Lecture 1

Introduction

Scribe: Thomas Vidick

1.1 Local Hamiltonians

In this course we study *local Hamiltonians*, and we study them through the lens of *computational complexity*. Informally, a local Hamiltonian is a collection of local constraints on a physical system. For us, the physical system will always be the space of n qubits, $\mathcal{H} = \mathbb{C}^n$. A Hamiltonian on $\mathcal{H}$ is a Hermitian matrix H on $\mathcal{H}$. Through the singular value decomposition $H = \sum_i \lambda_i |u_i\rangle\langle u_i|$, H assigns an energy $\lambda_i \in \mathbb{R}$ to the state $|u_i\rangle$. More generally, a state $|\psi\rangle$ is assigned energy $\langle\psi|H|\psi\rangle$. The eigenvector associated with the smallest eigenvalue is called the *ground state* and we generally denote it $|\Gamma\rangle$. We mostly consider "static" problems about H, such as estimating its smallest eigenvalue (ground energy), representing its ground state, or computing expectation values of local observables $\langle\Gamma|O|\Gamma\rangle$ on it.

Let's look at some examples. First take $n = 1$ and let $H = Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$. This is Hermitian with singular value decomposition $H = |0\rangle\langle 0| - |1\rangle\langle 1|$. Therefore the smallest eigenvalue of H is -1 and the associated ground state is $|1\rangle$. Another Hamiltonian with the same ground state is $H' = |0\rangle\langle 0| = (H + \mathrm{Id})/2$. Sometimes writing Hamiltonians as positive semi-definite operators will be convenient because the smallest energy is then clearly non-negative.

If we now consider $n = 2$ we can introduce some more interesting examples. First of all, $H = |0\rangle\langle 0| \otimes |0\rangle\langle 0|$ is interesting because it has a degenerate ground space, spanned by $\{|00\rangle, |10\rangle, |11\rangle\}$. Another example of a degenerate Hamiltonian is $H = Z \otimes Z$. This has singular value decomposition

$$H = |00\rangle\langle 00| + |11\rangle\langle 11| - |01\rangle\langle 01| - |10\rangle\langle 10| \,,$$

so that the ground energy is -1 and the ground space is spanned by $\{|01\rangle, |10\rangle\}$, i.e. the odd-parity computational basis states.

Of course, we should not restrict ourselves to Hamiltonians that are diagonal in the computational basis, and for example $H = X \otimes X$, where X is the Pauli matrix $X = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, similarly has ground space spanned by the odd-parity strings in the Hadamard basis, i.e. $\{|+\rangle|-\rangle, |-\rangle|+\rangle\}$.

All Hamiltonians we saw so far have eigenstates that are product states. A more interesting example is $H = \frac{1}{2}(X \otimes X + Z \otimes Z)$. Because X and Z commute, $X \otimes X$ and $Z \otimes Z$ commute hence are simultaneously diagonalizable. Their joint eigenbasis is the *Bell basis* $|\phi_{a,b}\rangle = (I \otimes X^a Z^b)\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. One

can then verify that H has smallest eigenvalue -1 with a unique associated ground state

$$|\Gamma\rangle = |\phi_{1,1}\rangle = \frac{1}{\sqrt{2}}\big(|01\rangle - |10\rangle\big),$$

the *singlet*.

Now that we've seen some small examples, let's look at three families of n-qubit Hamiltonians.

1.1.1 The QSAT hamiltonians

We call a Hamiltonian k-*local* if it can be expressed as a linear combination of terms, such that each term acts as identity on all but at most k of the qubits. For example, $H = \frac{1}{n}\sum_{i=1}^{n} I \otimes \cdots I \otimes |0\rangle\langle 0| \otimes I \otimes \cdots \otimes I$ is a 1-local Hamiltonian whose ground state is the $|11\cdots 1\rangle$ (all other computational basis states are assigned an energy equal to n minus their Hamming weight).

