Lecture 8

Multiprover interactive proof systems

In lecture 5 we considered the class $\mathrm{IP}[\mathcal{P}]$ of languages that can be decided by a BPP verifier interacting with a prover in class $\mathcal{P}$. We showed how to construct verification protocols for all of BQP in this model: informally, the result of the previous three lectures is that $\mathrm{BQP} \in \mathrm{IP}[\mathrm{BQP}]$ under the Learning with Errors assumption.[1] What if we do not wish to make computational assumptions? A first motivation for this is that we might simply not wish to rely on relatively untested assumptions — after all, isn't it likely that a few decades of algorithmic research (and even less for quantum algorithms) have barely scratched the surface of the possible ways of approaching a problem such as LWE or even e.g. factoring? A second motivation is that we could aim for more: firstly, in terms of complexity — almost by definition the class $\mathrm{IP}[\mathrm{BQP}]$ lies in BQP (and giving any more power to the prover risks breaking the computational assumption); what if we are interested in languages outside BQP?[2] Secondly, in terms of structural characterizations — in particular, remember how we stopped short of showing that the prover in the Mahadev protocol "has n qubits": can we achieve such a characterization in a different model?

In the next three lectures we switch gears and replace the use of computational assumptions by an assumption of spatial isolation. In this new model the verifier has the ability to interact with two (or more) provers that are restricted to acting locally on their respective quantum systems. This is the model that we already encountered in Section **??**. As we will see this physical (and, once properly formalized, mathematical) limitation on prover strategies will allow us to go further along the two motivating directions outlined in the preceding paragraph.

In this lecture we first introduce the model from a complexity-theoretic standpoint, discuss the recent characterization $\mathrm{MIP}^* = \mathrm{RE}$, and examine some consequences. In the following two lectures we introduce techniques that build towards a proof of the equality $\mathrm{MIP}^* = \mathrm{RE}$ by developing efficient tests for increasing numbers of qubits and increasingly complex computations.

8.1 Multiprover interactive proofs with entangled provers

We start with the main complexity-theoretic definition. Recall that a *promise language*[3] $L = (L_{yes}, L_{no})$ is specified by a pair of disjoint subsets L_{yes}, L_{no} of $\{0,1\}^*$ and that a *complexity class* is a collection of

[1] To be precise we should state what variant of the LWE assumption the result relies on.

[2] The practicality of a protocol in which the honest prover lies outside of BQP is almost entirely besides the point — our goal here is to study the problem of verification *per se*, and exploring it in the "high-complexity" regime is certain to yield useful insights which, who knows, may eventually lead to practical consequences of their own.

[3] Here the *promise* refers to the fact that it is not required that $L_{yes} \cup L_{no} = \{0,1\}^*$

languages.

Definition 8.1. The class MIP* is the class of promise languages $L = (L_{yes}, L_{no})$ such that there is a classical polynomial-time Turing machine M that on input 1^n returns the description of classical circuits for the verifier V_n in an interactive protocol with *two* quantum provers A and B such that:

- (Completeness:) There is a family of quantum provers $\{A_n, B_n\}_{n\in\mathbb{N}}$ such that for all $x \in L_{yes}$ the interaction of $V_{|x|}$ and $A_{|x|}, B_{|x|}$ on common input x accepts with probability at least $\frac{2}{3}$.
- (Soundness:) For any family of quantum provers $\{A_n, B_n\}_{n\in\mathbb{N}}$, for all $x \in L_{no}$ the interaction of $V_{|x|}$ and $A_{|x|}, B_{|x|}$ on common input x accepts with probability at most $\frac{1}{3}$.

Some comments on the definition are in order. Following tradition we called the provers A and B rather than P_1 and P_2; A stands for "Alice" and B for "Bob", a personification that is inspired from cryptography.[4] In general one may allow interaction with more than two provers; however the two-prover setting is sufficiently intersting for our purposes. (Furthermore, it can be shown that in purely complexity-theoretic terms there is no gain to considering more than 2 provers.) The number of rounds of interaction is left implicit in the definition; since V_n is polynomial-size there can be at most polynomially many rounds of interaction. Soon we will restrict ourselves to single-round protocols, which consist of a message from the verifier to each prover followed by an answer from each prover; again both for our purposes and in terms of complexity-theoretic expressive power this is without loss of generality.

Note that we did not (and will not) restrict the computational power of the provers — in fact, we did not even precisely specify what collection of strategies they may employ. For the time being we stay with the informal prescription that the provers may employ any quantum strategy that can be implemented *locally*, in *finite dimension*, and *without communication* — typically, local operations augmented with measurements on a shared entangled state that may have been agreed on prior to the protocol execution. We will see later how to formalize this more precisely.

The goal in complexity theory is to relate different classes of languages . This is especially interesting when the classes are defined in very different terms, as relations between them can provide insights into different models of computation. A pertinent example is the famous equality IP = PSPACE due to [LFKN92, Sha92]. Among the two classes, PSPACE is the simplest to define: this is the class of all languages that can be decided using a polynomial amount of space, and arbitrary time. A complete problem for PSPACE is the *quantified Boolean formula* (QBF) problem, which is to decide if a formula of the form $\exists x_1 \forall x_2 \exists x_3 \cdots (x_1 \wedge x_2 \wedge \neg x_3) \vee (\cdots)$ is satisfiable. Clearly this can be done in polynomial space by trying out all possibilities; it is also possible to show that any problem that is solvable in PSPACE can be reduced to this one, and so we say that QBF is *complete* for PSPACE. The class IP is defined very differently: it is the class of languages L such that membership $x \in L_{yes}$ can be decided efficiently by a randomized polynomial-time verifier interacting with a single infinitely powerful prover (so this is the single-prover analogue of MIP*). While it is not too hard to show that IP $\subseteq$ PSPACE, the converse inclusion is not easy at all — to see why, try coming up with a verification protocol for the QBF problem, and keep in mind that the prover is not to be trusted!

Our goal is to characterize the complexity of MIP* in terms of other complexity classes, with the hope of gaining insights about computation, entanglement, and verification of quantum devices. Before we do this let's first review what is known about the classical analogue of MIP*, in which the provers are restricted

[4] Indeed the model of multi-prover interactive proof systems is first intoduced by a team of cryptographers [BOGKW19] motivated by the development of *zero-knowledge* proof systems.

to classical strategies. This restriction affects both the completeness and soundness requirements in Definition 8.1, and so generally any stipulation of the set of allowed strategies for the provers will lead to a different complexity class.

8.1.1 Classical multiprover interactive proof systems

The * in MIP* refers to the fact that provers are allowed to use entanglement. If we omit it we get the class MIP of languages that have classical multiprover interactive proof systems. It was shown by Babai, Fortnow and Lund in the early 1990s that MIP = NEXP. This was shown shortly after the aforementioned result IP = PSPACE, which characterizes the unexpectedly large verification power of single-prover interactive proof systems.

Let's recall how MIP = NEXP is shown. The inclusion of MIP $\subseteq$ NEXP is not hard to obtain. To show it we give a non-deterministic exponential time algorithm that exactly computes the maximum acceptance probability of the verifier in an MIP protocol. This algorithm can therefore, given an instance x and a description of the verifier $V_{|x|}$, determine whether $x \in L_{yes}$ (the maximum success probability is $\geq \frac{2}{3}$) or $x \in L_{no}$ (the maximum success probability is $\leq \frac{1}{3}$), promised that one of them is the case, and thus decide any language $L \in$ MIP; thus MIP $\subseteq$ NEXP follows. To devise such an algorithm first observe that in order to do so it suffices to consider the maximum over deterministic strategies, as for any randomized strategy there is a deterministic one that succeeds with at least the same probability. Now note that a deterministic strategy is specified by a list of answers to each possible question for each of the provers. There are at most exponentially many questions because the bit representation of each question must have polynomial length (since the verifier runs in polynomial time) and similarly for answers. Finally, the success probability of a deterministic strategy can be computed exactly in exponential time simply by executing the verification procedure on each possible tuple of questions, weighted by the probability of the question being asked. Therefore, a non-deterministic algorithm can, in exponential time and space, guess an optimal strategy and compute its success probability.

