Quantum Complexity Theory

Proof Strategy for the Quantum Cook-Levin Theorem

LOCAL HAMILTONIAN $\in$ QMA

To put LH in QMA, we challenge Merlin to send us the ground state, and then we check it using phase estimation. In the NO instance he can't cheat because of the variational principle.

LOCAL HAMILTONIAN is QMA-hard

Here we need to show that finding ground state energies can be as difficult as doing nondeterministic quantum computation. To do this we will map an arbitrary quantum circuit with a constrained output and an unconstrained input register (i.e. a QMA verifier) into the ground state of a local Hamiltonian in such a way that the ground state energy will be sensitive to the acceptance prob of the QMA verifier.

Like the Cook-Levin tableau, these ground states will record the history of a quantum computation in a way that allows us to check validity with local constraints. These are called Feynman-Kitaev history states.

Quantum Complexity Theory

Suppose the QMA verifier runs the sequence of local unitary gates $U_1, ..., U_T$ on the input and the witness. The history of the computational steps looks like this:

$$|x\rangle|\xi\rangle \quad , \quad U_1|x\rangle|\xi\rangle \quad , \quad U_2U_1|x\rangle|\xi\rangle \quad , \quad ... \quad , \quad U_T...U_1|x\rangle|\xi\rangle$$

In the classical Cook-Levin proof, we would put each time step on its own set of bits. If we did this in the quantum case, the time steps might look like this:

$$|\psi_{t=0}\rangle|\psi_{t=1}\rangle...|\psi_{t=T}\rangle$$

Again in the classical proof, each gate acts on a few input bits and a few output bits. We can check that the inputs match the correct outputs using a local constraint that only acts on those bits.

We want local constraints that distinguish the state $|\psi_t\rangle|\psi_{t+1}\rangle$ from some other state $|\psi_t\rangle|\psi'\rangle$. What is the problem with this if the $|\psi\rangle$'s are *n* qubit states and we check them with a *k*-local operator?

Quantum Complexity Theory

The problem occurs if $|\psi_t\rangle|\psi_{t+1}\rangle$ and $|\psi_t\rangle|\psi'\rangle$ are both highly entangled states (which is the generic case).

Local observables are only sensitive to local reduced density matrices. If both states are highly entangled, then in both cases the RDMs will be very close to maximally mixed, and so local observables tell us nothing.

Entanglement fundamentally breaks the proof strategy of the classical Cook-Levin theorem.

However, paraphrasing the comments of Mike and Ike on quantum cloning and QKD, whatever quantum takes away with one hand, it gives us something new and beautiful with the other.

In particular, the solution to locally checking the time steps of a quantum computation is to entangle them, instead of recording them on separate registers in a tensor product.

Quantum Complexity Theory

Distilling the previous example, even locally checking the identity gate is impossible to do across a tensor product. This would require distinguishing $|\alpha\rangle|\alpha\rangle$ from $|\alpha\rangle|\beta\rangle$ using a k-local operator, when $|\alpha\rangle, |\beta\rangle$ are arbitrary n-qubit quantum states.

But notice that if we have the state $\frac{1}{\sqrt{2}}(|0\rangle|\alpha\rangle + |1\rangle|\beta\rangle)$, then the RDM of the first qubit (a local RDM) tells us a lot about the relation of $|\alpha\rangle, |\beta\rangle$.

This gives us hope for local constraints if we entangle the time steps of the computation,

$$|\psi\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^{T} U_t \ldots U_1 |x, \xi\rangle |t\rangle$$

Which is the (baseline form) of what is called a **Feynman-Kitaev history state**.

Quantum Complexity Theory

The notation $|x, \xi\rangle$ refers to the input and the witness. We are also free to consider the history state of a circuit $U_T \dots U_1|0^n\rangle$,

$$|\psi\rangle = \frac{1}{\sqrt{T+1}}\sum_{t=0}^{T} U_t \dots U_1|0^n\rangle|t\rangle = \frac{1}{\sqrt{T+1}}\sum_{t=0}^{T}|\psi_t\rangle|t\rangle$$

Where $|\psi_t\rangle = U_t \dots U_1|0^n\rangle$ is the state of the circuit at time t. Note the Hilbert space is now $\mathcal{H} = \mathcal{H}_{\text{data}} \otimes \mathcal{H}_{\text{clock}}$, where $\mathcal{H}_{\text{clock}}$ is a qudit of dimension T + 1 called the clock.

