Interactive Proofs with Approximately Commuting Provers

Matthew Coudron[1(✉)] and Thomas Vidick[2]

[1] Massachusetts Institute of Technology, Cambridge, MA, USA
mcoudron@mit.edu
[2] California Institute of Technology, Pasadena, CA, USA
vidick@cms.caltech.edu

Abstract. The class MIP* of promise problems that can be decided through an interactive proof system with multiple entangled provers provides a complexity-theoretic framework for the exploration of the nonlocal properties of entanglement. Very little is known in terms of the power of this class. The only proposed approach for establishing upper bounds is based on a hierarchy of semidefinite programs introduced independently by Pironio et al. and Doherty et al. in 2006. This hierarchy converges to a value, the field-theoretic value, that is only known to coincide with the provers' maximum success probability in a given proof system under a plausible but difficult mathematical conjecture, Connes' embedding conjecture. No bounds on the rate of convergence are known.

We introduce a rounding scheme for the hierarchy, establishing that any solution to its N-th level can be mapped to a strategy for the provers in which measurement operators associated with distinct provers have pairwise commutator bounded by $O(\ell^2/\sqrt{N})$ in operator norm, where ℓ is the number of possible answers per prover.

Our rounding scheme motivates the introduction of a variant of quantum multiprover interactive proof systems, called MIP^*_δ, in which the soundness property is required to hold against provers allowed to operate on the same Hilbert space as long as the commutator of operations performed by distinct provers has norm at most δ. Our rounding scheme implies the upper bound $\mathrm{MIP}^*_\delta \subseteq \mathrm{DTIME}(\exp(\exp(\mathrm{poly})/\delta^2))$. In terms of lower bounds we establish that $\mathrm{MIP}^*_{2^{-\mathrm{poly}}}$ contains NEXP with completeness 1 and soundness $1 - 2^{-\mathrm{poly}}$. We discuss connections with the mathematical literature on approximate commutation and applications to device-independent cryptography.

1 Introduction

In a multiprover interactive proof system, a *verifier* with bounded resources (a polynomial-time Turing machine) interacts with multiple all-powerful but non-communicating *provers* in an attempt to verify the truth of a mathematical statement — the membership of some input x, a string of bits, in a language L, such as 3-SAT. The provers always collaborate to maximize their chances of making the verifier accept the statement, and their maximum probability of success in

M.M. Halldórsson et al. (Eds.): ICALP 2015, Part I, LNCS 9134, pp. 355–366, 2015.
DOI: 10.1007/978-3-662-47672-7_29

doing so is called the *value* $\omega = \omega(x)$ of the protocol. We will sometimes refer to a given protocol as an "interactive game" and call the provers "players". A proof system's *completeness* c is the smallest value of $\omega(x)$ over all $x \in L$, while its soundness s is the largest value of $\omega(x)$ over $x \notin L$; a protocol is sound if $s < c$.

The class of all languages that have multiprover interactive proof systems with $c \geq 2/3$ and $s \leq 1/3$, denoted MIP, is a significant broadening of its non-interactive, single-prover analogue MA, as is witnessed by the characterization MIP = NEXP [BFL91]. This result is one of the cornerstones on which the PCP theorem [AS98, ALM+98] was built, with consequences ranging from cryptography [BOGKW88] to hardness of approximation [FGL+96].

Quantum information suggests a natural extension of the class MIP. The laws of quantum mechanics assert that, in the physical world, a set of non-communicating provers may share an arbitrary entangled quantum state, a physical resource which strictly extends their set of strategies but provably does not allow them to communicate. The corresponding extension of MIP is the class MIP* of all languages that have multiprover interactive proof systems with entangled provers [KM03].

Physical intuition for the significance of the prover's new resource, entanglement, dates back to Einstein, Podolsky and Rosen's paradoxical account [EPR35] of the consequences of quantum entanglement, later clarified through Bell's pioneering work [Bel64]. To state the relevance of Bell's results more precisely in our context we first introduce the mathematical formalism used by Bell to model locality. With each prover's private space is associated a separate Hilbert space. The joint quantum state of the provers is specified by a unit vector $|\Psi\rangle$ in the tensor product of their respective Hilbert spaces. Upon receiving its query from the verifier, each prover applies a local measurement (a positive operator supported on its own Hilbert space) the outcome of which is sent back to the verifier as its answer. The supremum of the provers' probability of being accepted by the verifier, taken over all Hilbert spaces, states in their joint tensor product, and local measurements, is called the entangled value ω^* of the game. The analogue quantity for "classical" provers (corresponding to shared states which are product states) is denoted ω.

