Frontiers of Quantum complexity *Instructor: Anand Natarajan, Anurag Anshu*

Lecture 11: The Quantum PCP Conjecture

October 9, 2025 *Scribe: Marcus Gozon*

1 Quantum PCP Conjecture

Today will be the first of four lectures on the quantum PCP conjecture.

Conjecture 1.1 (Quantum PCP Conjecture)**.** There is a (relative) promise gap $c > 0$ with constants a and b satisfying $b - a \geq c$ such that given a $k = O(1)$-local Hamiltonian $H = \sum_\alpha \beta_\alpha P_\alpha$ with $l = \sum_\alpha |\beta_\alpha|$, deciding if

$$\frac{E_0(H)}{l} < a \quad \text{or} \quad \frac{E_0(H)}{l} > b$$

is QMA-complete, where $E_0(H)$ is the ground state energy (minimum eigenvalue) of H. In other words, the constant (relative) gap local Hamiltonian problem is QMA-complete. Denote this problem as $\text{qPCP}_{a,b}[k]$.

Note that $l \leq \|H\|_\infty$ as each Pauli matrix satisfies $\|P_\alpha\|_\infty = 1$, and typically $l = \Theta(n)$ since usually $\beta_\alpha = \Theta(1)$ as in the case of "extensive" energy, i.e. energy that scales with the number of particles. In particular, when $|\beta_\alpha| = O(1)$, $l = O(n^{\text{locality of } H})$.

Recall that the local Hamiltonian (LH) problem with gap $c = 1/\text{poly}(n)$ was shown to be QMA-complete by the Feynman-Kitaev construction, and so what the above conjectures is that the problem remains QMA-hard when c is a constant. In addition, the conjecture is also saying that approximating the ground state energy up to an extensive level $cl/2$ is QMA-hard since such an algorithm could be used to solve the gap LH problem above.

Definition 1.2 (Quantum PCP)**.** A quantum probabilistically checkable proof is a proof system (QMA protocol) in which the verifier can only check a "few" locations of the proof, i.e. $O(1)$ qubits.

The problem $\text{qPCP}_{a,b}[O(1)]$ has a quantum probabilistically checkable proof since the verifier can do the following for a given proof $|\psi\rangle$: pick α with probability $\frac{|\beta_\alpha|}{l}$ and measure with the POVM $\{\frac{I+P_\alpha}{2}, \frac{I-P_\alpha}{2}\}$. Then accept if the outcome is $-\text{sign}(\beta_\alpha)$ where we regard $\frac{I+P_\alpha}{2}$ as being $+1$ and $\frac{I-P_\alpha}{2}$ as -1. Thus, the probability of accepting is

$$\sum_\alpha \frac{|\beta_\alpha|}{l} \langle\psi| \frac{I - \text{sign}(\beta_\alpha)P_\alpha}{2} |\psi\rangle = \frac{1}{2} - \frac{1}{2l} \sum_\alpha \beta_\alpha \langle\psi| P_\alpha |\psi\rangle = \frac{1}{2} - \frac{1}{2l} \langle\psi| H |\psi\rangle .$$

In the yes case, the prover can send the ground state $|\psi\rangle$ with $\frac{E_0(H)}{l} = \frac{\langle\psi|H|\psi\rangle}{l} < a$ so that $\Pr[\text{accept}] > \frac{1}{2} - \frac{a}{2}$, and in the no case, $\frac{\langle\psi|H|\psi\rangle}{l} \geq \frac{E_0(H)}{l} > b$ so that $\Pr[\text{accept}] < \frac{1}{2} - \frac{b}{2}$.

Note that there exists a constant $c > 0$ so that $\text{qPCP}_{0,c}[O(1)]$ is NP-hard, which is just a restatement of the PCP theorem. Recall that this says that there exists a constant $c > 0$ such that given a classical Hamiltonian (CSP) $H_{cl} = \sum_\alpha C_\alpha$ (i.e. set C_α to be the matrix assigning an energy penalty of 1 to the assignments that violate the associated constraint with no penalty otherwise), deciding if all clauses are satisfied or at most a $(1-c)$ fraction can be satisfied is NP-hard. This has had profound consequences for hardness of approximation, e.g. in a seminal work of Håstad, he proved that it is NP-hard to determine whether a system of linear equations over F_2 is at least $(1-\epsilon)$ satisfiable v.s. at most $(1/2+\epsilon)$ satisfiable for all $\epsilon > 0$, where $1/2$ is achieved through the trivial random assignment.

