Frontiers of Quantum complexity *Instructor: Anand Natarajan, Anurag Anshu*

Lecture 10: The Guided LH problem and quantum algorithms

Oct 7, 2025 *Scribe: Alexander Zlokapa*

1 Low-energy tensor network states

Let $E_0(H)$ denote the ground state of an n-qubit Hamiltonian H; let $E_{\max}(H) \leq m$ be a bound on the spectral norm.

Theorem 1.1. *Given a local Hamiltonian $H = \sum_a b_a P_a$ with spectral gap $\Delta = 1/\text{poly}(n)$ there is a tensor network state (TN) with energy $\langle \text{TN} | H | \text{TN} \rangle \leq E_0(H) + 2^{-\Theta(n)}$ and with constant bond dimension.*

This theorem isn't too useful for quantum complexity since $|\text{TN}\rangle$ is generically hard to prepare (and requires post selection). However, it provides some information about the entanglement of low-energy states, since the entanglement entropy scales as the log of the bond dimension. A similar result holds if a spectral gap is replaced by a promise gap (in the promise problem setting), so a similar statement holds for a witness in a QMA setting.

The proof is based on an AGSP (approximate ground space projector). For some integer k, introduce matrix

$$K(H) = \left(I - \frac{H - E_0(H)I}{m - E_0(H)} \right)^k .$$

Observe that $K(E_0) = 1$ and $K(E_j) \leq \left(1 - \frac{\Delta}{m - E_0}\right)^k$ for any $j \geq 1$ due to $E_j - E_0 \geq \Delta$. We choose $k = 2mn/\Delta$ to ensure that $K(E_j) \leq 2^{-2n}$. We set $|\text{TN}\rangle = K(H) |\psi\rangle$ for a randomly chosen product state $|\psi\rangle = |\rho_1\rangle \otimes \cdots \otimes |\rho_n\rangle$ (e.g., a random computational basis state). With high probability, we claim that (1) $|\text{TN}\rangle$ has very low energy and (2) $|\text{TN}\rangle$ is a tensor network state with constant bond dimension.

1. *Low energy.* We expand $|\psi\rangle$ in the Hamiltonian eigenbasis as $|\psi\rangle = a_{E_0} |E_0\rangle + \sum_{E > E_0} a_E |E\rangle$. Then $K(H) |\psi\rangle = a_{E_0} |E_0\rangle + O(2^{-2n}) \sum_{E > E_0} a'_E |E\rangle$, where a'_E might be slightly modified from a_E. Since $a_{E_0} \sim 2^{-n/2}$ with high probability, the state is exponentially dominated by $|E_0\rangle$.

2. *Tensor network state.* A product state is trivially a tensor network state with bond dimension 1. To prove $|\text{TN}\rangle$ has constant bond dimension we use MBQC teleportation: we apply k times the non-unitary operator $O = 1 - \frac{\Delta}{m - E_0}$. Prepare the product state. Prepare k sets of EPR pairs on $2n$ qubits with O applied on one end; this is also a tensor network state of bond dimension $O(1)$ since H is local. (A trivial argument gives bond dimension at most m, but with an ancilla we can get away with constant bond dimension given the constant locality of the Hamiltonian.) Project onto an EPR pair between pairs of qubits in the product state and an EPR pair with an O attached; project between the loose end and another EPR pair with an O attached; etc. This gives state

$$|\text{EPR}\rangle^{\otimes nk} \otimes K(H) |\psi\rangle .$$

This says the ground state has some entanglement structure; for example, a Haar random state cannot be represented as a 2D tensor network with constant bond dimension. However, it doesn't say *too* much about the ground state, since we can't actually prepare this tensor network state.

2 Guiding states

A state $|u\rangle$ is a guiding state for H if $\| \langle u| \Pi_{E_0(H)} |u\rangle \|^2 = \Omega(1/\text{poly}(n))$ and $|u\rangle$ has a classically succinct description, e.g., a description of a poly-size quantum circuit that prepares $|u\rangle$ up to inverse polynomial error. (Similar access models work: for example, one can give query access to the amplitudes of $|u\rangle$.) This is unlike the tensor network states above, which cannot be efficiently accessed.