Bell's work and the extensive literature on Bell inequalities [CHSH69, Ara02] and quantum games [CHTW04] establishes that there are protocols, or interactive games, for which $\omega^* > \omega$. This simple fact has important consequences for interactive proof systems. First, a proof system sound with classical provers may no longer be so in the presence of entanglement. Cleve et al. [CHTW04] exhibit a class of restricted interactive proof systems, XOR proof systems, such that the class with classical provers equals NEXP while the same proof systems with entangled provers cannot decide any language beyond EXP. Second, the completeness property of a proof system may also increase through the provers' use of entanglement. As a result optimal strategies may require the use of arbitrarily large Hilbert spaces for the provers — no explicit bound on the dimension of these spaces is known as a function of the size of the game.

In fact no better upper bound on the class MIP* is known other than its languages being recursively enumerable: they may not even be decidable! This unfortunate state of affairs stems from the fact that, while the value ω^* may be approached from below through exhaustive search in increasing dimensions, there is no verifiable criterion for the termination of such a procedure.

Bounding entangled-Prover Strategies. The question of deriving algorithmic methods for placing upper bounds on the entangled value ω^* of a given protocol has long frustrated researchers' efforts. Major progress came in 2006 through the introduction of a hierarchy of relaxations based on semidefinite programming [DLTW08, NPA07] that we will refer to as the QCSDP hierarchy. These relaxations follow a similar spirit as e.g. the Lasserre hierarchy in combinatorial optimization [Lau03], and can be formulated using the language of sums of squares of *non-commutative* polynomials. In contrast with the commutative setting, this leads to a hierarchy that is in general infinite and need not converge at any finite level.

The limited convergence results that are known for the QCSDP hierarchy involve a formalization of locality for quantum provers which originates in the study of infinite-dimensional systems such as those that arise in quantum field theory. Here the idea is that observations made at different space-time locations should be represented by operators which, although they may act on the same Hilbert space, should nevertheless commute — a minimal requirement ensuring that the joint outcome of any two measurements made by distinct parties should be well-defined and independent of the order in which the measurements were performed.

For the case of finite-dimensional systems this seemingly weaker condition is equivalent to the existence of a tensor product representation [DLTW08]. In contrast, for the case of infinite-dimensional systems the two formulations are not known to be equivalent. This question, known as Tsirelson's problem in quantum information, was recently shown to be equivalent to a host of deep mathematical conjectures [SW08, JNP+11], in particular Connes' embedding conjecture [Con76] and Kirchberg's QWEP conjecture [Kir93]. The validity of these conjectures has a direct bearing on our understanding of MIP*. The QCSDP hierarchy is known to converge to a value called the *field-theoretic value* ω^f of the game, which is the maximum success probability achievable by commuting strategies of the type described above. A positive answer to Tsirelson's conjecture thus implies that $\omega^* = \omega^f$ and both quantities are computable. However, even assuming the conjecture and in spite of strong interest (the use of the first few levels of the hierarchy has proven extremely helpful to study a range of questions in device independence [BSS14, YVB+14] and the study of nonlocality [PV10]) absolutely no bounds have been obtained on the convergence rate of the hierarchy. It is only known that if a certain technical condition, called a rank loop, holds, then convergence is achieved [NPA08]; unfortunately the condition is computationally expensive to verify (even for low levels of the hierarchy) and, in general, may not be satisfied at any finite level.

Beyond the obvious limitations for practical applications, these severe computational difficulties are representative of the intrinsic difficulty of working with the model of entangled provers. Our work is motivated by this state of affairs: we establish the first quantitative convergence results for the quantum SDP hierarchy. Our main observation is that successive levels of the hierarchy place bounds on the value achievable by provers employing a relaxed notion of strategy in which measurements applied by distinct provers are allowed to *approximately commute*: their commutator is bounded, in operator norm, by a quantity that goes to zero with the level in the hierarchy.

In this abstract we describe our quantitative results, use them to motivate the introduction of a sub-class MIP^*_{ac} of MIP^* and prove non-trivial lower and upper bounds on that class. We discuss the relevance of the study of MIP^*_{ac} for that of MIP^* and closely related results from the mathematical literature. We refer to the full version for precise definitions as well as complete proofs of the results announced here.

