Week 7: The QMA-Completeness of Local Hamiltonian

Stacey Jeffery

March 19, 2025

7.1 Introduction

While we believe that quantum computers can solve some problems significantly faster than classical computers, they also have limits, and quantum complexity theory tries to understand these. We have already seen how we can use quantum query lower bounds to understand these limits. However, quantum query lower bounds are limited in the following sense. If the oracle is instantiated by an input *string* $x \in \{0,1\}^N$, then a lower bound of $q(N)$ on $\mathsf{Q}(\mathrm{F})$ is a real lower bound on the time complexity of F as a function of its input size, N, but it can't be larger than N, so we cannot prove super-linear lower bounds in the input size. On the other hand, if we instantiate the oracle with a function, $f_x : [N] \to \{0,1\}$ defined somehow by an input string x of a possibly much smaller length n, then a lower bound of $q(N) \leq N$ might be superlinear in n, the input size – it could even be exponential in n. However, the quantum query lower bound doesn't strictly apply in such a setting, because by restricting the size of x, we have restricted the number of possible inputs f_x, which introduces structure to the problem that might be exploited to circumvent the lower bound. This is why, even though we have a lower bound of $\Omega(\sqrt{N})$ on the quantum query complexity of black-box search, this doesn't imply a lower bound on satisfiability, which we can see as an instantiation of black-box search, letting the input oracle be defined by the input formula.

So we actually don't have good ways to prove non-trivial lower bounds on the quantum complexity of concrete problems, but we shouldn't feel too bad about that, since we also can't prove such lower bounds classically. Instead, we define complexity classes characterizing different sets of problems we think might have different difficulties, and make conjectures about these, and even occasionally prove things. In this chapter, we will see a very small taste of this, as we will consider the quantum analogue of the class NP, and a complete problem for it. These lecture notes are partly based on [dW19, Chapter 14].

7.2 Warm up: Satisfiability and NP

A *circuit* is similar to a formula, which we have already seen, except that a formula is always a tree, whereas the nodes of a circuit are allowed to have multiple "parents" (i.e. outgoing edges). The problem of circuit satisfiability is, given a circuit C as input, decide if there exists an input x to the circuit such that $C(x) = 1$.

Note that this problem is much more difficult than *circuit evaluation*, analogous to formula evaluation $\textsc{Eval}_C(x)$, which we have seen: in this problem, a family of circuits C is a parameter of the problem, and the input consists of an input x to the circuit, and the output is $C(x)$. In contrast, for circuit satisfiability, the input is C, and we seem to need to search over all possible strings $x \in \{0,1\}^n$ to find one such that $C(x) = 1$.

Complexity Classes Recall that a complexity class is a set of *languages*, which are subsets $L \subseteq \{0,1\}^*$. This is just another way of describing *total functions*. For example, P is the set of languages L for which there exists a polynomial-time Turing Machine $\mathcal{A}_L(\cdot)$ such that $\mathcal{A}_L(x)$ outputs 1 if and only if $x \in L$.

The class NP is defined as the class of problems for which there is a polynomial-time *verifier*. That is, it is the set of languages $L \subseteq \{0,1\}^*$ such that there exists a polynomial-time Turing Machine $\mathcal{V}_L(\cdot,\cdot)$ that takes two inputs, $x \in \{0,1\}^*$ and $w \in \{0,1\}^{n(|x|)}$ for some polynomial n, and such that:

Soundness If $x \in \{0,1\}^* \setminus L$, no witness can convince $\mathcal{V}_L$ to accept x: i.e. for all $w \in \{0,1\}^{n(|x|)}$, $\mathcal{V}_L(x,w) = 0$.

Completeness If $x \in L$, there is witness that will convince $\mathcal{V}_L$ to accept x: i.e. exists $w \in \{0,1\}^{n(|x|)}$ such that $\mathcal{V}_L(x,w) = 1$.

Equivalently, we could let $\mathcal{V}_L(x,\cdot)$ be a polynomial-time *circuit* that depends on x, and is *uniformly generated*, meaning that there is a polynomial-time Turing Machine that writes down the circuit $\mathcal{V}_L(x,\cdot)$ given input x.

The problem SAT is NP-complete, which means that it's in NP, and if you can solve it in polynomial time, then you can solve any problem in NP in polynomial time. To see that it's in NP, note that for an input C, if $\mathcal{V}(C,x)$ takes as input[1] the circuit C and an input x to the circuit, and outputs $C(x)$ – that is, it evaluates the circuit – then it is an NP-verifier for SAT. If there exists x such that $C(x) = 1$, that x is a witness that makes $\mathcal{V}$ accept. If no such x exists, then $\mathcal{V}$ will never accept.

Seeing that every problem in NP can be reduced to SAT is much more involved, but it's essentially because Turing Machines and Circuits are equivalent models of computation. If we have any NP problem, L and an input x, the verifier $\mathcal{V}_L(x,\cdot)$ is an instance of SAT.

However, we can actually show something stronger, which is that if we restrict SAT to circuits of a very specific form, the resulting problem is still NP-complete. A clause C on $\{x_1,\ldots,x_n\}$ is the dysjunction (or) of some subset of $\{x_1,\ldots,x_n,\bar{x}_1,\ldots,\bar{x}_n\}$, which are called literals. For example,

$$C = x_1 \vee x_6 \vee \bar{x}_7.$$

We say C is a k-clause if it is a dysjunction of k literals. For example, C above is a 3-clause. We say C is satisfied by an assignment x to the variables if $C(x) = 1$.

Problem: k-SAT

Input: Some description of $C = \bigwedge_{j=1}^{m} C_j$ where each C_j is a k-clause on $\{x_1,\ldots,x_n\}$.
Output: 1 if $\exists x \in \{0,1\}^n$ such that $C(x) = 1$ and 0 if $\forall x \in \{0,1\}^n$, $C(x) = 0$.