The reverse inclusion, NEXP $\subseteq$ MIP, is harder. To get a hint of how it is shown, consider the problem of verifying that an exponential-size graph is 3-colorable. Formally, an instance x of this problem is specified by a pair $x = (1^n, C)$ where 1^n denotes an integer n written in unary, and C is the description of a classical circuit $C : \{0,1\}^n \to \{0,1\}^n \to \{0,1\}$. Any x that does not take this form is neither in L_{yes} nor in L_{no}, and need not be considered any further.[5] The circuit C implicitly specifies a graph $G_x = (V_x, E_x)$ with vertex set $V_x = \{0,1\}^n$ and edge set E_x such that $(i,j) \in E_x$ if and only if $C(i,j) = 1$. Then L_{yes} (resp. L_{no}) is the set of all strings x such that G_x is well-defined and is 3-colorable (resp. not 3-colorable). It is known that the language $L = (L_{yes}, L_{no})$ is complete for NEXP; intuitively this is a "scaled-up" version of the result that 3-coloring of n-vertex graphs is NP-complete. Now consider the following description for the actions of the verifier in a candidate multiprover interactive proof system for L:

1. The verifier parses its input x as $x = (1^n, C)$.
2. The verifier selects a pair of vertices (i,j) uniformly at random in $\{0,1\}^n \times \{0,1\}^n$. She sends i to Alice and j to Bob.
3. Alice and Bob each reply with a color $a, b \in \{0,1,2\}$.
4. The verifier accepts if and only if any of the following conditions hold: $C(i,j) = 0$ (there is no edge); $i = j$ and $a = b$ (same color for identical vertices); $C(i,j) = 1$ and $a \neq b$ (different colors for

[5] As usual, we consider that circuits are represented in some given, fixed manner, e.g. as a list of gates and bits that they act on.

neighboring vertices).[6]

It is clear that this protocol has completeness 1: whenever G_x is 3-colorable there is a winning strategy for the provers. Moreover, a moment's thought will reveal that if G_x is not 3-colorable then there is no perfect winning strategy; hence the maximum probability of success in this case is at most $1 - 2^{-\Omega(n)}$ (because any strategy must fail on at least one question). While this is a separation between the two cases, it is not sufficient to establish soundness, which requires that the maximum probability of success for an $x \in L_{no}$ be at most $\frac{1}{3}$.

What the proof of the inclusion NEXP $\subseteq$ MIP shows is that there is in fact a much better verifier, somewhat more involved than the one that we described here, which is such that whenever the graph is not 3-colorable then the maximum success probability is at most $\frac{1}{3}$. Achieving such a protocol essentially entails finding an efficient method that, informally, maps any graph to another graph of polynomially related size such that graphs that are 3-colorable are mapped to graphs that remain 3-colorable, but graphs that are not 3-colorable are mapped to graphs that are *very far* from 3-colorable. Achieving this can be done using advanced tools from the theory of error-correcting codes; we will not be able to say more in this lecture and refer the interested reader to e.g. [AB09].[7]

8.1.2 Interactive proof systems with entangled provers

Our focus is the class MIP*. What does the characterization MIP = NEXP say about it? Not much! The most important point to realize here is that allowing the provers to use entanglement is a double-edged sword:

First, it can affect the soundness property by allowing the provers to "cheat", meaning achieve a higher success probability. We already saw a good example of this with the Magic Square game. While this game doesn't quite look like the 3-coloring protocol we introduced in the previous section, by transforming it it is possible to come up with explicit instances of the latter that are associated with non-colorable graphs but such that there nevertheless exists a quantum strategy which succeeds with probability 1; see for example [Ji13].

As a result we are unable to transfer the lower bound NEXP $\subseteq$ MIP in a "black-box" manner, and the only trivial lower bound on MIP* is PSPACE, as clearly the verifier can ignore all but one of the provers and execute any classical IP protocol with the remaining prover. In fact it is interesting to note that such a "collapse" to IP does take place when one allows even more power to the provers, in the form of arbitrary non-signaling strategies as defined in Section **??**. Indeed it is not hard to see that the non-signaling constraints are linear, so that it is possible to write the optimal success probability of non-signaling provers in a multiprover interactive proof system as an exponential-size linear program (LP). Using that linear programs can be solved in time polynomial in their size it can be shown that the class of interactive proof systems with non-signaling provers, denoted MIP^{ns}, lies in EXP. Furthermore, if the number of provers is fixed to 2 and the number of rounds of interaction to 1 then the class "collapses" even further to PSPACE, because the associated LP can be solved more efficiently than a general LP; see [Ito10].

Second, entanglement can also affect the completeness property by increasing the power of the provers in the "honest" case. If we start with a classical protocol for a problem in NEXP this is not so interesting, because we already know that the provers have a good strategy without entanglement — we are not making

[6] One may modify this protocol by having the verifier only send pairs (i, j) such that either $i = j$ or (i, j) is an edge, since the other case is an automatic "free ride" for the provers; we gloss over this point here.

[7] Technically such a reduction is not obviously necessary, because the definition of MIP allows more complicated protocols than the 3-coloring game described here. Nevertheless, using appropriate manipulations it is possible to show that any proof of NEXP $\subseteq$ MIP does imply such a reduction.

use of the fact that they can do even better with entanglement, and indeed this fact is a new nuisance that we have to deal with in order to establish the soundness property. But what if we start from a more complex problem, that does not necessarily lie in NEXP, and attempt to design a protocol such that completeness *requires* the use of entanglement?

To see how far one might hope to go in this direction we ought to think about *upper bounds* on MIP*. Recall from the previous section that for MIP we simply enumerated over all possible strategies. In the quantum setting it is not so direct: since we do not place a priori bounds on the complexity of the provers, it is unclear what dimension one should choose in order to find an optimal strategy. If one was able to show an upper bound on the dimension that is sufficient to approach the optimal success probability (as a function of the size of the protocol) then one would automatically get a corresponding upper bound on the complexity of MIP*. However, no such bound is known! The only upper bound on MIP* is the following folklore result:

Lemma 8.2. MIP$^* \subseteq$ RE*, the set of recursively enumerable languages.*

Proof. Recall that a language $L = (L_{yes}, L_{no})$ is recursively enumerable if there exists a Turing machine such that on input x, if $x \in L_{yes}$ then the Turing machine eventually halts and accepts, whereas if $x \in L_{no}$ then the Turing machine may either halt and reject, or it may never halt.

Consider the Turing machine M that on input x specifying a verifier $V_{|x|}$ searches in increasing dimension and with increasing accuracy for a good strategy in the associated protocol. Since we have not introduced a precise formalism for strategies in MIP* protocols — we will do so for two-prover one-round protocols in Section 8.2.1 — we cannot make this too precise. At present it is sufficient to think intuitively that each prover is specified by a dimension of the Hilbert space on which they act, and for each possible question they may receive, in any round, a POVM on their space that is used to determine an answer; these POVM act on an initial quantum state that lies in the tensor product of the prover's Hilbert spaces. (Any unitary actions the provers may take can be incorporated in the POVMs.) For any given dimension d and accuracy ε the space of strategies in dimension at most d can be discretized to a finite set such that the optimum success probability over elements of that set will be within an additive ε of the optimum over all strategies in dimension at most d.

