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

Lecture 4: Perturbation Gadgets and Physical Hamiltonians

September 16, 2025 *Scribe: Kevin Kurashima*

1 Wrapping up lecture 3 (Feynman-Kitaev clock encoding)

Last lecture we discussed the Feynman-Kitaev construction for the QMA verifier of witnesses for the Local Hamiltonian problem and sketched a proof for QMA-completeness of the problem. The idea was to introduce a clock register keeping track of which step of the quantum circuit was being simulated, and to define a Hamiltonian penalizing states with invalid histories. However, the particular choice of encoding for the clock was left unspecified.

As a reminder, the propagation term of the Hamiltonian is of the form

$$\mathbb{1} \otimes |t\rangle\langle t| + \mathbb{1} \otimes |t+1\rangle\langle t+1| - U_t \otimes |t+1\rangle\langle t| - U_t^\dagger \otimes |t\rangle\langle t+1|$$

for time t. Kitaev's insight is to use a unary encoding for the clock, also known as the domain wall encoding. So, for $t = 0, 1, ..., T$, the clock register has T qubits, with $|\hat{0}\rangle = |000...0\rangle$, $|\hat{1}\rangle = |100...0\rangle$, $|\hat{2}\rangle = |110...0\rangle$, and so on.

This encoding is nice for two reasons. First, we only need 3 qubits from the clock register to recognize a clock transition $t \leftrightarrow t+1$; for the 3-qubit neighborhood around the clock bit c_t, we have the identification

$$|t+1\rangle\langle t| \leftrightarrow |110\rangle\langle 100| \tag{1}$$

For any other time t' encoded in unary, the 3-qubit neighborhood around c_t would either be all zeros or all ones, so we can uniquely identify the clock transition.

Second, we can easily penalize invalid time encodings; since every "0" bit must be followed by another "0" bit, the following Hamiltonian on the clock register penalizes non-domain wall configurations:

$$H_{\text{format}} = \sum_t |01\rangle\langle 01|_{t,t+1} \tag{2}$$

This adds an energy penalty to any state with a bitstring containing the sequence "01".

Now, the clock control only requires three clock qubits, and the unitary U_t acts on at most two qubits. We can thus strengthen our statement on QMA-completeness of the Local Hamiltonian problem, which we could previously only state for $\lceil \log T \rceil$+2-local terms for a binary clock encoding:

Theorem 1.1. *The 5-Local Hamiltonian problem is QMA-complete.*

2 Physical Hamiltonians

There are a number of constraints and characteristics on physically realizable and practical Hamiltonians. Not all of these are strictly necessary, but typical physical Hamiltonians will have at least some of these properties outlined below.

1. Lattice geometry: many physical Hamiltonians are defined on a (not necessarily rectangular) lattice. One example is the toric code.

2. Translational invariance: under spatial translation of the system by a fixed vector, the Hamiltonian describing the dynamics of the system doesn't change.

3. Nearest-neighbor interactions: many physical Hamiltonians describing many-body systems have interactions only between physically adjacent particles, and thus have largely 2-local terms.

4. Some systems are described by specific models, such as the Bose-Hubbard model, Fermi-Hubbard model, and the Heisenberg model.

Using the Feynman-Kitaev clock construction, we have been able to show QMA-hardness for 5-local Hamiltonians. However, often these kinds of Hamiltonians don't describe physical systems very well, so we would like to also say something about the hardness for physical Hamiltonian problems. Hardness results for physical Hamiltonians would allow us to, for example, talk about which kinds of quantum simulation and ground-state problems are tractable. One tool allowing us to prove hardness results for these kinds of physical Hamiltonians are perturbation gadgets.

3 Perturbation Gadgets (Kempe-Kitaev-Regev)

Our goal with perturbation gadgets is to replace awkward k-local Hamiltonian terms with more realistic 2-local interactions on the system with additional ancillas such that the low-energy effective Hamiltonian on the system approximates the original Hamiltonian. Our main result is the following:

Theorem 3.1. *Given a k-local Hamiltonian H on a system S, there exists a $\lceil k/2 \rceil$-local Hamiltonian H_{sim} on the system S and ancillas A and a unitary (often shallow circuit) U such that $U^\dagger H_{sim} U \approx H$ on the zero-ancilla subspace. This unitary U is known as the Schrieffer-Wolff transform.*

Classical example/motivation

Suppose we have a 4-body Hamiltonian, $H_{\text{cl}} = -Z_1 \otimes Z_2 \otimes Z_3 \otimes Z_4$. We have a ground state when the parity of the four bits is even. We motivate the idea of perturbation gadgets by asking how we can reproduce this parity check using only 3-local terms.