It turns out that 3-SAT is NP-complete (or k-SAT for any $k \geq 3$), whereas 2-SAT can be solved in polynomial time, so it's in P (thus, it is not known to NP-complete, and it is NP-complete if and only if P = NP). We describe at a very high level how to reduce any NP problem instance to an instance of 3-SAT. Let $\mathcal{V}_L$ be an NP verifier for any NP problem L. Let

$$\Psi(x) = (\psi_0(x,w), \psi_1(x,w), \ldots, \psi_T(x,w))$$

be the states of $\mathcal{V}_L(x,w)$ at steps 0 (no computation performed yet) through T (last step of computation applied). Without loss of generality suppose the answer to the computation is encoded in the first bit of $\psi_T(x,w)$, meaning that this bit is 1 if and only if $\mathcal{V}_L(x,w)$ accepts.

Then $\Psi(x)$ is a string of length $n' = (T+1)S = \mathsf{poly}(n)$, where S is the number of bits used by $\mathcal{V}_L$. The idea is to define a set of 3-clauses on n' variables, organized into $T+1$ blocks of S variables, $\psi_t = \psi_t[1],\ldots,\psi_t[S]$ for $t = 0,\ldots,T$, such that all clauses are satisfied by the assignment $\Psi(x)$, and any assignment that satisfies all clauses must represent a run of the algorithm that accepts x.

[1] Note that here x is the witness rather than the input to the problem.

Thus, some of the clauses will check that $\psi_0, \ldots, \psi_T$ represents a valid run of $\mathcal{V}_L(x, \cdot)$. For example, if the t-th operation in $\mathcal{V}_L$ is to compute the OR of the first two bits and store the value in the third bit, we want to check that $\psi_t[3] = \psi_{t-1}[1] \vee \psi_{t-1}[2]$, where we interpret each ψ_t as the S-bit string representing the t-th block of variables, that we index into with square brackets. Since the expression $x = y \vee z$ only involves three variables, we can write it as a conjunction of 3-clauses. In particular, it is satisfied whenever:

$$(x \vee y \vee \bar{z}) \wedge (x \vee \bar{y} \vee z) \wedge (x \vee \bar{y} \vee \bar{z}) \wedge (\bar{x} \vee y \vee z).$$

Thus, including these 4 clauses with $(x, y, z) = (\psi_t[3], \psi_{t-1}[1], \psi_{t-1}[2])$, means an assignment to the variables will only be accepted if the 3-rd bit of the t-th block of S bits is obtained from the $(t-1)$-th block by taking the OR of its first two bits. Using sufficiently many constraints like this, we can ensure that an assignment to the variables only satisfies all constraints if it does indeed represent a run of the algorithm $\mathcal{V}_L$, and by adding a 1-variable constraint

$$\psi_T[1]$$

we can ensure that we only accept if ψ_T has a 1 in the first bit, meaning the verifier accepts.

7.3 QMA

There are a number of natural quantum analogues to NP [Gha23], but the most popular one (and probably the most well-justified) is QMA. This is actually the quantum analogue of a randomized variant of NP called MA, for *Merlin-Arthur*. MA and QMA are defined for promise problems, which generalize languages to disjoint pairs $L = (L_0, L_1)$ such that $L_0, L_1 \subseteq \{0,1\}^*$. In such a problem, we want to decide if $x \in L_0$ or $x \in L_1$, promised that one of these is true. Like NP, the definition of MA involves a verifier. We imagine an interaction between Arthur, who is honest, but computationally bounded, and Merlin, who is computationally unbounded, but can't be trusted. Arthur gives Merlin an input x, and Merlin's goal is to try to convince Arthur that $x \in L_1$ by sending him a witness w. Informally, L is in MA if there is a polynomial-time algorithm Arthur can run such that if $x \in L_1$, Merlin can convince him with high probability to accept x, but if $x \in L_0$, then Merlin cannot convince him with very high probability to accept. Formally, we can define MA as follows.

Definition 7.3.1. *Let $a, b \in [0,1]$ be such that $b > a$. Let $\mathsf{MA}_{a,b}$ be the class of problems $L = (L_0, L_1)$ such that there exists a uniformly generated polynomial-time circuit $\mathcal{V}_L(x, \cdot)$ that takes an input $w \in \{0,1\}^n$ for n polynomial in $|x|$, such that*

Soundness *If $x \in L_0$, no witness can convince $\mathcal{V}_L(x, \cdot)$ to accept, except with small probability: i.e. for all $w \in \{0,1\}^n$, $\Pr[\mathcal{V}_L(x, w) = 1] \leq a$.*

Completeness *If $x \in L_1$, there is a witness that will convince $\mathcal{V}_L(x, \cdot)$ to accept, with high probability: i.e. exists $w \in \{0,1\}^n$, such that $\Pr[\mathcal{V}_L(x, w) = 1] \geq b$.*

Define $\mathsf{MA} := \mathsf{MA}_{1/3,2/3}$.