If $x \in L_{yes}$ by definition there must exist a finite dimension d and a strategy in dimension d that succeeds with probability at least (say) $\frac{2}{3} - \frac{1}{100}$; eventually, taking into account discretization errors M will identify a strategy that succeeds with probability at least $\frac{2}{3} - \frac{2}{100}$ and halt with acceptance, having successfully ruled out the case that $x \in L_{no}$. However, in case $x \in L_{no}$ the Turing machine will never find a strategy with success larger than $\frac{1}{3} + \frac{1}{100}$ (where the $\frac{1}{100}$ accounts for possible discretization errors and can be made arbitrarily small), but it will not be able to rule out the existence of such a strategy either; indeed, for all it knows such a strategy may exist in "just one more dimension". □

For a long time it was unclear where the complexity of MIP* lies, between the two "trivial" extremes of IP and RE. In 2012 Ito and the author showed that NEXP $\subseteq$ MIP* by adapting the proof of NEXP $\subseteq$ MIP by Babai et al. In the past few years better lower bounds were obtained. Quite astonishingly, in 2018 Natarajan and Wright [NW19] showed that NEEXP $\subseteq$ MIP*. One reason that this is "astonishing" is because NEEXP is a strictly (unconditionally) larger class that NEXP, and so their result established unconditionally that the presence of entanglement *increases* the verifier's ability to verify languages, even though the latter's complexity has not changed at all (it remains classical polynomial-time)! Building on this result in 2020 Ji et al. [JNV$^+$20] obtained the following characterization.

Theorem 8.3. MIP$^* =$ RE.

A complete problem for the class RE is the *halting problem*: given the description of a Turing machine M as input, does M eventually halt? What Theorem 8.3 shows is that this problem, even though it is *not decidable*, can be efficiently *verified* by asking questions to two provers sharing entanglement. In purely complexity-theoretic terms this is an extremely surprising result in and for itself; note that RE contains *any* bounded time or space complexity class — and much more. The following two lectures will be devoted to a sketch of the main arguments that go in the proof of the theorem; these arguments involve the design of tests for multiple qubits as well as delegation protocols and so we will be on familiar terrain. Aside from the complexity theory it turns out that the characterization $\mathrm{MIP}^* = \mathrm{RE}$ has some interesting consequences in the foundations of quantum mechanics as well as in the theory of operator algebras which we discuss next.

8.2 Consequences

Theorem 8.3 is related to a problem in the foundations of quantum non-locality called *Tsirelson's problem*, itself connected to a problem in the theory of von Neumann's algebra usually referred to as *Connes' Embedding Problem* (CEP). Even though they have no bearing on the remainder of the course, for motivation in this section we explain those connections. We start by (re-)introducing the language of nonlocal games that we already encountered in lecture **??** and which we will generally use to talk about multiprover interactive proof systems.

8.2.1 Nonlocal games

As we will see the proof of the "hard" part of Theorem 8.3 shows that $\mathrm{RE} \subseteq \mathrm{MIP}^*(2,1)$, where the $(2,1)$ refers to verifiers that are restricted to interacting with two provers in a single round. From now on we only only consider protocols that fall in this category. In this case once an input x has been fixed the associated verifier $V_{|x|}$ together with x itself implicitly define a two-player one-round game in the sense of lecture **??**. To be fully explicit as well as set notation, a two-player one-round game is specified by a distribution π on $\mathcal{X} \times \mathcal{Y}$, where $\mathcal{X}, \mathcal{Y}$ are finite sets of *questions*, and a predicate $R : \mathcal{X} \times \mathcal{Y} \times \mathcal{A} \times \mathcal{B} \to \{0,1\}$, for finite sets of *answers* $\mathcal{A}$ and $\mathcal{B}$, usually written as $R(a,b|x,y)$. Using this terminology the specification of the verifiers $\{V_n\}$ in a one-round two-prover interactive proof system for a language $L = L_{yes} \cup L_{no}$ is equivalent to an implicit specification of a family of games $\{G_x\}_{x \in L_{yes} \cup L_{no}}$. Explicitly, for each x we have $G_x = (\pi_x, R_x)$ where the distribution π_x is the distribution according to which the interactive proof verifier V on input x selects questions to the players and R_x denotes the decision that V makes, accept or reject, as a function of the questions sent and the answers received.

Remark 8.4. It is known that any multiprover interactive proof system that decides a language L can be "parallelized" to a single round of interaction [KKMV09]. However, this transformation in general requires the addition of a prover. It is not known how to modify a proof system with more than two provers to one with only two provers that decides the same language; it is only known, as a consequence of the inclusions $\mathrm{MIP}^* \subseteq \mathrm{RE}$ (Lemma 8.2) and $\mathrm{RE} \subseteq \mathrm{MIP}^*(2,1)$ (which follows from the proof of Theorem 8.3), that such a transformation *exists*. It is an open question whether there is a simple, or efficient, such transformation.

In general given a game game $G = (\pi, R)$, a *strategy* S for G is specified by a family of distributions $\{p(\cdot,\cdot|x,y)\}_{(x,y) \in \mathcal{X} \times \mathcal{Y}}$ on $\mathcal{A} \times \mathcal{B}$. The *success probability* of the strategy S in the game G is

$$\omega(G;S) \;=\; \sum_{x,y} \pi(x,y) \sum_{a,b} R(a,b|x,y) p(a,b|x,y) \;.$$

Informally this quantity is the average, over the referee's choice of questions and the player's probabilistic strategy, that the players provide valid answers to the referee. If one fixes a collection of possible strategies

$\mathcal{S}$ then one can define an associated *value* $\omega(G;\mathcal{S})$ for the game, which is the supremum success probability achievable using strategies $S \in \mathcal{S}$:

$$\omega(G;\mathcal{S}) \;=\; \sup_{S\in\mathcal{S}} \omega(G;S)\;.$$

For example, if $\mathcal{S}$ is the set of classical local strategies, i.e. all those families of distributions that take the form (3.3), then $\omega(G;\mathcal{S})$ is called the *classical value* of the game and is usually denoted $\omega(G)$. If $\mathcal{S}$ is the set of (tensor) quantum strategies, i.e. all those families of distributions that take the form (3.5), then $\omega(G;\mathcal{S})$ is called the *entangled value* of the game and is usually denoted $\omega^*(G)$. Explicitly,

$$\omega^*(G) \;=\; \sup_{|\psi\rangle,\{A_a^x\},\{B_b^y\}} \sum_{x,y} \pi(x,y) \sum_{a,b} R(a,b|x,y) \langle\psi| A_a^x \otimes B_b^y |\psi\rangle\;, \tag{8.1}$$

where the supremum is taken over all quantum states $|\psi\rangle \in \mathcal{H}_\mathsf{A} \otimes \mathcal{H}_\mathsf{B}$ for finite-dimensional $\mathcal{H}_\mathsf{A}$ and $\mathcal{H}_\mathsf{B}$ and collections of POVMs $\{A_a^x\}$ and $\{B_b^y\}$ on $\mathcal{H}_\mathsf{A}$ and $\mathcal{H}_\mathsf{B}$ respectively, one POVM for each question x or y to Alice or Bob respectively. Using the language of games the question of characterizing the complexity of the class $\mathrm{MIP}^*(2,1)$ boils down to determining the complexity of approximating the optimum of the optimization problem (8.1). This is because for a language $L \in \mathrm{MIP}^*(2,1)$ the problem of determining if some input $x \in L_{yes}$ reduces to evaluating the value $\omega^*(G_x)$, where G_x is the game associated with x and the verifier in a protocol for L, with good enough approximation to differentiate between the cases where $x \in L_{yes}$ ($\omega^*(G_x) \geq \frac{2}{3}$) and $x \in L_{no}$ ($\omega^*(G_x) \leq \frac{1}{3}$).[8]