2 A Rounding Scheme for the QCSDP Hierarchy

Our main technical result is a rounding procedure for the QCSDP hierarchy of semidefinite programs [NPA07, DLTW08]. The procedure maps any feasible solution to the N-th level of the hierarchy to a set of measurement operators for the provers that approximately commute. For simplicity we state and prove our results for the case of a single round of interaction with two provers and classical messages only. Extension to multiple provers is straightforward; we expect generalizations to multiple rounds and quantum messages to be possible but leave them for future work.

Definition 1. *An (m, ℓ) strategy for the provers specified by two sets of m POVMS $\{A_x^a\}_{1\leq a\leq \ell}$ and $\{B_y^b\}_{1\leq b\leq \ell}$ with ℓ outcomes each, where $x, y \in \{1, \ldots, m\}$.*

A strategy is said to be δ-AC if for every x, y, a and b, $\|A_x^a B_y^b - B_y^b A_x^a\| \leq \delta$, where $\|\cdot\|$ denotes the operator norm.

Our results apply to the QCSDP hierarchy of semidefinite programs as defined in [NPA07].

Theorem 1. *Let G be a 2-prover one-round game with classical messages in which each player has ℓ possible answers, and $\omega^N_{QCSDP}(G)$ the optimum of the N-th level of the QCSDP hierarchy. Then there exists a $\delta = O(\ell^2/\sqrt{N})$ and a δ-AC strategy for the provers with success probability $\omega^N_{QCSDP}(G)$ in G.*[1]

[1] Due to the approximate commutation of the provers' strategies the success probability of δ-AC strategies may a priori depend on the order in which the measurement operators are applied. In our context the parameter δ will always be small enough that we can neglect this effect. Moreover, for the particular kind of strategies constructed in our rounding scheme the value will not be affected by the order.

Our result is the first to derive the condition that the *operator norm* of commutators is small. In contrast it is not hard to show that a feasible solution to the first level of the hierarchy already gives rise to measurement operators that exactly satisfy a commutation relation *when evaluated on the state* (corresponding to the zeroth-order vector provided by the hierarchy). While the latter condition can be successfully exploited to give an exact rounding procedure from the first level for the class of XOR games [CHTW04], and an approximate rounding for the more general class of unique games [KRT10], we do not expect it to be sufficient in general. In particular, even approximate tightness of the first level of the hierarchy for three-player games would imply EXP = NEXP [Vid13]. We will furtherore show that the problem of optimizing over strategies which approximately commute, to within sufficiently small error and in *operator norm*, is NEXP-hard (see Section 3 for details).

The proof of Theorem 1 is constructive: starting from any feasible solution to the N-th level of the QCSDP hierarchy we construct measurement operators for the provers with pairwise commutators bounded by δ in operator norm, and which achieve a value in the game that equals the objective value of the N-th level SDP. Recall that this SDP has $O(m\ell)^N$ vector variables indexed by strings of length at most N over the formal alphabet $\{P_x^a, Q_y^b\}$ containing a symbol for each possible (question,answer) pair to any of the provers. Our main idea is to introduce a "graded" variant of the construction in [NPA08] (which was used to show convergence under the rank loop constraint). Rather informally, the rounded measurement operators, $\{\tilde{P}_x^a\}$ for the first prover and $\{\tilde{Q}_y^b\}$ for the second, can be defined as follows:

$$\tilde{P}_x^a \equiv \frac{1}{N-1}\sum_{i=1}^{N-1} \Pi_{\leq i}\Pi_{P_x^a}\Pi_{\leq i} \quad \text{and} \quad \tilde{Q}_y^b \equiv \frac{1}{N-1}\sum_{j=1}^{N-1} \Pi_{\leq j}\Pi_{Q_y^b}\Pi_{\leq j}.$$

Here $\Pi_{P_x^a}$ and $\Pi_{Q_y^b}$ are projectors as defined in [NPA08], i.e. as the projection onto vectors associated with strings ending in the formal label P_x^a, Q_y^b of the corresponding operator. The novelty is the introduction of the $\Pi_{\leq i}$, which project onto the subspace spanned by all vectors associated with strings of length at most i. Thus $\tilde{P}_x^a$ itself is not a projector, and it gives more weight to vectors indexed by shorter strings.

