CMSC 858L: Quantum Complexity

Instructor: Daniel Gottesman

Spring 2023

17 More QMA-Complete Problems

17.1 References

A survey of QMA-complete problems (as of 2012) is given in Bookatz, "QMA-complete problems," arXiv:1212.6312.

Hamiltonians in 2 spatial dimensions are discussed in Aharonov, van Dam, Kempe, Landau, Lloyd, and Regev, "Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation," quant-ph/0405098 and Oliveira and Terhal, "The complexity of quantum spin systems on a two-dimensional square lattice," quant-ph/0504050. Hamiltonians in 1 spatial dimension are discussed in Aharonov, Gottesman, Irani, and Kempe, "The power of quantum systems on a line," arXiv:0705.4077 and Hallgren, Nagaj, and Narayanaswami, "The Local Hamiltonian problem on a line with eight states is QMA-complete," arXiv:1312.1469. Translationally-invariant Hamiltonians are discussed in Gottesman and Irani, "The Quantum and Classical Complexity of Translationally Invariant Tiling and Hamiltonian Problems," arXiv:0905.2419.

Consistency of local density matrices was proven QMA-complete in Liu, "Consistency of Local Density Matrices is QMA-complete," quant-ph/0604166. (This was proved for Karp reductions in Broadbent and Grilo, "QMA-hardness of Consistency of Local Density Matrices with Applications to Quantum Zero-Knowledge," arXiv:1911.07782. Inequivalence of circuits was proven QMA-complete by Janzing, Wocjan, and Beth, "Non-identity check is QMA-complete," quant-ph/0305050.

17.2 Further variations of LOCAL HAMILTONIAN

We now a full classification of the complexity of the k-LOCAL HAMILTONIAN problem for different values of k. But because k-local doesn't correspond to geometric locality, this is still far from the more physically interesting Hamiltonians. What about Hamiltonians that are constrained to short-range or nearest-neighbor interactions in a small number of spatial dimensions?

It turns out that these problems can be QMA-complete as well. To prove this, we need to continue to work on the clock. We can use a checking circuit that can be laid out in 1 dimension without much difficulty — use SWAP gates to move qubits around as needed and then back, which only introduces polynomial overhead. But if all qubits need to interact with the clock, which would be stored at a particular location, that would introduce violations of geometric locality. Instead, the usual approach is to find some way of encoding the clock in the spatial position of some counter that moves around between the qubits and performs the appropriate gates at the appropriate "time." Typically, the strategy is to come up with some sort of classical process that can be implemented via geometrically local transition rules and geometrically local constraints and then encode the history of that with a geometrically local Hamiltonian.

This usually involves increasing the local dimension — the size of the individual qudits in the system. In 2 spatial dimensions, the local dimension can be reduced again to 2 (qubits) using perturbation theory gadgets, showing that 2-LOCAL HAMILTONIAN is QMA-complete using Hamiltonians with nearest-neighbor interactions in 2 spatial dimensions. In 1 spatial dimension, there is not enough room to apply the perturbation theory gadgets, so the problem is open for 2-local qubit Hamiltonians in 1 spatial dimension, but we know that 2-local Hamiltonians using large enough qudits in 1 spatial dimension are QMA-complete. (I believe that the best current result is for a local dimension of 8.)

Note that for classical SAT-type problems (including MAX-k-SAT and variations with more than two possible states), the 1-dimensional problem is in P. (It can be solved by dynamic programming.) The 2-dimensional version is NP-complete again. The superposition in a quantum state in some ways takes the place of a spatial dimension.

Even so, these Hamiltonians are not that physically realistic since they involve highly structured but different terms in different locations. One can go even further and focus on translationally-invariant Hamiltonians instead, which use the same Hamiltonian term at every site, translated appropriately throughout the system. However, there is an immediate complexity issue, which is that specifying a single Hamiltonian term is just constant size. The only thing that gives us a family of problems of different sizes is the number of qudits in the system. This is the sort of thing that complexity theorists would normally want to specify in unary. However, in that case, there is just one problem instance of any given size, so the problem is automatically solvable in P/poly since the advice for a given size can just specify the answer to the single instance of that size.