8.2.2 Computing upper bounds on $\omega^*(G)$

Let's put Theorem 8.3 aside for a moment and aim instead to contradict it by devising an algorithm that approaches the maximum success probability of quantum provers sharing entanglement in a two-prover one-round interactive proof system with verifier V. Equivalently, suppose that given an explicit game G we aim to approximate the value $\omega^*(G)$ defined in (8.1). In the proof of Lemma 8.2 we already saw an algorithm, let's call it algorithm A, that returns an increasing sequence of lower bounds

$$v_1 \leq v_2 \leq \cdots \leq v_k \leq \cdots \leq \omega^*(G)$$

by enumerating strategies in increasing dimension and with increasing level of accuracy. Using the definition (8.1) of the entangled value it is clear that $v_k \to_{k\to\infty} \omega^*(G)$. To make algorithm A into an actual approximation algorithm we need to have a sense of when to stop, e.g. when can we guarantee that $|v_k - \omega^*(G)| \leq \frac{1}{100}$?[9] A natural approach is to construct a companion algorithm B that constructs a decreasing sequence of *upper* bounds

$$w_1 \geq w_2 \geq \cdots \geq w_k \geq \cdots \geq \omega^*(G)\;.$$

Given algorithms A and B consider a third algorithm C that given a game G as input runs both algorithms in an interleaved fashion, computing v_1, w_1, v_2, w_2, etc., halts whenever $|v_k - w_k| \leq \frac{1}{100}$ and returns "YES" if and only if $\frac{1}{2}(v_k + w_k) > \frac{1}{2}$. Now suppose that both (v_k) and (w_k) converge to $\omega^*(G)$. Then C

[8] The correspondence is not entirely exact because complexity is measured as a function of the input size; for MIP* protocols the input is directly x, whereas for games G we think of the input as an explicit description of the underlying distribution π and predicate R. In particular it is possible that the description length of G_x is exponential in the description length of $V_{|x|}$, since the latter only specifies π and V implicitly through a circuit that computes them.

[9] The bound $\frac{1}{100}$ is arbitrary; we want it to be small enough to guarantee that the algorithm can eventually distinguish $\omega^*(G) \geq \frac{2}{3}$ from $\omega^*(G) \leq \frac{1}{3}$, so any bound $< \frac{1}{6}$ would do.

always terminates. Moreover, if $\omega^*(G) \geq \frac{2}{3}$ then $w_k \geq \frac{2}{3}$ for all k and so the value returned is at least $\frac{1}{2}((\frac{2}{3} - \frac{1}{100}) + \frac{2}{3}) = \frac{2}{3} - \frac{1}{50} > \frac{1}{2}$, whereas if $\omega^*(G) \leq \frac{1}{3}$ it is at most $\frac{1}{2}(\frac{1}{3} + (\frac{1}{3} + \frac{1}{100})) = \frac{1}{3} + \frac{1}{50} < \frac{1}{2}$. Thus C correctly distinguishes between the two cases.

So how do we determine such a sequence of upper bounds (w_k)? A general approach to finding an upper bound on the optimum of some optimization problem is to consider *relaxations* of the problem, i.e. optimization problems whose optimum is easier to find and is guaranteed to be at least as large as the original optimum. For example, consider the following relaxation

$$\begin{aligned}\omega^*(G) &= \sup_{|\psi\rangle,\{A_a^x\},\{B_b^y\}} \sum_{x,y} \pi(x,y) \sum_{a,b} R(a,b|x,y) \langle\psi| A_a^x \otimes B_b^y |\psi\rangle \\ &\leq \sup_{|u_a^x\rangle,|v_b^y\rangle} \sum_{x,y} \pi(x,y) \sum_{a,b} R(a,b|x,y) \langle u_a^x | v_b^y \rangle \ , \end{aligned} \tag{8.2}$$

where the supremum on the second line is over all families of vectors $|u_a^x\rangle, |v_b^y\rangle \in \mathcal{H}_\mathsf{A} \otimes \mathcal{H}_\mathsf{B}$ such that for every x, the $\{|u_a^x\rangle\}_a$ are orthogonal and $\sum_a \||u_a^x\rangle\|^2 = 1$; similarly for the $|v_b^y\rangle$. The inequality (8.2) is verified by setting $|u_a^x\rangle = A_a^x \otimes \mathrm{Id}\,|\psi\rangle$ and $|v_b^y\rangle = \mathrm{Id} \otimes B_b^y |\psi\rangle$. So (8.2) is a relaxation of (8.1). What did we gain in the process? Crucially, since the objective function in (8.2) only depends on the inner products between the vectors, without loss of generality we can restrict the vectors to lie in a Hilbert space $\mathcal{H}$ such that $\dim(\mathcal{H}) \leq \min(|\mathcal{X}||\mathcal{A}|, |\mathcal{Y}||\mathcal{B}|)$; this is true even if the original $\mathcal{H}_\mathsf{A}$ and $\mathcal{H}_\mathsf{B}$ were much larger. This means that by exhaustive search we can find arbitrarily good approximations to the optimum (8.2), without having to go beyond a certain fixed dimension that is determined by the size of the game. In fact, (8.2) is an optimization problem that falls in the class of *semidefinite programs* (informally, linear optimization problems over affine sections of the positive semidefinite cone) and can be solved in time polynomial in its size (as opposed to exponential for exhaustive search).

So the optimum (8.2) *can* be determined efficiently. How useful is it, i.e. how good is the inequality (8.1) $\leq$ (8.2)? Unfortunately, in general there can be an arbitrarily large (multiplicative) gap between the two [JP11], and in particular it can be that $\omega^*(G) \leq \frac{1}{3}$ but (8.2)$\geq \frac{2}{3}$.[10] The relaxation we have devised is thus too coarse for us to obtain a good algorithm right away. But maybe we can do better? What we did so far consists in adding a vector variable to represent $A_a^x \otimes \mathrm{Id}\,|\psi\rangle$ and $\mathrm{Id} \otimes B_b^y |\psi\rangle$. Each of these can be thought of as a degree-1 monomial in the matrix variables $\{A_a^x, B_b^y\}$, evaluated on $|\psi\rangle$. Considering vectors obtained from higher-degree monomials would allow us to impose more constraints, as for example we could require that

$$\big(\langle\psi| A_a^x \otimes \mathrm{Id}\,\big) \cdot \big(A_a^x \otimes B_b^y |\psi\rangle\big) \ = \ \big(\langle\psi| \,\mathrm{Id} \otimes \mathrm{Id}\,\big) \cdot \big(A_a^x \otimes B_b^y |\psi\rangle\big) \ ,$$

due to $\{A_a^x\}_a$ being projective. It is not hard to think of other such constraints. For any integer $k \geq 1$ let's define

$$w_k \ = \ \sup_{\Gamma^{(k)} \geq 0} \sum_{x,y} \pi(x,y) \sum_{a,b} R(a,b|x,y) \Gamma^{(k)}_{xa,yb} \ , \tag{8.3}$$

where the supremum is taken over all positive semidefinite matrices $\Gamma^{(k)}$ of dimension $\binom{|\mathcal{X}||\mathcal{A}|+|\mathcal{Y}||\mathcal{B}|}{k} \times \binom{|\mathcal{X}||\mathcal{A}|+|\mathcal{Y}||\mathcal{B}|}{k}$. Here we think of the entries of $\Gamma^{(k)}$ as being labeled by sequences $(z_1, c_1), \ldots, (z_k, c_k)$ where $z_i \in \mathcal{X} \cup \mathcal{Y}$ and $c_i \in \mathcal{A} \cup \mathcal{B}$, and $\Gamma^{(k)}$ is the Gram matrix of the associated vectors

$$|u_{(z_i,c_i)_{1\leq i\leq k}}\rangle \ = \ C_{c_k}^{z_k} \cdots C_{c_1}^{z_1} |\psi\rangle \ ,$$