In the quantum version of this class, QMA, we allow the verifier to be a quantum circuit, and the witness to be a quantum state. Without loss of generality, we can assume the verifier accepts if the first qubit is $|1\rangle$, and otherwise rejects, so we will

> let Π_{acc} be the projector onto having $|1\rangle$ in the first qubit register. **[1]**

Note that, for any circuit, since its action is unitary, there are states that the circuit will map to having $|1\rangle$ in the first qubit. Thus, we allow the algorithm to use some auxiliary space, which starts in the state $|0^{n_1}\rangle$. Formally, we have the following:

Definition 7.3.2. *Let $a, b \in [0,1]$ be such that $b > a$. Let $\mathsf{QMA}_{a,b}$ be the class of problems $L = (L_0, L_1)$ such that there exists a uniformly generated polynomial-time quantum circuit $\mathcal{V}_L(x)$ defined by 2-local gates $U_1, \ldots, U_T$ on $\mathbb{C}^{2^{n_1+n_2}}$ for T, n_1, n_2 all polynomial in $|x|$, such that*

Soundness *If $x \in L_0$, no witness can convince $\mathcal{V}_L(x)$ to accept, except with small probability: i.e. for all $|w\rangle \in \mathbb{C}^{2^{n_2}}$, $\|\Pi_{\text{acc}}\mathcal{V}_L(x)|0^{n_1}\rangle|w\rangle\|^2 \leq a$.*

Completeness *If $x \in L_1$, there is a witness that will convince $\mathcal{V}_L(x)$ to accept, with high probability: i.e. exists $|w\rangle \in \mathbb{C}^{2^{n_2}}$, such that $\|\Pi_{\text{acc}}\mathcal{V}_L(x)|0^{n_1}\rangle|w\rangle\|^2 \geq b$.*

Define $\mathsf{QMA} := \mathsf{QMA}_{1/3,2/3}$.

How important is the choice of a and b to the definition of QMA? It follows from the definition that whenever $a < a' < b' < b$,

$$\mathsf{QMA}_{a,b} \subseteq \mathsf{QMA}_{a',b'}. \tag{7.1}$$

On the other hand, you can take any QMA verifier, repeat it $\Theta(r)$ times, and take the majority, to get a verifier[2] with soundness 2^{-r} and completeness $1 - 2^{-r}$. If $r = \mathsf{poly}(|x|)$, then the verifier is still polynomial. Combined with (7.1), we have that for any polynomial r,

$$\mathsf{QMA} = \mathsf{QMA}_{2^{-r},1-2^{-r}}.$$

This is convenient, because it allows us to assume the verifier has smaller than constant error, which we will take advantage of in the next section. On the other side, for any polynomial r, we also have

$$\mathsf{QMA} = \mathsf{QMA}_{\frac{1}{2}-\frac{1}{r},\frac{1}{2}+\frac{1}{r}}.$$

An interesting case is when $b = 1$. For $b = 1$ and $a = 1/3$ (or some other constant), the resulting class is called QMA_1. It is not known whether $\mathsf{QMA} = \mathsf{QMA}_1$. On the one hand, the analogous classical complexity classes are equal. On the other hand, QMA_1 doesn't even have a unique definition, because it depends on the gate set used. That's because mapping between different complete quantum gate sets introduces small errors. So it's not even clear it makes sense to talk about "one-sided" quantum complexity classes.

7.4 Local Hamiltonian

As we have seen, a Hamiltonian can define a quantum computation: simulating that Hamiltonian. In theory, every quantum system (including any computation) is described by a Hamiltonian. Consider this in analogy to a *circuit*, which describes a classical computation. Just as we restricted our attention to circuits consisting of k-clauses, for characterizing QMA, we can restrict our attention to k-local Hamiltonians. The definition of locality for a Hamiltonian is identical to the definition we already know for unitaries, but we restate it precisely below.

Definition 7.4.1 (Local Hamiltonian)**.** *For a permutation $\pi \in S_n$ on n elements, let $S(\pi)$ be the unitary on $\mathbb{C}^{2^n}$ that permutes the n qubit registers according to π. A Hamiltonian H on $\mathbb{C}^{2^n}$ is* k-local *if there is a permutation $\pi \in S_n$ and a Hermitian matrix H' on $\mathbb{C}^{2^k}$ such that*

$$S(\pi^{-1})HS(\pi) = H' \otimes I_{2^{n-k}}.$$

In other words, H acts non-trivially on all but k qubits.

We can define the following for $a, b \in [0, m]$ such that $a < b$.

[2] This is not as trivial as it sounds, since the verifier doesn't get to have multiple copies of the witness. For details see [MW05] or [Reg06].

Problem: k-LocalHam$_{a,b}$

Input: Some classical description of $H = \sum_{j=1}^{m} H_j$ where each H_j is a k-local positive semidefinite[a] matrix on $\mathbb{C}^{2^n}$ with $\|H\| \leq 1$.
Output: 1 if $\exists$ a unit vector $|\psi\rangle \in \mathbb{C}^{2^n}$ such that $\|H|\psi\rangle\| \leq a$ and 0 if for all $|\psi\rangle \in \mathbb{C}^{2^n}$, $\|H|\psi\rangle\| \geq b$.

[a] Eigenvalues are real and non-negative.

The smallest $\|H|\psi\rangle\|$ for any $|\psi\rangle$ is just the smallest eigenvalue of H, also called the *ground energy*, so this problem is essentially asking for an estimate of the ground energy of H. Since the eigenvalues of each term H_j are in $[0, 1]$, the eigenvalues of H are in $[0, m]$.

This problem is in QMA whenever $b - a \geq \frac{1}{\text{poly}(n)}$. Without going into too many details, given an appropriate classical description of a Hamiltonian that is the sum of k-local terms, it is possible to simulate that Hamiltonian. A witness is a state $|\psi\rangle$ such that $\|H|\psi\rangle\| \leq a$.

We can relax the restriction that the terms H_j are positive semidefinite, and merely require them to be Hermitian, without changing the difficulty of the problem. On the other hand, as we will see in the following section, even when each H_j is a projector, this problem is QMA-complete for appropriate choices of k, a and b.

