As we have discussed, we expect that quantum computers can compute the ground state energy of local Hamiltonians in cases where the problem is hard classically; this may be an important application for quantum computers. On the other hand, we believe that in some cases computing the ground state energy of a local Hamiltonian is still hard even for quantum computers. Let us try to understand more deeply the reason for this belief.

Physicists are often interested in translation-invariant geometrically local Hamiltonians, in which all qubits interact in the same way with their neighbors (except for the qubits at the boundary of the sample), because such Hamiltonians provide good models of some real materials. But Hamiltonians that are not translationally invariant are also useful in physics (for example, when modeling a material with "disorder" due to dirt and other inperfections in the sample. If the Hamiltonian is not translation invariant, then we can formulate an instance of an n-qubit local Hamiltonian problem by specifying how the Hamiltonian varies from site to site in the system. Physicists sometimes refer to such (not translationally-invariant) systems as "spin glasses".

Even in the classical case, where the variable at each site is a bit rather than a qubit, finding the ground state energy of a spin glass to constant accuracy can be an NP-hard problem. Therefore, we don't expect classical and quantum computers to be able to solve the problem in general (unless NP is contained in BQP, which seems unlikely).

Let's first understand why the classical spin-glass problem can be NP-complete. Then we'll discuss the hardness of the quantum problem. We'll see that in the quantum case finding the ground state energy of a local Hamiltonian is QMA-complete. (Recall that QMA is the quantum analogue of NP: the class of problems such that the solution can be verified efficiently with a quantum computer if a suitable "quantum witness" is provided.)

For the classical case, we'll recall the notion of a "reduction" of one computational problem to another (B reduces to A if a machine that solves A can be used to solve B), and then we'll consider this sequence of reductions:

1) Any problem in NP reduces to CIRCUIT-SAT (already discussed previously); i.e., CIRCUIT-SAT is NP-complete.
2) CIRCUIT-SAT reduces to 3-SAT (3-SAT is NP-complete).
3) 3-SAT can be formulated as the problem of finding the ground state energy of a classical 3-local Hamiltonian to constant accuracy.
4) MAX 2-SAT is also NP-complete and can be formulated as the problem of finding the ground state energy of a classical 2-local Hamiltonian to constant accuracy.
5) The classical 2-local Hamiltonian problem is still hard in the case where the Hamiltonian is geometrically local, in three or even in two dimensions (cases of interest for describing real spin glasses).

(5) implies that a spin glass will not be able to relax to its ground state efficiently in any realistic physical process (which is part of what physicists mean by the word "glass").

Language: Recall (as discussed earlier) that if f is a uniform family of Boolean functions with variable input size , f: {0,1}* -> {0,1}, then the set of input strings accepted by f is called a *language*:

L = { x in {0,1}* : f(x) =1 } .

NP: We say that a language is in NP if there is a polynomial-size uniform classical circuit family (the *verifier* V(x,y)) such that:

If x in L, then there exists a "witness" y such that V(x,y)=1 (completeness).
If x not in L, then, for all y, V(x,y)=0 (soundness).

Reduction: We say that B reduces to A if there is a polynomial-size uniform classical circuit family R mapping x to R(x) such that B accepts x if and only if A accepts R(x). This means we can hook up R to a machine that solves A to construct a machine that solves B.

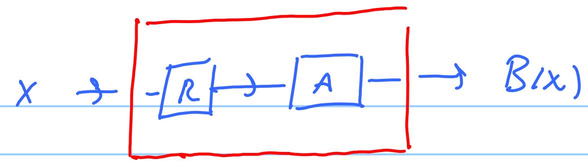

An important problem in NP is CIRCUIT-SAT. The input to the problem is a Boolean circuit C (with an n-bit input and G= poly(n) gates), and we are to evaluate the Boolean function f(C), where

f(C)= 1 if there is an input x such that C(x)=1,
f(C)= 0 otherwise.

CIRCUIT-SAT is in NP because we can simulate the circuit C. Given as a witness the value of x that C accepts, we can verify efficiently that C(x)=1.

Furthermore, CIRCUIT-SAT is NP-complete (any problem in NP reduces to CIRCUIT-SAT), as we discussed previously. If V(x, .) is the verifier for an NP problem with a fixed instance x, we may think of V(x, .) as a circuit whose input is the witness y. Solving the CIRCUIT-SAT problem for this Boolean circuit tells us *whether* there exists a witness that the verifier accepts, and therefore solves the NP problem.

Now we come to a further reduction that we did not discuss previously: CIRCUIT-SAT reduces to 3-SAT, and therefore 3-SAT, too, is an NP-complete problem (the Cook-Levin theorem).

For the SAT problem, the input is a "Boolean formula" with n variables, where each variable is a bit. The formula is a conjunction of clauses, and the formula is true if and only if every clause is true. In the k-SAT problem, each clause depends on at most k of the variables, where k is a constant. (In some formulations of k-SAT, each clause is required to be a disjunction of k "literals" (variables or their negations), but that is not an important requirement, since any formula, and in particular any k-bit formula, can be expressed in conjunctive normal form.). If f is a Boolean formula, the SAT function is:

SAT(f) = 1 if there exists x such that f(x)=1
SAT(f) = 0 otherwise.