Quantum k-SAT is a general class of k-local Hamiltonians where the only restriction is that each of the local terms is required to be a projection. So, $H = \frac{1}{m}\sum_{i=1}^{m} \Pi_i$, where Π_i takes the form $\Pi_i = \mathrm{Id}_{\overline{S_i}} \otimes P_i$, P_i a projection acting on the qubits in S_i only, with $S_i \subseteq [n]$ of size $|S_i| \leq k$.

The only observation we make about this class at the moment is that it allows us to model classical k-CSPs such as k-SAT. Take for example a 3-SAT formula with clauses such as $x_1 \vee \overline{x_2} \vee \overline{x_3}$. This can be represented by the local constraint $|1\rangle\langle 1|_1 \otimes |0\rangle\langle 0|_2 \otimes |1\rangle\langle 1|_3$, which assigns an energy penalty of 1 to the only non-satisfying configuration. To any 3-SAT formula φ we can in this way associate a 3-QSAT Hamiltonian H_φ such that the smallest eigenvalue of H_φ is exactly the fraction of clauses that are violated in an optimal assignment to φ. As a consequence, the problem of computing the smallest eigenvalue is at least as hard as the problem of deciding if a 3-SAT formula is satisfiable or not, i.e. NP-hard. (Note that for this statement to hold, it is important that the reduction $\varphi \mapsto H_\varphi$ can be performed in polynomial time, which is indeed the case.)

1.1.2 The Heisenberg model

The Heisenberg model is a general class of Hamiltonians that is studied in Physics. A Hamiltonian from this class is parametrized by a graph $G = (V, E)$ and takes the form

$$H = \frac{1}{|E|}\sum_{\{i,j\}\in E} \left(\alpha X_i \otimes X_j + \beta Y_i \otimes Y_j + \gamma Z_i \otimes Z_j\right).$$

Here, X_i denotes a Pauli X matrix acting on the i-th qubit, etc. The summation is over undirected edges in G. The coefficients $\alpha, \beta, \gamma \in \{0, 1\}$ vary depending on the sub-family of Hamiltonians one is considering:

- If $\alpha = \beta = 0$, $\gamma = 1$ then $H = \frac{1}{m}\sum_{i,j} Z_i \otimes Z_j$, where $m = |E|$, is diagonal in the computational basis. Each basis state $|x\rangle$ can be interpreted as a cut in the graph G, and the energy that H associates to this cut is equal to the number of edges that are not cut, minus the number of edges that are cut; equivalently, the total number of edges minus twice the number of cut edges (normalized by m). Thus, up to an additive shift of the number of edges, the smallest eigenvalue problem for H is exactly the MAXCUT problem in the graph G. In particular, we recover once more that this problem is NP-hard.
- If $\alpha = \beta = 1$, $\gamma = 0$ then we obtain the so-called *anti-ferromagnetic XY model*. For this model, the local terms take the form $X_i \otimes X_j + Y_i \otimes Y_j$. This term has smallest eigenvalue -2 with associated eigenstate the singlet.

- Finally we consider the case where $\alpha = \beta = \gamma = 1$. The local terms are $X_i \otimes X_j + Y_i \otimes Y_j + Z_i \otimes Z_j$, with smallest eigenvalue -3 with associated eigenvector the singlet. The only other eigenvalue is $+1$ with associated eigenspace the space spanned by the remaining three Bell states. In this case we may rescale the Hamiltonian to take the form

$$H = \frac{1}{m} \sum_{\{i,j\} \in E} \frac{1}{4} \Big(\mathrm{Id} - \alpha X_i \otimes X_j - \beta Y_i \otimes Y_j - \gamma Z_i \otimes Z_j \Big) .$$

  The advantage is that now each term, associated with an edge $\{i, j\}$ is a projection on the singlet state on the corresponding vertices i and j. Note that flipping signs turned the energy minimization problem into a maximization problem. This class is referred to as *quantum MAXCUT*. Informally speaking, the task is to determine how many singlet states can be "packed" along the edges of a graph. Note that contrary to MAXCUT, even when the graph is bipartite this problem is not trivial, as due to the phenomenon of monogamy of entanglement a given vertex can in general not be in a singlet state with more than one other vertices. We will return to this class of Hamiltonians in greater detail later on in the class.