Suppose $|u\rangle$ can be prepared on a quantum computer; then phase estimation starting from $|u\rangle = a_{E_0} |E_0\rangle + \sum_{E>E_0} a_E |E\rangle$ has success probability $1/|a_{E_0}|^2$ (and amplitude amplification can improve this to $1/|a_{E_0}|$), which succeeds with probability $\Omega(1/\text{poly}(n))$ by assumption.

Examples of guiding states.

1. *Random sparse Hamiltonians.* [Chen et al., "Sparse Random Hamiltonians Are Quantumly Easy"] Let H be a $2^n \times 2^n$ matrix with poly(n) entries per row. For sparse random Hamiltonians

$$H = \sum_{a=1}^{\text{poly}(n)} b_a P_a \tag{1}$$

   for random weight-n Paulis P_a and Bernoulli coefficients b_a, the spectrum is a semicircle law and thus the maximally mixed state has $1/\text{poly}(n)$ support on the subspace $\Pi_{\leq E_0+1/\text{poly}(n)}$. Note that the Hamiltonian has $O(n)$ locality and is thus less physically motivated than $O(1)$-local Hamiltonians. It is an open problem if certifying the ground state energy up to $1/\text{poly}(n)$ error is also efficient classically.

2. *Random local Hamiltonians.* Let $H = \sum_a b_a P_a$ for $O(1)$-local H. This spectrum has a Gaussian tail, making it hard to find a guiding state with good overlap on the low-energy subspace.

3. *Quantum chemistry.* There is numerical evidence that the ground states of various strongly correlated chemistry molecules can be approximated by a Matrix Product State. See Table I, Page 32 here. However, there quantities are extrapolated from small computations.

4. *Quartic speedups.* [Hastings, "Classical and Quantum Algorithms for Tensor Principal Component Analysis"] Guiding states may be able to give end-to-end polynomial speedups based on the Kikuchi method. A quadratic speedup is obtained by amplitude amplification, and further speedups are obtained by improved overlap compared to the maximally mixed state.

3 Guided local Hamiltonian problem

The guided local Hamiltonian was first stated in terms of a guiding state $|u\rangle$ that has "sampling access" (as in the dequantization work), which gives access to u_j with probability $|u_j|^2/\|u\|$ as well as an estimate of $\|u\|$. A simpler access model would simply provide a quantum circuit that prepares $|u\rangle$; this does not change the result.

Theorem 3.1 (Gharibian and Le Gall, "Dequantizing the Quantum Singular Value Transformation")**.** *Given a local Hamiltonian H and sampling access to a state $|u\rangle$ such that support of $|u\rangle$ on the ground state satisfies $\|\Pi_{E_0}|u\rangle\|^2 = 1/\mathrm{poly}(n)$, decide if $E_0(H) < a$ or $E_0(H) > b$ for $b - a = \Theta(1/\mathrm{poly}(n))$. This problem is BQP-complete.*

The guided local Hamiltonian is obviously in BQP since you can prepare $|u\rangle$ and apply quantum phase estimation to obtain the ground state (and eigenvalue) with probability $1/\mathrm{poly}(n)$. Our goal will be to show that it is BQP-hard; note that it is not BQP-hard when the precision is reduced to constant instead of inverse polynomial.

Warm up: We consider a simpler version of the guided local Hamiltonian problem. We add a promise that the ground state $|E_0\rangle$ itself can be prepared by a polynomial-time quantum circuit Q (whose gates can be efficiently determined from H) up to $1 - 1/\mathrm{poly}(n)$ fidelity. To show BQP-hardness, we use the clock Hamiltonian. For simplicity, take the clock register to be unary over $T + 1$ qubits. Recall that for the clock Hamiltonian constructed from

$$H_{\text{in}} = (I - |x, 0\rangle\langle x, 0|) \otimes |0\rangle\langle 0| \tag{2}$$

$$H_{\text{prop}} = \frac{1}{2}\sum_{t=1}^{T} I \otimes |t\rangle\langle t| + I \otimes |t-1\rangle\langle t-1| - U_t \otimes |t\rangle\langle t|\, t - 1 - U_t^\dagger \otimes |t-1\rangle\langle t-1|\, t \tag{3}$$

$$H_{\text{out}} = |0\rangle\langle 0| \otimes |T\rangle\langle T| \tag{4}$$

the history state

$$|\eta\rangle = \frac{1}{\sqrt{T+1}}\sum_{t=0}^{T} U_t \cdots U_1 |x\rangle |0\rangle \otimes |t\rangle \tag{5}$$