[10] This fact is not obvious, and constructing such "bad examples" is quite difficult. For some restricted types of games, such as XOR games or unique games, the inequality can be shown to be exact or close to exact respectively.

where $C_c^z = A_c^z \otimes \mathrm{Id}$ if $z \in \mathcal{X}$ and $c \in \mathcal{A}$, $C_c^z = \mathrm{Id} \otimes B_c^z$ if $z \in \mathcal{Y}$ and $c \in \mathcal{B}$, and $C_c^z = 0$ otherwise. In addition, we add any linear constraint on the entries of Γ that follows from the facts that $\{A_a^x\}$ and $\{B_b^y\}$ are projective measurements for all x, y, and that they act on different tensor factors and hence commute.

With this definition we can verify that $w_1 = (8.2)$; this follows since any positive semidefinite matrix Γ has a factorization as a matrix of inner produces. Moreover, $w_1 \geq w_2 \geq \cdots \geq w_k \geq \omega^*(G)$ since each successive level in the "hierarchy" consists in adding additional variables and constraints. Finally, using standard algorithms for semidefinite programs the optimization problem at the k-th level can be solved in time polynomial in its size, i.e. time $(|\mathcal{X}||\mathcal{A}| + |\mathcal{Y}||\mathcal{B}|)^{O(k)}$. Let's call Algorithm B the algorithm that on input k returns w_k.[11]

8.2.3 The commuting value and Tsirelson's problem

Unfortunately it is not the case that $w_k \to_{k\to\infty} \omega^*(G)$. Indeed, if this were the case algorithm C would return arbitrarily good approximations to $\omega^*(G)$ and thus contradict Theorem 8.3. Nevertheless, since (w_k) is non-increasing and larger than $\omega^*(G)$ the sequence must converge to some value. Interestingly, this value is a natural quantity that is referred to as the *commuting value* of the game and defined as

$$\omega^{com}(G) = \sup_{|\psi\rangle, \{A_a^x\}, \{B_b^y\}} \sum_{x,y} \pi(x,y) \sum_{a,b} R(a,b|x,y) \langle \psi | A_a^x B_b^y | \psi \rangle \ , \tag{8.4}$$

where the supremum is taken over all states $|\psi\rangle \in \mathcal{H}$ where $\mathcal{H}$ is a (possibly infinite-dimensional) separable Hilbert space and families of projective measurements $\{A_a^x\}$ and $\{B_b^y\}$ on $\mathcal{H}$ such that for all x, y, a, b, A_a^x and B_b^y commute. Since $A \otimes \mathrm{Id}$ and $\mathrm{Id} \otimes B$ always commute it always holds that $\omega^*(G) \leq \omega^{com}(G)$. The hierarchy of values (w_k) is introduced in [NPA08], where they show the following convergence result.

Lemma 8.5. *For any game G it holds that* $\lim_{k\to\infty} w_k = \omega^{com}(G)$.

Proof. First note that by definition $\omega^{com}(G) \leq \lim_{k\to\infty} w_k$, since none of the constraints imposed on the definition (8.3) of w_k makes use of the tensor product structure other than to say that $A_a^x \otimes \mathrm{Id}$ and $\mathrm{Id} \otimes B_b^y$ commute.

The remainder of the proof shows the reverse inequality. For any $k \geq 1$ fix a feasible solution $\Gamma^{(k)}$ to the optimization problem (8.3). The entries of $\Gamma^{(k)}$ are indexed by pairs of monomials m in non-commutative variables $\{A_a^x, B_b^y\}$ of degree at most k. Crucially, the constraints on the optimization problem require that (i) $\Gamma^{(k)} \geq 0$, and (ii) this matrix satisfies $\Gamma^{(k)}_{m_1,m_2} = \Gamma^{(k)}_{n_1,n_2}$ whenever both entries are well-defined and $m_1 m_2^* = n_1 n_2^*$ as monomials in $\{A_a^x, B_b^y\}$, because by definition any such constraint is imposed on the optimization problem.

For any monomial m and integer k at least as large as the degree of m let $\tau_k(m) = \Gamma^{(k)}_{m,1}$. Extend τ_k to a linear form on all non-commutative polynomials by setting $\tau_k(m) = 0$ if m has degree larger than k and extending by linearity. Since $|\tau_k| \leq 1$ for each k (this can be verified because the diagonal entries of $\Gamma^{(k)}$ are all constrained to equal 1, so using (i) all entries of $\Gamma^{(k)}$ must have modulus at most 1) by the Banach-Aleoglu theorem the sequence $(\tau_k)_{k\geq 1}$ admits a pointwise convergent subsequence $(\tau_{k_i})_{k_1 \leq k_2 \leq \ldots}$; let τ be the pointwise limit. Now crucially we observe that τ is a positive linear form. Indeed, for any polynomial

[11] Technically we need to allow B to return an approximation to w_k. Since well-behaved semidefinite programs such as (8.3) can be solved in time polynomial in their size and in the logarithm of the desired accuracy we could e.g. require that B returns an additive approximation of w_k that is within error at most 2^{-k}; this will suffice for our purposes.

$p = \sum_m \alpha_m m$ where m ranges over monomials we have

$$\begin{aligned}\tau(p^*p) &= \lim_i \ \tau_{k_i}(p^*p) \\ &= \lim_i \ \sum_{m,m'} \alpha_m^* \alpha_{m'} \ \tau_{k_i}(m^* m') \\ &= \lim_i \ \alpha^\dagger \Gamma^{(k_i)} \alpha \\ &\geq 0 \ ,\end{aligned}$$

where for the first line we used linearity of τ_{k_i}, for the second line we used the definition of τ_{k_i} (the equality holds for all i such that $k_i \geq \deg(p)$), for the third line we let $\alpha = (\alpha_m)$ and used property (ii), and for the last we used property (i).

At this point we may conclude in a single abstract step by invoking the GNS construction from C^*-algebra theory: for any positive linear functional τ on a C^*-algebra $\mathcal{A}$ there is a $*$-representation π of $\mathcal{A}$ on a Hilbert space $\mathcal{H}$ and a unit vector $|\xi\rangle \in \mathcal{H}$ such that

$$\forall a \in \mathcal{A} \ , \qquad \tau(a) \ = \ \langle \xi | \pi(a) | \xi \rangle \ . \tag{8.5}$$