Now we'll show that CIRCUIT-SAT reduces to 3-SAT. For a given circuit C (the input to CIRCUIT-SAT), how do we construct the corresponding Boolean formula R(C) (the input to 3-SAT)?

Suppose that the gates in the circuit C are chosen from the universal set (AND, OR, NOT), or any other gate set such that each gate has at most two-input bits and one output bit. We introduce a variable for the output of each gate, and we include in the formula R(C) a clause corresponding to each gate.

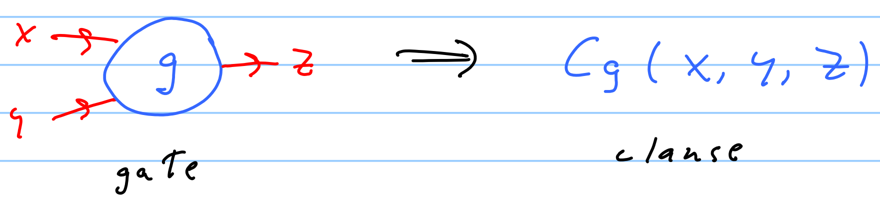

Here the three-variable clause $C_g(x,y,z)$ is true iff z is a valid output of the gate g when the inputs are (x, y). The circuit may also have inputs that are constants rather than variables; then, e.g. a gate with two input bits, one of which is a constant, becomes a two-variable clause, determined by the gate

$0, x \to g \to y \implies C_{g,0}(x,y)$ and the value of the constant.

$0 \to x \implies C_0(x)$

$1 \to y \implies C_1(x)$

or equivalently, we can regard inputing a constant as a gate with a one-bit output; the corresponding one-bit clause is true if x has the right value.

$x \to g$

the circuit also has an output bit, which becomes a 1-bit clause $C_g(x)$ which is true iff $x=1$ (that is, if and only if the output is accepted)

The formula R(C) has as variables the input x to the circuit C, and also additional variables corresponding to the outputs of all gates in the circuit C. R(C) has been constructed so that an assignment that satisfies every clause in C corresponds to a *valid history* of the computation of the circuit C acting on input x, where the input is accepted. If there is an input x that is accepted by the circuit C, then there will be a satisfying assignment for the 3-SAT formula R(C), and conversely if there is no input that C accepts, then there will be no satisfying assignment for R(C).

The key idea we have exploited to reduce CIRCUIT-SAT to 3-SAT is that the witness for 3-SAT is a valid history of the whole computation C(x) that accepts the input x. We can check the history efficiently because the circuit C has polynomial size and it is easy to check each of the poly(n) gates in the execution of the circuit. Later on, we will extend this idea --- that a valid history of the computation is an efficiently checkable witness --- to the quantum setting.

Notice that we may think of the clauses in the formula f as the terms in a 3-local classical Hamiltonian

$$H(x) = \sum_c H_c(x_{c1}, x_{c2}, x_{c3})$$

here the sum is over the clauses in the formula, where

$$H_c(x_{c1}, x_{c2}, x_{c3}) = \begin{cases} 0 & \text{if clause } C \text{ is true for assignment } x_{c1}, x_{c2}, x_{c3}, \\ 1 & \text{otherwise.} \end{cases}$$

Then $\min_x H(x) = 0$ if there is an assignment that satisfies every clause, while $\min_x H(x) \geq 1$ if there is no satisfying assignment (the number of violated clauses is ≥ 1 for any assignment).

We conclude that finding the minimum value of a 3-local classical Hamiltonian is NP-hard: if we could do it we could solve 3-SAT, and hence any problem in NP. This conclusion implies, as asserted earlier, that finding the ground state energy of a 3-local classical Hamiltonian to constant accuracy must be hard in general for quantum computers, unless NP $\subseteq$ BQP

In fact, finding the ground state energy to constant accuracy is NP-hard even for a 2-local classical Hamiltonian

$$H(x) = \sum_c H_c(x_{c1}, x_{c2})$$

Although 2-SAT (deciding whether a 2-SAT formula can be satisfied) is easy (there is a poly time algorithm), MAX-2SAT is an NP-hard problem. MAX-2SAT is the problem of finding the minimum number of violated clauses for any assignment, which is equivalent to minimizing the Hamiltonian function H.

Furthermore, we can make the 2-local Hamiltonian geometrically local without losing hardness. An example is the "Ising spin-glass model" in three dimensions. Suppose the binary variables are "spins" sitting at the sites of a cubic lattice, taking values $Z_i \in \{\pm 1\}$ at site i in the lattice. Consider the Hamiltonian

$$H = -\sum_{\langle ij \rangle} J_{ij} Z_i Z_j$$

Here $\langle ij \rangle$ labels the edge in the lattice that connects two nearest-neighbor sites with labels i and j

$J_{ij} \in \{\pm 1\}$ encodes the instance of the problem. If $J_{ij} = +1$, then we say the edge $\langle ij \rangle$ is <u>ferromagnetic</u>; it is energetically favorable for the neighboring spins to align (both +1 or both −1)

If $J_{ij}=-1$, then we say the edge $\langle ij\rangle$ is antiferromagnetic; it is energetically favorable for the neighboring spins to antialign (either +1 and −1 or −1 and +1).