One straightforward solution is to introduce an ancilla bit; call the parity check operator on the ancilla Z_a. Then, our 3-local terms are $H_1 = -Z_1 \otimes Z_2 \otimes Z_a$ and $H_2 = -Z_3 \otimes Z_4 \otimes Z_a$, and we have that $H_{\text{sim}} = H_1 + H_2$. Then, H_1 is minimized when Z_a measures the same parity as $Z_1 \otimes Z_2$, and H_2 is likewise minimized when Z_a measures the same parity as $Z_3 \otimes Z_4$. By minimizing H_{sim} over the ancilla bit, we see that the ground state exhibits even parity among the four original bits, which is what we desired for approximating H_{cl}.

Actually, for the classical case, our effective Hamiltonian matches the target exactly up to a constant. The main insight from the classical case was to couple the ancilla to the existing bits and impose constraints making the ancilla equal to the combined parity of bit groupings, thus forcing the ground state of the resulting sum of 3-local terms to agree with the original. We now want to move on to the quantum case; we can no longer minimize over the ancilla bits however, so our approach will be different.

Quantum attempt

Again, suppose we have a 4-body Hamiltonian, $H = -Z_1 \otimes Z_2 \otimes Z_3 \otimes Z_4 - X_1 \otimes X_2 \otimes X_3 \otimes \mathbb{1}_4$. We introduce two ancilla qubits, measured by Z_5 and X_6 respectively.

We then attempt to build up the 3-local Hamiltonian:

$$\begin{aligned} H_{\text{sim}} &= H_1^Z + H_2^Z + H_3^X + H_4^X \\ &= -(Z_1 \otimes Z_2 \otimes Z_5) - (Z_3 \otimes Z_4 \otimes Z_5) - (X_3 \otimes \mathbb{1}_4 \otimes X_6) - (X_1 \otimes X_2 \otimes X_6) \end{aligned}$$

This approach fails, however! The ground states of H_{sim} are entangled between the original qubits and the ancilla qubits, so we cannot maintain the ground state if we project to the zero-ancilla subspace, as we will see with a quick calculation.

Consider 3-local terms acting on the third qubit $H_2^Z = -Z_3 \otimes Z_4 \otimes Z_5$ and $H_3^X = -X_3 \otimes \mathbb{1}_4 \otimes X_6$. For simplicity, Since $H_2^Z + H_3^X$ commutes with Z_4, we can simplify our analysis to $-Z_3 \otimes Z_5 - X_3 \otimes X_6$. These terms anticommute, so there is no state of qubit 3 which can be a simultaneous eigenstate of both terms. Therefore, the ground state must entangle the third qubit with the ancilla qubits.

What we need to do is to overcome this problem is to force the ancilla qubits into a fixed state to avoid entanglement. We can do this by imposing a penalty term on the ancilla to, roughly speaking, keep them in a fixed state.

Introducing perturbation gadgets

This time, as we build up H_{sim}, we begin with a large penalty term on the ancilla qubits:

$$\Delta \left|1\right\rangle \left\langle 1\right|_5 + \Delta \left|1\right\rangle \left\langle 1\right|_6$$

We then perturb the penalty by $\delta H \equiv J(H_1 + H_2 + H_3 + H_4)$, with $\Delta \gg J$, to obtain

$$H_{\text{sim}} = \Delta \left(\left|1\right\rangle \left\langle 1\right|_5 + \left|1\right\rangle \left\langle 1\right|_6\right) + J(H_1 + H_2 + H_3 + H_4) \tag{3}$$

with $H_1 = Z_1 \otimes Z_2 \otimes X_5$, $H_2 = Z_3 \otimes Z_4 \otimes X_5$, $H_3 = X_3 \otimes X_4 \otimes X_6$, $H_4 = X_1 \otimes X_2 \otimes X_6$.

Why does this work? First, note that if one of the ancilla bits is in the excited state, we incur a penalty of Δ, and a penalty of 2Δ if both are excited. Therefore, for the unperturbed simulation Hamiltonian, the ground state is achieved when both ancilla bits are in their ground states.