The solution to still have an interesting complexity problem is to instead specify the number of qudits in binary, meaning an n-qudit instance has size $\log n$. However, measuring the energy takes a time polynomial in n, which is exponential in $\log n$; the problem is no longer in QMA. Instead, it is in a class $\mathrm{QMA}_{\mathrm{EXP}}$, which is like QMA, but has exponentially large checking circuits. In fact, using similar ideas to those we have discussed above, one can show that 1-dimensional translationally-invariant Hamiltonians with large enough qudits are $\mathrm{QMA}_{\mathrm{EXP}}$-complete. The idea is to use a quantum Turing machine (which has translation-invariant transition rules) and encode the instance in the number of qudits n, which becomes the input to the Turing machine. Unfortunately, this whole process uses incredibly large (but constant size) qudits, so the existing constructions are still not at all realistic.

Many other ground state properties are QMA-complete or QMA-hard, basically anything that is measurable on a ground state is a good candidate. In some cases, the problems can be even harder. For instance, it turns out to be undecideable to determine if a system has a spectral gap as the number of particles goes to infinity (i.e., if the limit of the difference between the lowest two eigenvalues is non-zero).

17.3 Density Matrix Consistency

Another related QMA-complete problem is *density matrix consistency* and variations.

Definition 1. *Let $\mathcal{S}$ be a set of subsets S_μ, where each S_μ is a set of at most k qubits. If ρ is a density matrix, then let $\rho(\mu) = \mathrm{Tr}_{i \notin S_\mu} \rho$ (i.e., ρ on only the qubits in S_μ). Then k-CONSISTENCY is a language consisting of instances $(\mathcal{S}, \{\rho_\mu\}, \epsilon)$ (with all quantities specified to polynomially many bits of precision) with the promise that either:*

1. *$\exists$ density matrix ρ such that $\rho(\mu) = \rho_\mu$ ("yes" instances)*
2. *$\forall$ density matrices σ, $\|\sigma(\mu) - \rho_\mu\|_1 \geq \epsilon$ for some μ. ("no" instances).*

In other words, the k-CONSISTENCY problem asks if there is a *global* density matrix consistent with a set of local density matrices we are provided with. It might not be clear why this is even a question, let alone why it might be hard. Certainly if $k = 1$, so all the local density matrices are single-qubit ones, then there is always a global density matrix consistent with them, the tensor product of the local values. Once $k \geq 2$, however, the local density matrices can be specified on overlapping sets of qubits, and that offers a possibility for inconsistency. Sometimes this is trivial and obvious, such as if $\rho_{12} = |00\rangle\langle 00|$ and $\rho_{23} = |11\rangle\langle 11|$: Qubit 2 can't be both $|0\rangle$ (as required by ρ_{12}) and $|1\rangle$ (as required by ρ_{23}. Sometimes it is less obvious, depending on the global structure of the density matrices. For instance, suppose $\rho_{i,i+1} = \frac{1}{2}(|01\rangle\langle 01| + |10\rangle\langle 10|)$, and qubit N is the same as qubit 0. That is, adjacent qubits on a ring of N qubits have the opposite value. This is possible when N is even but not when N is odd. Even more subtle global structures are possible.

However, this problem is in QMA. The basic idea is that the witness is the global density matrix ρ that is consistent with all ρ_μ. To verify, choose a set S_μ at random and make a random measurement Π on $\rho(\mu)$. Accept if the measurement outcome is 1; reject otherwise. If the witness is a good one, it will be accepted

with probability $\operatorname{Tr}\Pi\rho_\mu$, a known quantity, and if it is not a good witness (in the sense of condition 2 of the definition of k-CONSISTENCY), there is a polynomially small chance we have picked a μ such that $\|\rho(\mu) - \rho_\mu\|_1 \geq \epsilon$, in which case, there is a constant probability that we pick a measurement for which $\operatorname{Tr}\Pi\rho_\mu$ is significantly different from $\operatorname{Tr}\Pi\rho(\mu)$. Overall, there is thus a polynomially small difference in acceptance probability between the "yes" and "no" instances. This isn't the 2/3 vs. 1/3 difference we require for the usual definition of QMA, but it can be turned into that through the amplification techniques we will discuss later.

I won't prove QMA-completeness of this problem, but we can understand qualitatively the relationship with LOCAL HAMILTONIAN: Let ρ be some ground state of a k-local Hamiltonian H (either a pure ground state or a mixed state which is a mixture of degenerate ground states). Since H is k-local, it is a sum $\sum_\mu H_\mu$, where H_μ acts on the subset S_μ of qubits. Then the ground state energy of H is

$$\operatorname{Tr}(\rho H) = \sum_\mu \operatorname{Tr}(H_\mu \rho) = \sum_\mu \operatorname{Tr}(H_\mu \rho(\mu)). \tag{1}$$

That is, the ground state energy depends on only the local density matrices $\rho(\mu)$. One way we might try to find the ground state energy is by separately optimizing $\operatorname{Tr}(H_\mu\rho_\mu)$ for each μ, but once we've done that, we need to make sure that the separate ρ_μ we have found are all consistent with a single global ρ — the k-CONSISTENCY problem.