2 Quantum PCP Motivation

Since qPCP is already NP-hard, one might be tempted to believe that quantum computers can't solve NP-hard problems and so the problem is not in BQP. So why care about the qPCP conjecture? Rather than thinking of the conjecture in terms of BQP, one should think of it in terms of the structure of low-energy states since many NP-hard problems can be solved relatively fast in practice e.g. through SAT solvers. In particular, assuming QMA $\neq$ NP, the conjecture would rule out low-energy states with low "operational description complexity" i.e. having $\langle\psi| O |\psi\rangle$ be easy to compute classically for a local Hamiltonian O. An example would be $|\psi\rangle = U|0^n\rangle$ for a constant depth circuit U, or more generally, $|\psi\rangle = U|\phi\rangle$ where $|\phi\rangle$ is a stabilizer state. In these cases, the prover can just send a classical description of the circuit U with the verifier classically computing out $\langle\psi| O |\psi\rangle = \langle\phi| U^\dagger O U |\phi\rangle$ efficiently: since O is local and U is constant depth, by a lightcone argument (which we will see several times in the next few lectures), $U^\dagger O U$ will also be local. Then by the Gottesman-Knill Theorem, we can simulate $|\phi\rangle$ classically and also measure in any Pauli basis, and so we can expand out $U^\dagger O U$ into a sum of $O(\mathrm{poly}(n))$ Pauli operators $c_i P_i$ to compute (exactly[1]) each term $c_i \langle\phi| P_i |\phi\rangle$ and take their sum. Thus, such low-energy and low-complexity states would not exist for the gapped LH problem if proven to be QMA-hard as they would give us an NP protocol contradicting the assumption that QMA $\neq$ NP.

In much of the effort in solving quantum chemistry problems, it's not about finding such states but rather whether they exist since if you know there are nice states in the low-energy state regime, then you could use SAT solvers for the corresponding NP problem. The conjecture would show that assuming SAT solvers or ML algorithms are good at solving problems in practice would not be enough to explore the low-energy regime to estimate the ground energy. Such a consequence has been important enough to have been conjectured [FH14] and resolved in one sense:

Theorem 2.1 (No Low-Energy Trivial States (NLTS), [ABN23]). *There is a family of $O(1)$ local Hamiltonians $\{H_n\}_n$ and a constant $c > 0$ such that $\forall n$ and $\forall |\psi\rangle : \langle\psi_n| H_n |\psi_n\rangle \leq E_0 + cn$, the minimum circuit depth of $|\psi_n\rangle$ is $\Omega(\log n)$.*

In the above theorem, H_n are good quantum codes i.e. constant rate and relative distance[2] in which all the ground states are stabilizer states. Given the preceding discussion, this naturally leads to the following conjecture:

Conjecture 2.2 (No Low-Energy Trivial Magic States (NLTM)). There is a family of $O(1)$ local Hamiltonians $\{H_n\}_n$ and a constant $c > 0$ such that $\forall n$ and $\forall |\psi\rangle : \langle\psi_n| H_n |\psi_n\rangle \leq E_n + cn$, the minimum circuit depth of $|\psi_n\rangle$ starting from a stabilizer state is $\omega(1)$.

[1] Recall that for a stabilizer state $|\phi\rangle$ and a Pauli string P, $\langle\phi| P |\phi\rangle = \pm 1$ if $\pm P$ is in the stabilizer group of $|\phi\rangle$ and 0 otherwise.

[2] Subsequent work has shown that one can take quantum codes with large distance and low rate [GK24].

This conjecture would rule out all the ways we currently know of simulating quantum states classically since the two main classes of states in which the expectation value $\langle\psi| O |\psi\rangle$ is easy to compute are the constant depth circuits and the stabilizer states. This conjecture can be thought of as a practical version of the quantum PCP conjecture. A potential candidate for proving it is the quantum double model, which would generalize the good quantum codes used for the resolution of the NLTS conjecture. A barrier to the above conjecture is on how to prove stabilizer lower bounds, but there has been recent work in this direction [Par25], [WL25].

