CMSC 858L: Quantum Complexity

Instructor: Daniel Gottesman

Spring 2023

14 Local Hamiltonians and QMA

14.1 References

For the local Hamiltonian problem, see chapter 14 of Kitaev, Shen, and Vyalyi, *Classical and Quantum Computation*.

14.2 Project

Since we're at the halfway point of the class, it is time to start thinking about the project if you haven't already. To recap: The project is 35% of the grade, and should be performed in a group of 1–5 people. You should read and boil down the results from some research papers and write a written report of 10–30 pages. Each group will also give a brief presentation in the last week of class (May 9 and 11). The length of the presentation will depend on how many groups we end up with. The written report will also be due on May 11.

The first step is to form your groups and pick a topic. I would like everyone to email me your group and topic by April 4, 3 weeks from today. That will give you 5 weeks to do the project. Acceptable topics include basically anything relating to quantum complexity that we haven't actually done in class. Even the things we are doing in class generally are not exhaustive examinations of the topic and there is much more known about them that could be discussed. If you need help coming up with ideas, have questions about what topics we will cover in class or if something is acceptable, or need help finding suitable papers on the topic of your choice, let me know.

14.3 Local Hamiltonians

Working our way up to larger complexity classes, our next stop is QMA. Recall the definition of QMA:

Definition 1. *Fix a particular efficiently computable approximately universal gate set $\mathcal{G}$. Let QMA be the set of languages L such that there exists a uniform family of polynomial-size quantum circuits $Q_{n,m}$ (with gates drawn from $\mathcal{G}$) that takes two inputs $(x, |\psi\rangle)$ with the following properties: For any instance x,*

1. *If $x \in L$, then exists an n-qubit state $|\psi\rangle_x$ with $n = O(\mathrm{poly}(|x|))$ and $\mathrm{Prob}(M_{Q_{|x|,n}}(|x\rangle \otimes |\psi\rangle_x) = 1) \geq 2/3$.*
2. *If $x \notin L$, then for any n-qubit state $|\psi\rangle_x$ with $n = O(\mathrm{poly}(|x|))$, $\mathrm{Prob}(M_{Q_{|x|,n}}(|x\rangle \otimes |\psi\rangle_x) = 0) \geq 2/3$.*

In the first case, w_x is the witness *for $x \in L$.*

As the most natural quantum analogue of NP, a natural question to ask about it is what does a QMA-complete problem look like? The first one was found by Kitaev, and that is the *local Hamiltonian* problem. A *Hamiltonian* is a Hermitian (i.e. self-adjoint) operator which represents the energy of a system. Hamiltonians have an important role in physics because they determine almost all the properties of the system in question, including how they evolve in time given an initial state, their behavior at thermal equilibrium,

and properties of excitations. The spectrum of the Hamiltonian (the list of eigenvalues) already says a lot about a Hamiltonian, even without information about the eigenvectors themselves. At very low energies, the lowest eigenvalues are the most important, and the *ground state*, the lowest-energy state, most of all. (However, this to an extent dependent on the energies of other low-lying states.)

But what is a *local* Hamiltonian?

Definition 2. *A* k-local Hamiltonian *is a Hermitian operator* H *which can be written in the form*

$$H = \sum_S H_S \otimes I, \tag{1}$$

where the sum is taken over subsets S *of at most* k *qubits and* H_S *is a Hamiltonian acting only on the qubits in* S*.*

Thus, a k-local Hamiltonian is a Hamiltonian that is a sum of terms each of which acts on at most k qubits. Here, k is usually assumed to be a constant independent of the number of qubits in the system.

Note that we are not making any constraints on where these k qubits are located in physical space. They can be far apart from each other. If we want a Hamiltonian where all terms only interact qubits that are physically near each other, we will call it a *geometrically local* Hamiltonian or just one that is *local in* D *dimensions*, where D is the geometric dimension of the space containing the qubits' physical locations. I want to use the term *geometric dimension* here (or *spatial dimension* because sometimes we also want to talk about local Hamiltonians using qudits which have a Hilbert space of dimension d. Then the Hilbert space dimension d of each qudit is called the *local dimension*. The local Hamiltonian problem is unfortunately full of terminology like this that can be interpreted two ways. There is actually one more case of confusing terminology that we will get to soon.

The local Hamiltonian problem is then to determine the ground state energy (the lowest eigenvalue) of a local Hamiltonian. However, because the energy is a real number, we are going to want to limit the precision we need:

Definition 3. *The language* k-LOCAL HAMILTONIAN *works with instances of the form* (H, E, Δ) *with the promise that* $0 < \Delta = \Omega(1/\mathrm{poly}(n))$ *(where* n *is the size of the instance) and the lowest eigenvalue (*ground state energy*) of* H *is either* $\leq E$ *or* $\geq E + \Delta$*. A "yes" instance is* (H, E, Δ) *such that the ground state energy of* H *is* $\leq E$*.*

Δ is essentially the precision that we need to specify the Hamiltonian's ground state energy to. k-LOCAL HAMILTONIAN is a decision problem, but clearly there is a related search problem: Find the ground state energy to precision Δ. If we have an algorithm for the decision problem, we can solve the search problem by trying different values of E until we narrow down the ground state energy. Note that there is a potential hiccup depending on what the algorithm does when the promise is not satisfied (i.e., the ground state energy is between E and $E + \Delta$, but we can deal with this by trying a few different (E, Δ) pairs at each step before updating our window on E.