The intuition behind this rounding scheme is as follows. The winning probability is unchanged because it is determined by the action of the measurement operators on the subspace $\mathrm{Im}(\Pi_{\leq 1})$. On the other hand, the rounded operators approximately commute in the operator norm because the original operators commuted exactly on the subspace $\mathrm{Im}(\Pi_{\leq N-1})$, and we have now shifted the weight of the operators so that they are supported on that subspace. Furthermore, while truncating the operators abruptly at level $N-1$ (by conjugating by $\Pi_{\leq N-1}$ for example) could result in a large commutator, we perform a "smooth" truncation across vectors indexed by strings of increasing length.

3 Interactive Proofs with Approximately Commuting Provers

Motivated by the rounding procedure ascertained in Theorem 1 we propose a modification of the class MIP* in which the assumption that isolated provers must perform perfectly commuting measurements is relaxed to a weaker condition of *approximately commuting* measurements.

Definition 2. *Let* $\mathrm{MIP}^*_\delta(k, c, s)$ *be the class of promise problems* (L_{yes}, L_{no}) *that can be decided by an interactive proof system in which the verifier exchanges a single round of classical messages with k quantum provers* $P_1, \ldots, P_k$ *and such that:*

- *If the input* $x \in L_{yes}$ *then there exists a perfectly commuting strategy for the provers that is accepted with probability at least c,*
- *If* $x \in L_{no}$ *then any* δ*-AC strategy is accepted with probability at most s.*

Note that the definition of MIP^*_δ requires the completeness property to be satisfied with perfectly commuting provers; indeed we would find it artificial to seek protocols for which optimal strategies in the "honest" case would be required to depart from the commutation condition. Instead, only the soundness condition is relaxed by giving *more* power to the provers, who are now allowed to apply any "approximately commuting" strategy. The "approximately" is quantified by the parameter δ,[2] and for any $\delta' \leq \delta$ the inclusions $\mathrm{MIP}^*_\delta \subseteq \mathrm{MIP}^*_{\delta'} \subseteq \mathrm{MIP}^*$ trivially hold. It is important to keep in mind that while δ can be a function of the size of the protocol it must be independent of the dimension of the provers' operators, which is unrestricted.

δ-AC strategies were previously considered by Ozawa [Oza13] in connection with Tsirelson's problem. Ozawa proposes a conjecture, the "Strong Kirchberg Conjecture (I)", which if true implies the equality $\mathrm{MIP}^* = \cup_{\delta>0} \mathrm{MIP}^*_\delta$. We state and discuss the conjecture further as Conjecture 1 below. Unfortunately the conjecture seems well beyond the reach of current techniques (Ozawa himself formulates doubts as to its validity). However, in our context less stringent formulations of the conjecture would still imply conclusive results relating MIP^*_δ to MIP*; we discuss such variants in Section 4.

Further motivation for the definition of MIP^*_δ may be found by thinking operationally — with e.g. cryptographic applications in mind, how does one ascertain that "isolated" provers indeed apply commuting measurements? The usual line of reasoning applies the laws of quantum mechanics and special relativity to derive the tensor product structure from space-time separation. However, not only is strict isolation virtually impossible to enforce in all but the simplest experimental scenarios, but the implication "separation $\Longrightarrow$ tensor product" may itself be subject to questioning — in particular it may not be a testable prediction, at least not to precision that exceeds the number of measurements, or observations, performed. Relaxations of the tensor product condition

[2] As a first approximation the reader may think of δ as a parameter that is inverse exponential in the input length $|x|$. In terms of games, this corresponds to δ being inverse polynomial in the number of questions in the game, which is arguably the most natural setting of parameters.

have been previously considered in the context of device-independent cryptography; for instance Silman et al. [SPM13] require that the joint measurement performed by two isolated devices be close, in operator norm, to a tensor product measurement. Our approximate commutation condition imposes a weaker requirement, and thus our convergence results on the hierarchy also apply to their setting; we discuss this in more detail in Section 4.2.

A computationally Tractable Class? Theorem 1 can be interpreted as evidence that the hierarchy converges at a polynomial rate to the maximum success probability for MIP^*_{ac} provers. More formally, it implies the inclusion $\mathrm{MIP}^*_\delta \subseteq \mathrm{TIME}(\exp(\exp(\mathrm{poly})/\delta^2))$ for any $\delta > 0$, thereby justifying our claim that the class MIP^*_δ is tractable. This stands in stark contrast with $\mathrm{MIP}^* = \mathrm{MIP}^*_0$, for which no upper bound is known.