Before we see that k-LocalHam is QMA-complete, let us compare it with k-Sat. Let $C = \bigwedge_{j=1}^{m} C_j$ for some clauses C_j on n variables. For each $j \in [m]$, we can define an associated projector (which is a Hermitian matrix) on $\text{span}\{|x\rangle : x \in \{0,1\}^n\}$:

$$H_{C_j} = \sum_{x \in \{0,1\}^n : C_j(x) = 0} |x\rangle\langle x|.$$

If C_j is a k-clause, then this will be k-local. We can see this through an example. Suppose $C_j = x_1 \vee \bar{x}_2 \vee x_3$. Then $C_j(x) = 0$ when $x_1 = 0$, $x_2 = 1$ and $x_3 = 0$ (and the other bits of x take any values). Thus:

$$H_{C_j} = |0\rangle\langle 0| \otimes |1\rangle\langle 1| \otimes |0\rangle\langle 0| \otimes I_{2^{n-3}},$$

which is clearly 3-local. If each clause of C is a k-clause, then

$$H_C = \sum_{j=1}^{m} H_{C_j},$$

is an instance of k-local Hamiltonian. What does its ground energy have to do with the satisfiability of C? You will work this out in the following exercise.

Exercise 7.4.1. *Show that the ground energy of* H_C *is* $\min_{x \in \{0,1\}^n} |\{j \in [m] : C_j(x) = 0\}|$.

So C is satsifiable if and only if the ground energy of C is 0. More generally, this special case of k-LocalHam allows you to distinguish between formulas in which $\geq m - a$ clauses can be simultaneously satisfied, and formulas in which it is not possible to simultaneously satisfy more than $m - b$ clauses. Even for $b - a = \Omega(m)$, this problem is NP-hard.

We summarize the correspondance between classical satisfiability and quantum local Hamiltonian problems in Figure 7.1.

7.5 QMA-Completeness of Local Hamiltonian

We will show that 5-LocalHam$_{a,b}$ is QMA-complete, for some choice of a and b such that $b - a = \frac{1}{\text{poly}(n)}$.

Classical	Quantum
n variables	n qubits
m clauses	m Hamiltonian terms
assignment to the variables, x	quantum state, $\|\psi\rangle$
number of clauses satisfied by x	energy $\|\|H\|\psi\rangle\|\|$

Figure 7.1: Correspondance between classical satisfiability problems and local Hamiltonian.

Fix any QMA problem $L = (L_0, L_1)$, and let $\mathcal{V}_L(x)$ be a verifier for L, whose action is defined by 2-local unitaries $U_1, \ldots, U_T$. Let $|\psi_t(x)\rangle$ be the algorithm states of $\mathcal{V}(x)$ when given as witness a state $|w_x\rangle$ that maximizes its acceptance probability:

$$|\psi_0(x)\rangle = |0^{n_1}\rangle|w_x\rangle$$
$$\forall t \in [T],\ |\psi_t(x)\rangle = U_t|\psi_{t-1}(x)\rangle.$$

We will assume that this verifier has soundness $\frac{1}{T}$ and completeness $1 - \frac{1}{5T}$:

$$\forall x \in L_0,\ \|\Pi_{\text{acc}}|\psi_T(x)\rangle\|^2 \leq \frac{1}{T}$$
$$\forall x \in L_1,\ \|\Pi_{\text{acc}}|\psi_T(x)\rangle\|^2 \geq 1 - \frac{1}{5T}$$

In order to reduce L to k-LocalHam, we need to define a local Hamiltonian instance $H = H(x)$ where the ground energy is low whenever $x \in L_1$, and high whenever $x \in L_0$. To this end, we will define a local Hamiltonian instance where the ground state is the (normalized) *history state*:

$$|\Psi(x)\rangle := \frac{1}{\sqrt{T+1}} \sum_{t=0}^{T} |t\rangle_{\mathcal{C}}|\psi_t(x)\rangle \in \text{span}\{|t\rangle_{\mathcal{C}}|z\rangle : t \in \{0, \ldots, T\}, z \in \{0,1\}^{n_1+n_2}\} =: \mathcal{H}. \tag{7.2}$$

We have labelled the *clock* register, with a $\mathcal{C}$, for clarity. *Any* unit vector in $\mathcal{H}$ has the form $\frac{1}{\sqrt{T+1}} \sum_t |t\rangle|\phi_t\rangle$ for some (possibly unnormalized, possibly zero) vectors $|\phi_t\rangle$. We will define local Hamiltonians – in fact, each will be a projector – on $\mathcal{H}$ that "penalize" any state of this form that is not a history state, ensuring it ends up having too high energy (i.e. $\|H|\psi\rangle\|$ is too big), and thus ensuring only a history state can be a low energy state. But we will also have a local Hamiltonian, H_{final}, that penalizes any history state if the final state $|\psi_T(x)\rangle$ does not accept, ensuring that the energy is higher if the acceptance probability is too low.

7.5.1 Reduction to $(\log(T) + 2)$-local Hamiltonian

We first show a simpler reduction from L to $(\log(T)+2)$-$\text{LocalHam}_{a,b}$. This contains the main ideas of the reduction, and in Section 7.5.2, we use some tricks to modify the Hamiltonian terms so that they are 5-local.

First, define the following projectors, for $i \in [n_1]$:

$$H_{\text{init}}(i) = |0\rangle\langle 0|_{\mathcal{C}} \otimes I_2^{\otimes(i-1)} \otimes |1\rangle\langle 1| \otimes I_2^{\otimes(n_1-i)} \otimes I_{2^{n_2}}. \tag{7.3}$$

Clearly, $H_{\text{init}}(i)$ is $(\log(T)+1)$-local, since $\mathcal{C}$ is a $\log(T)$-qubit register, and $H_{\text{init}}(i)$ acts as the identity on all but $\mathcal{C}$ and one additional qubit. Together, these terms ensure that at time 0 – the only time in which the clock register is set to $|0\rangle_{\mathcal{C}}$ – the first n_1 qubits are set to $|0\rangle$ – any state not satisfying this will be mapped to something bigger than 0 by at least one of these projectors, resulting in an "energy penalty".