If all edges were ferromagnetic, it would be easy to minimize the energy — all spins would align. But antiferromagnetic edges can generate frustration. This means it is not possible to minimize $-J_{ij}Z_iZ_j$ for all edges simultaneously.

For purposes of visualization, it is convenient to represent spins by lattice cells — i.e., by squares on the 2D square lattices or by cubes in the 3D cubic lattice

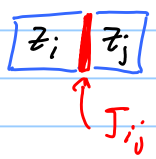

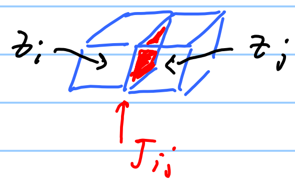

There is a $-J_{ij}Z_iZ_j$ coupling neighboring spins associated with each edge where two squares meet in 2D, or each face where two cubes meet in 3D

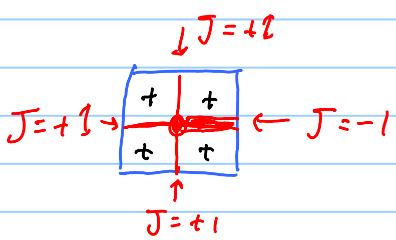

Consider a site in 2D where 4 edges meet. If one of these edges is antiferro and the other three are ferro, then there must be an odd number of "excited" edges meeting at this site

More generally, if the no. of antiferro edges is odd, then there must be an odd number of excited edges and if the no of antiferro edges is even, then there must be an even number of excited edges. If the number of $J=-1$ edges meeting at a site is odd, we say that there is an "Ising vortex" at that site. For any spin config, then, there are "domain walls" of excited edges, where the walls end on Ising vortices

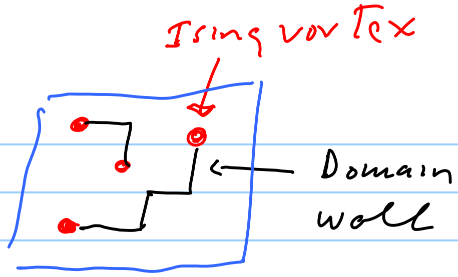

Minimizing the energy, then is equivalent in 2D to finding the minimum "chain" of excited edges with boundary points at the positions of the Ising vortices. There is a poly-time classical algorithm that finds the min chain

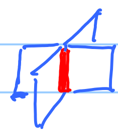

In 3D, there is an Ising vortex on an edge if there are an odd number of $J=-1$ faces that meet at that edge. These vortices form closed loops, and each spin configuration has a "domain wall" of excited faces bounded by the vortex loops. To minimize the energy, then, we find the minimum area surface with a specified 1D boundary — this problem is NP-hard

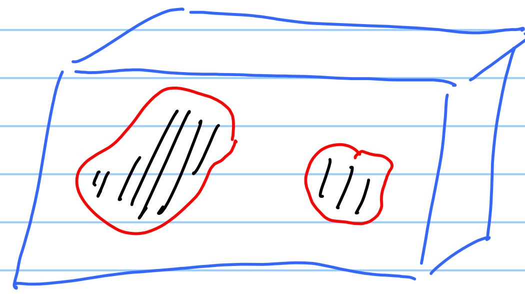

Finding the minimum energy configuration is hard because there are many ways for the domain walls to be pinned — stuck at local minima of the energy, such that many spins need to flip at once to find a lower energy configuration. "Local searching" for the global energy minimum fails.

In fact, there are also NP-hard spin glass problems in 2D, if we introduce local "magnetic field" terms in H as well as antiferromagnetic terms. For example, on a square lattice consider

$$H = -\sum_{\langle ij\rangle} J_{ij} Z_i Z_j - \sum_i h_i Z_i$$

where $J_{ij} \in \{1, 0, -1\}$ and $h_i \in \{1, 0, -1\}$

The local magnetic field $\{h_i\}$ compounds the frustration: Each spin wants to align with the local field, but by doing so the edge connecting the spins might become excited.

So we see that minimizing the energy of a geometrically 2-local classical Hamiltonian can be NP-hard because of frustration — there is no way to satisfy all the "clauses" and there are many local minima of the energy that are not global minima. In the quantum local Hamiltonian problem, we have $H = \sum_a H_a$

where the $\{H_a\}$ might not be mutually commuting; hence we might expect the problem to be even harder — that the H_a's cannot be simultaneously diagonalized compounds the frustration even further. Indeed, the ground state could be highly entangled, with no succinct classical description. Let's try to characterize the hardness of the quantum problem.

The quantum k-local Hamiltonian problem.

Let $H = \sum_a H_a$ be an n-qubit Hamiltonian that is <u>k-local</u> — each H_a acts nontrivially on at most k qubits, where k is a constant, and $\|H_a\| \le h$ = constant. The $2^k \times 2^k$ matrix H_a is specified to poly(n) bits of precision.