For us $\mathcal{A}$ is the algebra of non-commutative polynomials in $\{A_a^x, B_b^y\}$ with complex coefficients satisfying the POVM and commutation conditions, and so the image $\tilde{A}_a^x = \pi(A_a^x)$, $\tilde{B}_b^y = \pi(B_b^y)$, together with the state $|\xi\rangle$, immediately gives us a commuting strategy for G with value $\lim_k w_k$:

$$\begin{aligned}\lim_{k\to\infty} w_k = \lim_{i\to\infty} \ w_{k_i} &= \lim_{i\to\infty} \ \sum_{x,y} \pi(x,y) \sum_{a,b} R(a,b|x,y) \Gamma^{(k_i)}_{xa,yb} \\ &= \lim_{i\to\infty} \ \sum_{x,y} \pi(x,y) \sum_{a,b} R(a,b|x,y) \Gamma^{(k_i)}_{(xa,yb)} \\ &= \lim_{i\to\infty} \ \sum_{x,y} \pi(x,y) \sum_{a,b} R(a,b|x,y) \tau_{k_i}((xa,yb)) \\ &= \sum_{x,y} \pi(x,y) \sum_{a,b} R(a,b|x,y) \tau((xa,yb)) \\ &= \sum_{x,y} \pi(x,y) \sum_{a,b} R(a,b|x,y) \langle \xi | \pi(xa,yb) | \xi \rangle \\ &= \sum_{x,y} \pi(x,y) \sum_{a,b} R(a,b|x,y) \langle \xi | \pi(xa) \pi(yb) | \xi \rangle \\ &= \sum_{x,y} \pi(x,y) \sum_{a,b} R(a,b|x,y) \langle \xi | \tilde{A}_a^x \tilde{B}_b^y | \xi \rangle \ ,\end{aligned}$$

where the first line is by definition of w_{k_i}, the second line by the linear constraints (ii), the third by definition of τ_{k_i}, the fourth by definition of τ, the fitfth by (8.5), the sixth because π is a representation and the last by definition of $\tilde{A}_a^x$ and $\tilde{B}_b^y$.

It is also possible to finish the construction more concretely by defining an infinite-dimensional matrix $\Gamma = \lim_i \Gamma^{(k_i)}$, where for the limit to make sense we embed each $\Gamma^{(k_i)}$ as the top left corner of an infinite-dimensional matrix by padding with zeroes. Since all finite minors of Γ are positive semidefinite, it is positive semidefinite and therefore admits a factorization $\Gamma_{m,m'} = \langle m | m' \rangle$ for some $\{|m\rangle\}$ in a Hilbert space $\mathcal{H}$. We can then define $\tilde{A}_a^x$ as the projection on the span of all $|m\rangle$ such that $m = A_a^x m'$ for some m', i.e. the first variable of monomial m is A_a^x. Using the relations satisfied by the inner products between the vectors $|m\rangle$ (i.e. condition (ii) above) it is possible to verify that the $\tilde{A}_a^x$ together with analogously defined

$\tilde{B}_b^y$ and $|\psi\rangle = |1\rangle$ satisfy the required conditions for a commuting strategy, and that the associated value is once again $\lim_k w_k$. □

The two values $\omega^*(G)$ and $\omega^{com}(G)$ were introduced by Tsirelson in a series of papers laying the foundations for the mathematical study of non-locality [Tsi93]. Rather than using the language of games (which at the time was not much in use yet), Tsirelson directly studied the underlying *correlation sets* defined as

$$C^*(n,k) = \Big\{ \big(\langle \psi, A_a^x \otimes B_b^y \psi\rangle\big)_{a,b,x,y} \;:\; \mathcal{H}_\mathsf{A}, \mathcal{H}_\mathsf{B} \text{ Hilbert spaces, } \psi \in \mathcal{H}_\mathsf{A} \otimes \mathcal{H}_\mathsf{B},\ \|\psi\| = 1, \\ \forall (x,y) \in \{1,\dots,n\}^2,\ \{A_a^x\}_{a\in\{1,\dots,k\}}, \{B_b^y\}_{b\in\{1,\dots,k\}} \text{ POVM on } \mathcal{H}_\mathsf{A}, \mathcal{H}_\mathsf{B} \text{ resp.} \Big\}, \quad (8.6)$$

$$C^{com}(n,k) = \Big\{ \big(\langle \psi, A_a^x B_b^y \psi\rangle\big)_{a,b,x,y} \;:\; \mathcal{H} \text{ Hilbert space, } \psi \in \mathcal{H},\ \|\psi\| = 1, \\ \forall (x,y) \in \{1,\dots,n\}^2,\ \{A_a^x\}_{a\in\{1,\dots,k\}}, \{B_b^y\}_{b\in\{1,\dots,k\}} \text{ PVM on } \mathcal{H} \\ \text{s.t. } [A_a^x, B_b^y] = 0\ \forall (a,b) \in \{1,\dots,k\}^2 \Big\}.^{12} \quad (8.7)$$

By taking direct sums of POVMs and scaled vectors it is not hard to see that both sets are convex subsets of $[0,1]^{n^2k^2}$. Note that in the definition of $C^*(n,k)$ we did not restrict the dimension of $\mathcal{H}_\mathsf{A}$ and $\mathcal{H}_\mathsf{B}$ to be finite. This is to match Tsirelson's presentation; for our purposes the distinction is not important as it is not hard to see that allowing infinite-dimensional strategies in the definition of the entangled value $\omega^*(G)$ does not change the supremum.[13][14] However, in case the Hilbert spaces in *both* definitions are taken to be finite-dimensional then the two sets can be shown to coincide. (This fact essentially follows from von Neumann's Double Commutant Theorem, though it can also be shown directly; we skip the proof.) In his paper Tsirelson states as "fact" the claim that $C^*(n,k) = C^{com}(n,k)$ for arbitrary separable Hilbert spaces and all $n,k \geq 1$. Having realized that a proof of the claim seemed elusive (with the inclusion $C^*(n,k) \subseteq C^{com}(n,k)$ that we already observed being the only obvious one), in a subsequent note[15] Tsirelson reformulates the "fact" as an open problem and, realizing that the answer may be negative, formulates as an "even more important" problem the question of whether the closure $\overline{C^*(n,k)} = C^{com}(n,k)$. (Here the overline designates closure in the usual topology for $\mathbb{R}^{n^2k^2}$. It is not hard to verify that C^{com} is closed.) Two and a half decades after its introduction Tsirelson's first problem was solved by Slofstra [Slo19], who used techniques from the theory of nonlocal games to show the existence of finite n,k such that $C^*(n,k) \neq C^{com}(n,k)$. Until the proof of Theorem 8.3, an apparently purely complexity-theoretic result, Tsirelson's "even more important problem" remained open. However, we can now observe the following corollary to Theorem 8.3.

Corollary 8.6. *There exists finite $n,k \geq 1$ such that $\overline{C^*(n,k)} \subsetneq C^{com}(n,k)$.*

Proof. Suppose for contradiction that $\overline{C^*(n,k)} = C^{com}(n,k)$ for all $n,k \geq 1$. As an immediate consequence, for any game G it holds that $\omega^*(G) = \omega^{com}(G)$. Therefore, algorithm C described in Section 8.2.2 always converges in finite time to a correct answer. This contradicts Theorem 8.3, which implies that the problem "Given a game G, is $\omega^*(G) \geq \frac{2}{3}$ or $\omega^*(G) \leq \frac{1}{3}$?" is undecidable. □

[13] To show this, observe that any state $|\psi\rangle \in \mathcal{H}_\mathsf{A} \otimes \mathcal{H}_\mathsf{B}$, even in infinite dimensions, always has a Schmidt decomposition $|\psi\rangle = \sum_i \lambda_i |u_i\rangle|u_i\rangle$ such that $\sum_i \lambda_i^2 = 1$. $|\psi\rangle$ can be arbitrarily well approximated in finite dimension by truncating the coefficients; using that the restriction of a POVM to a subspace is a POVM we find arbitrarily good approximations to the game value in finite dimension.

[14] It does change the definition of the set however: as shown in [CS18] some elements of $C^*(n,k)$ cannot be represented in finite dimensions.

[15] "Bell inequalities and operator algebras", available at `https://www.tau.ac.il/~tsirel/download/bellopalg.pdf`.

Note how indirect the proof of Lemma 8.6 is! In particular, while it asserts the existence of n, k there is no obvious way to determine what these integers are, or even upper bounds on them, from the proof. In fact it is possible to tweak the argument to get an explicit construction; we refer to [JNV$^+$20] for more.

