UCSD Summer school notes

Multi-prover interactive proofs and the quantum PCP conjectures

1 Introduction

In a previous lecture we introduced the quantum "Hamiltonian" PCP conjecture, which is a conjecture on the hardness of approximating the smallest eigenvalue ("ground state energy") of a local Hamiltonian. The conjecture is inspired by the classical PCP theorem, which makes an analogous statement about the hardness of approximating the largest possible fraction of clauses that can be satisfied in an instance of a classical constraint satisfaction problem.

Readings. Chapter 6 of the survey [VW$^+$16] discusses quantum multi-prover interactive proof systems. The survey [AAV13] discusses the quantum PCP conjecture(s).

1.1 Three formulations of the PCP theorem

The PCP theorem has three main formulations. The formulations are equivalent, but they reflect three different points of view that one can take on the theorem.

(i) The first formulation justifies the theorem's acronym, PCP for "probabilistically checkable proof": all languages in NP have polynomial-size proofs of membership that can be verified by an efficient (polynomial-time) probabilistic procedure that only queries a constant number of locations of the proof in any execution. Such a procedure necessarily has a small chance of making the wrong decision, and we require that all valid statements (i.e. $x \in L$) have proofs that are accepted with probability at least $\frac{2}{3}$, while for any wrong statement (i.e. $x \notin L$), and any purported proof, the proof will be accepted by the verification procedure with probabiliy at most $\frac{1}{3}$.

(ii) A second formulation is directly in terms of hardness of approximation for constraint satisfaction problems. A particularly refined result, due to Hastad [Hås01], states that for any $\varepsilon > 0$, the following problem is NP-complete: given a 3SAT formula φ, distinguish between φ being satisfiable and at most a fraction $(\frac{7}{8} + \varepsilon)$ of clauses of φ being satisfiable. Another result of the same flavor due to Hastad is that, for all small enough $\varepsilon > 0$, the following problem is NP-complete: given a set of bipartite parity equations of the form $x_i \oplus y_j = b_{ij}$ over variables $\{x_i, y_j | i, j \in \{1, \ldots, n\}\}$, for some constants $b_{ij} \in \{0, 1\}$, distinguish between the cases when there exists an assignment satisfying a fraction at least $\frac{12}{16}$ of the equations, or no assignment satisifies more than a fraction $\frac{11}{16} + \varepsilon$ of the clauses.

(iii) The third formulation is in terms of interactive proof systems, or games.[1] A first statement, in fact weaker than the PCP theorem (and proven before it) is the characterization MIP = NEXP due to Babai et al. [BFL90]. The formulation that is equivalent to the PCP theorem is in terms of multiplayer games. It states that the following problem is NP-complete: given an explicit description of the referee in a 2-player 1-round game, decide between the cases where there is a strategy for the players that has success probability at least $\frac{2}{3}$ in the game, or no strategy for the players can succeed with probability larger than $\frac{1}{3}$.

The equivalence between the first two formulations is immediate: to go from (i) to (ii), treat proof entries as variables, and each of the verifier's possible tests (for each possible fixing of his random bits) as a constraint. Conversely, given a 3SAT formula that is hard to approximate, the verifier can check a purported assignment with low error probability by looking up the variables that are part of just one clause chosen at random. (Error amplification can be performed in a straightforward way.) The implication from (iii) to (i) is also direct, as the proof can list a strategy for each player, and the verification procedure similate the referee in the game. The implication from (ii) to (iii) uses the technique of oracularization, that we will review a little later.

1.2 Quantum formulations

Each of the three formulations for the classical PCP theorem can be turned into a "quantum PCP conjecture". You already saw the quantum "constraint satisfaction" (or, "local Hamiltonian") formulation, that is analogue to (ii). The proof-checking formulation (i) also has a natural quantum analogue, in terms of a quantum verification procedure performing a measurement on a constant number of qubits of a quantum proof. These two quantum formulations can be shown equivalent in a similar way as in the classical case, though one has to be a little more careful [AALV09]. In particular, the reduction becomes a quantum reduction, and it is open whether there could exist a classical reduction between the two formulations.

Deciding what is the right formulation for the quantum analogue of (iii) is more delicate. A first question is what kind of games one should consider. To discuss known results that will help us make the right conjecture, we'll adopt the language of interactive proofs for a bit, as this is how many results are more easily formulated. A first modification then would be to allow the verifier to be a quantum polynomial-time procedure, and to exchange qubits with the provers. However, it has been known ever since the corresponding class, QMIP, was introduced by Kobayashi and Matsumoto [KM03], that this does not lead to an interesting class: indeed, MIP $\subseteq$ QMIP is trivial, and QMIP $\subseteq$ NEXP can be shown by arguing that the quantum provers do not need more than a polynomial number of qubits of private space, so that optimal actions for the provers can be guessed in non-deterministic exponential time. As a result, QMIP = MIP: nothing particularly exciting there.

A more interesting modification consists in allowing the provers to share entanglement. Depending on whether the verifier remains classical or is allowed to be quantum as well, this leads to the class MIP* or QMIP*. Luckily for us, it was shown by Reichardt et al. [RUV12] that QMIP* = MIP*, so that we can focus our attention on either class. The class MIP* with a classical verifier is convenient because we can also talk about the scaled-down version, multiplayer games in which the (classical) verifier is specified explicitly, but for which we allow strategies for the players that use an arbitrary entangled state. A natural conjecture

[1] The language of interactive proof systems is generally used to describe protocols where the verifier's actions are specified implicitly, e.g. as a family of polynomial-size circuits. The language of multiplayer games is generally used to describe interactions where the verifier's actions are specified explicitly, e.g. as a truth table listing all possible questions, answers, and whether the answer is valid or not.

would then be that $\mathrm{QMIP}^* = \mathrm{QMA}_{\mathrm{EXP}}$, where $\mathrm{QMA}_{\mathrm{EXP}}$ is the quantum-proof analogue of NEXP. Alternatively, we could conjecture that approximating the quantum (entangled) value of a multiplayer quantum game is QMA-hard. Let's do it:

Conjecture 1 (Quantum games PCP)**.** *Approximating the maximum success probability of entangled players in a multiplayer game specified in explicit form is* QMA*-hard.*

However, it is not a priori clear that this is the right conjecture! Note that, for the case of the local Hamiltonian problem, two inclusions are given without any work: first, constant-factor approximations for the local Hamiltonian problem are at least as hard as NP (since 3SAT is a special case), and second, they are no harder than QMA (since even the problem for inverse-polynomial approximation is in QMA).

Amazingly, neither inclusion holds a priori for the problem of approximating the value of entangled-player games. First, due to the introduction of entanglement, the problem could have become *easier*: after all, if the players always have a perfect winning strategy, then the problem is trivial (in BPP!). Second, there is no easy upper bound, because we do not know how much entanglement can be useful to the players. In fact, the problem of deciding between the value being exactly 1, or strictly less than 1, has recently shown to be undecidable [Slo17] (whereas the analogous problem classically is still in NP). Although this result no longer goes through once one considers approximations, it suggests that the class MIP^* may be more powerful than basic intuition suggests — a fact to keep in mind.

