CS286.2 Lecture 7: Quantum PCP conjectures

Scribe: Brenden Roberts

Last time, we saw Kitaev's proof that the $O(\log n)$-LH problem is QMA-complete. By using unary notation for the clock the locality can be improved to 5, and with more work Kitaev, Kempe and Regev showed that 2-LH is QMA-complete.

Theorem 1 (Quantum Cook-Levin)**.** *Given a set of n qubits and m local quantum constraints, or Hamiltonians, $0 \leq H_i \leq \mathrm{Id}$, acting on $q = O(1)$ qubits each, it is QMA-hard to distinguish between ($H = \sum_i H_i$):*

- $E_0(H) = \min_{|\psi\rangle} \langle\psi|H|\psi\rangle \leq a$
- $E_0(H) \geq b$,

for $a = 2^{-\operatorname{poly}(n)}$ and $b = a + 1/\operatorname{poly}(n)$.

A few comments on the parameters a and b. The fact that a can be taken exponentially small in the number of qubits follows from *sequential* error amplification for QMA as shown by Marriott and Watrous. Note that parallel error reduction wouldn't do, as it decreases the error but increases the number of qubits. The best known results for 5-LH have $b - a = \Omega(\frac{1}{T^2})$, where T is the number of gates in the QMA verifier circuit. Since the number of Hamiltonians is linear in T, this essentially gives a gap that is quadratic in $1/m$.

The classical Cook-Levin theorem states that 3SAT is NP-hard. This implies that 3-LH$_{a,b}$ is NP-hard for $a = 0, b = 1$. The PCP theorem implies that 3-LH$_{a,b}$ is NP-hard for $a = 0, b = \gamma m$ for some $\gamma \geq 0$ (e.g. $\gamma = 1/8 - \varepsilon$ for 3SAT, see notes from Lecture 1), leading to a gap that scales linearly with the number of constraints.

The CSP formulation of QPCP

Conjecture 2 (QPCP, CSP variant)**.** *There exists $\gamma > 0$ and q such that given an instance of q-LH it is QMA-hard to distinguish between):*

- $E_0(H) \leq a$
- $E_0(H) \geq b = a + \gamma m$.

This is the most immediate generalization of the PCP theorem to the quantum setting. Just as we saw different statements for the PCP theorem, however, we may introduce corresponding variants of the QPCP conjecture. Before proceeding we explore the physical implications of (QPCP, CSP variant) by introducing yet another complexity class.

The complexity class QCMA

QCMA stands for "quantum-classical Merlin-Arthur": QCMA is the set of languages that can be verified by a *quantum* polynomial-time machine given access to a *classical* proof as advice.

Definition 3. *A language* $L \in$ QCMA *if there exists a classical polynomial time mapping* $x \in \{0,1\}^* \mapsto V_x$*, a quantum circuit on* $n + p = poly(|x|)$ *qubits, such that for all* x*:*

- $x \in L \Rightarrow \exists\, w \in \{0,1\}^p$ *such that* $\Pr(V_x \text{ accepts } |0^n\rangle|w\rangle) \geq \frac{2}{3}$
- $x \in L \Rightarrow \forall\, w,\ \Pr(V_x \text{ accepts } |0^n\rangle|w\rangle) \leq \frac{1}{3}$.

Note that in the definition we wrote $|w\rangle$ as a quantum state, but it is really just a quantum representation of the classical string in the computational basis $\{|0\rangle, |1\rangle\}^{\otimes p}$, and it is neither more nor less useful to the verifier than the classical string w — these are two notation for the same object.

Clearly,

$$\text{MA} \subseteq \text{QCMA} \subseteq \text{QMA},$$

as quantum verifiers are (potentially) more powerful than classical verifiers, and quantum proofs (potentially) more useful than classical ones. It is not known, but strongly believed, that the subset relations are strict (*i.e.*, all three classes are distinct). There are oracle separations between all three classes, and this is essentially the best we can hope for since we don't even know that MA $\subsetneq$ QCMA! For the sake of the foregoing discussion we'll assume that QCMA $\neq$ QMA. Since the QPCP conjecture makes the assertion that q-LH$_{a,b}$ is QMA-complete, it in particular asserts that the problem is *not* in QCMA (as otherwise the two classes would collapse).