8.2.4 Connes Embedding Problem

Quantum mechanics and the theory of operator algebras have been intertwined since their origin. In the 1930s [VN32] von Neumann laid the foundations for the theory of (what are now known as) von Neumann algebras with the explicit goal of establishing Heisenberg's matrix mechanics on a rigorous footing (quoting from the preface, in the translation by Beyer: "The object of this book is to present the new quantum mechanics in a unified representation which, so far as it is possible and useful, is mathematically rigorous"). Following the initial explorations of Murray and von Neumann the new theory took on a life of its own, eventually leading to multiple applications unrelated to quantum mechanics, such as to free probability or noncommutative geometry.

In his 1976 paper completing the classification of injective von Neumann algebras [Con76] Connes made a casual remark that has become a central conjecture in the theory of operator algebras. Since we do not have the mathematical language to express it precisely, I will paraphrase Connes' remark as the comment that "any finite von Neumann algebra *ought to* be well-approximated by finite-dimensional matrix algebras." (In more formal terms, CEP states that every von Neumann algebra type II_1 factor embeds into an ultrapower of the hyperfinite II_1 factor.) Although this conjecture may at first seem rather specific (and in fact as far as I know Connes himself did not pursue the question any further than the remark made in his paper), in the two decades that followed the problem rose to prominence thanks to the work of other mathematicians, such as Kirchberg and Voiculescu, who gave equivalent reformulations of the conjecture in operator algebras and free probability. (See e.g. [Cap15] for more reformulations.) Kirchberg's formulation is closest to us: Kirchberg showed that CEP is equivalent to the *QWEP conjecture* about the equivalence of the minimal and maximal tensor products on the full group C^* algebra of a nonabelian free group [Kir93].[16] Informally, the minimal and maximal tensor products of two $C*$ algebras provide two ways to define the closure of the algebraic tensor product, with respect to two different norms, the minimal norm

$$\|x\|_{min} = \sup_{\pi_A,\pi_B} \left\|\pi_A \otimes \pi_B(x)\right\|$$

where the supremum ranges over all pairs of representations $\pi_A : C^*(\mathbb{F}_2) \to \mathcal{B}(\mathcal{H}_\mathsf{A})$ and $\pi_B : C^*(\mathbb{F}_2) \to \mathcal{B}(\mathcal{H}_\mathsf{B})$, whereas

$$\|x\|_{max} = \sup_{\pi} \|\pi(x)\|$$

where here $\pi : C^*(\mathbb{F}_2) \otimes C^*(\mathbb{F}_2) \to \mathcal{B}(\mathcal{H})$ is any representation that is such that $\pi(a \otimes b) = \pi_A(a)\pi_B(b)$ where $\pi_A, \pi_B : C^*(\mathbb{F}_2) \to \mathcal{B}(\mathcal{H})$ are representations with commuting range. Clearly, $\|x\|_{min} \leq \|x\|_{max}$ always, and these two norms can be seen to be the smallest and largest "reasonable" norms that one may put on the tensor product of two C^*-algebras.

With this reformulation it may not be surprising that Kirchberg's QWEP is directly related to Tsirelson's problem, and indeed building on work of Fritz [Fri12] and Junge et al. [JNP$^+$11] Ozawa [Oza13a] showed that Tsirelson's "even more important" problem is equivalent to CEP. This brings us to a second corollary of Theorem 8.3.

[15] The brief discussion in this section is adapted from [Vid19].

[16] Concretely, a C^* algebra can always be represented as a sub-algebra of the algebra of bounded linear operators on a Hilbert space that is closed under taking adjoints, and closed under the norm topology. A von Neumann algebra is further restricted to be closed under the weak toperator topology.

Corollary 8.7. *CEP has a negative answer, i.e. there exists a von Neumann algebra that is not hyperfinite.*

For more background on the relation between Tsirelson's problem and Kirchberg's conjecture, presented in an accessible way, I recommend [Fri12]. For additional results and the connection to CEP, presented in a less accessible way, I recommend [Oza13b].

Bibliography

[Aar10] Scott Aaronson. Bqp and the polynomial hierarchy. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 141–150. ACM, 2010.

[AB09] Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.

[ACGK17] Scott Aaronson, Alexandru Cojocaru, Alexandru Gheorghiu, and Elham Kashefi. On the implausibility of classical client blind quantum computing. *arXiv preprint arXiv:1704.08482*, 2017.

[AG17] Dorit Aharonov and Ayal Green. A quantum inspired proof of $\mathrm{P}^{\setminus \#\mathrm{p}} \subseteq$ IP. *arXiv preprint arXiv:1710.09078*, 2017.

[AGV09] Adi Akavia, Shafi Goldwasser, and Vinod Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In *Theory of Cryptography Conference*, pages 474–495. Springer, 2009.

[AV12] Dorit Aharonov and Umesh Vazirani. Is quantum mechanics falsifiable? A computational perspective on the foundations of quantum mechanics. *arXiv preprint arXiv:1206.3686*, 2012.

[AV13] Dorit Aharonov and Umesh Vazirani. *Is quantum mechanics falsifiable? A computational perspective on the foundations of quantum mechanics*. Computability: Turing, Gödel, Church, and Beyond. MIT Press, 2013.

[BCM+18] Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 320–331. IEEE, 2018.

[Bel64] John S Bell. On the einstein podolsky rosen paradox. *Physics Physique Fizika*, 1(3):195, 1964.

[BOGKW19] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: How to remove intractability assumptions. In *Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali*, pages 373–410. 2019.

[Bre06] Frédéric Brechenmacher. *Histoire du théorème de Jordan de la décomposition matricielle (1870-1930). Formes de représentation et méthodes de décomposition.* PhD thesis, 2006.

[BV14] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) lwe. *SIAM Journal on Computing*, 43(2):831–871, 2014.

[Cap15] Valerio Capraro. *Connes' Embedding Conjecture*, pages 73–107. Springer International Publishing, Cham, 2015.

[CCKW19] Alexandru Cojocaru, Léo Colisson, Elham Kashefi, and Petros Wallden. Qfactory: classically-instructed remote secret qubits preparation. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 615–645. Springer, 2019.

[CLS17] Richard Cleve, Li Liu, and William Slofstra. Perfect commuting-operator strategies for linear system games. *Journal of Mathematical Physics*, 58(1):012202, 2017.

[CM14] Richard Cleve and Rajat Mittal. Characterization of binary constraint system games. In *International Colloquium on Automata, Languages, and Programming*, pages 320–331. Springer, 2014.

[CM16] Toby Cubitt and Ashley Montanaro. Complexity classification of local hamiltonian problems. *SIAM Journal on Computing*, 45(2):268–316, 2016.

[Con76] Alain Connes. Classification of injective factors cases II_1, II_∞, III_λ, $\lambda \neq 1$. *Annals of Mathematics*, pages 73–115, 1976.

[CR20] Rui Chao and Ben W Reichardt. Quantum dimension test using the uncertainty principle. *arXiv preprint arXiv:2002.12432*, 2020.

[CRSV17] Rui Chao, Ben W Reichardt, Chris Sutherland, and Thomas Vidick. Overlapping qubits. *arXiv preprint arXiv:1701.01062*, 2017.

[CRSV18] Rui Chao, Ben W Reichardt, Chris Sutherland, and Thomas Vidick. Test for a large amount of entanglement, using few measurements. *Quantum*, 2:92, 2018.

[CS17] Andrea Coladangelo and Jalex Stark. Robust self-testing for linear constraint system games. *arXiv preprint arXiv:1709.09267*, 2017.