Having shown that the new class has "reasonable" complexity, it is natural to ask whether the additional power granted to the provers might actually make the class trivial — could provers that are δ-AC be no more useful than a single quantum prover, even for very small δ? We show this is not the case by establishing the inclusion $\mathrm{NEXP} \subseteq \mathrm{MIP}^*_{2^{-\mathrm{poly}}}(2, 1, 1 - 2^{-\mathrm{poly}})$. This is a direct analogue of the same lower bound for MIP^* [IKM09], and is proven using the same technique. We conjecture that the inclusion $\mathrm{NEXP} \subseteq \mathrm{MIP}^*_{2^{-\mathrm{poly}}}(3, 1, 2/3)$ also holds, and that this can be derived by a careful extension of the results in [IV12, Vid13].

4 Discussion

Our introduction of MIP^*_{ac} is motivated by a desire to develop a framework for the study of quantum multiprover interactive proof systems that is both computationally tractable and relevant for typical applications of such proof systems. Our main technical result, Theorem 1, demonstrates the first aspect. In this section we discuss the relevance of the new model, its connection with the standard definition of MIP^*, and possible applications to quantum information.

4.1 Commuting Approximants: Some Results, Limits, and Possibilities

While we believe MIP^*_{ac} is of interest in itself, we do not claim that approximately commuting provers are more natural than commuting provers, or provers in tensor product form; the main goal in introducing the new class is to shed light on its thus-far-intractable parent MIP^*. In light of the results from Section 2 the relationship between the two classes seems to hinge on the general mathematical problem of finding exactly commuting approximants to approximately commuting matrices.

Limits for Commuting Approximants. The main objection to the existence of a positive answer for the "commuting approximants" question is revealed

by a beautiful construction of Voiculescu who exhibits a surprisingly simple scenario in which commuting approximants provably do not exist [Voi83]. The following is a direct consequence of Voiculescu's result.

Theorem 2 (Voiculescu). *For every $d \in \mathbb{N}$ there exists a pair of unitary matrices $U_1, U_2 \in \mathbb{C}^{d\times d}$ with $\|[U_1, U_2]\| \leq O(\frac{1}{d})$, such that for any pair of complex matrices $A, B \in \mathbb{C}^{d\times d}$ satisfying $[A, B] = 0$, $\max(\|U_1 - A\|, \|U_2 - B\|) = \Omega(1)$.*

In Voiculescu's example U_1 is a d-dimensional cyclic permutation matrix, and U_2 is a diagonal matrix whose eigenvalues are the d^{th} roots of unity. The proof draws on a connection to homology, in particular using a homotopy invariant to establish the lower bound on distance to commuting approximants. A succinct and elementary proof of the result is given by Exel and Loring [EL89].

In the context of non-local entangled strategies one is most concerned with Hermitian matrices representing measurements, rather than unitaries. However, as a consequence of Theorem 2 we see that if one considers the Hermitian operators $M_k^j = \frac{(-i)^j}{2}(U_k + (-1)^j U_k^\dagger)$ ($j \in \{0,1\}$) we have that $\|[M_1^j, M_2^{j'}]\| \leq O(\frac{1}{d})$, and yet any exactly commuting set of matrices must be a constant distance away in the operator norm. Thus Theorem 2 rules out the strongest form of a "commuting approximants" statement, which would ask for approximants in the same space as the original matrices, and with a commutator bound that does not depend on the dimension of the matrices.

Theorem 2 invites us to refine the "commuting approximants" question and distinguish the ways in which it may avoid the counter-example.

Ozawa's Conjecture. Motivated by the study of Tsirelson's problem and the relationship with Tsirelson's conjecture, Ozawa [Oza13] introduces two equivalent conjectures, the "Strong Kirchberg Conjecture (I)" and "Strong Kirchberg Conjecture (II)" respectively, which conjecture the existence of commuting approximants to approximately commuting sets of POVM measurements and unitaries respectively. The novelty of these conjectures, which allows them to avoid the immediate pitfall given by Voiculescu's example, is that Ozawa considers approximants in a larger Hilbert space than the original approximately commuting operators. Precisely, his Strong Kirchberg Conjecture (I) states the following:

