathsf{r}^* + 1$, rather than $\mathsf{r}^* + 2$.

[7] Since the players are allowed the use of entanglement, we cannot hope for a single function that underlies their strategy. Indeed, the players could measure their entangled state to obtain shared randomness and select a random R according to which they answer.

[8] Intuitively, a single point in the encoded oracle can summarize a large amount of information from the original oracle (e.g., very large linear combinations).

These are the ideas that underlie our algebraic construction of an unconditional (perfect) zero knowledge low-degree IPCP for $\mathbf{NEXP}$.

The BFL protocol, and why it leaks: Oracle 3SAT (O3SAT) is the following $\mathbf{NEXP}$-complete problem: given a boolean formula B, does there exist a boolean function A (a witness) such that

$$B(z, b_1, b_2, b_3, A(b_1), A(b_2), A(b_3)) = 0$$
$$\text{for all } z \in \{0,1\}^r, b_1, b_2, b_3 \in \{0,1\}^s \ ?$$

The BFL protocol is an IPCP for O3SAT that is then (generically) converted to an MIP. In the BFL protocol, the honest prover first sends a PCP oracle $\hat{A}\colon \mathbb{F}^s \to \mathbb{F}$ that is the unique multilinear extension (in some finite field $\mathbb{F}$) of a valid witness $A\colon \{0,1\}^s \to \{0,1\}$. The verifier must check that (a) $\hat{A}$ is a boolean function on $\{0,1\}^s$, and (b) $\hat{A}$'s restriction to $\{0,1\}^s$ is a valid witness for B. To do these checks, the verifier arithmetizes the formula B into an arithmetic circuit $\hat{B}$, and reduces the checks to conditions that involve $\hat{A}$, $\hat{B}$, and other low-degree polynomials. A technique in [39] allows the verifier to "bundle" all of these conditions into a single low-degree polynomial f such that (with high probability over the choice of f) the conditions hold if and only if f sums to 0 on $\{0,1\}^{r+3s+3}$. The verifier checks that this is the case by engaging in a sumcheck protocol with the prover.[9]

We observe that the BFL protocol is *not* zero knowledge for two reasons: (i) the verifier has oracle access to $\hat{A}$ and, in particular, to the witness A; (ii) the prover's messages during the sumcheck protocol leak further information about A (namely, hard-to-compute partial sums of f, which itself depends on A).

A blueprint for zero knowledge: We now describe the "blueprint" for an approach to achieve zero knowledge in the BFL protocol. The prover does not send $\hat{A}$ directly, but instead a *commitment* to it. After this, the prover and verifier engage in a sumcheck protocol with suitable zero knowledge guarantees; at the end of this protocol, the verifier needs to evaluate f at a point of its choice, which involves evaluating $\hat{A}$ at three points. Now the prover reveals the requested values of $\hat{A}$, without leaking any information beyond these, so that the verifier can perform its check. We explain how these ideas motivate the need for certain algebraic tools, which we later develop and use to instantiate our approach.

(1) Randomized low-degree extension: Even if the prover reveals only three values of $\hat{A}$, these may still leak information about A. We address this problem via a *randomized low-degree extension*. Indeed, while the prover in the BFL protocol sends the *unique* multilinear extension of A, one can verify that *any* extension of A of sufficiently low degree also works. We exploit this flexibility as follows: the prover randomly samples $\hat{A}$ in such a way that any three evaluations of $\hat{A}$ do not reveal any information about A. Of course, if any of these evaluations is within the systematic part $\{0,1\}^s$, then no extension of A has this property. Nevertheless, during the sumcheck protocol, the prover can ensure that the verifier chooses only evaluations outside of $\{0,1\}^s$ (by aborting if the verifier deviates), which incurs only a small increase in the soundness error.[10] With this modification in place, it suffices for the prover to let $\hat{A}$ be a random degree-4 extension of A: by a dimensionality argument, any 3 evaluations outside of $\{0,1\}^s$ are now independent and uniformly random in $\mathbb{F}$. We are thus able to reduce a claim about A to a claim which contains *no information* about A.

(2) Algebraic commitments: As is typical in zero knowledge protocols, the prover will send a *commitment* to $\hat{A}$, and then selectively reveal a limited set of evaluations of $\hat{A}$. The challenge in our setting is that this commitment must *also* be a low-degree polynomial, since we require a low-degree oracle. For this, we devise a new algebraic commitment scheme based on the sumcheck protocol; we discuss this in Section II-D2.

