Lecture 8

The Quantum PCP conjecture

Scribe: Victor Braun
Reviewer: Pengxiang Wang

In the last lecture, we stated two variants of the *"Classical PCP theorem"*. Furthermore, we also proved a weaker version to give an intuition on the stronger result. After introducing the quantum analog of different classical complexity theory notions, it is natural to wonder if there is a quantum equivalent to the mentioned PCP theorem.

In this and the following lectures, we will thus define the *QPCP conjecture* and explain the different challenges that proving such conjecture induce. For completeness, we will also introduce some positive and negative results that were found while trying to solve this problem in recent years.

8.1 Constraint Satisfaction Problem (CSP) variant

Let us first recall that given a 3-local Hamiltonian $H = \frac{1}{m}\sum_{i=1}^{m} H_i$, $0 \leq H_i \leq I$, separating the cases when $\lambda_{\min}(H) \leq a$ or $\lambda_{\min}(H) \geq b$ is QMA-hard whenever $a = 2^{-n}$ and $b = \Omega(\frac{1}{n^2})$.

It appears that the result whenever $a - b \geq \gamma$ with γ being a fixed constant, is not known to be true or not, and is equivalent to the QPCP conjecture. Before formally stating the conjecture, let us introduce (we could do it earlier) a more general definition of a local Hamiltonian.

Definition 8.1. Let $m, q, d : \mathbb{N} \to \mathbb{N}$, we say that H is $[m, q]_d$-local if $H = \frac{1}{m}\sum_{i=1}^{m} H_i$, where each H_i acts non-trivially on at most q qudits (quantum objects lying in $\mathbb{C}^d$).

Remark 8.2. If not specified, we usually set $m(n) = \text{poly}(n)$, $q(n) = O(1)$ and $d = 2$.

We are now ready to formally define our first variant of the QPCP conjecture:

Conjecture 8.3. (**QPCP conjecture, CSP variant**)
There exists a constant $q \in \mathbb{N}$, $\gamma > 0$, $a, b : \mathbb{N} \to [0, 1]$ with $a(n) - b(n) \geq \gamma$ for all n such that given any q-LH H on n qubits, it is QMA*-hard to distinguish between $\lambda_{\min}(H) \leq a$ and $\lambda_{\min}(H) \geq b$.*

This variant can be also seen as the *"hardness of approximation"* variant, because proving it would imply that, unless BQP $=$ QMA, there is no quantum polynomial-time algorithm that can approximate $\lambda_{\min}(H)$.

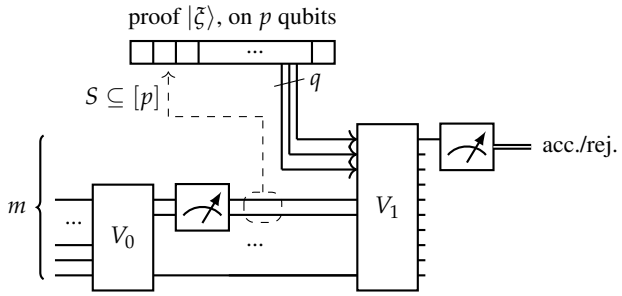

Figure 8.1: *Non-adaptive QPCP verifier*

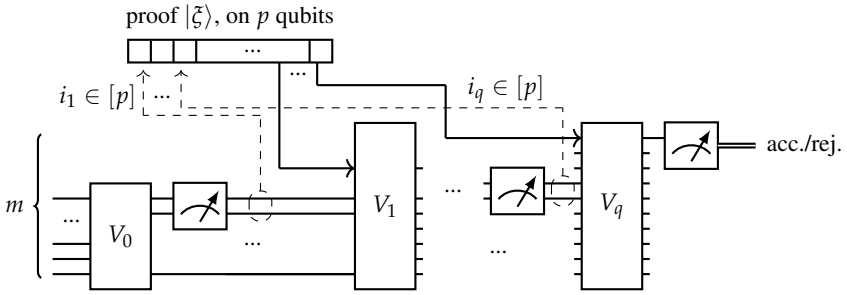

Figure 8.2: *Adaptive QPCP verifier*

8.2 Proof Checking variant

Analogously to the classical case, the QPCP conjecture comes more naturally with the definition of what is a QPCP verifier itself. This machine doesn't need any explicit randomness parameter r, as this behavior is automatically inherited from the probabilistic quantum behavior of the verifier.