Conjecture 1 (Ozawa). Let $m, \ell \geq 2$ be such that $(m, \ell) \neq (2,2)$ [3]. For every $\kappa > 0$ there exists $\varepsilon > 0$ such that, if $\dim \mathcal{H} < \infty$ and (P_i^k), (Q_j^l) is a pair of m projective ℓ-outcome POVMs on $\mathcal{H}$ satisfying $\|[P_i^k, Q_j^l]\| \leq \varepsilon$, then there is a finite-dimensional Hilbert space $\tilde{H}$ containing $\mathcal{H}$ and projective POVMs $\tilde{P}_i^k$, $\tilde{Q}_j^l$

[3] The case $(m, \ell) = (2,2)$ is the only nontrivial setting for which we have some understanding. In particular nonlocal games with two inputs and two outputs per party can be analyzed via an application of Jordan's lemma [Mas05].

on $\tilde{\mathcal{H}}$ such that $\|[\tilde{P}_i^k, \tilde{Q}_j^l]\| = 0$ and $\|\Phi_{\mathcal{H}}(\tilde{P}_i^k) - P_i^k\| \leq \kappa$ and $\|\Phi_{\mathcal{H}}(\tilde{Q}_j^l) - Q_j^l\| \leq \kappa$. Here $\Phi_{\mathcal{H}}$ denotes the compression to $\mathcal{H}$, defined by $\Phi_{\mathcal{H}}(M) \equiv P_{\mathcal{H}} M P_{\mathcal{H}}$, where $P_{\mathcal{H}}$ is the projection onto $\mathcal{H}$.

Ozawa gives an elegant proof of a variant of the conjecture that applies to just two approximately commuting unitaries, thereby establishing that extending the Hilbert space can allow one to avoid the complications in Voiculescu's example. He also establishes that the conjecture is *stronger* than Kirchberg's conjecture (itself equivalent to Tsirelson's problem and Connes' embedding conjecture), casting doubt, if not on its validity, at least on its approachability.

Nevertheless, we can mention the following facts. First, Conjecture 1 implies the equality $\mathrm{MIP}^*_{ac} = \mathrm{MIP}^*$; in fact it implies that $\mathrm{MIP}^*_{\delta} = \mathrm{MIP}^*$ for small enough δ, depending on how the parameter ε in Conjecture 1 depends on κ, m and d. For this it suffices to verify that a state ρ optimal for a strategy based on POVMs P_i^k and Q_j^l in a given protocol can be lifted to a state $\tilde{\rho}$ on $\tilde{\mathcal{H}}$ such that the correlations exhibited by performing the POVMs $\tilde{P}_i^k$, $\tilde{Q}_j^l$ on $\tilde{\rho}$ approximately reproduce those generated by P_i^k, Q_j^l on ρ; this is easily seen to be the case provided κ is small enough.

Second, Conjecture 1 can be weakened in several ways without losing the implication that $\mathrm{MIP}^*_{ac} = \mathrm{MIP}^*$. For instance, it is not necessary for the exactly commuting $\tilde{P}_i^k$, $\tilde{Q}_j^l$ to approximate the P_i^k, Q_j^l in operator norm — in the context of interactive games, only the correlations obtained by measuring a particular state need to be preserved, and this does not in general imply an approximation as strong as that promised in Conjecture 1.

Dimension Dependent Bounds. An alternative relaxation for the "commuting approximants" question is to allow the approximation error to depend explicitly on the dimension of the matrices. A careful analysis of the rounding scheme from Theorem 1 shows that it produces d-dimensional POVM elements with an $O(1/\sqrt{\log(d)})$ bound on the commutators (this is because the dimension of the subspace $\mathrm{Im}(\Pi_{\leq N-1})$ is exponential in N). Unfortunately, Voiculescu's result (Theorem 2) shows that one can only hope for good approximants in the operator norm if the commutator bound is $o(1/d)$. It remains instructive to find *any* explicit existence result for commuting approximants in the general case, regardless of dimension dependence. Concretely, we conjecture that Conjecture 1 may be true with a parameter κ that scales with the dimension d of the operators $\{P_i^k, Q_j^l\}$ as $\kappa = \varepsilon^c \,\mathrm{poly}(d)^{(ml)^2}$ for some constant $0 < c \leq 1$.