(3) Sumcheck in zero knowledge: We need a sumcheck protocol where the prover's messages leak little information about f. The prior work in [37] achieves an IPCP for sumcheck that is "weakly" zero knowledge: any verifier learns at most one evaluation of f for each query it makes to the PCP oracle. If the verifier could evaluate f by itself, as was the case in that paper, this guarantee would suffice for zero knowledge. In our setting, however, the verifier *cannot* evaluate f by itself because f is (necessarily) hidden behind the algebraic commitment.

One approach to compensate would be to further randomize $\hat{A}$ by letting $\hat{A}$ be a random extension of A of some well-chosen degree d. Unfortunately, this technique is incompatible with our low-degree IPCP to MIP* transformation: such a low-degree extension is at most d-wise independent, whereas our lifting lemma, and more generally low-degree testing, requires zero knowledge against any $\Omega(d^2)$ queries.

We resolve this by relying on more algebraic techniques, achieving an IPCP for sumcheck with a much stronger zero knowledge guarantee: any malicious verifier that makes polynomially-many queries to the PCP oracle learns only a *single* evaluation of f. This suffices for zero knowledge in our setting: learning one evaluation of f implies learning only three evaluations of $\hat{A}$, which can be made "safe" if $\hat{A}$ is chosen to be a random extension of A of sufficiently high degree. Our sumcheck protocol uses as building blocks both our algebraic commitment scheme and the "weak" zero knowledge sumcheck in [37]; we summarize its construction in Section II-D3.

[9]The soundness of the sumcheck protocol depends on the PCP oracle being the evaluation of a low-degree polynomial, and so the verifier in [8] checks this using a low-degree test. In our setting of *low-degree* IPCPs a low-degree test is not necessary.

[10]The honest verifier will be defined so that it always chooses evaluations *outside* of $\{0,1\}^s$, so completeness is unaffected.

2) Algebraic commitments from algebraic query complexity lower bounds: We provide a high-level description of an information-theoretic commitment scheme in the low-degree IPCP model (i.e., a low-degree *interactive locking scheme* [36]). See full version [28] for details.[11]

In this scheme, the prover commits to a message by sending to the verifier a PCP oracle that perfectly hides the message; subsequently, the prover can reveal positions of the message by engaging with the verifier in an interactive proof, whose soundness guarantees statistical binding.

Committing to an element: We first consider the simple case of committing to a single element a in $\mathbb{F}$. Let k be a security parameter, and set $N := 2^k$. Suppose that the prover samples a random B in $\mathbb{F}^N$ such that $\sum_{i=1}^{N} B_i = a$, and sends B to the verifier as a commitment. Observe that any $N-1$ entries of B do not reveal any information about a, and so any verifier with oracle access to B that makes fewer than N queries cannot learn any information about a. However, as B is unstructured it is not clear how the prover can later convince the verifier that $\sum_{i=1}^{N} B_i = a$.

Instead, we can consider imbuing B with additional algebraic structure. Namely, the prover views B as a function from $\{0,1\}^k$ to $\mathbb{F}$, and sends its unique multilinear extension $\hat{B}\colon \mathbb{F}^k \to \mathbb{F}$ to the verifier. Subsequently, the prover can reveal a to the verifier, and then engage in a sumcheck protocol for the claim "$\sum_{\vec{\beta}\in\{0,1\}^k} \hat{B}(\vec{\beta}) = a$" to establish the correctness of a. The soundness of the sumcheck protocol protects the verifier against cheating provers and hence guarantees that this scheme is binding.

However, giving B additional structure calls into question the hiding property of the scheme. Indeed, surprisingly, a result of Juma et al. [31] shows that this new scheme is in fact *not* hiding (in fields of odd characteristic): it holds that $\hat{B}(2^{-1},\ldots,2^{-1}) = a \cdot 2^{-k}$ for any choice of B, so the verifier can learn a with only a single query to $\hat{B}$!