Next, define the following projectors, for $t = \{0, \ldots, T-1\}$:

$$\begin{aligned} H_t &= \frac{1}{2} \sum_{z \in \{0,1\}^{2^{n_1+n_2}}} (|t\rangle|z\rangle - |t+1\rangle U_{t+1}|z\rangle) \left(\langle t|\langle z| - \langle t+1|\langle z|U_{t+1}^\dagger\right) \\ &= \frac{1}{2}\left(|t\rangle\langle t| \otimes I_{2^{n_1+n_2}} + |t+1\rangle\langle t+1| \otimes I_{2^{n_1+n_2}} - |t+1\rangle\langle t| \otimes U_{t+1} - |t\rangle\langle t+1| \otimes U_{t+1}^\dagger\right) \end{aligned}$$

since $\sum_{z \in \{0,1\}^{2^{n_1+n_2}}} |z\rangle\langle z| = I_{2^{n_1+n_2}}$. This acts as the identity on every qubit except the $\log(T)$ qubits in $\mathcal{C}$, and the 2 qubits acted non-trivially on by the 2-local gate U_{t+1}. Thus it is $(\log(T)+2)$-local. This term gives an energy penalty to any state of the form $|t\rangle|\phi_t\rangle + |t+1\rangle|\phi_{t+1}\rangle$ where $|\phi_{t+1}\rangle \neq U_{t+1}|\phi_t\rangle$, so together, these terms ensure that the ground state is some history state, for some run of $\mathcal{V}(x)$, starting from some initial state.

Finally, define

$$H_{\text{final}} = |T\rangle\langle T|_{\mathcal{C}} \otimes |0\rangle\langle 0| \otimes I_{2^{n_1+n_2-1}}.$$

This is clearly $(\log(T)+1)$-local. This term gives an energy penalty if the final state has $|0\rangle$ in the answer register.

Together, these give:

$$H = \sum_{i=1}^{n_1} H_{\text{init}}(i) + \sum_{t=0}^{T-1} H_t + H_{\text{final}}.$$

Theorem 7.5.1 (Completeness)**.** *Let $a = \frac{1}{\sqrt{5T(T+1)}}$. Then for all $x \in L_1$, there is a unit vector $|\Psi\rangle$ such that $\|H|\Psi\rangle\| \leq a$.*

Proof. We will let $|\Psi\rangle = |\Psi(x)\rangle$ be the history state, defined in (7.2). We first compute:

$$\begin{aligned} H_{\text{init}}(i)|\Psi(x)\rangle &= \frac{1}{\sqrt{T+1}}|0\rangle_{\mathcal{C}} \otimes \left(I_2^{\otimes(i-1)} \otimes |1\rangle\langle 1| \otimes I_2^{\otimes(n_1-i)} \otimes I_{2^{n_2}}\right)|\psi_0(x)\rangle \\ &= \frac{1}{\sqrt{T+1}}|0\rangle_{\mathcal{C}} \otimes \left(I_2^{\otimes(i-1)} \otimes |1\rangle\langle 1| \otimes I_2^{\otimes(n_1-i)}\right)|0^{n_1}\rangle \otimes |w_x\rangle = 0 \end{aligned}$$

for any $i \in [n_1]$, since $|\psi_0(x)\rangle = |0^{n_1}\rangle|w_x\rangle$ has $|0\rangle$ in the first n_1 qubit registers. Next,

$$\left(\langle t|\langle z| - \langle t+1|\langle z|U_{t+1}^\dagger\right)(|t\rangle|\psi_t(x)\rangle + |t+1\rangle|\psi_{t+1}(x)\rangle) = \langle z|\psi_t(x)\rangle - \langle z|U_{t+1}^\dagger|\psi_{t+1}(x)\rangle = 0,$$

for any $t \in \{0, \ldots, T-1\}$ and $z \in \{0,1\}^{n_1+n_2}$, since $\langle z|U_{t+1}^\dagger|\psi_{t+1}(x)\rangle = \langle z|U_{t+1}^\dagger U_{t+1}|\psi_t(x)\rangle = \langle z|\psi_t(x)\rangle$. Thus we have

$$H_t|\Psi(x)\rangle = H_t\left(|t\rangle|\psi_t(x)\rangle + |t+1\rangle|\psi_{t+1}(x)\rangle\right) = 0.$$

Thus,

$$\begin{aligned} \|H|\Psi(x)\rangle\|^2 &= \frac{1}{T+1}\|H_{\text{final}}|\psi(x)\rangle\|^2 = \frac{1}{T+1}\left\||T\rangle_{\mathcal{C}} \otimes (|0\rangle\langle 0| \otimes I_{2^{n_1+n_2}})|\psi_T(x)\rangle\right\|^2 \\ &= \frac{1}{T+1}\|(I - \Pi_{\text{acc}})|\psi_T(x)\rangle\|^2, \end{aligned}$$

since we assume $\mathcal{V}_L(x)$ accepts if and only if there is a 1 in the first qubit register. By assumption that $|w_x\rangle$ is a witness that maximizes acceptance probability, and since $x \in L_1$, we get:

$$\|H|\psi(x)\rangle\|^2 = \frac{1}{T+1}\left(1 - \|\Pi_{\text{acc}}|\psi_T(x)\rangle\|^2\right) \leq \frac{1}{5T(T+1)}.$$

□

Theorem 7.5.2 (Soundness)**.** *Let* $b = \frac{1}{\sqrt{4T(T+1)}}$*. Then for all* $x \in L_0$*, and any unit vector* $|\Psi\rangle$*,* $\|H|\Psi\rangle\| \geq b$*.*

Proof. Suppose $|\Psi\rangle$ is a unit vector that minimizes $\|H|\Psi\rangle\|^2$. Then

$$\|H|\Psi\rangle\|^2 = \langle\Psi|H|\Psi\rangle = \langle\Psi|\sum_{i=1}^{n_1} H_{\text{init}}(i)|\Psi\rangle + \sum_{t=0}^{T-1}\langle\Psi|H_t|\Psi\rangle + \langle\Psi|H_{\text{final}}|\Psi\rangle. \tag{7.4}$$