We are promised that the ground state energy E_0, the lowest eigenvalue of H, satisfies either

(i) $E_0 \le E_{low}$, or

(ii) $E_0 > E_{high}$

where $E_{high} - E_{low} > \frac{1}{\text{poly}(n)}$.

We are to output: $f(H) = \begin{cases} 1 & \text{if } E_0 \le E_{low} \\ 0 & \text{if } E_0 > E_{high} \end{cases}$

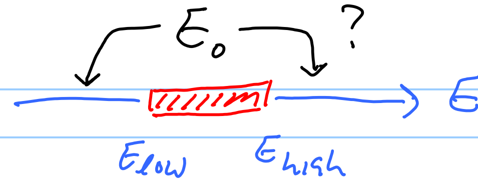

Note that in the formulation of the problem there is a "promise gap" so that we can answer the yes/no question by determining E_0 to $1/\mathrm{poly}(n)$ accuracy.

We already know that this problem is NP-hard, even in the special case where H is classical (all H_a commute). Also recall that there is a class analogous to NP for randomized computation (MA) and quantum computation (QMA):

A language L is in QMA if there exists a poly-size uniform quantum circuit family (the <u>verifier</u> V) and a single-qubit measurement $\{E_0, E_1\}$ such that

- If $x \in L$ there is a witness $|\psi_x\rangle$ such that
  $\langle \Psi | E_1 | \Psi \rangle \geq \frac{2}{3}$ where $|\Psi\rangle = V\left(|\psi_x\rangle \otimes |x\rangle \otimes |0\rangle^{*}\right)$

- If $x \notin L$, then $\langle \Psi | E_1 | \Psi \rangle \leq \frac{1}{3}$ for all $|\psi_x\rangle$

We claim: the k-local Hamiltonian problem is QMA-complete for $k \geq 5$ (Kitaev; see Chap. 14 of KSV)

The result can be improved to geometrically 2-local H in 2D (for qubits) or geometrically 2-local H in 1D (for higher dimensional qudits, with $d \geq 12$).

We need to show:

(1) The k-local Hamiltonian problem is in QMA.

(2) Any problem in QMA is reducible to k-local Hamiltonian. We'll show this for $k = 5$ and without geometrical constraints, as in Kitaev's original discovery.

We have already shown part (1). We have seen that, if the ground state $|\psi_0\rangle$ is provided as a witness, then we can compute E_0 to $1/\mathrm{poly}(n)$ accuracy with $\mathrm{poly}(n)$ accuracy using the phase estimation algorithm. But how to achieve the reduction (2)?

For the reduction, we'll follow the strategy used to show that 3SAT is NP complete: For any problem in QMA, we'll construct a witness that encodes the whole history of the computation performed by the verifier, and a Hamiltonian H such that computing the ground state energy of H amounts to checking that each step in the computation is valid.

For a given problem in QMA, suppose that the verifier circuit V_T has T gates: $U_1, U_2, \ldots, U_T$ chosen from a universal set, where each U_t acts on at most two qubits. For the corresponding k-local Hamiltonian problem, we'll suppose that Merlin provides the history state encoding the computation performed by the verifier:

$$|\eta\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^{T} |\psi(t)\rangle \otimes |t\rangle$$

where $|\psi(t)\rangle = (U_t U_{t-1} \cdots U_1)|\psi_{wit}\rangle$;

here $|\psi_{wit}\rangle$ is the witness for the given QMA problem, and $|\psi(t)\rangle$ is the state obtained after the first t steps of the verifier circuit. The state $|t\rangle$ is the state of a clock register that records the time $t \in \{0, 1, \ldots, T\}$; since $\langle t|s\rangle = \delta_{ts}$, the $T+1$ states appearing in superposition is state $|\eta\rangle$ are mutually orthogonal, so $|\eta\rangle$ is properly normalized: $\langle\eta|\eta\rangle = 1$.

Now we want to choose our Hamiltonian H so that it locally "check" the history encoded in $|\eta\rangle$;

That is, H will impose an energetic penalty if a step in the circuit is invalid.

We will choose H to be of the form

$$H = H_{in} + H_{out} + H_{prop} + H_{clock}$$

The purpose of H_{in} is to enforce that the verifier circuit's input qubits (other than the witness itself) are set to the right initial values. E.g. for each scratch qubit that should be in the state $|0\rangle$ at time $t=0$, we include in H_{in} the term

$$H_{in}^{(j)} = \left(|1\rangle\langle 1|\right)^{(j)} \otimes I^{(else)} \otimes \left(|0\rangle\langle 0|\right)^{(clock)}$$

There is an energy penalty of 1 if the scratch qubit is set to $|1\rangle$ rather than $|0\rangle$ as the execution of the verifier circuit begins (time $t=0$). Same for the bits of the input x.

The purpose of H_{out} is to impose a penalty if the verifier circuit for the given QMA problem fails to accept

$$H_{out} = \left(|0\rangle\langle 0|\right)^{(output)} \otimes I^{(else)} \otimes \left(|T\rangle\langle T|\right)^{(clock)}$$

There is an energy cost of 1 if the output qubit has the value $|0\rangle$ rather than $|1\rangle$ after the verifier circuit is executed (time $t=T$).