Δ is sometimes called the *promise gap* of the instance. This should be distinguished from the *spectral gap* of the Hamiltonian, which is the difference between the lowest two eigenvalues. (Yes, this is the other example of confusing terminology I was mentioning.) The spectral gap has an important impact on the physics of the system, but the promise gap is just something somewhat artificial introduced to make this a problem in QMA rather than some much higher class. (Although precision is a perfectly reasonable physics concept too.) Spectral gap vs. promise gap is a bit more confusing generally than the terminology "local" and "dimension" because they both refer to something about the energies. In some example instances, the two are actually the same, but that is not generally true.

Now, what is the complexity of k-LOCAL HAMILTONIAN? The first thing we can say is:

Theorem 1. *k-LOCAL HAMILTONIAN is in QMA.*

Proof sketch. There exist efficient quantum algorithms to simulate the dynamics of any k-local Hamiltonian on an initial state provided to us. Thus, if Merlin claims that an instance of k-LOCAL HAMILTONIAN is

a "yes" instance, he can provide us with a witness which is a copy of the ground state $|\psi\rangle$. We then create an ancilla in the state $\sum_t |t\rangle$ and conditionally run the simulation of the Hamiltonian H on $|\psi\rangle$ for a time t. If $|\psi\rangle$ is indeed an eigenstate of energy E_0, then the Hamiltonian gives a phase $e^{-iE_0t/\hbar}$ after a time t. The ancilla thus ends up unentangled with the witness and in the state $\sum_t e^{-iE_0t/\hbar}|t\rangle$. The quantum Fourier transform will then let us measure E_0. There are some details to fill out here (which I will skip), but we don't get E_0 exactly, but merely some approximation $\tilde{E}_0$ to E_0 whose accuracy is inverse polynomial in the size of the simulation and ancilla. We pick our circuits so that the accuracy is less than $\Delta/2$ with probability $\geq 2/3$.

If we find $\tilde{E}_0 < E + \Delta/2$, then we can accept the witness. In the case that Merlin is being honest and has provided us with a valid witness with energy at most E, we will accept it with high probability. On the other hand, if this is a "no" instance and Merlin provided us a different eigenstate, then $E_0 \geq E + \Delta$, so with probability at least 2/3, $\tilde{E}_0 > E + \Delta/2$, and we will reject the witness.

But what if we have a "no" instance and Merlin gives us something which is not an eigenstate? We can expand an arbitrary state $|\psi\rangle = \sum_i \alpha_i |\psi_i\rangle$, where $|\psi_i\rangle$ run over an orthonormal basis of eigenstates with energies E_i. Since it is a "no" instance, $E_i \geq E + \Delta$ for all i. What happens when we run our simulation? We now get the state

$$\sum_i \alpha_i \sum_t e^{-iE_it/\hbar}|t\rangle|\psi_i\rangle. \tag{2}$$

This is entangled between the ancilla and system. Since the eigenstates are orthogonal to each other, if we trace out the witness, the state of our ancilla is a mixture over pure states $\sum_t e^{-iE_it/\hbar}|t\rangle$ with probability $|\alpha_i|^2$. When we measure, we get an estimate of E_i.

However, there is still one more complication: The Fourier transform doesn't give us a direct estimate of E_i, but rather of a multiple of $1/E_i$; you need multiple such estimates to actually learn E_i. If we ask for multiple copies of the witness, we could measure it multiple times, but we would be getting different E_i each time, making it hard to extract the answer. However, we can just measure it repeatedly on the same copy: In the case where the witness is genuinely an eigenstate, this doesn't disturb the state (modulo imperfections in the simulation), and in the case where it is superposition we get

$$\sum_i \alpha_i \left(\sum_t e^{-iE_it/\hbar}|t\rangle\right) \otimes \left(\sum_t e^{-iE_it/\hbar}|t\rangle\right) \cdots \otimes |\psi_i\rangle. \tag{3}$$

That is, in this case, we get the *same* E_i for each ancilla used. We can put them together to get a single value of E_i out. Which one is random, but they are all at least $E + \Delta$, so will always fail the witness check with probability at least 2/3. □

Note that the fact that the Hamiltonian is local is playing an important role here, as is the accuracy Δ. If the Hamiltonian is arbitrary, we don't have an efficient algorithm to simulate the Hamiltonian and therefore have no way to estimate the energy. If Δ is too small, we don't have a way to efficiently learn the energy to sufficient accuracy to distinguish between the yes and no instances. Also note that we really need a quantum simulation here, unless the ground state has an efficient classical description and we have an efficient classical algorithm to compute the energy with this Hamiltonian of a state provided. This problem is not generally in NP or MA.