Furthermore, it is allowed, in the same way as its classical analog, to read a predefined number of qubits. We now define more formally how such a machine works:

Definition 8.4. (**QPCP**)
$L \in \mathsf{QPCP}_{c,s}[p,q]$ if there exists a poly-time computable function $V : x \mapsto V_x$, where V_x is a quantum circuit on $m(|x|)$ ancilla bits (it's implicit that $m(n) = O(\text{poly}(n))$ and on $p(|x|)$ proof qubits such that V_x queries less than $q(|x|)$ proof qubits and:

- If $x \in L \Rightarrow \exists |\xi\rangle \in (\mathbb{C}^2)^{\otimes p}$, $\Pr(V_x \text{ accepts } |0^m\rangle|\xi\rangle) \geq c$.
- If $x \notin L \Rightarrow \forall |\xi\rangle \in (\mathbb{C}^2)^{\otimes p}$, $\Pr(V_x \text{ accepts } |0^m\rangle|\xi\rangle) \leq s$.

It's a nice exercise to see how we can "query" qubits in a quantum circuit. For intuition, we give two ways (one non-adaptive and one adaptive) of seeing this process.

In both cases, the idea is that we first perform some measurement whose outcome will specify the proof qubit(s) we want to read. We can then imagine that by some process, the corresponding quantum objects (say, particles) are given to our quantum computer and placed in one or more "empty" slots of the q circuit qubits. The "new" qubits are then used as regular circuit qubits and the computation continues normally. Visually, the processes are shown in (8.1) and (8.2).

With this definition in mind, we are ready to state the second variant of the QPCP conjecture as follows:

Conjecture 8.5. (**QPCP conjecture, proof checking variant**)
$\mathsf{QMA} \subseteq \mathsf{QPCP}_{1/3,\,2/3}[\mathrm{poly}(n), O(1)]$.

It is worth noting that both in the classical and quantum version, the proof length is bounded by $O(\mathrm{poly}(n))$, with the only difference that in the quantum version, this bound is explicit.

We will show below that this variant is equivalent to the previous one. It is quite counterintuitive as the CSP variant seems like a more "natural" extension to the known QMA-hardness of the 3-LH problem; whereas on the other side, the second variant states something surprising, as it would induce that no matter the quantity of entanglement in the quantum proof, reading a constant number q of its qubits is in a way "enough" for the verifying process.

8.3 Equivalence of the variants

As the two variants define the same conjecture, the following lemma is natural:

Lemma 8.6. *For the QPCP conjecture, CSP variant $\Leftrightarrow$ Proof Checking variant.*

Proof. Starting with the easiest, the $(\Rightarrow)$ direction, we know that if $L \in \mathsf{QMA}$, there exist, by the CSP variant QPCP conjecture, a computable function $H : x \mapsto H_x$ and a constant γ such that if $x \in L$, $\lambda_{\min}(H) \leq a$ and if $x \notin L$, $\lambda_{\min}(H) \geq b$, with $a - b \geq \gamma$. The quantum verifier V_x simply selects H_i uniformly at random and measures it, exactly the same way we did to prove that k-LH$\in$ QMA. Because the difference between completeness and soundness is constant, we can boost it by repeating the process until having $\frac{1}{3}$ and $\frac{2}{3}$ as desired.

The $(\Leftarrow)$ direction is more tricky, as it uses notions of quantum tomography, we sketch the proof as follows. We start with our QPCP verifier V_x and would like to compute a corresponding H_x in poly-time. We could simply define it by seeing that the querying process is q-local, however, the challenge is that the q positions are not predefined, and so doing it in such say would force H_x to operate on all possible proof qubits. We rather create $H_x = \binom{p}{q}^{-1} \sum_{S\subseteq[p]:|S|=q} H_S$, now how to find H_S? Let's say that for a run of V_x, the positions of the proof qubits queried by V_x is a set $S \subseteq [p]$, then all the circuit can be interpreted as a big POVM $(\Pi_S^{acc.}, \Pi_S^{rej.})$, the probability that the state is accepted is given by $\Pr(V_x \text{ accepts } |0^m\rangle|\xi\rangle) = ||\Pi_S^{acc.}|\tilde{\xi}\rangle_S||^2$, where $|\tilde{\xi}\rangle_S$ is the proof restricted to the queried qubits. Our mission now is thus to find $\Pi_S^{acc.}$ for each S. This is done by performing tomography: we run V_x repeatedly on a proof that contains $|0\rangle$ except on the proof qubits indexed by S, which are initialized to an arbitrary state $|u_1\rangle, |u_2\rangle, ..., |u_q\rangle$. If we select only the runs where V_x actually queries exactly the qubits in S, we can estimate $|\Pi_S^{acc.}|u_1\rangle|u_2\rangle \cdots |u_q\rangle\|^2$. Repeating this a polynomial number of times, over a family of states $|u_i\rangle$ that form an ε-net over all single-qubit states, will allow us to estimate all matrix coefficients of $\Pi_S^{acc.}$ to within inverse polynomial accuracy. □