Nevertheless, for various reasons, some of which will become clear later, it is not unreasonable to take Conjecture 1 as our quantum analogue to formulation (iii) of the PCP theorem. Multiple questions immediately arise: Can we give evidence for the conjecture? For example, can we show that the approximation problem is at least NP-hard? Or that it is QMA-hard for inverse-polynomial approximation (this would be an analogue of the Cook-Levin theorem stating QMA-hardness of inverse-polynomial approximations for the local Hamiltonian problem)? Can we relate Conjecture 1 to the other quantum formulations of the PCP conjecture? Note that the fact that the local Hamiltonian problem is trivially NP-hard, but the same statement for multiplayer entangled games is non-obvious, indicates that any reduction between the two formulations is a non-trivial statement.

The goal of this lecture is to explore these questions. In Section 2 we formally introduce entangled-prover interactive proof systems, and make a few observations about them. In Section 3 we introduce a "toy" version of the PCP theorem, the so-called "exponential-size PCP" based on the Blum-Luby-Runbinfeld linearity test. In Section 4 we introduce, and sketch a proof of, a quantum analogue of this result — an "exponential quantum PCP". This will also give us the opportunity to discuss the relation between Conjecture 1 and the "regular" Hamiltonian QPCP conjecture. Finally, in Section 5 we briefly consider applications to the problem of delegated computation.

2 Multi-prover interactive proofs with entangled provers

We review a few basic facts about classical multi-prover interactive proof systems, and discuss their quantum counterparts.

Definition 2. Let k, r be polynomially bounded functions of n, and $c > s$ functions of n such that $c - s > \mathrm{poly}^{-1}(n)$. The class $(\mathrm{Q})\mathrm{MIP}^*_{c,s}(k,r)$ is the class of promise languages $L = L_{yes} \cup L_{no}$ such that for every n there is a polynomial-size (quantum) verifier in a k-round interactive proof system with r provers such that, if $x \in L_{yes}$ there is a strategy for the provers that is accepted with probability at least c, and if $x \in L_{no}$ then no strategy for the provers is accepted with probability greater than s.

The * in the definition refers to the fact that provers are allowed to use entanglement. The Q in QMIP refers to the use of quantum messages. If neither is present, we get MIP, the class of languages that have classical multi-prover interactive proof systems. It is known that $\text{MIP} = \text{MIP}_{1,2^{-n}}(2,1) = \text{NEXP}$: multi-prover interactive proof systems can be parallelized to a single-round interaction with two provers only, and can be made to have completeness 1 and soundness that is exponentially small. Most of these facts have easy proofs. The completeness 1 requires a little more astuteness, but it is not too hard. The soundness amplification is much more technical, as achieving it while keeping the number of rounds to 1 requires a parallel repetition theorem. Of course, while the inclusion $\text{MIP} \subseteq \text{NEXP}$ is immediate, the converse inclusion $\text{NEXP} \subseteq \text{MIP}$ requires much more work. A key ingredient in the proof is an encoding of a satisfying assignment as a multilinear function, a simpler version of which we will discuss later.

What about the quantum case? We can add the Q, or the *, or both. As already mentioned, adding the Q only does not change the class. This is not only the case for multiprover interactive proofs, but even with a single prover: $\text{QIP} = \text{IP} = \text{PSPACE}$ [JJUW11]. While $\text{QMIP} = \text{MIP}$ is not too hard, $\text{QIP} = \text{IP}$ requires more work, and in fact we know of no direct proof of this fact: instead, the proof in [JJUW11] argues that $\text{QIP} \in \text{PSPACE}$ by devising a clever algorithm for deciding the maximum acceptance probability in a protocol; since $\text{PSPACE} \in \text{IP}$ (itself a highly non-trivial statement!) the equality $\text{QIP} = \text{IP}$ follows using the trivial inclusion $\text{IP} \subseteq \text{QIP}$.

As already mentioned as well, adding just the * or both the * and the Q yields the same class, $\text{QMIP}^* = \text{MIP}^*$ — though once again, this requires a highly non-trivial proof [RUV12], which employs many of the "rigidity" techniques we discussed in a previous lecture. So this is the one class we would like to study, $\text{MIP}^*_{c,s}(k,r)$. In terms of structural properties, parallelization to a single round is known, but only if one allows one extra prover, and moreover quantum messages [KKMV09] (the only known proof of the inclusion $\text{QMIP}^* \subseteq \text{MIP}^*$ requires a protocol with polynomially many rounds). In terms of number of provers, little is known, and it could be the case that $\text{MIP}^*(2,\text{poly}) \subsetneq \text{MIP}^*(3,\text{poly}) \subsetneq \cdots$. We do know that perfect completeness can always be achieved [KKMV09], and that parallel repetition holds with sufficient generality to amplify soundness: $\text{MIP}^*_{1,s}(k,1) = \text{MIP}^*_{1,2^{-n}}(k,1)$ for any s bounded away from 1 by an inverse polynomial.

The question we are most interested in, of course, is the question of the complexity of the class. The most important point to realize here is that allowing the provers to use entanglement is a double-edged sword. First, it can allow them to "cheat", meaning have a higher success probability. And this does not only arise in contrived examples. A good example of this is the Magic Square game. This game can be thought of as the "clause-vs-variable" game for a simple instance of MAX-3XOR, with 9 variables and 6 equations. The instance is unsatisfiable, so classical provers cannot win with certainty. But quantum provers can. As a result 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.

The example of XOR games demonstrates that this is not an isolated phenomenon: indeed, entanglement can "collapse" complexity classes! Define $\oplus\text{MIP}_{c,s}$ as the class of languages having 2-prover 1-round interactive proof systems, where the provers answer a single bit, and the verifier bases his decision on the parity of the two answer bits only. Then it is known that with classical provers, $\oplus\text{MIP}_{c,s} = \text{NEXP}$, for some choice of constants $s < c$ [Hås01]. But in the case of provers sharing entanglement, $\oplus\text{MIP}^* \subseteq \text{PSPACE}$! The inclusion of the class in EXP follows from the formulation of the quantum value of an XOR game as a semidefinite program (SDP) that we saw earlier. And in fact, the specific kind of SDP that arises can be solved even more efficiently, in PSPACE.

Remark 3. One way to understand this "collapse" is by observing that it happens in an even more systematic way when one generalizes to "non-signaling strategies". We will not define this formally here, but intuitively

a non-signaling strategy is any strategy for the provers that respects the non-signaling principle, i.e. is such that the marginal distribution of answers for any set of players is independent from the questions to the other players. It is not hard to see that the non-signaling constraints are linear, thus the optimal value of a multi-player non-signaling game can be expressed as a linear program (LP). Therefore the corresponding class $\mathrm{MIP}^{ns}(k,r) \subseteq \mathrm{EXP}$. Furthermore, if $k = 2$ and $r = 1$ (one round of interaction with two players), the LP can be solved more efficiently, so $\mathrm{MIP}^{ns}(1,2) \subseteq \mathrm{PSPACE}$ [Ito10].