is the unique ground state of $H_{\text{in}} + H_{\text{prop}}$. If we are in the YES case (i.e., $\Pr[Q \text{ outputs } 1 \text{ on } x] \geq 1 - 2^{-n}$), then the history state is exponentially close to the ground state of the clock Hamiltonian. We also note that the spectral gap of $H_{\text{in}} + H_{\text{prop}}$ is $\gamma \geq 2/(T+1)^2$. Moreover, we observe that the expectation of H_{out} on $|\eta\rangle$ is

$$\langle\eta| H_{\text{out}} |\eta\rangle = \frac{1 - \Pr[Q \text{ outputs } 1 \text{ on } x]}{T+1}, \tag{6}$$

and $|\eta\rangle$ can be prepared efficiently by a polynomial-sized quantum circuit via controlled unitaries on the uniform superposition over the clock register. However, in the NO case (i.e., $\Pr[Q \text{ outputs } 1 \text{ on } x] \leq 2^{-n}$), the history state is no longer the ground state. To show that the ground state of the NO case can be efficiently prepared, we will modify the clock Hamiltonian and show that $|\eta\rangle$ has good overlap with its ground state. We set

$$H' = \Delta(H_{\text{in}} + H_{\text{prop}}) + H_{\text{out}} \tag{7}$$

for $\Delta = 100T^2$. We want to lower bound the overlap between $|\eta\rangle$ (i.e., the ground state of $H_0 = \Delta(H_{\text{in}} + H_{\text{prop}})$) and the ground state $|\phi\rangle$ of H' (i.e., H_0 perturbed by H_{out}). We will expand the ground state as $|\psi\rangle = \alpha |\eta\rangle + \sqrt{1-\alpha^2} |\eta_\perp\rangle$ for $\alpha > 0$. Since $|\eta\rangle$ is the ground state of H_0 and since $\langle\eta_\perp| H_0 |\eta_\perp\rangle \geq \gamma$, we have

$$\langle\psi| H' |\psi\rangle \geq \Delta\gamma(1-\alpha^2) + \langle\psi| H_{\text{out}} |\psi\rangle . \tag{8}$$

We expand the second term as

$$\langle\psi| H_{\text{out}} |\psi\rangle = \alpha^2 \langle\eta| H_{\text{out}} |\eta\rangle + (1-\alpha^2) \langle\eta_\perp| H_{\text{out}} |\eta_\perp\rangle + 2\Re(\alpha\sqrt{1-\alpha^2} \langle\eta| H_{\text{out}} |\eta_\perp\rangle) \tag{9}$$

and use bounds

$$| \langle\eta| H_{\text{out}} |\eta_\perp\rangle | \leq \|H_{\text{out}}\| \leq 1, \quad \langle\eta| H_{\text{out}} |\eta\rangle \geq 0, \quad \langle\eta_\perp| H_{\text{out}} |\eta_\perp\rangle \geq 0. \tag{10}$$

This gives

$$\langle\psi| H' |\psi\rangle \geq \Delta\gamma(1-\alpha^2) - 2\alpha\sqrt{1-\alpha^2}. \tag{11}$$

Applying the immediate bound

$$\langle\eta| H' |\eta\rangle = \langle\eta| H_{\text{out}} |\eta\rangle \leq 1 \tag{12}$$

from $\|H_{\text{out}}\| \leq 1$, we obtain

$$\Delta\gamma(1-\alpha^2) - 2\alpha\sqrt{1-\alpha^2} \leq 1 \implies \alpha \geq 1 - \frac{1}{2(\Delta\gamma)} \geq 1 - \frac{1}{50}. \tag{13}$$

Since $\alpha = |\langle\eta|\psi\rangle|$, we ensure that phase estimation from $|\eta\rangle$ yields the ground state of H' with constant probability.

Guided local Hamiltonian problem: Returning to the original guided local Hamiltonian problem, the proof idea is to modify the above construction in a manner that the history state itself is close to a simple state. We pad the beginning of the circuit with polynomially many identities. Then the history state is close to $\sum_{t=1}^{M} |x\rangle |0\rangle \otimes |t\rangle$ with fidelity $1 - M/(M+T)$, which is chosen as the 'guiding state'. By similar reasoning as above, the ground state can be efficiently prepared by phase estimation.