If someone hands you this state, then how would you check the output of the circuit? To check the output we could collapse the **clock register**, and hope we get $t = T$, which happens with probability 1/T. Then if we look at the computational register we'll see the state of the circuit at $t = T$.

How could we modify the history state to increase the probability of seeing the output of the circuit? Pad the end of the circuit with identity gates…

Quantum Complexity Theory

Given the history state corresponding to a verifier circuit, the probability of acceptance (output of $|1\rangle$) is given by the expectation value,

$$|\psi\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^{T} |\psi_t\rangle |t\rangle \qquad O_{\text{accept}} = |1\rangle\langle 1|_1 \otimes |T\rangle\langle T|$$

$$p_{\text{accept}} = \langle\psi|O_{\text{accept}}|\psi\rangle$$

These equations are correct, but when we map a QMA verifier to a local Hamiltonian we want to assign a higher energy to inputs that the verifier rejects. Therefore we include an energy penalty for rejection:

$$H_{\text{out}} = |0\rangle\langle 0|_1 \otimes |T\rangle\langle T|$$

Quantum Complexity Theory

Similarly, the Hamiltonian in our reduction will include terms that enforce the input of the computation. If the input is $x = x_1 ... x_n$, then we include

$$H_{\text{in}} = \sum_{i=1}^{n} |\neg x_i\rangle\langle\neg x_i|_i \otimes |0\rangle\langle 0|$$

Which assigns higher energy to states that do not have the intended input bit at t = 0. This covers the case when the string x is input to the verifier. If instead we define QMA verifiers in terms of circuits that are efficiently computable from x, the input constraint may just check for ancillas in a standard state like $|0\rangle$,

$$H_{\text{in}} = \sum_{i=1}^{n} |1\rangle\langle 1|_i \otimes |0\rangle\langle 0|$$

In any case, just as in the Cook-Levin proof the local constraint terms do not act on the registers that hold the witness at t = 0. Rather, we will design the overall Hamiltonian so that is has lower energy iff an acceptable witness exists.

Quantum Complexity Theory

So far we have written down a few local terms that check the inputs and outputs of a history state. But what about the main problem of creating ground states that look like history states?

Our solution to this problem will be closely related to a much easier problem, which is to construct a Hamiltonian on just $\mathcal{H}_{\text{clock}}$ with a unique ground state given by

$$|\psi\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^{T} |t\rangle$$

Which is a uniform superpositon state of a particle on a line with T + 1 sites. From physics, this looks like a low energy state of a particle hopping on a line (in a higher energy state we would expect the magnitude and phase of the wave function to oscillate rapidly across space). This propagation Hamiltonian is:

$$H_{\text{prop}} = \sum_{t=1}^{T} H_{\text{prop}}(t) \quad , \quad H_{\text{prop}}(t) = \frac{1}{2}(|t\rangle - |t-1\rangle)(\langle t| - \langle t-1|)$$

Quantum Complexity Theory

The terms $H_{\text{prop}}(t)$ are most useful as projectors, but we can expand them out as

$$H_{\text{prop}}(t) = \frac{1}{2}(|t\rangle - |t-1\rangle)(\langle t| - \langle t-1|) = \frac{1}{2}\left(|t\rangle\langle t| + |t-1\rangle\langle t-1| - |t\rangle\langle t-1| - |t-1\rangle\langle t|\right)$$

Consider an arbitrary state $|\alpha\rangle = \sum_{t=0}^{T} \alpha_t |t\rangle$, for which the expectation value is

$$\langle\alpha|H_{\text{prop}}(t)|\alpha\rangle = |\alpha_t - \alpha_{t-1}|^2$$

And so the states that minimize H_{prop} have amplitude distributions that are as "flat as possible." This is one of many ways to see that the uniform superposition is the ground state of H_{prop}.