But second, entanglement also increases 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 don't care that they can do even better with entanglement. But what if we do not start from a classical protocol — what if we start from a quantum problem, and design a protocol such that completeness *only* holds by using entanglement? What is amazing is that at the moment, the sky is the limit! There are no reasonable (or even unreasonable) upper bounds on MIP^*.

Open Question 4. Show that languages in MIP^* are decidable.

As mentioned earlier, it is known that the problem of deciding between cases when there exists a perfect strategy, or there does not, is undecidable. A possible approach to resolving the above question in the affirmative would be to prove the validity of Connes' embedding conjecture in the theory of von Neumann algebras. Let's hope there might be an easier way.

3 Exponential-size PCPs for NP

We will not have time to go over the more advanced lower bounds on the class MIP^*. We will introduce some of the most important ideas by considering the proof of a "toy" version of the PCP theorem, the so-called "exponential-size PCP". Consider the following theorem.

Theorem 5 (Exponential PCP)**.** *For any $L \in \mathrm{NP}$ there is a probabilistic verification procedure V such that V flips a polynomial number of random coins, looks up a constant number of locations (depending on the coin flips) of a proof π, and decides to accept or reject. V makes the correct decision with probability at least $\frac{2}{3}$.*

Note that, even though the proof is pretty long, the theorem is not trivial, because the amount of "information" that V gets from π is only constant. Note also that, as usual, the proof is not trusted, and it is not at all obvious at first that a proof, as long as it may be, can be "checked" with good probability while only making a constant number of queries to it.

Since our goal is to formulate quantum analogues of classical statements about multiplayer games, we can first reformulate Theorem 5 as a result about games.

Theorem 6 (Exponential PCP)**.** *For any $L \in \mathrm{NP}$ there is an efficient procedure $x \mapsto V_x$, where V_x is a description for a verifier in a two-player one-round game with questions of $\mathrm{poly}(n)$ bits and answers of $O(1)$ bits, such that if $x \in L$ then $\omega(V_x) \geq \frac{2}{3}$ (i.e. there is a good strategy for the players), whereas if $x \notin L$ then $\omega(V_x) \leq \frac{1}{3}$ (there is no good strategy for the players).*

The equivalence between Theorem 5 and Theorem 6 is shown using the technique of oracularization, as discussed in the introduction.

How do we prove Theorem 6? We have to start by fixing an NP-complete language. To simplify the presentation we will cheat a little bit, and make the following (provably false) assumption: given a system

of 3XOR equations, it is NP-hard to determine whether all equations are satisfiable. Of course, this is not actually a hard problem (Gaussian elimination!), and in "real life" one needs to consider a problem that is actually hard; the problem that is most often used is similar but involves quadratic equations.

The advantage of this formulation is that amplification is very easy. In particular, assume that you were given access to an exponentially long "proof" π, indexed by strings $u \in \{0,1\}^n$, such that the u-th entry of the proof contains the bit $u \cdot x$, where x is a claimed satisfying assignment to φ. Then I claim that such a proof could be easily verified, with high probability, using a single uniformly distributed query. How? Write each equation present in φ in the form $v_j \cdot x = b_j$, where $v_j \in \{0,1\}^n$ has Hamming weight exactly 3, and $b_j \in \{0,1\}$. The verifier chooses a random vector $w \in \{0,1\}^m$ and forms the equation $v = \sum_j w_j v_j \bmod 2$. If x satisfies all the equations, then $v \cdot x = \sum w_j b_j$. If, however, x does not satisfy just one (or more) of the equations, then the probability, over the choice of w, that $v \cdot x = \sum w_j b_j$ is exactly $\frac{1}{2}$. Thus the verifier can determine which of the two cases hold by looking up a single uniformly distributed entry of the proof.

This is a good start, but how do we construct a game out of all this? The construction has two steps:

(i) Distribute the proof among the players, in a way that makes it possible to "look up" any particular entry by sending an appropriate question to each player;

(ii) Develop a scheme to verify that the players answer question in (i) according to a proof that is properly encoded.

The first point is pretty easy: we simply require each player to obtain a separate copy of the proof $\pi = (\pi_v)_{v \in \{0,1\}^n}$. If we want to look up a specific entry v, we can ask either player. We can even ask both, and make sure they give the same answer, so we know that they have the same proof in mind.

The second point is more delicate. In general each player's answer to questions of the form $v \in \{0,1\}^n$ is a bit π_v (by the above, we can assume it's the same bit from each player), but there is no guarantee that there is a hidden x such that $\pi_v = v \cdot x$ for each v. Note that a small slack we can allow ourselves, is that we only need this to hold for "most" v: the reason is that in the scheme for verifying equations described earlier, the final question v is distributed uniformly at random.

Checking the right format is precisely the goal of the linearity test. Here is a standard formulation.

Theorem 7 (BLR linearity test)**.** *Let $\pi \in \{0,1\}^n$ be such that on average over uniformly random $u, v \in \{0,1\}^n$, $\pi_u + \pi_v = \pi_{u+v} \bmod 2$ with probability at least $1 - \varepsilon$, for some $\varepsilon \geq 0$. Then there exists an $x \in \{0,1\}^n$ such that $\pi_v = v \cdot x$ with probability at least $1 - O(\varepsilon)$, over the choice of a uniformly random v.*

The theorem gives us precisely the kind of test we want. The verifier in the game can pick u, v uniformly at random. It would like to query locations u, v and $w = u + v$ in the proof, but it only has access to two players. We could add a third player and be done. Or, we can ask both u and v to the first player, and $u + v$ to the second. To make sure the players don't cheat, a trick is that some fraction of the time, we still send u and v to the first player, but we send u *or* v to the second. The second player has a single question, so we know it replies π_u or π_v. This way we can cross-check the first player's answers, and make sure that the answer to u (resp. v), on question (u, v), only depends on u and not on v (resp. only on v and not on u).

Putting everything together, and modulo our silly assumption on linear equations being NP-hard, we have proven Theorem 6. To prove the real theorem, we need to consider another language, that involves quadratic, instead of linear equations. The main idea is the same, but the details are a bit more tedious, so we'll skip the actual implementation.

4 Exponential-size QPCPs for QMA

Here is a weaker formulation of Conjecture 1, a quantum extension of the exponentially-large PCP considered in the previous section.

Theorem 8. *For every language in* QMA *there is an efficient procedure $x \mapsto V_x$, where V_x is a description for a verifier in a multi-player one-round game with questions of* $\mathrm{poly}(n)$ *bits and answers of* $O(1)$ *bits, such that if $x \in L$ then $\omega^*(V_x) \geq \frac{2}{3}$, whereas if $x \notin L$ then $\omega^*(V_x) \leq \frac{1}{3}$.*

The main difference with Theorem 6 is the consideration of the entangled value, instead of the classical value, of the game, and the claim of QMA-hardness, not NP-hardness. Note also the mention of a "multi-player" game. It is known how to prove the theorem with a game with 5 players, but it is open if one can get it down to 2 players only.