There is also additional motivation from quantum simulation. Recall that for any quantum circuit, one has the associated Feynman-Kitaev Hamiltonian. Thus, the ground state "encodes" the computation, but the issue is the brittleness of the ground energy regime as the gap shrinks to 0 as $n \to \infty$. Then the possibility of transforming the Hamiltonian to have a large gap could be more robust to errors. This is particularly useful for adiabatic simulation where one evolves a simple Hamiltonian to a target Hamiltonian with the hope that the ground state of the original Hamiltonian is transformed to an approximate ground state of the target Hamiltonian.

3 Quantum PCP Nonexamples

Here are a few examples of local Hamiltonians which cannot be used to prove the quantum PCP conjecture:

- Lattice Hamiltonians (i.e. the interaction graph has an edge between qubits i and j when the there is a Pauli operator whose support contains both of them): for every $\epsilon > 0$, there is a classical algorithm that approximates the ground energy λ, i.e. $E_0(H) \leq \lambda \leq E_o(H) + \epsilon n$ for a $\sqrt{n} \times \sqrt{n}$ lattice in classical time $2^{O(1/\epsilon^2)}$: Chop up the lattice with block length/width of size $2/\epsilon$ (see figure). Throw away edges/Pauli's for all qubits lying on the lines and call the new Hamiltonian H' and so

  $$\|H' - H\|_\infty \leq \frac{\epsilon n}{2} + \frac{\epsilon n}{2} \leq \epsilon n.$$

  Then brute force solve for the ground energy $E_0(H_{sq})$ within each of the squares in $2^{O(1/\epsilon^2)}$ time, choosing $\lambda = \sum_{sq} E_0(H_{sq})$ since after removing the associated Pauli's, the systems associated to each square are now independent. More generally, for d dimensions, the argument would give $2^{O(1/\epsilon^d)}$, and so lattices of constant dimension wouldn't be useful for the quantum PCP conjecture. However lattices of dimension $d = \Omega(\log\log n)$ would still be possible candidates. Thus, we would expect codes (on lattices) in something like at least $d = \Omega(\log\log n)$ dimension to be used in a successful proof. The above setting wouldn't hold for a little bit of long-range interaction that often exists in the real world.

- 2-local Hamiltonians on high-degree graphs [BH13] - for a d regular graph, one can estimate the ground energy up to an $O(n/\sqrt[3]{d})$ error, which we will see more of in the next few lectures.

4 Positive Results for Quantum PCP

Recall that the Feynman-Kitaev construction had a gap of at most $O(1/n^3)$. It turns out that one can make the gap as large as $O(n^{-\epsilon})$ through a simple observation made by Lijie Chen:

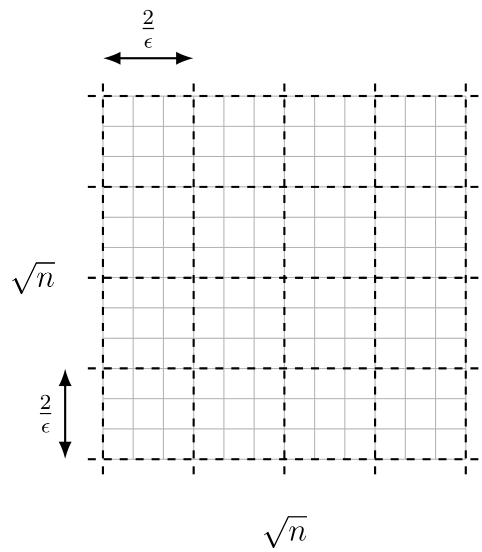

Figure 1: Chopping argument on the interaction graph of the lattice Hamiltonian

Proposition 4.1. $\mathrm{qPCP}_{e^{-n^c}, n^{-\epsilon}}[5]$ is QMA complete for any constant $c > 0$.