Quantum Complexity Theory

To go from a single particle to the history state, define a unitary:

$$R = \sum_{t=0}^{T} U_t ... U_1 \otimes |t\rangle\langle t|$$

Append a register of $|0^n\rangle$ to our single particle ground state, and note that

$$R|0^n\rangle\left(\frac{1}{\sqrt{T+1}}\sum_{t=0}^{T}|t\rangle\right) = \frac{1}{\sqrt{T+1}}\sum_{t=0}^{T}|\psi_t\rangle|t\rangle$$

Is a ground state of $RH_{\text{prop}}R^\dagger$. The rotated propagation terms have the form

$$RH_{\text{prop}}(t)R^\dagger = \frac{1}{2}\Big(I \otimes |t\rangle\langle t| + I \otimes |t-1\rangle\langle t-1| - U_t \otimes |t\rangle\langle t-1| - U_t^\dagger \otimes |t-1\rangle\langle t|\Big)$$

Quantum Complexity Theory

Because of the unitary equivalence, we can move freely between the single hopping particle Hamiltonian and the propagation Hamiltonian that enforces history state ground states.

To reduce the notation, redefine $H_{\text{prop}} = \sum_{t=1}^{T} H_{\text{prop}}(t)$ to act on $\mathcal{H} = \mathcal{H}_{\text{data}} \otimes \mathcal{H}_{\text{clock}}$, with

$$H_{\text{prop}}(t) = \frac{1}{2}\Big(I \otimes |t\rangle\langle t| + I \otimes |t-1\rangle\langle t-1| - U_t \otimes |t\rangle\langle t-1| - U_t^\dagger \otimes |t-1\rangle\langle t|\Big)$$

By itself, H_{prop} only enforces the correct propagation of the gates in the circuit, but it does not check the input. Therefore H_{prop} is 2^n- fold degenerate. Combining it with

$$H_{\text{in}} = \sum_{i=1}^{n} |1\rangle\langle 1|_i \otimes |0\rangle\langle 0|$$

Singles out an input, and enforces $|\psi\rangle = \dfrac{1}{\sqrt{T+1}} \sum_{t=0}^{T} U_t \ldots U_1 |0^n\rangle |t\rangle$ to be the *unique* ground state.

Quantum Complexity Theory

The ground energy of $H_{\mathrm{in}} + H_{\mathrm{prop}}$ is 0. We now want to add H_{out} and show that the ground energy changes by an amount that is related to the acceptance probability of the circuit.

Standard perturbation theory would tell us that the first order shift in the ground energy caused by perturbing the Hamiltonian with H_{out} corresponds to the expectation of H_{out} in the original ground state, and is therefore related to the acceptance probability.

But if we leave the part of the input register containing the witness to be unconstrained, $H_{\mathrm{in}} + H_{\mathrm{prop}}$ would have a degenerate ground space (spanned by possible witness inputs). This exponential degeneracy is too much for perturbation theory to handle.

Quantum Complexity Theory

We begin with a QMA verifier that has completeness and soundness amplified exponentially close to 1 and 0. In the YES instance, we can check variationally that the ground energy of $H = H_{\text{in}} + H_{\text{prop}} + H_{\text{out}}$ is very near zero. In the NO instance, we need to show the ground energy is pushed up by some amount that is at least $1/\text{poly}(n)$.

To solve this problem and analyze the ground energy of $H = H_{\text{in}} + H_{\text{prop}} + H_{\text{out}}$ in a NO instance, Kitaev proved the following "geometrical lemma."

Lemma. Let A, B be positive semi-definite operators, each with a zero eigenspace $\ker(A), \ker(B)$ satisfying $\ker(A) \cap \ker(B) = \{0\}$. Let Δ_A, Δ_B denote the minimum non-zero eigenvalues of A and B, then

$$A + B \succeq \left(\min\{\Delta_A, \Delta_B\} 2 \sin^2 \theta/2\right) I$$

Where θ is the angle between $\ker(A), \ker(B)$.

Quantum Complexity Theory

In our setting, we take $A = H_{\text{in}} + H_{\text{out}}, B = H_{\text{prop}}$. Therefore

$$\min\{\Delta_A, \Delta_B\} = \Omega(T^{-2})$$

To compute the angle between the kernels, we use the fact that the cosine of the angle is the maximum inner product between vectors in the respective kernels.

$$\cos\theta = \max_{\eta \in \ker(A)} \max_{\xi \in \ker(B)} |\langle \xi | \eta \rangle|$$

This question asks, “what is the maximum overlap between valid history states, and states that are accepted at time t = T, given that this is a NO instance?” The answer is (1/T) times 1 – soundness.