Note that in a sense this result "trivially" follows from known results. Indeed, it is known that QMA $\subseteq$ PSPACE $=$ IP, and it is not hard to parallelize an IP protocol to a single round of interaction with multiple provers. However, if the players are allowed to use entanglement this may no longer be sound. In addition, to obtain constant answer size one would have to apply a technique similar to the classical linearity test. So, it wouldn't be that trivial! More importantly, in the protocol from the theorem the players can be implemented in BQP, if they are given access to the right encoding of the quantum witness for the QMA-complete problem; this is in contrast to going through the proof of IP $=$ PSPACE, which requires PSPACE-complete provers. This is important for applications to delegation that we will discuss later.

We won't give a complete proof of the theorem here; you can find details in [NV16]. The structure is the same as in the classical case. What we need to figure out are:

(i) A method for distributing a *quantum* witness among the players, in a way that makes it possible for the verifier to "check" the witness, provided it is correctly formatted, by asking "simple" questions, with constant-size answers;

(ii) A scheme to verify that the players share a valid encoding of a witness.

Note that contrary to the classical case, where (i) is achieved just by giving each player a copy of the witness, in the case of a quantum witness it is no longer so obvious! Indeed, it is not clear at all how to "split" a quantum proof $|\psi\rangle$ (presumably, the ground state of an instance of the QMA-complete local Hamiltonian problem) between two or more players. We will discuss this problem first, in the next section. Then we will turn to point (ii).

4.1 Distributing quantum proofs

The difficulty in distributing a quantum proof between multiple players is related to the fact that we don't know if the "Hamiltonian" QPCP implies the "Games" QPCP, Conjecture 1. If we had a straightforward way to do this (of course, the "splitting" would need to have many "nice" properties), then presumably the implication would follow using a similar argument as the proof of the same implication in the classical setting, via oracularization.

Remark 9. The other direction of the implication, from "Games" QPCP to "Hamiltonian" QPCP, is also not known. We won't discuss it, because we (or at least I) don't have any idea how to even get started on showing it. Recall that in the classical case, starting from a game we construct a CSP by associating variables to the player's strategies, and writing constraints for the verifier's actions. But in the quantum case, a strategy is not simply a list of bits, it is given by measurement operators. And these measurement operators may be of arbitrary dimension. How can we construct a Hamiltonian from them?

In this section we will sketch a method for distributing quantum proofs to multiple players. The method won't quite achieve a reduction from "Games" QPCP to "Hamiltonian" QPCP, but we will still be able to use it to prove our exponential-size QPCP, Theorem 8.

Before we get started let's get a sense of why the classical technique of oracularization fails. Interestingly, the problem arises already with the completeness property. Indeed, the oracularization technique relies on the ability to compare the values of variables returned by Alice or Bob. In the quantum case, the "assignment" is a quantum state, and all its "variables" (qubits) may be arbitrarily entangled with each other. What this means is that, if we give Alice and Bob each an identical, but distinct, copy of the same entangled state, and they send back the same qubit of that state to the verifier, in general that qubit will be totally mixed, so that the verifier has no way to check whether the qubits are "identical" or not. Observe also that such a splitting doesn't use any entanglement between the players, which is a bit suspicious.[2]

To make our life easier, just as in the classical case we started with a simple system of linear equations, let's also imagine we start with a really simple Hamiltonian: for example, this is a 2-local Hamiltonian on a line, with only nearest-neighbor interactions of a particularly nice kind, let's say just XX or ZZ terms (by this we mean a projection of the form $\frac{1}{2}(\mathrm{Id} \pm \sigma_X(i) \otimes \sigma_X(j))$ or $\frac{1}{2}(\mathrm{Id} \pm \sigma_Z(i) \otimes \sigma_Z(j))$, where $i, j \in \{1, \ldots, n\}$ are neighbors on the n-qubit line. Let's even assume that, in the case of a positive instance, the Hamiltonian is frustration-free, so it has a ground state $|\Gamma\rangle$ with energy zero.[3] This Hamiltonian has a ground state, and our goal is to devise a scheme for splitting the ground state between two or more players in a way that the players can convince us that they have a low-energy ground state, while only showing us a constant number of qubits at a time. Of course, in case H does not have a low-energy ground state, then the players shouldn't be able to convince us either.

Basic splitting does not work. A first idea is to split the qubits. For example, one player gets even-numbered qubits, and another gets odd-numbered qubits. When the verifier wants to evaluate the energy of a local term acting on neighboring qubits i and $i+1$, it asks each qubit from the player that holds it, and measures it. (This gives us a game with quantum answers. In any case the game won't work out, so let's not worry about it.) Let's consider this transformation on the following "EPR Hamiltonian". The Hamiltonian has a term

$$\frac{1}{4}\big(2\,\mathrm{Id} + (\sigma_X(i) \otimes \sigma_X(i+1) + \sigma_Z(i) \otimes \sigma_Z(i+1))\big)$$

on any two nearest neighbors $(i, i+1)$, for $i \in \{1, \ldots, n\}$ and $n+1 \equiv 1$ (i.e. the line "wraps around"). This Hamiltonian is asking the ground state to be in an EPR pair in-between any two neighboring qubits. This is impossible, because a qubit can't be maximally entangled with both its left neighbor and its right neighbor. So, the Hamiltonian is highly frustrated, and the minimal energy is around $\frac{n}{2}$, with n the number of qubits. But the associated game is easy to win. The verifier is going to chose a term at random, acting on qubits i and j. It'll request qubits i from Alice, j from Bob (if i is even, and the other way round otherwise), receive one qubit from each player, and estimate the energy by projecting onto an EPR pair. There is a perfect winning strategy: Alice and Bob share a single EPR pair, and whatever qubit the verifier asks them for, they always reply the same thing!

To summarize the difficulty, the problem we are facing is that a quantum strategy for the players, in which whenever they are asked for "qubit i" they provide a qubit, does not mean that, in the "background",

[2] Technically, it is possible, since the class of games in Theorem **??** contains NEXP. But using that theorem would defeat our goal of giving a simple, "quantum" proof for Theorem 8.

[3] Technically this is an instance of Quantum 2-SAT, ant it has an efficient (classical) algorithm [Bra11]. We could allow 3-local terms, in which case it is not known if the problem is in NP, but we will keep them 2-local for simplicity.

they do have an n-qubit state from which the i-th qubit is taken. In the classical case there is no such difficulty: every list of answers, a_i, b_j to questions i, j, can be put together in a "proof" $(a_1, \ldots, a_n, b_1, \ldots, b_n)$. But in the quantum case, we can have qubits ρ_{ij} such that there does not exist an $2n$-qubit state $|\psi\rangle$ that has ρ_{ij} as its reduced density on the i-th and $(n+j)$-th qubits — a good example is the EPR Hamiltonian discussed above.