The purpose of H_{clock} is to impose a penalty if the clock register is not properly encoded (we'll return to this issue later).

The purpose of H_{prop} is to impose a penalty if the state $|\psi(t)\rangle$ does not have the form $U_t|\psi(t-1)\rangle$; i.e., was not obtained by faithfully executing step t of the verifier circuit. Hence we write:

$$H_{prop} = \sum_{t=1}^{T} H_{prop}(t) \quad , \text{ where}$$

$$H_{prop}(t) = \frac{1}{2}\Big[I \otimes |t\rangle\langle t| + I \otimes |t-1\rangle\langle t-1| - U_t \otimes |t\rangle\langle t-1| - U_t^{\dagger} \otimes |t-1\rangle\langle t| \Big]$$

The action of $H_{prop}(t)$ on the relevant part of the valid history state $|\eta\rangle$ is

$$|\psi(t-1)\rangle \otimes |t-1\rangle \longmapsto \frac{1}{2}\big[\psi(t-1) \otimes |t-1\rangle - \psi(t) \otimes |t\rangle \big]$$

$$|\psi(t)\rangle \otimes |t\rangle \longmapsto \frac{1}{2}\big[\psi(t) \otimes |t\rangle - \psi(t-1) \otimes |t-1\rangle \big].$$

Acting on $|\eta\rangle$, the $I \otimes |t\rangle\langle t|$ and $-U_t\,|t\rangle\langle t-1|$ terms give cancelling contributions. Hence

$H_{prop}|\eta\rangle = 0$ if $|\eta\rangle$ is a valid history state

Therefore a valid history state (where the state at time t is obtained from the state at time $t-1$ by applying the proper gate) such that the initial state is also valid is a null vector of both H_{prop} and H_{in}. Furthermore, if the quantum verifier accepts the input with probability $1-\varepsilon$, then

$\langle \eta | H_{out} | \eta \rangle = \dfrac{\varepsilon}{T+1}$ (the $t=T$ portion of the state has amplitude $1/\sqrt{T+1}$, and

We therefore conclude that the ground state energy E_0 of

$H = H_{in} + H_{out} + H_{prop}$ is $E_0 \le \langle \eta | H_{out} | \eta \rangle \le \dfrac{\varepsilon}{T+1}$

Note also that it is possible to amplify the success probability by repeating the verification on on multiple copies of the witness. Actually, the amplification is a little bit subtle: Merlin might try to fool Arthur. Instead of sending a product state $|\psi_x\rangle^{\otimes m}$ (m copies of the witness) he might send an entangled state instead. But the amplification still works. Each copy sent by Merlin may be a mixed state (obtained from the partial trace over the other copies), a mixture of a state the verifier accepts with prob $\geq 2/3$ and of another state that it might reject. But Merlin cannot fool Arthur into accepting after many trials unless there is some state that is accepted with high probability in each trial. Therefore — we may assume ε is exponentially small, and therefore:

$$E_0 \leq 2^{-\Omega(n)}$$

So now we have seen that for a problem in QMA such that the verifier accepts with high prob., the corresponding Hamiltonian H has an eigenvector with eigenvalue close to zero. There are two things left to show

- If the verifier rejects with high probability, then

$$E_0 \geq 1/\mathrm{poly}(n)$$

(Then we can choose the promise gap of size $1/\mathrm{poly}(n)$, such that $E_0 \leq E_{low}$ for a YES answer and $E_0 \geq E_{high}$ for a NO answer)

- So far H_{prop} is not local! It acts on the clock register, which is a $(T+1)$-dimensional system. We need to show we can encode the clock using qubits such that H_{prop} is k-local.

We'll come back to the issue of encoding the clock later. First let's try to understand the spectrum of $H = H_{in} + H_{out} + H_{prop}$.

Start by considering H_{prop}. We've seen that a valid history is a null vector of H_{prop} (has eigenvalue zero)

What are the other eigenspaces and eigenvalues of H_{prop}?

It's easier to compute the spectrum of H_{prop} by transforming to a "rotating frame" basis that "freezes the motion" of the state $|\psi(t)\rangle$. That is, let

$$V_t = U_t U_{t-1} \cdots U_1$$ (the unitary applied after the first t steps of the circuit)

and consider $V = \sum_{t=0}^{T} V_t \otimes |t\rangle\langle t|$

Then the history state $\eta = \frac{1}{\sqrt{T+1}} \sum_{t=0}^{T} V_t |\psi(0)\rangle \otimes |t\rangle\langle t|$

is mapped to a "constant" by $V^\dagger$

$$V^\dagger|\eta\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^{T} V_t^\dagger V_t |\psi(0)\rangle \otimes |t\rangle\langle t| = |\psi(0)\rangle \otimes \sum_{t=0}^{T} |t\rangle\langle t|$$

and the Hamiltonian transforms as

$$H'_{prop} = V^\dagger H_{prop} V = \sum_{t=1}^{T} \tfrac{1}{2}\Big[I \otimes |t\rangle\langle t| + I \otimes |t-1\rangle\langle t-1| - V_t^\dagger U_t V_{t-1} \otimes |t\rangle\langle t-1| - V_{t-1}^\dagger U_t^\dagger V_t \otimes |t-1\rangle\langle t| \Big]$$

$$= I \otimes \sum_{t=1}^{T} \tfrac{1}{2}\big[|t\rangle\langle t| + |t-1\rangle\langle t-1| - |t\rangle\langle t-1| - |t-1\rangle\langle t| \big]$$