In fact, this is not good enough, since usually separately optimizing the terms $\operatorname{Tr}(H_\mu\rho_\mu)$ *doesn't* give us a globally valid state. There might be multiple solutions for each term, leaving ambiguity about which set of choices we should make (and there are exponentially many sets of choices if each subset of qubits has 2 or more choices), or it might be that the Hamiltonian is frustrated and it is instead a question of compromising by violating each term as little as possible. However, the overall goal, to minimize $\sum_\mu \operatorname{Tr}(H_\mu\rho_\mu)$, is a *convex optimization* problem, and there are good algorithms for that. The wrinkle here is that we must do so under the constraint that the ρ_μ are globally consistent, but if we somehow have an oracle for the k-CONSISTENCY problem, we are still able to run this convex optimization algorithm and find the local density matrices of that Hamiltonian's ground state. Thus, if we can solve k-CONSISTENCY, we can solve k-LOCAL HAMILTONIAN and therefore any problem in QMA when $k \geq 2$.

Note that this reduction is a *Turing reduction* rather than the *Karp reductions* we discussed before: we run a general program which makes use of the hypothetical subroutine to solve k-CONSISTENCY, as opposed to converting instances of k-LOCAL HAMILTONIAN to instances of k-CONSISTENCY. However, recently a different reduction of essentially the same problem has been found which is a Karp reduction.

17.4 Circuit Non-Equivalence

Another type of QMA-complete problem are properties of quantum circuits. One good example is *circuit non-equivalence*, in which we are given two circuits and wish to determine if they compute the same unitary or not:

Definition 2. *The CIRCUIT NON-EQUIVALENCE problem consists of instances of the form (C_1, C_2, a, b) where C_1 and C_2 are polynomial-size quantum circuits on n qubits (involving efficiently computable gates) computing the unitaries U_1 and U_2 respectively, and $b - a \geq \Omega(1/\mathrm{poly}(n))$ with the following promise: Either*

1. *$\forall\phi$, $\|U_1 - e^{i\phi}U_2\|_{sup} \geq b$ ("yes" instances)*
2. *$\exists\,\phi$ such that $\|U_1 - e^{i\phi}U_2\|_{sup} \leq a$ ("no" instances).*

It is not surprising that this is in QMA: The witness is simply a few copies of a state $|\psi\rangle$ that distinguishes the two unitaries, one for which $\|U_1|\psi\rangle - e^{i\phi}U_2|\psi\rangle\| \geq b$. There is in this case a measurement that distinguishes the states $U_i|\psi\rangle$, so to verify the witness, just run the two circuits on $|\psi\rangle$ and make the measurement to see difference results. The complication is again that you really need multiple copies to run both circuits and see the difference in statistics, but this can be dealt with.

This problem is QMA-complete. In fact, the special case NON-IDENTITY CHECK, where one of the circuits is the identity, is also QMA-complete. It is not hard to show this directly from the definition of QMA: Let V be the verification circuit for some QMA language, amplified so that the success probability is very high. Then we run V, copy the output qubit to an ancilla, then run $V^\dagger$. This is C_1. C_2 is just the identity. For a "yes" instance, there exists some input state (a valid witness) such that the overall effect is to flip the ancilla qubit, which is definitely not the identity. For the "no" instance, the ancilla isn't changed and then $V^\dagger$ inverts V, returning the system to its initial state, an overall identity operation.

17.5 QCMA

Recall that QCMA is like QMA, but with the difference that the witness must be a classical string instead of a quantum state. (However, the verification circuit is still quantum, which is what distinguishes QCMA from MA.) I think it is fair to say that most people believe that QCMA $\neq$ QMA, but there is not overwhelming evidence of this. There are known oracle separations between QCMA and QMA, but of course that is far from definitive. Mostly I would say this belief is based on the understanding that quantum states can be very complicated and providing one is potentially much more powerful than providing simply classical information.

If we take a QMA-complete problem and add extra promises to make sure that the witness is classically describable, we tend to get a QCMA-complete problem. For instance, if we consider k-local Hamiltonians with the additional promise that the ground state can be created by a polynomial-size quantum circuit, then it is a QCMA-complete problem. Or we could use NON-IDENTITY CHECK with the extra constraint that U_1 acts differently from the identity on some basis state input.