Using an error-correcting code. Here is a better idea. One way to think of the classical "replication" scheme (both players get a copy of the proof) is as an encoding, bit by bit, of the proof, using an error-correcting code: the 2-bit repetition code. (This doesn't correct any error, but it corrects a deletion, which, if you think about it, is all that is needed.) So consider a quantum error-correcting code. For concreteness, let's focus on the 7-qubit Steane code $\mathcal{C}$. This code encodes one qubit and corrects one error. Furthermore, the code has the nice property of being a CSS code that is "self-dual". What this means is that the code has stabilizers $\sigma_X(x)$ and $\sigma_Z(z)$ where x, z each lie in the same subspace of $\mathbb{F}_2^7$: for this code, the space is spanned by the vectors (00001111), (0110011), (101010101). Finally, it holds that the logical X and Z operators for the code are specified by a tensor product of σ_Z and σ_Z observables on a subset of the qubits: for the 7-qubit code the logical operators are $\sigma_X(x)$ and $\sigma_Z(z)$ for $x = z = (1111111)$.

Let $|\Gamma\rangle$ be the ground state of the local Hamiltonian H that the verifier wants to check, using the players. Encode each of the qubits of $|\Gamma\rangle$ using $\mathcal{C}$, yielding a $7n$-qubit state $|\tilde{\Gamma}\rangle$. For example, if $H = \frac{1}{4}(2\,\mathrm{Id} + \sigma_X \otimes \sigma_X + \sigma_Z \otimes \sigma_Z)$ then $|\Gamma\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ and

$$|\tilde{\Gamma}\rangle \;=\; \frac{1}{\sqrt{2}}\left(|\bar{0}\rangle_7|\bar{0}\rangle_7 + |\bar{1}\rangle_7|\bar{1}\rangle_7\right) ,$$

where we wrote $|\bar{b}\rangle_7$ for the 7-qubit encoding of bit b. Finally, give one share of each qubit to each of the players.

Why is this useful? Recall that the main difficulty we faced earlier is that we had no way to "piece together" the one- or two-qubit answers from the players into a consistent (encoding of) an n-qubit state. The code provides us a means to do this, by using a subset of the players as an "anchor" for the remaining players. Specifically, the code is some form of a "commitment scheme": if we ask all but one players for their share of qubit i, then we can ask for the remaining player not only for qubit i, but also qubit j; there is no way the player can return qubits that depend jointly on i and j, because there is a unique way it can return a qubit that will complete the qubits from the first players into a valid codestate for the i-th qubit.

This is the intuition, but in practice there are many difficulties that arise when one attempts to implement the idea. First of all, we still have a game in which the players send back qubits, instead of classical bits. Second, are we sure that we have completely overcome the kind of cheating strategy that we discussed when considering the EPR Hamiltonian? This needs to be shown!

We will not give all details here; the technique is developed in the papers [FV15, Ji16]. For the remainder of this section we focus on just one of the issues, which turns out to be the more delicate one. Consider the problem of testing that the 7 players share a valid $7n$-qubit codestate. An even simpler problem is testing that 2 players share n EPR pairs.[4] Our goal is to achieve this while using classical questions of polynomial length, and classical answers of constant length. Can we do it?

[4]The latter problem is similar to the former, as any codestate of the 7-qubit code (or any error-detection code) is locally equivalent to an EPR pair when one bundles 6 of the parties together.

4.2 A quantum linearity test

In a previous lecture we developed a lot of machinery, and applied it to obtain a rigidity theorem for the CHSH game. A consequence of that result is that a high success probability in the game can be used as a "witness", or "certificate", for one qubit of entanglement (an EPR pair). To obtain interesting tests for complexity or cryptography we need to scale things up. This is the purpose of the "quantum linearity test" introduced in this section. The test resembles very much the Blum-Luby-Rubinfeld linearity test from property testing, that was used in our proof of the classical exponential PCP (Theorem 6). It will provide us with the means to verify that two entangled players share an n-qubit maximally entangled state; equivalently, a tensor product of n EPR pairs.

Remark 10. If you are familiar with techniques in complexity theory, you may wonder, given that we have a good test for a single qubit, why not consider its parallel repetition? There are two obstacles to achieving our aims in this way. The first is that parallel repetition will have answer sizes that scale with the number of repetitions, whereas we aim to keep the answer size constant. The second is that we don't know how to show that parallel repetition achieves the desired result![5] It is known that the quantum value of the parallel repeated CHSH game goes down exponentially with the number of repetitions. But here we aim for something different, which is to certify properties of strategies that achieve close to the optimal success probability. This is analogous to the "direct product tests" in complexity theory, and it seems muh more tricky to achieve in the quantum case: showing that a player performs measurements that are in tensor product form, on a tensor product of qubits, is more delicate than showing that the player applies a "direct product function".

To get started, the following lemma gives a "robust" characterization of the maximally entangled state in dimension $d = 2^n$ as the unique common eigenvalue-1 eigenvector of all operators of the form $\sigma_P \otimes \sigma_P$, where σ_P ranges over the Pauli matrices. It generalizes a statement used at the end of the proof of the CHSH rigidity theorem.

Lemma 11. *Let $\varepsilon \geq 0$, n an integer, $d = 2^n$, and $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ a unit vector such that*

$$\mathrm{E}_{a,b} \langle\psi| \big(\sigma_X(a)\sigma_Z(b) \otimes \sigma_X(a)\sigma_Z(b)\big) |\psi\rangle \geq 1 - \varepsilon \,, \tag{1}$$

where the expectation is uniform over all $a, b \in \{0,1\}^n$. Then $|\langle\psi|\phi_{2^n}\rangle|^2 \geq 1 - \varepsilon$. In particular, $|\psi\rangle$ has Schmidt rank at least $(1-\varepsilon)2^n$.

Proof. Consider the case $n = 1$. The "swap" matrix

$$S = \frac{1}{4}\big(\sigma_I \otimes \sigma_I + \sigma_X \otimes \sigma_X + \sigma_Z \otimes \sigma_Z + \sigma_W \otimes \sigma_W\big)$$

squares to identity and has a unique eigenvalue-1 eigenvector, the vector $|\phi_2\rangle = \frac{1}{\sqrt{2}}(|1\rangle|1\rangle + |2\rangle|2\rangle)$. Thus $\langle\psi|S|\psi\rangle \geq 1 - \varepsilon$ implies $|\langle\psi|\phi\rangle|^2 \geq 1 - \varepsilon$. The same argument for general n shows $|\langle\psi|\phi_{2^n}\rangle|^2 \geq 1 - \varepsilon$. Any unit vector $|u\rangle$ of Schmidt rank at most r satisfies $|\langle u|\phi_{2^n}\rangle|^2 \leq r2^{-n}$, concluding the proof. □

Following the intuition we gained from analyzing the CHSH game, Lemma 11 provides a clear roadmap for establishing that a bipartite system is (close to) a maximally entangled state:

(i) Verify that the players apply observables that are isometric to the Pauli $\sigma_P = \pm\sigma_X(a)\sigma_Z(b)$ by testing the group relations for the associated group $H^{(n)}$;

[5] At least, not with good enough error bounds. "Rigidity" with a polynomial scaling in n is known; see e.g. [CN16].