– the transformed Hamiltonian H'_{prop} acts nontrivially only on the clock. It is a sum of overlapping 2×2 blocks:

$$H'_{prop} = \sum_{t=1}^{T} \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{pmatrix}_{t-1,t}$$ meaning it acts on the space spanned by $\{|t-1\rangle, |t\rangle\}$

Because of the overlaps, H'_{prop} is actually $\begin{pmatrix} 1 & -\frac{1}{2} \\ -\frac{1}{2} & 1 \end{pmatrix}_{t-1,t}$ in each block, except in the spaces spanned by $\{|0\rangle, |1\rangle\}$ and by $\{|T-1\rangle, |T\rangle\}$, where it is:

$$\begin{pmatrix} \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & 1 \end{pmatrix}_{0,1} \quad \text{and} \quad \begin{pmatrix} 1 & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{pmatrix}_{T-1,T}$$

That is, H'_{prop} acts on the clock as a $(T+1)\times(T+1)$ matrix

$$H'_{prop} = \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} & & & & \\ -\frac{1}{2} & 1 & -\frac{1}{2} & & & \\ & -\frac{1}{2} & 1 & \ddots & & \\ & & \ddots & & & \\ & & & & 1 & -\frac{1}{2} \\ & & & & -\frac{1}{2} & \frac{1}{2} \end{pmatrix}$$

That is -- it is $(-\frac{1}{2}, -\frac{1}{2}, \dots -\frac{1}{2})$ just above and below the diagonal, $(\frac{1}{2}, 1, 1, \dots, 1, \frac{1}{2})$ on the diagonal, and 0 elsewhere.

We may express it as $H'_{prop} = I - \frac{1}{2}M$, where

$$M = \begin{pmatrix} 1 & 1 & & & & \\ 1 & 0 & 1 & & & \\ & 1 & 0 & 1 & & \\ & & 1 & \ddots & & \\ & & & & 0 & 1 \\ & & & & 1 & 1 \end{pmatrix}$$

That is

$M: |t\rangle \mapsto |t+1\rangle + |t-1\rangle$ for $t \in \{1, 2, \dots, T-1\}$

$M: |0\rangle \mapsto |0\rangle + |1\rangle$

$M: |T\rangle \mapsto |T-1\rangle + |T\rangle$

Diagonalizing M will also diagonalize H'_{prop}. The eigenvectors of M are linear combinations of the (unnormalized) vectors

$$|\omega\rangle = \sum_{t=0}^{T} e^{i\omega t}|t\rangle$$

In the expansion of $M|\omega\rangle$, the coefficient of $|t\rangle$ for $t \in \{1, 2, \dots, T-1\}$ is

$$e^{i\omega(t-1)} + e^{i\omega(t+1)} = (e^{i\omega} + e^{-i\omega})e^{i\omega t}$$
$$= 2\cos\omega\,(e^{i\omega t})$$

The action on $|t\rangle$ is multiplication by $2\cos\omega$, which does not depend on the sign of ω. To construct an eigenstate of M, then, we may consider a linear combination of $|\omega\rangle$ and $|-\omega\rangle$, where the value of $|\omega\rangle$ is chosen so that M acts properly on the states at the "boundary" — i.e. on $|t=0\rangle$ and $|t=T\rangle$.

The coefficient on $|0\rangle$ in $M|\omega\rangle$ is $1+e^{i\omega}$. Therefore, the coefficient of $|0\rangle$ in

$$M\left(e^{i\omega/2}|\omega\rangle + e^{-i\omega/2}|-\omega\rangle\right) \text{ is}$$

$$e^{i\omega/2}\left(1+e^{i\omega}\right) + e^{-i\omega/2}\left(1+e^{-i\omega}\right)$$

$$= \left(e^{i\omega}+e^{-i\omega}\right)\left(e^{i\omega/2}+e^{-i\omega/2}\right) = 2\cos\omega\left(e^{i\omega/2}+e^{-i\omega/2}\right)$$

This action on $\left(e^{i\omega/2}+e^{-i\omega/2}\right)|0\rangle$ is also consistent with an eigenstate with eigenvalue $2\cos\omega$, for any value of ω.