Sending an extension of B has created a new problem: querying the extension outside of $\{0,1\}^k$, the verifier can learn information that may require many queries to B to compute. Indeed, this additional power is precisely what underlies the soundness of the sumcheck protocol. To resolve this, we need to understand what the verifier can learn about B given some low-degree extension $\hat{B}$. This is precisely the setting of *algebraic query complexity* [30].[12]

Indeed the foregoing theory suggests a natural approach for overcoming the problem created by the extension of B: instead of considering the multilinear extension, we can let $\hat{B}$ be chosen uniformly at random from the set of degree-d extensions of B, for some $d > 1$. It is not hard to see that if d is very large (say, $|\mathbb{F}|$) then 2^k queries are required to determine the summation of $\hat{B}$ on $\{0,1\}^k$. However, we need d to be small to achieve soundness. Fortunately, a result of [31] shows that $d = 2$ suffices: given a random multiquadratic extension $\hat{B}$ of B, one needs 2^k queries to $\hat{B}$ to determine $\sum_{\vec{\beta}\in\{0,1\}^k} \hat{B}(\vec{\beta})$.[13]

Committing to a polynomial: The prover in our zero knowledge protocols needs to commit not just to a single element but rather to the evaluation of an m-variate polynomial Q over $\mathbb{F}$ of degree $d > 1$. We extend our ideas to this setting. We follow a similar general approach, however, arguing the hiding property now requires a *stronger* algebraic query complexity lower bound than the one proved in [31]. Not only do we need to know that the verifier cannot determine $Q(\vec{\alpha})$ for a particular $\vec{\alpha} \in \mathbb{F}^m$, but we need to know that the verifier cannot determine $Q(\vec{\alpha})$ for *any* $\vec{\alpha} \in \mathbb{F}^m$, or even *any linear combination of any such values*. We prove that this stronger guarantee holds in the same parameter regime: if $d > 1$ then 2^k queries are both necessary and sufficient. See the discussion full version [28] for a more detailed overview.

Decommitting in zero knowledge: To use our commitment scheme in zero knowledge protocols, we must ensure that, in the decommitment phase, the verifier cannot learn any information beyond the value $a := Q(\vec{\alpha})$, for a chosen $\vec{\alpha}$. To decommit, the prover sends the value a and has to convince the verifier that the claim "$\sum_{\vec{\beta}\in\{0,1\}^k} \hat{B}(\vec{\alpha},\vec{\beta}) = a$" is true. However, if the prover and verifier simply run the sumcheck protocol on this claim, the prover leaks partial sums $\sum_{\vec{\beta}\in\{0,1\}^{k-i}} \hat{B}(\vec{\alpha}, c_1,\ldots,c_i,\vec{\beta})$, for $c_1,\ldots,c_i \in \mathbb{F}$ chosen by the verifier, which could reveal additional information about Q. Instead, the prover and verifier run on this claim the IPCP for sumcheck of [37], whose "weak" zero knowledge guarantee ensures that this cannot happen. (Thus, in addition to the commitment, the honest prover also sends the evaluation of a random low-degree polynomial as required by the IPCP for sumcheck of [37].)

3) A zero knowledge sumcheck protocol: We describe the "strong" zero knowledge variant of the sumcheck protocol that we use in our construction. The protocol relies on the algebraic commitment scheme described in the previous section. We first cover some necessary background, and then describe our protocol.

Previous sumcheck protocols: The sumcheck protocol [29] is an IP for claims of the form "$\sum_{\vec{\alpha}\in H^m} F(\vec{\alpha}) = 0$", where H is a subset of a finite field $\mathbb{F}$ and F is an m-variate polynomial over $\mathbb{F}$ of small individual degree. The protocol has m rounds: in round i, the prover sends the univariate polynomial $g_i(X_i) := \sum_{\vec{\alpha}\in H^{m-i}} F(c_1,\ldots,c_{i-1},X_i,\vec{\alpha})$, where $c_1,\ldots,c_{i-1} \in \mathbb{F}$ were sent by the verifier in previous

[11] We use the commitment scheme perspective to illustrate the key ideas in our construction. In the technical sections, we prove the zero knowledge property directly using algebraic query complexity lower bounds, without explicitly using any commitment scheme.

[12] Interestingly, in [30] a connection between algebra and zero knowledge is also exhibited. Namely, to show that the result $\mathbf{NP} \subseteq \mathbf{CZK}$ [2] algebrizes, it is necessary to exploit the algebraic structure of the oracle to design a zero knowledge protocol for verifying the existence of certain sets of query answers.