(ii) Verify that the players' σ_P stabilize their entangled state by playing a ''consisteny test", in which both players are sent the same question, and checked for identical answers.

Step (i) is easier said than done. To explain how the step can be implemented we first need to understand the underlying group structure; this is done in Section 4.2.1. Then we also need to provide "operational tests" that can be used to certify the group relations for $H^{(n)}$; this is done in Section 4.2.2.

4.2.1 The Weyl-Heisenberg group

The Weyl-Heisenberg group H is the fancy name for the "Pauli group modulo complex conjugation", which is the 8-element group $\{\pm\sigma_I, \pm\sigma_X, \pm\sigma_Z, \pm\sigma_W\}$ whose multiplication table matches that of the 2×2 matrices

$$\sigma_X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad \sigma_Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{2}$$

$\sigma_I = \sigma_X^2 = \sigma_Z^2$ and $\sigma_W = \sigma_X\sigma_Z = -\sigma_Z\sigma_X$. As we saw in the proof of the CHSH rigidity theorem, this group has four 1-dimensional representations, uniquely specified by the image of σ_X and σ_Z in $\{\pm 1\}$, and a single irreducible 2-dimensional representation, given by the matrices defined above.

To obtain a test for n-qubit maximally entangled states we are led to consider the "n-qubit Weyl-Heisenberg group" $H^{(n)}$, the matrix group generated by n-fold tensor products of the 8 matrices identified above. The group $H^{(n)}$ has cardinality $2 \cdot 4^n$, and elements of $H^{(n)}$ have a unique representative of the form $\pm\sigma_X(a)\sigma_Z(b)$ for $a, b \in \{0,1\}^n$.

The irreducible representations of $H^{(n)}$ are easily computed from those of H; for us the only thing that matters is that the only irreducible representation which satisfies $\rho(-I) = -\rho(I)$ has dimension 2^n and is given by the defining matrix representation (in fact, it is the only irreducible representation in dimension larger than 1).

With the upcoming application to rigidity in mind, we state a version of the Gowers-Hatami theorem tailored to the group $H^{(n)}$ and a specific choice of presentation for the group relations.

Corollary 12. *Let n, d be integer, $\varepsilon \geq 0$, $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ a permutation-invariant state, σ the reduced density of $|\psi\rangle$ on either system, and $f : \{X, Z\} \times \{0,1\}^n \to U(\mathbb{C}^d)$. For $a, b \in \{0,1\}^n$ let $X(a) = f(X, a)$, $Z(b) = f(Z, b)$, and assume $X(a)^2 = Z(b)^2 = I_d$ for all a, b (we call such operators, unitaries with eigenvalues in $\{\pm 1\}$,* observables*). Suppose that the following inequalities hold: consistency*

$$\mathrm{E}_a \langle\psi|\big(X(a) \otimes X(a)\big)|\psi\rangle \geq 1 - \varepsilon\,, \qquad \mathrm{E}_b \langle\psi|\big(Z(b) \otimes Z(b)\big)|\psi\rangle \geq 1 - \varepsilon\,, \tag{3}$$

linearity

$$\mathrm{E}_{a,a'} \big\|X(a)X(a') - X(a+a')\big\|_\sigma^2 \leq \varepsilon\,, \qquad \mathrm{E}_{b,b'} \big\|Z(b)Z(b') - Z(b+b')\big\|_\sigma^2 \leq \varepsilon\,, \tag{4}$$

and anti-commutation

$$\mathrm{E}_{a,b} \big\|X(a)Z(b) - (-1)^{a\cdot b}X(a)Z(b)\big\|_\sigma^2 \leq \varepsilon\,. \tag{5}$$

Then there exists a $d' \geq d$, an isometry $V : \mathbb{C}^d \to \mathbb{C}^{d'}$, and a representation $g : H^{(n)} \to U_{d'}(\mathbb{C})$ such that $g(-I) = -I_{d'}$ and

$$\mathrm{E}_{a,b} \big\|X(a)Z(b) - V^* g(\sigma_X(a)\sigma_Z(b))V\big\|_\sigma^2 = O(\varepsilon).$$

Note that the conditions (4) and (5) in the corollary are very similar to the conditions required of an approximate representation of the group $H^{(n)}$; in fact it is easy to convince oneself that their exact analogue suffice to imply all the group relations. The reason for choosing those specific relations is that they can be checked using multiplayer games; see the next section for this. Condition (3) is necessary to derive the conditions for the application of the Gowers-Hatami theorem from (4) and (5), and is also testable; see the proof.

Remark 13. Corollary 12 can be seen as an extension of the BLR linearity test. The latter makes a similar statement, but for the commutative group $\{\pm\sigma_X(a) |\, a \in \{0,1\}^n\}$.

Proof. To apply the Gowers-Hatami theorem we need to construct an (ε,σ)-representation f of the group $H^{(n)}$. Using that any element of $H^{(n)}$ has a unique representative of the form $\pm\sigma_X(a)\sigma_Z(b)$ for $a,b \in \{0,1\}^n$, we define $f(\pm\sigma_X(a)\sigma_Z(b)) = \pm X(a)Z(b)$. Next we need to verify that f is an approximate representation. Let $x,y \in H^{(n)}$ be such that $x = \sigma_X(a_x)\sigma_Z(b_x)$ and $y = \sigma_X(a_y)\sigma_Z(b_y)$ for n-bit strings (a_x,b_x) and (a_y,b_y) respectively. Up to phase, we can exploit successive cancellations to decompose $(f(x)f(y)^* - f(xy^{-1})) \otimes I$ as

$$\begin{aligned}
&\big(X(a_x)Z(b_x)X(a_y)Z(b_y) - (-1)^{a_y\cdot b_x}X(a_x+a_y)Z(b_x+b_y)\big)\otimes I \\
&\quad = X(a_x)Z(b_x)X(a_y)\big(Z(b_y)\otimes I - I\otimes Z(b_y)\big) \\
&\qquad + X(a_x)\big(Z(b_x)X(a_y) - (-1)^{a_y\cdot b_x}X(a_y)Z(b_x)\big)\otimes Z(b_y) \\
&\qquad + (-1)^{a_y\cdot b_x}\big(X(a_x)X(a_y)\otimes Z(b_y)\big)\big(Z(b_x)\otimes I - I\otimes Z(b_x)\big) \\
&\qquad + (-1)^{a_y\cdot b_x}\big(X(a_x)X(a_y)\otimes Z(b_y)Z(b_x) - X(a_x+a_y)\otimes Z(b_x+b_y)\big) \\
&\qquad + (-1)^{a_y\cdot b_x}\big(X(a_x+a_y)\otimes I\big)\big(I\otimes Z(b_x+b_y) - Z(b_x+b_y)\otimes I\big).
\end{aligned}$$