Now consider the other boundary, at $t=T$. The coefficient of $|T\rangle$ in $e^{i\omega/2}|\omega\rangle + e^{-i\omega/2}|-\omega\rangle$ is $e^{i\omega/2}e^{i\omega T} + e^{-i\omega/2}e^{-i\omega T}$, while the coefficient of $|T\rangle$ in $M\left[e^{i\omega/2}|\omega\rangle + e^{-i\omega/2}|-\omega\rangle\right]$ is

$$\left(1+e^{-i\omega}\right)e^{i\omega/2}e^{i\omega T} + \left(1+e^{i\omega}\right)e^{-i\omega/2}e^{-i\omega T}$$

we are to choose ω so that

$$\left(1+e^{-i\omega}\right)e^{i\omega/2}e^{i\omega T} + \left(1+e^{i\omega}\right)e^{-i\omega/2}e^{-i\omega T}$$

$$= \left(e^{i\omega}+e^{-i\omega}\right)\left(e^{i\omega/2}e^{i\omega T} + e^{-i\omega/2}e^{-i\omega T}\right).$$

After some cancellations, this condition becomes

$$e^{i\omega/2}e^{i\omega T} + e^{-i\omega/2}e^{-i\omega T} = e^{3i\omega/2}e^{i\omega T} + e^{-3i\omega/2}e^{-i\omega T}$$

or $\quad \cos\left[\omega\left(T+\tfrac{1}{2}\right)\right] = \cos\left[\omega\left(T+\tfrac{3}{2}\right)\right]$

The condition $\cos\left[\omega(T+1) - \frac{\omega}{2}\right] = \cos\left[\omega(T+1) + \frac{\omega}{2}\right]$
is satisfied provided that

$$\omega(T+1) = \pi K \quad \text{where } K = \text{integer}.$$

Therefore we have established that

$\{2\cos\omega_K\}$ are $T+1$ distinct eigenvalues of M

where $\omega_K = \frac{\pi}{T+1} K, \quad K \in \{0, 1, 2, \ldots, T\}$

For $K = T+1 \Rightarrow$ $\omega_K = \pi$, we have $|\psi_K\rangle = 0$ – not an eigenstate

The corresponding eigenvectors are

$$|\psi_K\rangle = N_K \sum_{t=0}^{T} \cos\left[\omega_K(t+\tfrac{1}{2})\right] |t\rangle$$

and the eigenvalue of H'_{prop} is $E_K = 1 - \cos\omega_K$.

These $T+1$ eigenstates are a complete basis, and the two smallest eigenvalues are:

$$E_0 = 0$$
$$E_1 = 1 - \cos\left(\frac{\pi}{T+1}\right) = 2\sin^2\left(\frac{\pi}{2(T+1)}\right) \approx \frac{\pi^2}{2(T+1)^2}.$$

If $T = \text{poly}(n)$, then the eigenvalue gap for H_{prop} is $E_1 - E_0 \geq \frac{1}{\text{poly}(n)}$.

But we want to consider the spectrum for the full Hamiltonian $H = H_{in} + H_{out} + H_{prop}$

For $H_{in} + H_{out}$, the null space is spanned by all vectors that have a valid input at $t=0$ and answer YES (are accepted) at $t = T$:

$$(H_{in} + H_{out})\,|\text{valid input, accepted output}\rangle = 0,$$

while $\langle H_{in} + H_{out}\rangle \geq 1$ for all vectors orthogonal to this null space. Thus $H_{in} + H_{out}$ has eigenvalue gap $= 1$. Now, if the verifier accepts the input with probability one, then there is a simultaneous null eigenvector of H_{prop} and of $H_{in} + H_{out}$. That is, there is a valid history with a valid input where the output is accepted. But if the acceptance probability is small, that means the angle between these two null spaces cannot be too small.

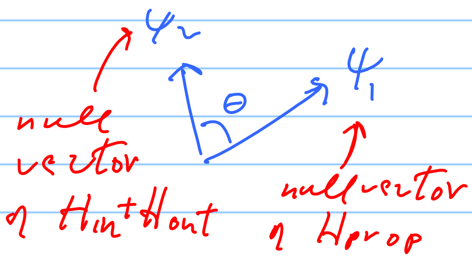

We can relate this angle to the ground state energy of the full Hamiltonian $H = H_{in} + H_{out} + H_{prop}$.

In general, suppose that H_1 and H_2 are two hermitian operators, each with lowest eigenvalue zero, and eigenvalue gap at least Δ. Then

$H_1 \geq \Delta(I - \Pi_1)$ where Π_1 is projection onto null space of H_1

$H_2 \geq \Delta(I - \Pi_2)$ where Π_2 is projection onto null space of H_2

Thus $\quad H_1 + H_2 \geq \Delta(2I - \Pi_1 - \Pi_2)$

and $\quad \langle H_1 + H_2\rangle \geq 2\Delta - \Delta\langle \Pi_1 + \Pi_2\rangle$

But suppose $|\psi_1\rangle$ and $|\psi_2\rangle$ are two vectors such that $|\langle\psi_1|\psi_2\rangle| = \cos\theta$ for $0 \leq \theta \leq \pi/2$. With suitable phase conventions we may choose a basis in the two-dim space spanned by $|\psi_1\rangle$ and $|\psi_2\rangle$ such that

$$|\psi_1\rangle = \begin{pmatrix} \cos\theta/2 \\ \sin\theta/2 \end{pmatrix}, \quad |\psi_2\rangle = \begin{pmatrix} \cos\theta/2 \\ -\sin\theta/2 \end{pmatrix} \Longrightarrow$$

$$|\psi_1\rangle\langle\psi_1| + |\psi_2\rangle\langle\psi_2| = \begin{pmatrix} 2\cos^2\theta/2 & 0 \\ 0 & 2\sin^2\theta/2 \end{pmatrix}$$

and therefore, in any state

$$\langle\, |\psi_1\rangle\langle\psi_1| + |\psi_2\rangle\langle\psi_2| \,\rangle \le 2\cos^2\frac{\theta}{2} = 1+\cos\theta$$