[CS18] Andrea Coladangelo and Jalex Stark. Unconditional separation of finite and infinite-dimensional quantum correlations. *arXiv preprint arXiv:1804.05116*, 2018.

[DFPR14] Vedran Dunjko, Joseph F Fitzsimons, Christopher Portmann, and Renato Renner. Composable security of delegated quantum computation. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 406–425. Springer, 2014.

[EPR35] Albert Einstein, Boris Podolsky, and Nathan Rosen. Can quantum-mechanical description of physical reality be considered complete? *Physical review*, 47(10):777, 1935.

[Fri12] Tobias Fritz. Tsirelson's problem and Kirchberg's conjecture. *Reviews in Mathematical Physics*, 24(05):1250012, 2012.

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, volume 9, pages 169–178, 2009.

[GKK19] Alexandru Gheorghiu, Theodoros Kapourniotis, and Elham Kashefi. Verification of quantum computation: An overview of existing approaches. *Theory of computing systems*, 63(4):715–808, 2019.

[GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N Rothblum. Delegating computation: interactive proofs for muggles. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 113–122. ACM, 2008.

[GKW17] Rishab Goyal, Venkata Koppula, and Brent Waters. Lockable obfuscation. In *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, pages 612–621. IEEE, 2017.

[GKW18] Rishab Goyal, Venkata Koppula, and Brent Waters. Collusion resistant traitor tracing from learning with errors. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 660–670. ACM, 2018.

[GMR85] Shafi Goldwasser, Silvio Micali, and Ronald Rivest. A "paradoxical" solution to the signature problem. In *Advances in Cryptology*, pages 467–467. Springer, 1985.

[GV19] Alexandru Gheorghiu and Thomas Vidick. Computationally-secure and composable remote state preparation. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1024–1033. IEEE, 2019.

[GVW01] Oded Goldreich, Salil Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. In *International Colloquium on Automata, Languages, and Programming*, pages 334–345. Springer, 2001.

[GVW15] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. *Journal of the ACM (JACM)*, 62(6):45, 2015.

[Ito10] Tsuyoshi Ito. Polynomial-space approximation of no-signaling provers. In *International Colloquium on Automata, Languages, and Programming*, pages 140–151. Springer, 2010.

[Ji13] Zhengfeng Ji. Binary constraint system games and locally commutative reductions. *arXiv preprint arXiv:1310.3794*, 2013.

[JNP+11] Marius Junge, Miguel Navascues, Carlos Palazuelos, David Perez-Garcia, Volkher B Scholz, and Reinhard F Werner. Connes' embedding problem and Tsirelson's problem. *Journal of Mathematical Physics*, 52(1):012102, 2011.

[JNV+20] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP*= RE. *arXiv preprint arXiv:2001.04383*, 2020.

[JP11] Marius Junge and Carlos Palazuelos. Large violation of bell inequalities with low entanglement. *Communications in Mathematical Physics*, 306(3):695, 2011.

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

[KKMV09] Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, and Thomas Vidick. Using entanglement in quantum multi-prover interactive proofs. *Computational Complexity*, 18(2):273–307, 2009.

[KMW17] Elham Kashefi, Luka Music, and Petros Wallden. The quantum cut-and-choose technique and quantum two-party computation. *arXiv preprint arXiv:1703.03754*, 2017.

[KRR14] Yael Tauman Kalai, Ran Raz, and Ron D Rothblum. How to delegate computations: the power of no-signaling proofs. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 485–494. ACM, 2014.

[LFKN92] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM (JACM)*, 39(4):859–868, 1992.

[Mah18] Urmila Mahadev. Classical verification of quantum computations. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 259–267. IEEE, 2018.

[Mer90] N David Mermin. Simple unified form for the major no-hidden-variables theorems. *Physical review letters*, 65(27):3373, 1990.

[Mer93] N David Mermin. Hidden variables and the two theorems of john bell. *Reviews of Modern Physics*, 65(3):803, 1993.

[MF16] Tomoyuki Morimae and Joseph F Fitzsimons. Post hoc verification with a single prover. *arXiv preprint arXiv:1603.06046*, 2016.

[Mor18] Tomoyuki Morimae. Blind quantum computing can always be made verifiable. *arXiv preprint arXiv:1803.06624*, 2018.

[MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 700–718. Springer, 2012.

[MYS12] Matthew McKague, Tzyh Haur Yang, and Valerio Scarani. Robust self-testing of the singlet. *Journal of Physics A: Mathematical and Theoretical*, 45(45):455304, 2012.

[NC02] Michael A Nielsen and Isaac Chuang. Quantum computation and quantum information, 2002.

[NPA08] Miguel Navascués, Stefano Pironio, and Antonio Acín. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New Journal of Physics*, 10(7):073013, 2008.

[NW19] Anand Natarajan and John Wright. Neexp is contained in mip. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 510–518. IEEE, 2019.

[Oza13a] Narutaka Ozawa. About the Connes embedding conjecture. *Japanese Journal of Mathematics*, 8(1):147–183, 2013.

[Oza13b] Narutaka Ozawa. About the connes embedding conjecture. *Japanese Journal of Mathematics*, 8(1):147–183, 2013.

[P$^+$16] Chris Peikert et al. A decade of lattice cryptography. *Foundations and Trends® in Theoretical Computer Science*, 10(4):283–424, 2016.

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):34, 2009.

[RRR16] Omer Reingold, Guy N Rothblum, and Ron D Rothblum. Constant-round interactive proofs for delegating computation. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 49–62. ACM, 2016.

[RT19] Ran Raz and Avishay Tal. Oracle separation of BQP and PH. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 13–23, 2019.

[RUV13] Ben W Reichardt, Falk Unger, and Umesh Vazirani. Classical command of quantum systems. *Nature*, 496(7446):456–460, 2013.

[Sha92] Adi Shamir. IP= PSPACE. *Journal of the ACM (JACM)*, 39(4):869–877, 1992.

[Slo19] William Slofstra. The set of quantum correlations is not closed. In *Forum of Mathematics, Pi*, volume 7. Cambridge University Press, 2019.

[SW87] Stephen J Summers and Reinhard Werner. Maximal violation of bell's inequalities is generic in quantum field theory. *Communications in Mathematical Physics*, 110(2):247–259, 1987.

[Tsi93] Boris S Tsirelson. Some results and problems on quantum Bell-type inequalities. *Hadronic Journal Supplement*, 8(4):329–345, 1993.

[Unr16] Dominique Unruh. Computationally binding quantum commitments. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 497–527. Springer, 2016.

[Vid19] Thomas Vidick. From operator algebras to complexity theory and back. *Notices of the American Mathematical Society*, 66(10), 2019.

[Vid20] Thomas Vidick. Verifying quantum computations at scale: A cryptographic leash on quantum devices. *Bulletin of the American Mathematical Society*, 57(1):39–76, 2020.

[VN32] J Von Neumann. Mathematische grundlagen der quantenmechanik. 1932.

[VW16] Thomas Vidick and John Watrous. Quantum proofs. *Foundations and Trends® in Theoretical Computer Science*, 11(1-2):1–215, 2016.

[VZ20] Thomas Vidick and Tina Zhang. Classical proofs of quantum knowledge. *arXiv preprint arXiv:2005.01691*, 2020.

[WBMS16] Xingyao Wu, Jean-Daniel Bancal, Matthew McKague, and Valerio Scarani. Device-independent parallel self-testing of two singlets. *Physical Review A*, 93(6):062121, 2016.

[WZ17] Daniel Wichs and Giorgos Zirdelis. Obfuscating compute-and-compare programs under LWE. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 600–611. IEEE, 2017.