(It is worth staring at this sequence of equations for a little bit. In particular, note the "player-switching" that takes place in the 2nd, 4th and 6th lines; this is used as a means to "commute" the appropriate unitaries, and is the reason for including (3) among the assumptions of the corollary.) Evaluating each term on the vector ψ, taking the squared Euclidean norm, and then the expectation over uniformly random a_x,a_y,b_x,b_y, the inequality $\|AB\psi\| \leq \|A\|\|B\psi\|$ and the assumptions of the theorem let us bound the overlap of each term in the resulting summation by $O(\varepsilon)$. Using $\|(A\otimes I)\psi\| = \|A\|_\sigma$ by definition, we obtain the bound

$$\mathrm{E}_{x,y}\left\|f(x)f(y)^* - f(xy^{-1})\right\|_\sigma^2 = O(\varepsilon).$$

We are now in a position to apply the Gowers-Hatami theorem, which gives an isometry V and exact representation g such that

$$\mathrm{E}_{a,b}\left\|X(a)Z(b) - \frac{1}{2}V^*\big(g(\sigma_X(a)\sigma_Z(b)) - g(-\sigma_X(a)\sigma_Z(b))\big)V\right\|_\sigma^2 = O(\varepsilon). \tag{6}$$

Using that g is a representation, $g(-\sigma_X(a)\sigma_Z(b)) = g(-I)g(\sigma_X(a)\sigma_Z(b))$. It follows from (6) that $\|g(-I)+I\|_\sigma^2 = O(\varepsilon)$, so we may restrict the range of V to the subspace where $g(-I) = -I$ without introducing much additional error. □

4.2.2 Testing the Weyl-Heisenberg group relations

Corollary 12 makes three assumptions about the observables $X(a)$ and $Z(b)$: that they satisfy approximate consistency (3), linearity (4), and anti-commutation (5). To complete our test, we need to show how these relations can be "certified" in a two-playe game. There are multiple ways this can be done; here is one.

Linearity test:

(a) The referee selects $W \in \{X, Z\}$ and $a, a' \in \{0,1\}^n$ uniformly at random. He sends (W, a, a') to one player and (W, a), (W, a'), or $(W, a + a')$ to the other.

(b) The first player replies with two bits, and the second with a single bit. The referee accepts if and only if the player's answers are consistent.

As always in this section, the test treats both players symmetrically. As a result we can assume that the players' strategy is symmetric, and is specified by a permutation-invariant state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ and a measurement for each question: an observable $W(a)$ associated to questions of the form (W, a), and a four-outcome measurement $\{W^{a,a'}\}$ associated with questions of the form (W, a, a').

The linearity test described above is almost identical to the BLR linearity test, except for the use of the basis label $W \in \{X, Z\}$. The following lemma states conditions that a strategy must satisfy in order to succeed with high probability in the test.

Lemma 14. *Suppose that a family of observables $\{W(a)\}$ for $W \in \{X, Z\}$ and $a \in \{0,1\}^n$, generates outcomes that succeed in the linearity test with probability $1 - \varepsilon$, when applied on a symmetric bipartite state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ with reduced density matix σ. Then the following hold: approximate consistency*

$$\mathrm{E}_a \langle\psi| \big(X(a) \otimes X(a)\big) |\psi\rangle \;=\; 1 - O(\varepsilon), \qquad \mathrm{E}_b \langle\psi| \big(Z(b) \otimes Z(b)\big) |\psi\rangle \;\geq\; 1 - O(\varepsilon),$$

and linearity

$$\mathrm{E}_{a,a'} \big\| X(a)X(a') - X(a+a') \big\|_\sigma^2 = O(\varepsilon), \qquad \mathrm{E}_{b,b'} \big\| Z(b)Z(b') - Z(b+b') \big\|_\sigma^2 \;=\; O(\varepsilon).$$

Exercise 15. Prove the lemma. (In the case of classical strategies, the conditions are an immediate reformulation of the test. The proof for quantum strategies is not much harder.)

Testing anti-commutation can be done using the CHSH game. Unfortunately, this game does not have quantum value 1, and this makes it hard to use it in a modular way. To avoid this we use instead the Magic Square game. This is a fascinating game, but for lack of time we only only recall the part of the analysis of the game that is useful for us (see e.g. the paper [WBMS16] for a description of the game and a proof of Lemma 16 below).

Lemma 16 (Magic Square)**.** *The Magic Square game is a two-player game with nine possible questions (with binary answers) for one player and six possible questions (with two-bit answers) for the other player which has the following properties. The distribution on questions in the game is uniform. Two of the questions to the first player are labelled X and Z respectively. For any strategy for the players that succeeds in the game with probability at least $1 - \varepsilon$ using a bipartite state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ and observables X and Z for questions X and Z respectively, it holds that*

$$\big\| \big((XZ + ZX) \otimes I_d\big) |\psi\rangle \big\|^2 \;=\; O\big(\sqrt{\varepsilon}\big). \tag{7}$$

Moreover, there exists a strategy which succeeds with probability 1 in the game, using $\psi = \phi_4$ and Pauli observables $\sigma_X \otimes I_2$ and $\sigma_Z \otimes I_2$ for questions X and Z respectively.

Based on the Magic Square game we devise the following "anti-commutation test".

Anti-commutation test:

(a) The referee selects $a, b \in \{0,1\}^n$ uniformly at random under the condition that $a \cdot b = 1$. He plays the Magic Square game with both players, with the following modifications: if the question to the first player is X or Z he sends (X, a) or (Z, b) instead; in all other cases he sends the original label of the question in the Magic Square game together with both strings a and b.

(b) Each player provides answers as in the Magic Square game. The referee accepts if and only if the player's answers would have been accepted in the game.

Using Lemma 16 it is straightforward to show the following.

Lemma 17. *Suppose a strategy for the players succeeds in the anti-commutation test with probability at least $1-\varepsilon$, when performed on a symmetric bipartite state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ with reduced density matrix σ. Then the observables $X(a)$ and $Z(b)$ applied by the player upon receipt of questions (X, a) and (Z, b) respectively satisfy*

$$\mathrm{E}_{a,b:\, a\cdot b=1} \left\| X(a)Z(b) - (-1)^{a\cdot b} Z(b)X(a) \right\|_\sigma^2 \;=\; O\big(\sqrt{\varepsilon}\big). \tag{8}$$

4.2.3 Application: a robust test for high-dimensional entangled states

We are ready to put all the pieces together and prove the soundness of our entanglement test. Although arguably a direct "quantization" of the BLR result, it is the only test known which achieves constant robustness, and was an important step forward in quantum complexity — all prior n-qubit tests required success that is inverse polynomially (in n) close to the optimum in order to provide any meaningful conclusion.

n-qubit Pauli braiding test: With probability $1/2$ each,

(a) Execute the linearity test;

(b) Execute the anti-commutation test.