Without loss of generality, we can assume

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

for some vectors $|\psi_t\rangle$ (not necessarily normalized, possibly 0). While we can't assume that the vectors $|\psi_t\rangle$ have anything to do with a run of the verifier, or even that they're unit vectors, we will see that if $\|H|\Psi\rangle\|^2$ is small, they must actually be pretty close. Let

$$|\tilde{\psi}_0\rangle = \frac{(|0^{n_1}\rangle\langle 0^{n_1}| \otimes I_{2^{n_2}})|\psi_0\rangle}{\|(|0^{n_1}\rangle\langle 0^{n_1}| \otimes I_{2^{n_2}})|\psi_0\rangle\|},$$

and for $t > 0$, let $|\tilde{\psi}_t\rangle = U_t|\tilde{\psi}_{t-1}\rangle$. We will formalize that:

1. $|\tilde{\psi}_0\rangle$ must be close to $|\psi_0\rangle$, or $\langle\Psi|\sum_{i=1}^{n_1} H_{\text{init}}(i)|\Psi\rangle$ will be large;
2. $|\psi_t\rangle$ must be close to $U_t|\psi_{t-1}\rangle$, or $\langle\Psi|H_t|\Psi\rangle$ will be large, meaning that $|\tilde{\psi}_t\rangle$ must be close to $|\psi_t\rangle$;
3. $|\psi_T\rangle$ must have high weight on a 1 in the answer register, or $\langle\Psi|H_{\text{final}}|\Psi\rangle$ will be large. But since $|\tilde{\psi}_T\rangle$ is close to $|\psi_T\rangle$, $|\tilde{\psi}_T\rangle$ must have high weight on a 1 in the answer register.

Since $|\tilde{\psi}_0\rangle = |0^{n_1}\rangle|w\rangle$ for some $|w\rangle$, together, these imply that $|w\rangle$ is a witness that is accepted by the verifier with high probability. If "high probability" is too high, we get a contradiction. Thus, assume towards a contradiction that

$$\|H|\Psi\rangle\|^2 \leq \frac{1}{4T(T+1)}. \tag{7.5}$$

For $z \in \{0,1\}^{n_1}$, let $\Pi_z = |z\rangle\langle z| \otimes I_{2^{n_2}}$.

$$\begin{aligned}(T+1)\langle\Psi|\sum_{i=1}^{n_1} H_{\text{init}}(i)|\Psi\rangle &= \langle\psi_0|\sum_{i=1}^{n_1}\sum_{z\in\{0,1\}^{n_1}:z_i=1}\Pi_z|\psi_0\rangle \geq \langle\psi_0|\sum_{z\in\{0,1\}^{n_1}\setminus\{0\}}\Pi_z|\psi_0\rangle \\ &= \langle\psi_0|(I-\Pi_{0^{n_1}})|\psi_0\rangle = 1-\langle\psi_0|\Pi_{0^{n_1}}|\psi_0\rangle = 1-\left\||\tilde{\psi}_0\rangle\right\|^2 \\ &= \left\||\psi_0\rangle - |\tilde{\psi}_0\rangle\right\|^2.\end{aligned}$$

Then by (7.5), since all terms of (7.4) are non-negative, each term is at most $\frac{1}{4T(T+1)}$, so we have:

$$\left\||\psi_0\rangle - |\tilde{\psi}_0\rangle\right\|^2 \leq \frac{1}{4T}. \tag{7.6}$$

The next caclulation we need is a straightforward exercise.

Exercise 7.5.1. *Show that for any* $t \in \{0,\ldots,T-1\}$*,* $\langle\Psi|H_t|\Psi\rangle = \frac{1}{2(T+1)}\left\|U_{t+1}|\psi_t\rangle - |\psi_{t+1}\rangle\right\|^2$*.*

From this we can use Cauchy-Schwarz to derive:

$$\begin{aligned}(T+1)\sum_{t=0}^{T-1}\langle\Psi|H_t|\Psi\rangle &= \frac{1}{2}\sum_{t=0}^{T-1}\|U_{t+1}|\psi_t\rangle - |\psi_{t+1}\rangle\|^2 \\ &\geq \frac{1}{2T}\left(\sum_{t=0}^{T-1}\|U_{t+1}|\psi_t\rangle - |\psi_{t+1}\rangle\|\right)^2 \\ &= \frac{1}{2T}\left(\sum_{t=0}^{T-1}\|U_T\ldots U_{t+1}|\psi_t\rangle - U_T\ldots U_{t+2}|\psi_{t+1}\rangle\|\right)^2,\end{aligned}$$

where we use the fact that any unitary, including $U_T\ldots U_{t+2}$, doesn't change the norm. We can continue by using the triangle inequality (twice):

$$\begin{aligned}(T+1)\sum_{t=0}^{T-1}\langle\Psi|H_t|\Psi\rangle &\geq \frac{1}{2T}\left\|\sum_{t=0}^{T-1}(U_T\ldots U_{t+1}|\psi_t\rangle - U_T\ldots U_{t+2}|\psi_{t+1}\rangle)\right\|^2 \\ &= \frac{1}{2T}\|U_T\ldots U_1|\psi_0\rangle - |\psi_T\rangle\|^2 \\ &\geq \frac{1}{2T}\left(\left\|U_T\ldots U_1|\tilde{\psi}_0\rangle - |\psi_T\rangle\right\| - \left\|U_T\ldots U_1|\tilde{\psi}_0\rangle - U_T\ldots U_1|\psi_0\rangle\right\|\right)^2 \\ &= \frac{1}{2T}\left(\left\||\tilde{\psi}_T\rangle - |\psi_T\rangle\right\| - \left\||\tilde{\psi}_0\rangle - |\psi_0\rangle\right\|\right)^2 \geq \frac{1}{2T}\left(\left\||\tilde{\psi}_T\rangle - |\psi_T\rangle\right\| - \frac{1}{\sqrt{4T}}\right)^2\end{aligned}$$

by (7.6). Again, using the fact that each term of (7.4) must be at most $\frac{1}{4T(T-1)}$, we get:

$$\begin{aligned}\frac{1}{2T}\left(\left\||\tilde{\psi}_T\rangle - |\psi_T\rangle\right\| - \frac{1}{\sqrt{4T}}\right)^2 &\leq \frac{1}{4T} \\ \left\||\tilde{\psi}_T\rangle - |\psi_T\rangle\right\| &\leq \frac{1}{\sqrt{2}} + \frac{1}{\sqrt{4T}}.\end{aligned} \tag{7.7}$$