[13] This is the main reason why our application to constructing MIP* protocols requires low-degree test against entangled provers, rather than just a multilinearity test, as was used in [11].

rounds; the verifier checks that $\sum_{\alpha_i \in H} g_i(\alpha_i) = g_{i-1}(c_{i-1})$ and replies with a uniformly random challenge $c_i \in \mathbb{F}$. After round m, the verifier outputs the claim "$F(c_1, \ldots, c_m) = g_m(c_1, \ldots, c_m)$". If F is of sufficiently low degree and does not sum to a over the space, then the output claim is false with high probability. Note that the verifier does not need access to F.

The "weak" zero knowledge IPCP for sumcheck in [37] modifies the above protocol as follows. The prover first sends a PCP oracle that (allegedly) equals the evaluation of a random "masking" polynomial R; the verifier checks that R is (close to) low degree. Subsequently, the prover and verifier conduct the following interactive proof. The prover sends $z \in \mathbb{F}$ that allegedly equals $\sum_{\vec{\alpha} \in H^m} R(\vec{\alpha})$, and the verifier responds with a uniformly random challenge $\rho \in \mathbb{F}^*$. The prover and verifier now run the (standard) sumcheck protocol to reduce the claim "$\sum_{\vec{\alpha} \in H^m} \rho F(\vec{\alpha}) + R(\vec{\alpha}) = \rho a + z$" to a claim "$\rho F(\vec{c}) + R(\vec{c}) = b$", for a random $\vec{c} \in \mathbb{F}^m$. The verifier queries R at $\vec{c}$ and then outputs the claim "$F(\vec{c}) = \frac{b - R(\vec{c})}{\rho}$". If $\sum_{\vec{\alpha} \in H^m} F(\vec{\alpha}) \neq a$, then with high probability over ρ and the verifier's messages in the sumcheck protocol, this claim is false.

A key observation is that if the verifier makes no queries to R, then the prover's messages are identically distributed to the sumcheck protocol applied to a random polynomial Q. When the verifier does make queries to R, simulating the resulting conditional distribution involves techniques from Algebraic Complexity Theory, as shown in [37]. Given Q, the verifier's queries to $R(\vec{\alpha})$, for $\vec{\alpha} \in \mathbb{F}^m$, are identically distributed to $Q(\vec{\alpha}) - \rho F(\vec{\alpha})$. Thus, the simulator need only make at most one query to F for every query to R; that is, any verifier making q queries to R learns no more than it would learn by making q queries to F alone.

As discussed, this zero knowledge guarantee does not suffice for the application that we consider: in the $\mathbf{NEXP}$ protocol, the polynomial F is defined in terms of the $\mathbf{NEXP}$ witness. In this case the verifier can learn enough about F to break zero knowledge by making only $O(\deg(F))$ queries to R.

Our sumcheck protocol: The "strong" zero knowledge guarantee that we aim for is the following: any polynomial-time verifier learns no more than it would by making *one* query to F, regardless of its number of queries to the PCP oracle.

The main idea to achieve this guarantee is the following. The prover sends a PCP oracle that is an *algebraic commitment* Z to the aforementioned masking polynomial R. Then, as before, the prover and verifier run the sumcheck protocol to reduce the claim "$\sum_{\vec{\alpha} \in H^m} \rho F(\vec{\alpha}) + R(\vec{\alpha}) = \rho a + z$" to a claim "$\rho F(\vec{c}) + R(\vec{c}) = b$" for random $\vec{c} \in \mathbb{F}^m$.

We now face two problems. First, the verifier cannot simply query R at $\vec{c}$ and then output the claim "$F(\vec{c}) = \frac{b - R(\vec{c})}{\rho}$", since the verifier only has oracle access to the commitment Z of R. Second, the prover could cheat the verifier by having Z be a commitment to an R that is far from low degree, which allows cheating in the sumcheck protocol.

The first problem is addressed by the fact that our algebraic commitment scheme has a decommitment sub-protocol that is zero knowledge: the prover can reveal $R(\vec{c})$ in such a way that no other values about R are also revealed as a side-effect. As discussed, this relies on the protocol of [37], used as a subroutine.

The second problem is addressed by the fact that our algebraic commitment scheme is "transparent" to low-degree structure; that is, the algebraic structure of the scheme implies that if the commitment Z is a low-degree polynomial (as in a *low-degree* IPCPs), then R must also be low degree (and vice versa).