We will see this more precisely, but using Schrieffer-Wolff transformations to treat H_{sim} terms with perturbation theory, we can get the effective Hamiltonian on the low-energy subspace (that is, the zero-ancilla subspace) to well approximate H. The perturbative terms H_1 and H_2 induce transitions $\left|0\right\rangle_5 \to \left|1\right\rangle_5 \to \left|0\right\rangle_5$ due to the X_5 operator, and the perturbative terms H_3 and H_4 induce the same kind of transition on the "6" ancilla qubit. Therefore, the effective Hamiltonian looks like

$$H_{\text{eff}} = -(Z_1 \otimes Z_2 \otimes Z_3 \otimes Z_4) - (X_1 \otimes X_2 \otimes X_3 \otimes X_4)$$

which is what we desire.

3.1 Introducing the Schrieffer-Wolff Transformation

To show how the perturbation theory expansion works, we restrict our attention to a single ancilla, and consider involutory operators A and B on disjoint parts of the system so that $H = A \otimes B$. Then, our simulated Hamiltonian is constructed by perturbing the penalty term on the ancilla:

$$H_{\text{sim}} = \Delta \left|1\right\rangle \left\langle 1\right|_a + J(A \otimes \mathbb{1} + \mathbb{1} \otimes B) \otimes X_a \tag{4}$$

again with $\Delta \gg J > 0$. Note that since A and B act on disjoint parts of the system, $[A, B] = 0$.

We can visualize H_{sim} as a block-diagonal matrix in the computational basis of the ancilla $\{|0\rangle_a, |1\rangle_a\}$:

$$H_{\text{sim}} = \begin{pmatrix} 0 & J(A+B) \\ J(A+B) & \Delta \mathbb{1} \end{pmatrix} \tag{5}$$

which can now be diagonalized. All of the blocks commute with each other, so our eigenvalues are exactly

$$\lambda_{\pm} = \frac{\Delta}{2} \pm \sqrt{\left(\frac{\Delta}{2}\right)^2 + J^2(A+B)^2} \tag{6}$$

Since $\Delta \gg J$, $A^2 = \mathbb{1}$, and $B^2 = \mathbb{1}$, we can do a perturbative expansion on the eigenvalues, and our diagonalized Hamiltonian $\Lambda = U H_{\text{sim}} U^\dagger$ is

$$\Lambda = \begin{pmatrix} \frac{-J^2(A+B)^2}{\Delta} & 0 \\ 0 & \Delta \mathbb{1} + \frac{J^2(A+B)^2}{\Delta} \end{pmatrix} + O\left(\frac{J^4}{\Delta^3}\right) \tag{7}$$

for unitary U. Choosing $J = \sqrt{\Delta}$, we get that the low-energy term is now $-(A+B)^2$, the high-energy term is now $\Delta \mathbb{1} + (A+B)^2$, and the higher-order terms are $O(1/\Delta)$. It turns out that the unitary U has a nice form as well:

$$U \equiv e^S = e^{-i\frac{J}{\sqrt{\Delta}}(A+B)\otimes Y_a} = e^{-i(A+B)\otimes Y_a} \tag{8}$$

Finally, since A and B commute with each other and are involutory, $(A+B)^2 = 2\mathbb{1} + 2AB$ up to a constant. So, the lower-energy part of the system, H_{eff}, matches with the desired Hamiltonian closely in the zero-ancilla subspace.

In the example above, A and B correspond to H_1 and H_2 respectively to get the term $(Z_1 \otimes Z_2 \otimes Z_3 \otimes Z_4)$ in the effective Hamiltonian. They similarly correspond to H_3 and H_4 to get the $(X_1 \otimes X_2 \otimes X_3 \otimes X_4)$ term in H_{eff}.

3.2 Proof Sketch of Theorem 3.1

We can generalize the approach for perturbation gadget treatments to Hamiltonians with multiple k-local terms. For k-local Hamiltonian $H = \sum_\alpha b_\alpha P_\alpha$, our simulated Hamiltonian has the form

$$H_{\text{sim}} = \sum_\alpha b_\alpha \left(P_\alpha^{(1)} + P_\alpha^{(2)}\right) + \Delta \sum_g |1\rangle\langle 1|_g \tag{9}$$

Then, treating it with the Schrieffer-Wolff transformation, we have that the effective Hamiltonian in the zero-ancilla subspace is given by $e^{iS} H_{\text{sim}} e^{-iS}$, with the relation being approximate when S is a shallow circuit.