For our final calculation, let $\Pi_{\text{rej}} = I - \Pi_{\text{acc}}$ denote the projector onto a 0 in the answer register.

$$\begin{aligned}(T+1)\langle\Psi|H_{\text{final}}|\Psi\rangle &= \langle\psi_T|\Pi_{\text{rej}}|\psi_T\rangle = \|\Pi_{\text{rej}}|\psi_T\rangle\|^2 \\ &\geq \left(\left\|\Pi_{\text{rej}}|\tilde{\psi}_T\rangle\right\| - \left\|\Pi_{\text{rej}}|\psi_T\rangle - \Pi_{\text{rej}}|\tilde{\psi}_T\rangle\right\|\right)^2 \geq \left(\left\|\Pi_{\text{rej}}|\tilde{\psi}_T\rangle\right\| - \left\||\psi_T\rangle - |\tilde{\psi}_T\rangle\right\|\right)^2,\end{aligned}$$

where we have once again used the triangle inequality. Combining this with (7.7), and the fact that each term of (7.4) must be at most $\frac{1}{4T(T+1)}$, we get:

$$\begin{aligned}\left(\left\|\Pi_{\text{rej}}|\tilde{\psi}_T\rangle\right\| - \left(\frac{1}{2} + \frac{1}{\sqrt{4T}}\right)\right)^2 &\leq \frac{1}{4T} \\ \left\|\Pi_{\text{rej}}|\tilde{\psi}_T\rangle\right\| &\leq \frac{1}{\sqrt{4T}} + \frac{1}{\sqrt{2}} + \frac{1}{\sqrt{4T}} = \frac{1}{\sqrt{2}} + \frac{1}{\sqrt{T}}.\end{aligned}$$

Thus, whenever T is greater than some sufficiently large constant,

$$\left\|\Pi_{\text{rej}}|\tilde{\psi}_T\rangle\right\|^2 \leq \frac{1}{2} + o(1) \quad \text{and so} \quad \left\|\Pi_{\text{acc}}|\tilde{\psi}_T\rangle\right\|^2 \geq \frac{1}{2} - o(1).$$

This brings us to a contradiction, because, it implies that if we give the verifier input

$$\frac{|\tilde{\psi}_0\rangle}{\left\||\tilde{\psi}_0\rangle\right\|} = |0^{n_1}\rangle\frac{|w\rangle}{\left\||\tilde{\psi}_0\rangle\right\|},$$

and measure it's final state $\frac{|\tilde{\psi}_T\rangle}{\||\tilde{\psi}_0\rangle\|}$, we will get a 1 with probability

$$\geq \frac{\frac{1}{2} - o(1)}{\left\||\tilde{\psi}_0\rangle\right\|} = \Omega(1),$$

by the following exercise:

Exercise 7.5.2. *Show that* $\left\||\tilde{\psi}_0\rangle\right\| = O(1)$.

This is a contradiction because $x \in L_0$, so the acceptance probability should be at most $1/T$. Thus, our assumption in (7.5) must be false. □

7.5.2 Reduction to 5-local Hamiltonian

The reason our previous construction was $\log(T)$-local was because of the clock register. The main change we make to facilitate making our Hamiltonians more local, is that we will encode the register $\mathcal{C}$ in *unary* encoding:

$$|t\rangle_{\mathcal{C}} = |1\rangle_{\mathcal{C}_1} \otimes \cdots \otimes |1\rangle_{\mathcal{C}_t} \otimes |0\rangle_{\mathcal{C}_{t+1}} \otimes \cdots \otimes |0\rangle_{\mathcal{C}_T}.$$

Above, each register $\mathcal{C}_j$ is a qubit register. Encoding a number t by a string of t 1s (followed by 0s) is called *unary*. This is not very space-efficient, but that's not really a priority here. We will ensure that the register $\mathcal{C}$, which is now a T-qubit register, does in fact encode a number in unary – we will refer to this as the *unary subspace* – by including the following additional terms in our Hamiltonian, for each $t \in [T-1]$:

$$H_{\text{unary}}(t) = I_{\mathcal{C}_1} \otimes I_{\mathcal{C}_{t-1}} \otimes |01\rangle\langle 01|_{\mathcal{C}_t\mathcal{C}_{t+1}} \otimes I_{\mathcal{C}_{t+2}} \otimes I_{\mathcal{C}_T} \otimes I_{2^{n_1+n_2}}.$$

This is clearly 2-local. Moreover, it introduces an energy penalty for any state that has a string in the clock register that does not encode a number in unary. That's because a string encodes a unary number precisely when it consists of zero or more 1s, followed by 0s; which is precisely when there is never a 0 followed by a 1. If there is a 0 in the t-th clock qubit, and a 1 in the $(t+1)$-th clock qubit, $H_{\text{unary}}(t)$ will not map such a state to 0, resulting in a non-zero "energy penality".