Overall, the only value that a malicious verifier can learn is $F(\vec{c})$, for $\vec{c} \in I^m$ of its choice (where I is some sufficiently large subset of $\mathbb{F}$, fixed in advance). More precisely, we prove the following theorem, which shows a strong zero knowledge sumcheck protocol.

Theorem II.3 (Informally stated, see full version [28])**.** *There exists a low-degree IPCP for sumcheck, with respect to a low-degree polynomial F, that satisfies the following zero knowledge guarantee: the view of any probabilistic polynomial-time verifier in the protocol can be perfectly and efficiently simulated by a simulator that makes only a single query to F.*

Our sumcheck protocol leaks a single evaluation of F. We stress that this limitation is *inherent*: the honest verifier always outputs a true claim about one evaluation of F, which it cannot do without learning that evaluation. Nevertheless, this guarantee is strong enough for our application, as we can ensure that learning a single evaluation of F does not harm zero knowledge.

We remark that our strong zero knowledge sumcheck protocol can be transformed into a standard IPCP, by the standard technique of adding a (classical) low-degree test to the protocol.

III. Discussion and open problems

The framework that we use to prove that $\mathbf{NEXP} \subseteq \mathbf{PZK\text{-}MIP}^*$ elucidates the role that algebra plays in the design of proofs systems with entangled provers. Namely, we show that a large class of algebraic protocols (low-degree IPCPs) can be transformed in a black box manner to MIPs with entangled provers. This abstraction decouples the mechanisms responsible for soundness against entangled adversaries from other classical components in the proof system. In turn, this allows us to focus our attention on designing proof systems with desirable properties (zero knowledge, in this work), without having to deal with the complications that arise from entanglement, and then derive MIP^* protocols from these classical protocols.

These ideas also enable us to re-interpret prior constructions of MIP* protocols at a higher level of abstraction. For example, the protocol in [11] can be viewed as applying our lifting lemma to the (low-degree) IPCP in [8]. As another example, one can start with *any* PCP for some language $\mathscr{L}$, low-degree extend the PCP, and then apply our lifting lemma to obtain a corresponding MIP* protocol for $\mathscr{L}$; in fact, the protocol in [12] can be viewed in this perspective.

In more detail, we say that a transformation from IPCP to MIP* is *black box* if it maps an IPCP protocol into an MIP* protocol whose verifier can be expressed as an algorithm that only accesses the queries and messages of the IPCP verifier, but does not access its input (apart from its length). The following corollary shows that any IPCP protocol can be transformed into an MIP* protocol via a black box transformation. While a proof of this fact is implicit in [12, 13], the framework developed in this paper allows us to crystallize its structure and give a compellingly short proof of it.

Corollary III.1. *There is a black box transformation that maps any* r*-round IPCP protocol for a language* $\mathscr{L}$ *to a* 2*-prover* $(\mathsf{r}+1)$*-round MIP* for* $\mathscr{L}$

The round complexity of $\mathsf{r}+1$ in Corollary III.1 is less than in our lifting lemma $(\mathsf{r}+2)$, because now we do not require that zero knowledge is preserved. We make the foregoing discussion precise in the full version [28].

We conclude this section by discussing several open problems.

In this work we show that there exist perfect zero knowledge MIP* protocols for all languages in $\mathbf{NEXP}$, with *polynomially-many* rounds. Since round complexity is a crucial resource in any interactive proof system, it is essential to understand whether zero knowledge MIP* protocols with low round complexity exist. (After all, without the requirement of zero knowledge, every language in $\mathbf{NEXP}$ has a MIP* protocol with just *one round* [12, 13].) We remark that the "oracularization" technique of Ito et al. [26] reduces the round complexity of any MIP* to one round, but this technique does *not* preserve zero knowledge.

Open Problem 1. Do there exist constant-round zero knowledge MIP* protocols for $\mathbf{NEXP}$?

At the beginning of this section, we reflected on the fact that known results that establish the power of MIP* protocols rely on *algebraic* structure, which enables classical-to-quantum black box transformations of protocols. But is algebraic structure inherently required, or does some combinatorial structure suffice?

Open Problem 2. Is there a richer class of classical protocols (beyond low-degree IPCPs) that can be black-box transformed into MIP* protocols?

For instance, could we replace low-degree polynomials with, say, error correcting codes with suitable local testability and decodability properties? One place to start would be to understand whether local testers for tensor product codes [40] are sound against entangled provers.