1.1.3 The Motzkin Hamiltonian

The Motzkin Hamiltonian is our first example of a Hamiltonian that acts on higher-dimensional particles, in this case qutrits instead of qubits. That is, the Motzkin Hamiltonian is a Hermitian matrix on $(\mathbb{C}^3)^{\otimes n}$. We use the labels $|L\rangle, |R\rangle, |0\rangle$ for an orthonormal basis of $\mathbb{C}^3$.

A Motzkin path is a well-parenthesized expression. For example, if $n = 1$ there is only 0, if $n = 2$ there is LR and 00, if $n = 3$ we have $0LR$, $L0R$, $LR0$ and 000, etc.

The Motzkin Hamiltonian has the property that its ground state $|\Gamma\rangle$ is a uniform superposition of all Motzkin paths. This is interesting for two reasons. Firstly, because of the "well-parenthesized" condition, this ground state has a fairly large amount of entanglement; namely, it is possible to show that the entanglement entropy across the middle cut scales as $\Omega(\log n)$. Secondly, it is possible to give a realization of the Motzkin Hamiltonian as a "1D" Hamiltonian, namely all local constraints are nearest-neighbor on a line.

To see how this is done, introduce the states

$$|\phi\rangle = \frac{1}{\sqrt{2}}(|00\rangle - |LR\rangle) , \qquad |\psi_L\rangle = \frac{1}{\sqrt{2}}(|L0\rangle - |0L\rangle) , \qquad |\psi_R\rangle = \frac{1}{\sqrt{2}}(|0R\rangle - |R0\rangle) .$$

The Motzkin Hamiltonian is then defined as follows:

$$H = |R\rangle\langle R|_1 + |L\rangle\langle L|_n + \sum_{j=1}^{n-1} \Big(|\phi\rangle\langle\phi|_{j,j+1} + |\psi_L\rangle\langle\psi_L|_{j,j+1} + |\psi_R\rangle\langle\psi_R|_{j,j+1} \Big) .$$

It can be shown that the smallest eigenvalue of H is 0, and the associated eigenvector is the Motzkin State $|M_n\rangle$, the uniform superposition over all Motzkin paths (see the problem sheet for a proof of this).

The Motzkin Hamiltonian is an example of a local Hamiltonian that has additional structure, in this case the Hamiltonian is *geometrically local* in $1D$. This restriction, of geometric locality, is sometimes imposed as it is "physically relevant": most quantum mechanical systems exist in one, two or three-dimensional Euclidean space. Another restriction that can be considered is that of *translation invariance*: this means that the Hamiltonian is defined on e.g. a $2D$ grid, and the local terms are invariant under translation. This constraint represents some kind of spatial homogeneity of the system. Another condition, of a different

nature, that arises at times, is that the Hamiltonian is *spectrally gapped*: the smallest and second smallest eigenvalues are separated by, say, an inverse polynomial in the number of qubits (note that for a Hamiltonian of bounded (polynomial) norm the "typical" gap between eigenvalues is exponentially small). Each of these restrictions may affect the complexity of problems associated with the local Hamiltonian, and we will see how in later lectures.

1.2 The class QMA

Having introduced our main object of study we now get ready to formulate the main complexity-theoretic question that we will ask about it. For this we introduce the complexity class QMA ("Quantum Merlin-Arthur"), which is the quantum analogue of the class NP. Before that, let's review the definition of the complexity class BQP ("Bounded-error Quantum Polynomial time").