An Alternative Norm. Another relaxation of the "commuting approximants" question, which would be sufficient to imply $\mathrm{MIP}^*_{ac} = \mathrm{MIP}^*$, is to allow for any set of commuting approximants which approximately preserves the winning probability of the game. For concreteness we include a precise version of a possible statement along these lines:

Conjecture 2. There exists a function $f(\varepsilon, k) : \mathbb{R}^+ \times \mathbb{N} \to \mathbb{R}^+$ satisfying $\lim_{\varepsilon \to 0} f(\varepsilon, k) = 0$ for all $k \in \mathbb{N}$, such that for every game G and (m, ℓ) strategy (A_x^a, B_y^b, ρ) which is δ-AC, there exists a 0-AC strategy $(\tilde{A}_x^a, \tilde{B}_y^b, \rho)$ for G satisfying

$$\left|\omega^*\big(((A_x^a, B_y^b, \rho); G\big) - \omega^*\big((\tilde{A}_x^a, \tilde{B}_y^b, \rho); G\big)\right| \leq f(\delta, m\ell).$$

4.2 Device-Independent Randomness Expansion and Weak Cross-Talk

A device-independent randomness expansion (DIRE) protocol is a protocol which may be used by a classical verifier to certify that a pair of untrusted devices are producing true randomness. Under the sole assumptions that the devices do not communicate with each other, and that the verifier has access to a small initial seed of uniform randomness, the protocol allows for the generation of much larger quantities of certifiably uniform random bits; hence the term "randomness expansion". This conclusion relies only on the assumption that the two devices do not communicate, and in particular does not require any limit on the computational power of the devices, as is typically the case in the study of pseudorandomness. The precise formalization of DIRE protocols is rather involved, and we direct the interested reader to the flourishing collection of works on the topic [CK11, PAM10, MS14].

Our definition of MIP^*_{ac} is directly relevant to the notion of devices with *weak cross-talk* introduced in [SPM13] as a model which relaxes the assumption that the devices must not communicate, leading to protocols that are more robust to leakage than the traditional model of device-independence. [SPM13] proposes the use of the QCSDP hierarchy in order to optimize over the set of "weakly interacting" quantum strategies that they introduce, but no bounds are shown on the rate of convergence. This is where MIP^*_{ac} becomes relevant. Our notion of δ-AC strategies is easily seen to be a relaxation of weak cross-talk, and thus the analogue of the approach in [SPM13] when performed with a δ-AC constraint is at least as robust as the weak cross-talk approach. Our rounding scheme for the QCSDP hierarchy thus provides a specific algorithm and complexity bound that applies to both δ-AC strategies and strategies with weak cross-talk.

References

[ALM+98] Arora, S., Lund, C., Motwani, R., Sudan, M., Szegedy, M.: Proof verification and the hardness of approximation problems. J. ACM **45**(3), 501–555 (1998)

[Ara02] Aravind, P.K.: The magic squares and Bell's theorem. Technical report (2002 arXiv:quant-ph/0206070

[AS98] Arora, S., Safra, S.: Probabilistic checking of proofs: A new characterization of NP. J. ACM **45**(1), 70–122 (1998)

[Bel64] John, S.: Bell. On the Einstein-Podolsky-Rosen paradox. Physics **1**, 195–200 (1964)

[BFL91] Babai, L., Fortnow, L., Lund, C.: Non-deterministic exponential time has two-prover interactive protocols. Comput. Complexity **1**, 3–40 (1991)

[BOGKW88] Ben-Or, M., Goldwasser, S., Kilian, J., Wigderson, A.: Multi-prover interactive proofs: How to remove intractability assumptions. In: Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC), pp. 113–131 (1988)

[BSS14] Bancal, J.-D., Sheridan, L., Scarani, V.: More randomness from the same data. New Journal of Physics **16**(3), 033011 (2014)

[CHSH69] Clauser, J.F., Horne, M.A., Shimony, A., Holt, R.A.: Proposed experiment to test local hidden-variable theories. Phys. Rev. Lett. **23**, 880–884 (1969)

[CHTW04] Cleve, R., Høyer, P., Toner, B., Watrous, J.: Consequences and limits of nonlocal strategies. In: Proc. 19th IEEE Conf. on Computational Complexity (CCC 2004), pp. 236–249. IEEE Computer Society (2004)

[CK11] Colbeck, R., Kent, A.: Private randomness expansion with untrusted devices. Journal of Physics A: Mathematical and ..., 1–11 (2011)

[Con76] Connes, A.: Classification of injective factors cases ii_1, ii_∞, iii_λ, $\lambda \neq 1$. Annals of Mathematics **104**(1), 73–115 (1976)