Open Problem 3. When suitably adapted to the multi-prover setting, is the random hyperplane test in [40] for tensor product codes sound against entangled provers?

ACKNOWLEDGMENT

We are grateful to Thomas Vidick for multiple technical and conceptual suggestions that greatly improved our results and their presentation. We thank Claude Crépeau for useful comments on an earlier version of this paper. We also thank Zeph Landau, Chinmay Nirkhe, and Igor Shinkar for helpful discussions.

REFERENCES

[1] S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof systems," *SIAM Journal on Computing*, vol. 18, no. 1, pp. 186–208, 1989, preliminary version appeared in STOC '85.

[2] O. Goldreich, S. Micali, and A. Wigderson, "Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems," *Journal of the ACM*, vol. 38, no. 3, pp. 691–729, 1991, preliminary version appeared in FOCS '86.

[3] R. Ostrovsky and A. Wigderson, "One-way functions are essential for non-trivial zero-knowledge," in *Proceedings of the 2nd Israel Symposium on Theory of Computing Systems*, ser. ISTCS '93, 1993, pp. 3–17.

[4] R. B. Boppana, J. Håstad, and S. Zachos, "Does co-NP have short interactive proofs?" *Information Processing Letters*, vol. 25, no. 2, pp. 127–132, 1987.

[5] L. Fortnow, "The complexity of perfect zero-knowledge (extended abstract)," in *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, ser. STOC '87, 1987, pp. 204–209.

[6] W. Aiello and J. Håstad, "Statistical zero-knowledge languages can be recognized in two rounds," *Journal of Computer and System Sciences*, vol. 42, no. 3, pp. 327–345, 1991, preliminary version appeared in FOCS '87.

[7] M. Ben-Or, S. Goldwasser, J. Kilian, and A. Wigderson, "Multi-prover interactive proofs: how to remove intractability assumptions," in *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, ser. STOC '88, 1988, pp. 113–131.

[8] L. Babai, L. Fortnow, and C. Lund, "Non-deterministic exponential time has two-prover interactive protocols," *Computational Complexity*, vol. 1, pp. 3–40, 1991, preliminary version appeared in FOCS '90.

[9] 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*, 2004, pp. 236–249.

[10] J. S. Bell, "On the Einstein Podolsky Rosen paradox," 1964.

[11] T. Ito and T. Vidick, "A multi-prover interactive proof for NEXP sound against entangled provers," in *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science*, ser. FOCS '12, 2012, pp. 243–252.

[12] T. Vidick, "Three-player entangled XOR games are NP-hard to approximate," *SIAM Journal on Computing*, vol. 45, no. 3, pp. 1007–1063, 2016.

[13] A. Natarajan and T. Vidick, "Two-player entangled games are np-hard," in *Proceedings of the 32nd Annual IEEE Conference on Computational Complexity*, ser. CCC '18, 2018, pp. 20:1–20:18.

[14] J. Watrous, "Limits on the power of quantum statistical zero-knowledge," in *Proceedings of the 43rd Symposium on Foundations of Computer Science*, ser. FOCS '02, 2002, p. 459.

[15] H. Kobayashi, "Non-interactive quantum perfect and statistical zero-knowledge," in *Proceedings of the 14th Algorithms and Computation International Symposium*, 2003, pp. 178–188.

[16] I. Damgård, S. Fehr, and L. Salvail, "Zero-knowledge proofs and string commitments withstanding quantum attacks," in *Proceedings of the 24th Annual International Cryptology Conference*, ser. CRYPTO '04, 2004, pp. 254–272.

[17] J. Watrous, "Zero-knowledge against quantum attacks," *SIAM Journal on Computing*, vol. 39, no. 1, pp. 25–58, 2009.

[18] H. Kobayashi, "General properties of quantum zero-knowledge proofs," in *Proceedings of the 5th Theory of Cryptography Conference*, ser. TCC '08, 2008, pp. 107–124.

[19] D. Unruh, "Quantum proofs of knowledge," in *Proceedings of the 31st Annual International Conference on Theory and Application of Cryptographic Techniques*, ser. EUROCRYPT '12, 2012, pp. 135–152.

[20] ——, "Non-interactive zero-knowledge proofs in the quantum random oracle model," in *Proceedings of the 34th Annual International Conference on Theory and Application of Cryptographic Techniques*, ser. EUROCRYPT '15, 2015, pp. 755–784.