As a side note, we can introduce a subclass of QMA, QCMA *"Quantum-Classical Merlin Arthur"* where the quantum verifier is quantum but has only access to a said "classical" proof (diagonal in the Z basis), we then define the following hierarchy:

Theorem 8.7. $\mathsf{NP} \subseteq \mathsf{MA} \subseteq \mathsf{QCMA} \subseteq \mathsf{QMA}$.

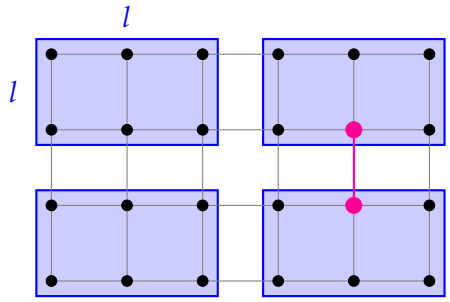

Figure 8.3: *Chopping qubit grid with $l \times l$ areas*

8.4 QPCP for geometrically local Hamiltonians

Surprisingly, when we restrict the Hamiltonians to a geometrical structure (such as 1D or even 2D Hamiltonians), the QPCP conjecture does not hold (unless of course, $\mathsf{P} = \mathsf{QMA}$). This is because, even though these versions of the LH problem are still QMA-hard, they are easy ways to obtain rough (constant factor) approximations to the ground state energy. Let's for example take a 2D Hamiltonian. Its interaction graph can be represented as a 2-dimensional grid. Now, let's choose a large constant l and let's chop the grid into $l \times l$ regions as showed in figure 8.3. We name each rectangle C_i and create the associated H_i that corresponds to H restricted to the qubits in C_i. Then we can write $H = \sum_i H_i + H_b$ where H_b corresponds to the interaction pairs that "cross" the rectangle, as shown in pink in the figure. The idea is that if we choose l large enough, the number of crossing terms will be negligible, $||H_b|| \leq 2n \cdot \frac{n}{l} \ll n^2 = m$. And so, we are only left to compute the minimal energy of each H_i, which is done in poly-time, and add the results up to obtain a good approximation, with small additive error, to the minimum energy of H.

To go even further, this idea of "chopping" the qubits can be generalized to any graph. We state this fact formally in the following theorem:

Theorem 8.8. *Let H be a 2-local Hamiltonian acting on qudits (H acts on $(\mathbb{C}^d)^{\otimes n}$), on a regular graph of degree D. Then for every $|\psi\rangle \in (\mathbb{C}^d)^{\otimes n}$ and for all partitions $\mathcal{X} = \bigcup_i X_i$ of $[n]$ such that $\forall i,\ |X_i| = r$:*

There exists a state $|\phi\rangle = |\phi_1\rangle \otimes ... \otimes |\phi_{n/r}\rangle$ with each $|\phi_i\rangle$ only acting on the qudits of X_i such that $|\langle\psi|H|\psi\rangle - \langle\phi|H|\phi\rangle| \leq c \cdot \left(\frac{d^6}{D} \cdot \frac{\mathbb{E}[S(X_i)_\psi]}{r}\right)^{\frac{1}{8}}$.

Because the entanglement entropy $S(.)$ is never larger than the number of qubits, the right part of the product is smaller than 1, meaning that if we take $D \gg d$, we can approximate the energy of H by the energy of H on a product state. The intuition comes from the "monogamy" of entanglement.

We can then see that this theorem implies that the QPCP conjecture is false on graphs with $D = \omega(1)$ unless $\mathsf{QCMA} = \mathsf{QMA}$. What is also interesting about this theorem is that it is not true classically, as SAT remains hard on graphs of high degree.