Theorem 18. *Suppose that a family of observables $W(a)$, for $W \in \{X, Z\}$ and $a \in \{0,1\}^n$, and a state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$, generate outcomes that pass the n-qubit Pauli braiding test with probability at least $1-\varepsilon$. Then $d = (1 - O(\sqrt{\varepsilon}))2^n$.*

As should be apparent from the proof it is possible to state a stronger conclusion for the theorem, which includes a characterization of the observables $W(a)$ and the state $|\psi\rangle$ up to local isometries. For simplicity we only recorded the consequence on the dimension of $|\psi\rangle$.

Proof. Using Lemma 14 and Lemma 17, success with probability $1-\varepsilon$ in the test implies that conditions (3), (4) and (5) in Corollary 12 are all satisfied, up to error $O(\sqrt{\varepsilon})$. (In fact, Lemma 17 only implies (5) for strings a, b such that $a \cdot b = 1$. The condition for string such that $a \cdot b = 0$ follows from the other conditions.) The conclusion of the corollary is that there exists an isometry V such that the observables $X(a)$ and $Z(b)$ satisfy

$$\mathrm{E}_{a,b} \left\| X(a)Z(b) - V^* g(\sigma_X(a)\sigma_Z(b)) V \right\|_\sigma^2 \;=\; O(\sqrt{\varepsilon}).$$

Using again the consistency relations (3) that follow from part (a) of the test together with the above we get

$$\mathrm{E}_{a,b} \, \langle\psi| (V\otimes V)^* \big(\sigma_X(a)\sigma_Z(b) \otimes \sigma_X(a)\sigma_Z(b)\big)(V\otimes V)|\psi\rangle \;=\; 1 - O(\sqrt{\varepsilon}).$$

Applying Lemma 11, $(V \otimes V)|\psi\rangle$ has Schmidt rank at least $(1 - O(\sqrt{\varepsilon}))2^n$. But V is a local isometry, which cannot increase the Schmidt rank. □

5 Application to delegation

In Section 4.2 we developed a test that forces two entangled provers to share a maximally entangled state, and to perform specific Pauli measurements on that state when asked to do so. In this section we sketch three different ways in which this can be leveraged to achieve the seemingly more complex task of delegating an arbitrary computation to entangled provers.

5.1 Computation by teleportation

This is the method used in [RUV12].[6] The first idea is to use the entanglement test to achieve a form of state and process tomography. For example, the verifier can ask Alice to perform a certain arbitrary two-qubit measurement on some of the qubits of the entanglement shared with Bob. He can also tell Bob to perform a regular Pauli measurement on the corresponding two qubits. Since Bob cannot not tell the difference with the entanglement test, he must perform the right measurements. By the properties of the maximally entangled state, this allows to check that Alice indeed performed the measurement she was asked to, even if it is not a Pauli measurement.

The next idea is to use computation by teleportation. In this model, the server only needs to apply at most one gate on each qubit, followed by a teleportation measurement — the next gate is applied on the qubit at the receiving end of the teleportation, etc. By using similar ideas to the state tomography explained above, it is possible to use one prover to verify that the other exactly implements the required steps of the computation. Working out the details of this is not trivial.

5.2 Computation using magic states

It is possible to execute a universal computation by only making applying Clifford gates, as long as magic states can be used. This is similar to the protocol from [ABOE08] that uses the polynomial-based authentication code. The main requirement is to ensure that the server receives the right authenticated qubits, be they computation qubits or magic states. For this an adaptation of the PBT can be used, where the observables tested are not only the single-qubit Pauli but also Clifford observables acting on a constant number of qubits.

5.3 Post-hoc verification with two servers

It is possible to adapt the Morimae-Fitzsimons protocol [MF16] that we saw in a previous lecture as follows. A first server is tasked with preparing the same encoded ground state as in the MF protocol. In addition, the two servers are required to share as many EPR pairs as there are qubits in the state. The first server performs measurements in the Bell basis to teleport the state to the second server. The measurement outcomes are communicated to the verifier, who interprets them as a one-time pad.

The second server is tasked with measuring a specific term of the Hamiltonian, selected by the verifier, on the teleported state. To verify that the right measurement is performed the PBT is played with constant probability with the two servers. Note that the behavior of the first server does not need to be checked: as long as the second server makes measurements that are isometric to the desired Hamiltonian, the optimal state on which to perform the measurements is the ground state — so it is in the interest of the first server to cooperate and indeed teleport a ground state. Upon receiving measurement outcomes, the verifier can correct for the one-time pad that results from teleportation.

[6] A more detailed explanation of this scheme is available in Section 4 of the "Week 10" notes on delegation.

References

[AALV09] Dorit Aharonov, Itai Arad, Zeph Landau, and Umesh Vazirani. The detectability lemma and quantum gap amplification. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 417–426. ACM, 2009.

[AAV13] Dorit Aharonov, Itai Arad, and Thomas Vidick. Guest column: the quantum pcp conjecture. *Acm sigact news*, 44(2):47–79, 2013.

[ABOE08] Dorit Aharonov, Michael Ben-Or, and Elad Eban. Interactive proofs for quantum computations. *arXiv preprint arXiv:0810.5375*, 2008.

[BFL90] László Babai, Lance Fortnow, and Carsten Lund. Nondeterministic exponential time has two-prover interactive protocols. In *Foundations of Computer Science, 1990. Proceedings., 31st Annual Symposium on*, pages 16–25. IEEE, 1990.

[Bra11] Sergey Bravyi. Efficient algorithm for a quantum analogue of 2-sat. *Contemporary Mathematics*, 536:33–48, 2011.

[CN16] Matthew Coudron and Anand Natarajan. The parallel-repeated magic square game is rigid. *arXiv preprint arXiv:1609.06306*, 2016.

[FV15] Joseph Fitzsimons and Thomas Vidick. A multiprover interactive proof system for the local hamiltonian problem. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*, pages 103–112. ACM, 2015.

[Hås01] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.

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

[Ji16] Zhengfeng Ji. Classical verification of quantum proofs. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 885–898. ACM, 2016.

[JJUW11] Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. Qip= pspace. *Journal of the ACM (JACM)*, 58(6):30, 2011.

[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.

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

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

[NV16] Anand Natarajan and Thomas Vidick. Robust self-testing of many-qubit states. *arXiv preprint arXiv:1610.03574*, 2016.

[RUV12] Ben W Reichardt, Falk Unger, and Umesh Vazirani. A classical leash for a quantum system: Command of quantum systems via rigidity of chsh games. *arXiv preprint arXiv:1209.0448*, 2012.

[Slo17] William Slofstra. The set of quantum correlations is not closed. *arXiv preprint arXiv:1703.08618*, 2017.

[VW+16] Thomas Vidick, John Watrous, et al. Quantum proofs. *Foundations and Trends® in Theoretical Computer Science*, 11(1-2):1–215, 2016.

[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.