[21] A. Broadbent, Z. Ji, F. Song, and J. Watrous, "Zero-knowledge proof systems for QMA," in *Proceedings of the 57th Annual Symposium on Foundations of Computer Science*, ser. FOCS '16, 2016, pp. 31–40.

[22] S. Menda and J. Watrous, "Oracle separations for quantum statistical zero-knowledge," *CoRR*, vol. abs/1801.08967, 2018.

[23] C. Crépeau, L. Salvail, J. Simard, and A. Tapp, "Two provers in isolation," in *Proceedings of the 17th International Conference on the Theory and Application of Cryptology and Information Security*, ser. ASIACRYPT '11, 2011, pp. 407–430.

[24] C. Dwork, U. Feige, J. Kilian, M. Naor, and S. Safra, "Low communication 2-prover zero-knowledge proofs for NP," in *Proceedings of the 11th Annual International Cryptology Conference*, ser. CRYPTO '92, 1992, pp. 215–227.

[25] 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*, ser. CCC '08, 2008, pp. 187–198.

[26] T. Ito, H. Kobayashi, and K. Matsumoto, "Oracularization and two-prover one-round interactive proofs against nonlocal strategies," in *Proceedings of the 24th IEEE Annual Conference on Computational Complexity*, ser. CCC '09, 2009, pp. 217–228.

[27] J. Kempe, H. Kobayashi, K. Matsumoto, B. Toner, and T. Vidick, "Entangled games are hard to approximate," *SIAM Journal on Computing*, vol. 40, no. 3, pp. 848–877, 2011.

[28] A. Chiesa, M. A. Forbes, T. Gur, and N. Spooner, "Spatial isolation implies zero knowledge even in a quantum world," *CoRR*, vol. abs/1803.01519, 2018.

[29] C. Lund, L. Fortnow, H. J. Karloff, and N. Nisan, "Algebraic methods for interactive proof systems," *Journal of the ACM*, vol. 39, no. 4, pp. 859–868, 1992.

[30] S. Aaronson and A. Wigderson, "Algebrization: A new barrier in complexity theory," *ACM Transactions on Computation Theory*, vol. 1, no. 1, pp. 2:1–2:54, 2009.

[31] A. Juma, V. Kabanets, C. Rackoff, and A. Shpilka, "The black-box query complexity of polynomial summation," *Computational Complexity*, vol. 18, no. 1, pp. 59–79, 2009.

[32] Y. Kalai and R. Raz, "Interactive PCP," in *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*, ser. ICALP '08, 2008, pp. 536–547.

[33] R. Raz and S. Safra, "A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP," in *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, ser. STOC '97, 1997, pp. 475–484.

[34] O. Goldreich and M. Sudan, "Locally testable codes and PCPs of almost-linear length," *Journal of the ACM*, vol. 53, pp. 558–655, July 2006, preliminary version in STOC '02.

[35] J. Kilian, E. Petrank, and G. Tardos, "Probabilistically checkable proofs with zero knowledge," in *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, ser. STOC '97, 1997, pp. 496–505.

[36] V. Goyal, Y. Ishai, M. Mahmoody, and A. Sahai, "Interactive locking, zero-knowledge PCPs, and unconditional cryptography," in *Proceedings of the 30th Annual Conference on Advances in Cryptology*, ser. CRYPTO'10, 2010, pp. 173–190.

[37] E. Ben-Sasson, A. Chiesa, M. A. Forbes, A. Gabizon, M. Riabzev, and N. Spooner, "Zero knowledge protocols from succinct constraint detection," in *Proceedings of the 15th Theory of Cryptography Conference*, ser. TCC '17, 2017, pp. 172–206.

[38] E. Ben-Sasson, A. Chiesa, A. Gabizon, and M. Virza, "Quasilinear-size zero knowledge from linear-algebraic PCPs," in *Proceedings of the 13th Theory of Cryptography Conference*, ser. TCC '16-A, 2016, pp. 33–64.

[39] L. Babai, L. Fortnow, L. A. Levin, and M. Szegedy, "Checking computations in polylogarithmic time," in *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*, ser. STOC '91, 1991, pp. 21–32.

[40] E. Ben-Sasson and M. Sudan, "Robust locally testable codes and products of codes," *Random Structures and Algorithms*, vol. 28, no. 4, pp. 387–402, 2006.