Implications of QPCP, assuming QCMA $\subsetneq$ QMA

Take H to be a hard instance for q-LH$_{a,a+\gamma m}$. Let $S_0(H) = \text{span}\{|\psi\rangle : |\psi\rangle \text{ an eigenvector of } H \text{ with eigenvalue} < a + \gamma m\}$. The problem of deciding between $E_0(H) \leq a$ and $E_0(H) \geq a + \gamma m$ is equivalent to distinguishing between

- $S_0(H) = \emptyset \quad (\iff E_0(H) \geq a + \gamma m)$
- $S_0(H) \neq \emptyset \quad (\iff E_0(H) \leq a)$

If this problem cannot be decided in QCMA, there can be no classical proof certifying which case holds and that can be verified efficiently. In particular, whenever $S_0(H) \neq \emptyset$ there cannot be any state $|\psi\rangle \in S_0(H)$ for which there would exist a classical representation efficiently certifying this fact (that $|\psi\rangle \in S_0(H)$). This rules out the possibility that local Hamiltonians always have low-energy state that are "simple": for instance, there cannot (in general) be a good product-state approximation to the ground state; in fact the ground state cannot even be well-approximated by applying a low-depth circuit to a product state. Indeed, if this were the case then a classical description of the product state, together with the circuit, would let us efficiently estimate the ground state energy of the Hamiltonian and hence solve LH in QCMA. Other kinds of representations, such as certain types of tensor networks (allowing for efficient energy computation), are also ruled out. Thus morally QPCP implies the following strong statement:

QPCP: Even "simple" (read, local) Hamiltonians can have very "complex" (read, highly entangled) low-energy states.

Remark 4. *There do exist efficient classical representations of these states; for example, "$|\psi\rangle$ is the smallest eigenvector of H." However, these are* not efficiently verifiable*. An interesting question is: how large can $S_0(H)$ be before it must contain a "simple" state, for instance one that is product, or close to product, across the qubits?*

Using the properties of the Gibbs distribution, one can argue that the equilibrium state of a physical system with Hamiltonian H at temperature $T > 0$ is supported on states $|\psi\rangle$ with $\langle\psi|H|\psi\rangle < a + \gamma m$, where γ is a parameter independent of n, m that goes to 0 as T goes to 0. That is, if $\rho_T \propto e^{-H/T}$ denotes the Gibbs state at temperature T, $\text{supp}(\rho_T) \in S_0(H)$ (up to exponentially small error). Returning to the discussion above, we see that for QPCP to hold no state in the support of ρ_T can have an efficient classical representation. Hence QPCP makes the very strong statement that, not only are ground states complex (this already follows from the quantum Cook-Levin theorem), but this also holds for low, but non-zero,-temperature equilibrium states. This is significant because, as much as one can debate the physical relevance of the ground state (how does one reach absolute 0 temperature anyways?), positive temperatures seem much more attainable.. This seems to indicate that it is hard to do any meaningful physics: how can one hope to make meaningful predictions if one cannot even compute the "simplest" quantity, the energy of the system?

There is an important caveat however, which is that "physical" Hamiltonians are special and tend to have additional structure:

- They are geometrically local (*e.g.*, the interaction graph is 2D or 3D-local), rather than the more abstract form of locality we have been considering thus far. In fact you can show (exercise!) that *QPCP is false for H local in any constant number of dimensions*, so that for the conjecture to hold we are forced to drop this important "physical" requirement. How does the largest achievable $b - a$ promise gap scale with the geometric lcoality of the constraint Hamiltonians?

- They have a great deal of symmetry, *e.g.*, translational invariance. It is known that the quantum Cook-Levin theorem holds even for 1D translationally invariant systems; since for QPCP we have to consider dimensions that increase with n it is not immediately clear what kind of symmetry restriction one could impose without jeopardizing the possible validity of the conjecture.