To give some intuition about the theorem, we can state the de Finetti theorem, stating that a n-exchangeable distribution must be close to a product distribution. Formally:

Theorem 8.9. *Let P be a distribution over $[d]^n$ such that P is n-exchangeable, i.e. $Pr_P(s) = Pr_P(\pi(s))$ for every permutation $\pi \in S_n$. Then, there exists a distribution μ on $[d]$ such that*

$$\forall k, ||tr_{n-k}(P) - \int Q^{\otimes k} d\mu||_1 \leq \min\left\{\frac{2kd}{n}, \frac{k(k-1)}{2}\right\}$$

We won't prove this theorem in the course, but we give a proof sketch for $d = 2$:

Proof. The exchangeability condition implies that P must be a convex combination of uniform distributions over strings of a given Hamming weight, e.g $\Pr(010) = \Pr(100) = \Pr(001)$. Let's now take one such distributions, P_i. The idea is that the marginal distribution on k variables is hypergeometric, and if we have $k \ll n$, then this distribution is very close to a product of binomial distributions. □

The quantum analog of this theorem is not enough to prove Theorem 8.8 but already gives us the flavor of it in some sense. Formally, it is stated as:

Theorem 8.10. (**Quantum De Finetti, Renner, ETH**)
Let ρ be a density matrix on $(\mathbb{C}^d)^{\otimes n}$ such that ρ is n-exchangeable, i.e ρ is unchanged after any permutation of the n qudits. Then, there exists a measure μ such that

$$\forall k, ||tr_{n-k}(\rho) - \int \sigma^{\otimes k} d\mu||_\infty \leq \frac{2k(d+k)}{n+d}$$

The achieved bound is stronger than in Theorem 8.8, however, it requires n-exchangeability, which is not the case there. Therefore, the theorem is not directly applicable; but it gives the flavor of what we need: a high degree of symmetry, or many constraints between "almost all" pairs of qudits, will force the entanglement to be very weak overall and thus a good product state approximation can be found.

8.5 Global entanglement

One of the implications of the QPCP conjecture is the existence of states that are *impossible* to create with a poly-time circuit.

More precisely, assume that QCMA $\neq$ QMA, the QPCP conjectures implies that there exist families of q-LH such that $\lambda_{\min}(H) \leq a$ and that $\forall|\psi\rangle, \langle\psi|H|\psi\rangle \leq b = a + \gamma_{cst}$, $|\psi\rangle$ does not have a poly-time quantum circuit representation. This comes from the fact that if $|\psi\rangle = C|0\rangle$, then C is a classical witness to the statement $\lambda_{\min}(H) \leq b$, i.e. H is a YES-instance.

To be more formal, we define the NLTS$_d$ conjecture as follows:

Definition 8.11. NTLS$_{d=\text{depth}}$ is true $\Leftrightarrow$ There exists a family of LH such that if $\langle\psi|H|\psi\rangle \leq b$ then $|\psi\rangle$ has no depth-d quantum circuit

The above statement can be formalized as: "if QCMA $\neq$ QMA and if the QPCP conjecture is true, then NLTS$_{\text{poly}(n)}$ is also true". Furthermore, we can state a weaker theorem:

Theorem 8.12. *There exists $d = \Omega(\log n)$ such that NLTS$_d$ is true.*

To give an intuition to this theorem, we can introduce the notion on globally entangled states:

Definition 8.13. $|\psi\rangle$ is said to be *non-trivial* if it has no constant depth circuit representation.

Definition 8.14. $|\psi\rangle$ is called *globally entangled* if there exists $|\psi'\rangle \perp |\psi\rangle$ such that for any local observable O, $\langle\psi|O|\psi\rangle = \langle\psi'|O|\psi'\rangle$.

In other words, such states $|\psi\rangle, |\psi'\rangle$ are very easy to distinguish as they are orthogonal, but as soon as we restrict them to a subset of qubits, it's *impossible* to differentiate the two states. Such states are for example the cat states $|GHZ_{\pm}\rangle = \frac{1}{\sqrt{2}}(|0^n\rangle \pm |1^n\rangle)$.

In the next lecture, we will continue to explore this statement by showing a lemma that basically says that *non-triviality* implies *global entanglement*. Thus, loosely speaking, Theorem 8.12 which we prove in the next couple lecture implies that there are local Hamiltonians such that all states with low energy are globally entangled.