Proof. Let H_{FK} be the Feynman-Kitaev Hamiltonian. Recall that this has a gap with parameters $a = e^{-n^c}$ and $b = 1/n^3$ where n is the number of particles and c is any constant, so either $E_0(H_{FK}) \leq na = ne^{-n^c}$ or $E_0(H_{FK}) \geq bn \geq 1/n^2$. Then define $H^M = H_{FK}^{(1)} + \cdots + H_{FK}^{(M)}$, and so $E_0(H^M) = M \cdot E_0(H_{FK})$. In the yes case, $E_0(H^M) \leq Mne^{-n^c}$, and in the no case, $E_0(H^M) \geq M/n^2$ where the number of particles is now $n' = Mn$. Then taking $M = n^t$, we would have $n' = n^{t+1}$, and $b' = n^{t-2}/n^t = n^{-2} = n'^{\frac{-2}{t+1}}$ and $a' = n^{t+1}e^{-n^c}/n^t = ne^{-n^c} = n'^{\frac{1}{t+1}} e^{-n'^{\frac{c}{t+1}}}$. Thus, first choosing t large enough so that $\frac{2}{t+1} < \epsilon$ and c sufficiently large will give the required bound. $\square$

4.1 Gap Amplification and Known Techniques

Conjecture 4.2 (Gap Amplification Conjecture)**.** For all k-local H_n, one can find a k-local $H'_{n'}$ efficiently such that

- if $E_0(H)/l \approx 0 \implies E_0(H')/l' \approx 0$
- if $E_0(H)/l \geq \delta \implies E_0(H')/l' \geq d\delta$ where $d > 1$

Note that this implies the qPCP conjecture by just applying the above $O(\log n)$ times. Likewise, the qPCP conjecture in some sense also implies this conjecture by the induced mapping of a given proof of the qPCP theorem. This conjecture naturally comes from the classical gap amplification procedure, which makes the following transformation: A CSP gets mapped as follows:

$$\underset{\substack{2-\text{local}\\ |\Sigma|=q \\ E_0(C)/l=\gamma}}{C} \longrightarrow \underset{\substack{2-\text{local}\\ |\Sigma|=q^t \\ E_0(C)/l=\Theta(t\gamma)}}{C'} \longrightarrow \underset{\substack{2-\text{local}\\ |\Sigma|=q \\ E_0(C)/l=\Theta(t\gamma)}}{C''}$$

There are two standard methods known for proving gap amplification:

- Powering method/tensoring
- Detectability lemma

Both are based on "AGSPs," or approximate ground state projections, which we will see in the upcoming lectures. As a simple example of the powering method, we can map

$$\frac{H}{l} \mapsto H' = \mathbb{1} - \left(\mathbb{1} - \frac{H}{l}\right)^t$$

If $E_0(H) = 0$, then $E_0(H') = 0$ (by the same ground state), and if $E_0(H)/l = \delta$, then $E_0(H') = 1 - (1-\delta)^t = \Omega(t\delta)$. Then the issue is that if H is k-local, H' would be kt-local which would ruin the locality with repeated applications of this transformation.

References

[ABN23] Anurag Anshu, Nikolas P. Breuckmann, and Chinmay Nirkhe. Nlts hamiltonians from good quantum codes. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, page 1090–1096, New York, NY, USA, 2023. Association for Computing Machinery.

[BH13] Fernando G. S. L. Brandão and Aram Wettroth Harrow. Product-state approximations to quantum ground states. In *Symposium on the Theory of Computing*, 2013.

[FH14] Michael H. Freedman and Matthew B. Hastings. Quantum systems on non-k-hyperfinite complexes: a generalization of classical statistical mechanics on expander graphs. *Quantum Info. Comput.*, 14(1–2):144–180, January 2014.

[GK24] Louis Golowich and Tali Kaufman. NLTS Hamiltonians and Strongly-Explicit SoS Lower Bounds from Low-Rate Quantum LDPC Codes. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, volume 287 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 54:1–54:23, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[Par25] Natalie Parham. Quantum circuit lower bounds in the magic hierarchy, 2025.

[WL25] Fuchuan Wei and Zi-Wen Liu. Long-range nonstabilizerness from quantum codes, orders, and correlations, 2025.