[DLTW08] Doherty, A.C., Liang, Y-C., Toner, B., Wehner, S.: The quantum moment problem and bounds on entangled multi-prover games. In: Proc. 23rd IEEE Conf. on Computational Complexity (CCC 2008), pp. 199–210 (2008)

[EL89] Exel, R., Loring, T.: Almost commuting unitary matrices. In: Proceedings of the American Mathematical Society 106(4), 913–915 (1989)

[EPR35] Einstein, A., Podolsky, B., Rosen, N.: Can quantum-mechanical description of physical reality be considered complete? Physical Review **47**, 777–780 (1935)

[FGL+96] Feige, U., Goldwasser, S., Lovász, L., Safra, S., Szegedy, M.: Interactive proofs and the hardness of approximating cliques. J. ACM **43**(2), 268–292 (1996)

[IKM09] Ito, T., Kobayashi, H., Matsumoto, K.: Oracularization and two-prover one-round interactive proofs against nonlocal strategies. In: Proc. 24th IEEE Conf. on Computational Complexity (CCC 2009), pp. 217–228. IEEE Computer Society (2009)

[IV12] Ito, T., Vidick, T., A multi-prover interactive proof for NEXP sound against entangled provers. In: Proc. 53rd FOCS, pp. 243–252 (2012)

[JNP+11] Junge, M., Navascues, M., Palazuelos, C., Perez-Garcia, D., Scholz, V.B., Werner, R.F.: Connes' embedding problem and tsirelson's problem. J. Math. Physics 52(1) (2011)

[Kir93] Kirchberg, E.: On non-semisplit extensions, tensor products and exactness of group C^*-algebras. Inventiones mathematicae **112**(1), 449–489 (1993)

[KM03] Kobayashi, H., Matsumoto, K.: Quantum multi-prover interactive proof systems with limited prior entanglement. Journal of Computer and System Sciences **66**(3), 429–450 (2003)

[KRT10] Kempe, J., Regev, O., Toner, B.: Unique games with entangled provers are easy. SIAM J. Comput. **39**(7), 3207–3229 (2010)

[Lau03] Laurent, M.: A comparison of the Sherali-Adams, Lovász-Schrijver, and Lasserre relaxations for 0–1 Programming. Mathematics of Operations Research **28**(3), 470–496 (2003)

[Mas05] Ll. Masanes. Extremal quantum correlations for n parties with two dichotomic observables per site. Technical report (2005). arXiv:quant-ph/0512100

[MS14] Miller, C.A., Shi, Y.: Robust protocols for securely expanding randomness and distributing keys using untrusted quantum devices. In: Proc. 46th STOC. ACM New York (2014)

[NPA07] Navascués, M., Pironio, S., Acín, A.: Bounding the set of quantum correlations. Phys. Rev. Lett. **98**, 010401 (2007)

[NPA08] Navascués, M., Pironio, S., Acín, A.: A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. New Journal of Physics, 10(073013) (2008)

[Oza13] Ozawa, N.: Tsirelson's problem and asymptotically commuting unitary matrices. Journal of Mathematical Physics 54(3) (2013)

[PAM10] Pironio, S., Acín, A., Massar, S.: Random numbers certified by Bell's theorem. Nature, 1–26 (2010)

[PV10] Pál, K.F., Vértesi, T.: Maximal violation of a bipartite three-setting, two-outcome Bell inequality using infinite-dimensional quantum systems. Phys. Rev. A **82**, 022116 (2010)

[SPM13] Silman, J., Pironio, S., Massar, S.: Device-independent randomness generation in the presence of weak cross-talk. Phys. Rev. Lett. **110**, 100504 (2013)

[SW08] Scholz, V.B., Werner, R.F.: Tsirelson's problem. Technical report (2008). arXiv:0812.4305v1 [math-ph]

[Vid13] Vidick, T.: Three-player entangled XOR games are NP-hard to approximate. In: Proc. 54th FOCS (2013)

[Voi83] Voiculescu, D.: Asymptotically commuting finite rank unitary operators without commuting approximants. Acta Sci. Math. (Szeged) **45**, 429–431 (1983)

[YVB+14] Yang, T.H., Vertesi, T., Bancal, J-D., Scarani, V., Navascues, M.: Robust and Versatile Black-Box Certification of Quantum Devices. Phys. Rev. Lett. 113(4), (July 22, 2014)