- They have a "spectral gap" $E_1(H) - E_0(H) = \Omega(1)$, where $E_1(H)$ is the second smallest eigenvalue of H (or at least a "low density of states" above the minimum energy, meaning that for "typical" families of local Hamiltonians the dimension of the span of eigenvectors with eigenvalue at most $E_0(H) + \alpha$ for some $\alpha > 0$ increases polynomially, rather than exponentially, with n). These restrictions imply that the ground state is in some sense "well-isolated", and this can lead to simpler ground states — we will see an example of this phenomenon when we discuss the area law.

The proof-checking formulation of QPCP

We now introduce a quantum analogue of the proof-checking variant of the PCP theorem.

Definition 5. *A language $L \in QPCP[q]$ if there exists a polynomial time mapping $x \in \{0,1\}^* \mapsto V_x$, a quantum circuit on $n + p = \mathrm{poly}(|x|)$ qubits where V_x acts nontrivially on only q qubits of the p-qubit proof, such that*

- $x \in L \Rightarrow \exists\ |\xi\rangle \in \mathbb{C}^{2^p}$ *such that* $\Pr(V_x \text{ accepts } |0^n\rangle|\xi\rangle) \geq \frac{2}{3}$
- $x \in L \Rightarrow \forall\ |\xi\rangle,\ \Pr(V_x \text{ accepts } |0^n\rangle|\xi\rangle) \leq \frac{1}{3}$

Conjecture 6 (QPCP, proof-checking variant)**.** *Membership in any language $L \in$ QMA can be decided using a verification procedure that receives a polynomial-length quantum proof, but only ever looks at a constant number of qubits of the proof: there exists a $q = O(1)$ such that* QMA $\subseteq$ *QPCP[q].*

Remark 7. *The wording "act on q qubits" used in the definition to is not very natural, because it forces one to think of V_x as a mixed classical/quantum procedure: what we really mean is that V_x is allowed to toss some classical random coins, and dependending on the outcome of these coin tosses, decide on at most q qubits of the proof that it wants to examine (using a quantum circuit, that also acts on the ancilla qubits). Can you find a more natural definition?*

We conclude by showing that the two variants of QPCP introduced thus far are equivalent.

Claim 8. *(QPCP, CSP variant) and (QPCP, proof-checking variant) are equivalent.*

Proof. We first show CSP $\Rightarrow$ proof-checking. Given an instance of q-LH$_{a,a+\gamma m}$ we describe a verification procedure of the form described in Definition 5. The verifier proceeds as follows: choose a local term H_i at random, measure the q qubits of the proof on which H_i acts using $\{H_i, \mathrm{Id} - H_i\}$, and accept if and only if the second outcome is obtained. Then if the proof given to V_x is the quantum state $|\xi\rangle$,

$$\Pr(\text{accept}) = 1 - \frac{\langle\xi|H_i|\xi\rangle}{m} \begin{cases} \geq 1 - a/m \\ \leq 1 - a/m - \gamma \end{cases},$$

based on whether the instance of LH is positive or not. Since γ is a constant the gap between these two probabilities can be amplified to the required $2/3 - 1/3$ separation through parallel repetition of the verification procedure. This will only require a constant number of copies of $|\xi\rangle$, and the total number of qubits accessed, a constant multiple of q, remains constant.

Next we show proof-checking $\Rightarrow$ CSP. By executing the quantum circuit V_x one eventually obtains the positions of q qubits that are to be measured in a certain basis. THe basis can be computed using a quantum algorithm which would simulate the verification procedure V_x by repeatedly running it sufficiently many times and doing tomography. Since q is a constant a polynomial number of repetitions will suffice to obtain a description of the measurement that is accurate to within polynomial precision. For every possible subset S of q qubits of the proof we obtain a measurement

(possibly trivial in case the verifier never observes those qubits) H_S; putting all the nontrivial H_S together gives an instance of LH that faithfully reproduces the constraints imposed by V_x.

Note that the reduction we just described requires a *quantum* algorithm. A classical algorithm would have to compute the exponential-size matrix that represents the polynomial-size quantum circuit used by V_x, and thus generically require exponential time. The equivalence between the two QPCP conjectures is then a "quantum equivalence". There is nothing particularly wrong with that, but it still raises the question: are the two conjectures also equivalent under classical polynomial-time reductions? □