The end result is that the ground energy is at least $\Omega(T^{-3})$ in a NO instance.

Quantum Complexity Theory

So far we have given a reduction from QMA verifier acceptance probabilities to the ground energy of a Hamiltonian. But is our Hamiltonian local?

If we represent the clock states $|t\rangle$ as binary strings, then we need at most $\mathcal{O}(\log n)$ qubits to represent the clock. Therefore terms like $U_t \otimes |t\rangle\langle t-1|$ are $\mathcal{O}(\log n)$-local.

This is where Kitaev's 1999 proof stopped. He showed that the $\mathcal{O}(\log n)$-local Hamiltonian problem is QMA-complete. (note that this also requires noting that the $\mathcal{O}(\log n)$-local Hamiltonian problem is in QMA).

Quantum Complexity Theory

A bit about history: the form of H_{prop} was introduced by Feynman around 1985 in the context of classical computers built from microscopic components. He called the clock a “pointer” and did not look at ground states, but rather time evolution generated by H_{prop} .

Feynman’s gates were classical and reversible; in part because quantum logic gates were not yet widely conceptualized.

None of the quantum computer scientists read Feynman’s old papers, so it was left to the Kitaev to find this gem.

As the story goes, Kitaev agreed to give a talk on “Quantum NP” at Hebrew University, and worked out the details (defining quantum NP and proving the quantum Cook-Levin theorem) on the flight to Israel where he gave the talk.

Richard Feynman

Alexei Kitaev

Quantum Complexity Theory

The reduction to a 5-local Hamiltonian came soon after Kitaev's original proof, and it is based on encoding the clock in unary:

$$|t\rangle \rightarrow |\underbrace{11...1}_{t\ \mathrm{times}}0...0\rangle$$

A new part of the Hamiltonian $H_{\mathrm{cons}} = \sum_{t=1}^{T} |01\rangle\langle 01|_{t-1,t}$ ensures that all clock qubit states with energy below 1 are valid unary encodings.

The propagation terms are now encoded as follows:

$$U_t \otimes |t\rangle\langle t-1| \rightarrow U_t \otimes |110\rangle\langle 100|_{t-1,t,t+1}$$

Which suffices to make them 5-local. See "Quantum NP: A survey" for more unary encoding details.

Quantum Complexity Theory

In 2016, Bausch and Crosson modified the Feynman-Kitaev construction to obtain a ground energy scaling like $\Omega(T^{-2})$ in the NO instance, instead of the $\Omega(T^{-3})$ from the geometrical lemma.

We joked about calling our paper, "One small trick for increasing the promise gap of the local Hamiltonian problem that Dr Kitaev doesn't want you to know."

It turned out the joke was on us: while we thought we improved the scaling compared to the standard Feynman-Kitaev construction, we subsequently proved the geometrical lemma is not tight. The ground energy for Kitaev's construction scales like $\Omega(T^{-2})$, by a more involved proof.

Then we proved that, of all possible (weighted, complex) graphs, the path graph assigns the best energy penalty to nonaccepting computations. So it is unlikely that anyone can improve the FK construction without totally altering the framework.

Quantum Complexity Theory

What are these weird Hamiltonian terms with clocks and projectors in them? What happened to Pauli operators?

$$-\left(U_t \otimes |110\rangle\langle 100|_{t-1,t,t+1} + U_t^\dagger \otimes |100\rangle\langle 110|_{t-1,t,t+1}\right)$$

There is a general method to reduce k-local Hamiltonians to r-local Hamiltonians with r < k using "perturbative gadgets." These gadgets are r-local Hamiltonians whose physics at low energy resembles the target k-local Hamiltonian. Some terms in the gadget have norm poly(n), while others have norm 1. The higher order interaction terms then appear at a higher order in perturbation theory.

Using these kinds of tricks, Biamonte and Love showed that the local Hamiltonian problem is QMA-complete for Hamiltonians of the form:

$$\begin{aligned} H_{\text{ZZXX}} &= \sum_i h_i \sigma_i^z + \sum_i \Delta_i \sigma_i^x + \\ &+ \sum_{i,j} J_{ij} \sigma_i^z \sigma_j^z + \sum_{i,j} K_{ij} \sigma_i^x \sigma_j^x . \end{aligned}$$