The complexity class BQP

The complexity class BQP is the quantum analogue of the class BPP. It consists of all languages that can be decided in quantum polynomial time. More formally,

Definition 1.1. A language $L \in$ BQP if there exists a classical polynomial time algorithm A that maps inputs $x \in \{0,1\}^*$ to quantum circuits C_x on $n = \text{poly}(|x|)$ qubits, where the circuit is considered as a sequence of unitary operators each on 2 qubits, i.e $C_x = U_T U_{T-1} ... U_1$ where each $U_i \in \text{L}\big(\mathbb{C}^2 \otimes \mathbb{C}^2\big)$, such that:

1. *Completeness*: $x \in L \Rightarrow \Pr(C_x \text{ accepts } |0_n\rangle) \geq \frac{2}{3}$
2. *Soundness*: $x \notin L \Rightarrow \Pr(C_x \text{ accepts } |0_n\rangle) \leq \frac{1}{3}$

We say that the circuit "C_x accepts $|\psi\rangle$" if the first output qubit measured in $C_x|\psi\rangle$ is 1. More specifically, letting $\Pi_1^{|1\rangle} = |1\rangle\langle 1|_1$ be the projection of the first qubit on state $|1\rangle$,

$$\Pr(C_x \text{ accepts } |\psi\rangle) = \| \ (\Pi_1^{|1\rangle} \otimes \text{Id}_{n-1}) C_x |\psi\rangle \ \|_2^2 \ .$$

The complexity class QMA

The complexity class QMA (or $BQNP$, as Kitaev originally named it) is the quantum analog of the class NP. More formally,

Definition 1.2. A language $L \in$ QMA if there exists a classical polynomial time algorithm A that maps inputs $x \in \{0,1\}^*$ to quantum circuits C_x on $n + q = \text{poly}(|x|)$ qubits, such that:

1. *Completeness*: $x \in L \Rightarrow \exists |\psi\rangle \in \mathbb{C}^{2^q}, \ \||\psi\rangle\|_2 = 1, \text{ such that } \Pr(C_x \text{ accepts } |0_n\rangle \otimes |\psi\rangle) \geq \frac{2}{3}$
2. *Soundness*: $x \in L \Rightarrow \forall |\psi\rangle \in \mathbb{C}^{2^q}, \ \||\psi\rangle\|_2 = 1, \ \Pr(C_x \text{ accepts } |0_n\rangle \otimes |\psi\rangle) \leq \frac{1}{3}$

Proposition 1.3. NP $\subseteq$ QMA

Proof. The circuit C_x will be the NP verifier "hardwired" with the input x, and it will measure every bit of $|\psi\rangle$ independently and use it as the classical proof π. □

Exercise 1.1. Let $\mathrm{QMA}_{c,s}$ be same as QMA but with completeness and soundness parameters c and s instead of $\frac{2}{3}$ and $\frac{1}{3}$. Show that $\mathrm{QMA}_{c,s} = \mathrm{QMA}$ as long as $c - s \in \frac{1}{O(\mathrm{poly}(|x|))}$. (The main issue to consider during repetition is the entanglement that could exist between multiple quantum proofs).

Exercise 1.2. Show the progressively stronger sequence of inclusions $\mathrm{NP} \subseteq \mathrm{QMA} \subseteq \mathrm{PP} \subseteq \mathrm{PSPACE} \subseteq \mathrm{EXP}$.

Remark 1.4. Note that QMA is a largest eigenvalue problem:

$$\begin{aligned}\Pr(C_x \text{ accepts } |0_n\rangle \otimes |\psi\rangle) &= \|(\Pi_1^{|1\rangle} \otimes \mathrm{Id})C_x|0_n\rangle \otimes |\psi\rangle\|_2^2 \\ &= \langle\psi| \otimes \langle 0_n|C_x^\dagger(\Pi_1^{|1\rangle} \otimes \mathrm{Id})C_x|0_n\rangle \otimes |\psi\rangle \\ &= \langle\psi|D|\psi\rangle \ ,\end{aligned}$$

where $D = \mathrm{Id} \otimes \langle 0_n|C_x^\dagger(\Pi_1^{|1\rangle} \otimes \mathrm{Id})C_x|0_n\rangle \otimes \mathrm{Id}$ is a $2^q \times 2^q$ matrix. Since $\langle\psi|\psi\rangle = 1$, finding $|\psi\rangle$ which maximizes $\langle\psi|D|\psi\rangle$ is exactly same as the maximum eigenvalue λ of the matrix D.