We slightly modify $H_{\text{init}}(i)$, H_t and H_{final}, using the fact that $\mathcal{C}$ now has a unary encoding to decrease the locality. Recall from (7.3) that $H_{\text{init}}(i)$ should ensure that when $\mathcal{C}$ contains $|0\rangle_{\mathcal{C}}$, the i-th qubit of the algorithm is in the state $|0\rangle$. Now we will exploit the fact that $|0\rangle_{\mathcal{C}}$ is the only state in unary encoding that has $|0\rangle$ in $\mathcal{C}_1$, so as long as we're in the subspace spanned by unary strings, to "check" that $\mathcal{C}$ contains 0, we need only look at the first qubit, $\mathcal{C}_1$. That is, for $i \in [n_1]$, define:

$$H_{\text{init}}(i) = |0\rangle\langle 0|_{\mathcal{C}_1} \otimes I_{\mathcal{C}_2\ldots\mathcal{C}_T} \otimes I_2^{\otimes(i-1)} \otimes |1\rangle\langle 1| \otimes I_2^{\otimes(n-i)}.$$

Clearly, $H_{\text{init}}(i)$ is 2-local.

Next, recall that we defined H_t to be the orthogonal projector onto states of the form:

$$\begin{aligned} &|t\rangle_{\mathcal{C}}|z\rangle - |t+1\rangle_{\mathcal{C}} U_{t+1}|z\rangle \\ &= |1^t\rangle_{\mathcal{C}_1\ldots\mathcal{C}_t}|0\rangle_{\mathcal{C}_{t+1}}|0^{T-t-1}\rangle_{\mathcal{C}_{t+2}\ldots\mathcal{C}_T}|z\rangle - |1^t\rangle_{\mathcal{C}_1\ldots\mathcal{C}_t}|1\rangle_{\mathcal{C}_{t+1}}|0^{T-t-1}\rangle_{\mathcal{C}_{t+2}\ldots\mathcal{C}_T} U_{t+1}|z\rangle. \end{aligned}$$

We can decrease the locality of H_t by projecting on states of the form (up to reordering of registers):

$$|\ell_1,\ldots,\ell_{t-1}\rangle_{\mathcal{C}_1\ldots\mathcal{C}_{t-1}}|\ell_{t+3},\ldots,\ell_T\rangle_{\mathcal{C}_{t+3}\ldots\mathcal{C}_T} \underbrace{\left(|1\rangle_{\mathcal{C}_t}|0\rangle_{\mathcal{C}_{t+1}}|0\rangle_{\mathcal{C}_{t+2}}|z\rangle - |1\rangle_{\mathcal{C}_t}|1\rangle_{\mathcal{C}_{t+1}}|0\rangle_{\mathcal{C}_{t+2}} U_{t+1}|z\rangle\right)}_{|\phi_{t,z}\rangle}$$

for any values $\ell_1, \ldots, \ell_{t-1}, \ell_{t+3}, \ldots, \ell_T \in \{0,1\}$. Restricting to the unary subspace of $\mathcal{C}$, this does the same thing, because as long as a number is encoded in unary, the unique location of its 10 encodes its value. Thus, we can define, for $t = \{0, \ldots, T-1\}$:

$$H_t = I_{\mathcal{C}_1 \ldots \mathcal{C}_{t-1} \mathcal{C}_{t+3} \ldots \mathcal{C}_T} \otimes \sum_{z \in \{0,1\}^{2^{n_1+n_2}}} |\phi_{t,z}\rangle\langle\phi_{t,z}|,$$

which is 5-local, because it acts non-trivillay on the qubits in $\mathcal{C}_t$, $\mathcal{C}_{t+1}$ and $\mathcal{C}_{t+2}$, as well as the two qubits acted on by the 2-local gate U_{t+1}.

Finally, similar to the case of $H_{\text{init}}(i)$, for H_{final}, as long as we're in the unary subspace, we can "check" if $\mathcal{C}$ contains T by simply "checking" if $\mathcal{C}_T$ is set to $|1\rangle$:

$$H_{\text{final}} = I_{\mathcal{C}_1 \ldots \mathcal{C}_{T-1}} \otimes |1\rangle\langle 1|_{\mathcal{C}_T} \otimes |0\rangle\langle 0| \otimes I_{2^{n_1+n_2-1}}.$$

Clearly this is 2-local. Then we define

$$H = \sum_{t=1}^{T-1} H_{\text{unary}}(t) + \sum_{i=1}^{n_1} H_{\text{init}}(i) + \sum_{t=0}^{T-1} H_t + H_{\text{final}},$$

which is 5-local, since all terms have locality at most 5.

Soundness and completeness can be shown similar to Section 7.5.1. We leave this as an exercise.

7.6 Extensions and Further Reading

We have seen that 5-LocalHam$_{a,b}$ is QMA-complete whenever $b - a = \frac{1}{\text{poly}(n)}$. Using more clever tricks, this can be improved to show that even 2-LocalHam$_{a,b}$ is QMA-complete [KKR06], which is in contrast to the classical world, where 2-Sat is in P. A long-standing open problem called the *quantum PCP conjecture* is to show that 2-LocalHam$_{a,b}$ is QMA-complete even when $b - a = \Omega(m)$, for m the number of local Hamiltonian terms.

Bibliography

[Gha23] Sevag Gharibian. The 7 faces of quantum np, 2023. arXiv: 2310.18010 3

[KKR06] Julia Kempe, Alexei Yu Kitaev, and Oded Regev. The complexity of the local hamiltonian problem. *SIAM Journal on Computing*, 35(5):1010–1097, 2006. 11

[MW05] Chris Marriott and John Watrous. Quantum Arthur–Merlin games. *Computational Complexity*, 14(2):122–152, 2005. arXiv: cs/0506068 4

[Reg06] Oded Regev. Witness-preserving amplification of qma, 2006. 4

[dW19] Ronald de Wolf. Quantum computing lecture notes. arXiv: 1907.09415v5, 2019. 1