More generally, if Π_1 and Π_2 are projectors, where the max overlap between spaces projected by Π_1 and Π_2 is $|\langle\psi_1|\psi_2\rangle| = \cos\theta$, then $\langle \Pi_1 + \Pi_2 \rangle \le 1+\cos\theta$

Thus $\langle H_1 + H_2 \rangle \geq 2\Delta - \Delta\langle \Pi_1 + \Pi_2\rangle$

$$\geq \Delta(1-\cos\theta) = 2\Delta\sin^2\theta/2 .$$

Now we need to estimate the angle between the null spaces of $H_1 = H_{in} + H_{out}$ and $H_2 = H_{prop}$. That is, we want to find

$$\cos^2\theta = \max|\langle\eta_1|\eta_2\rangle|^2$$

$$= \max\langle\eta_2|\Pi_1|\eta_2\rangle$$

where we maximize over η_1 in the null space of H_1 and η_2 in the null space of H_2 (Π_1 is the projector onto null space of H_1).

A vector in the null space of $H_2 = H_{prop}$ is a valid history state $|\eta_2\rangle = \frac{1}{\sqrt{T+1}}\sum_{t=0}^{T} |\psi(t)\rangle \otimes |t\rangle$. The projector onto null space of H_1 acts trivially on states with $t \in \{1, 2, \dots, T-1\}$, so, after transforming to rotating frame

$$\langle\eta_2'|\Pi_1|\eta_2'\rangle = \frac{T-1}{T+1} + \frac{1}{T+1}\langle\tilde{\eta}_2|\Pi_{in} + \Pi_{out}'|\tilde{\eta}_2\rangle$$

where $\tilde{\eta}_2$ is a state of the nonclock variables, Π_{in} projects onto valid input state, Π_{out}' projects onto

$V_T^\dagger\left(|1\rangle^{out}\otimes \text{anything}^{else}\right)$

because we have transformed to the rotating frame

We know that $\langle\tilde{\eta}_2|\Pi_{in}+\Pi'_{out}|\tilde{\eta}_2\rangle \le (1+\cos\phi)$

where ϕ is angle between spaces Π_{in} and Π'_{out} project onto, as we showed above. This angle ϕ is given by

$\cos^2\phi = \varepsilon$ where ε is the max. acceptance prob. That is, if the input is valid (in the support of Π_{in}) and the history is valid, then the prob of $|1\rangle^{out}$ is at most ε.

We have therefore shown that, if θ is angle between H_1 and H_2 eigenspaces, then

$$\cos^2\theta \le \frac{T-1}{T+1}+\frac{1}{T+1}\left(1+\sqrt{\varepsilon}\right)$$

$$= 1-\frac{1-\sqrt{\varepsilon}}{T+1} \qquad \text{if } \varepsilon \text{ is the max acceptance prob.}$$

Now $\Delta = 2\sin^2\left(\frac{\pi}{2(T+1)}\right)$ is a lower bound for the gap of H_1 and H_2

and $\langle H_1+H_2\rangle \geq 2\Delta\sin^2\frac{\theta}{2}$

where $\sin^2\theta = 1-\cos^2\theta \geq \dfrac{1-\sqrt{\varepsilon}}{T+1}$

and using $\sin^2\frac{\theta}{2} = \dfrac{\sin^2\theta}{4\cos^2\theta/2} \geq \frac{1}{4}\sin^2\theta$, we have

$$E_0 \geq 4\sin^2\left(\frac{\pi}{2(T+1)}\right)\times\frac{1}{4}\,\frac{1-\sqrt{\varepsilon}}{T+1} = \sin^2\frac{\pi}{2(T+1)}\times\frac{1-\sqrt{\varepsilon}}{T+1}$$

$$\geq \text{const}\times\left(1-\sqrt{\varepsilon}\right)\times 1/(T+1)^3 = \left(1-\sqrt{\varepsilon}\right)\frac{1}{\text{poly}(n)}$$

To summarize, we have shown that if acceptance prob is $\geq 1-\varepsilon$, then $E_0 \leq \frac{\varepsilon}{T+1}$

and if acceptance prob is $\leq \varepsilon$, then

$$E_0 \geq \text{const} \times (1-\sqrt{\varepsilon}) \frac{1}{(T+1)^3}.$$

With suitable amplification to make ε small (compared to $1/(T+1)^3$) in the case where the answer is YES, we have reduced the QMA problem to an instance of the Hamiltonian problem.

But -- we still need to see that the Hamiltonian can be made local. We'd like to encode the clock register using qubits. One way is to use a "unary" encoding with T qubits, where

$$|t=0\rangle = |000\cdots0\rangle$$
$$|t=1\rangle = |100\cdots0\rangle$$
$$|t=2\rangle = |1100\cdots0\rangle$$
etc

We can add a term to the Hamiltonian that imposes an energy penalty on the clock if its state is not validly encoded.

The encoding is valid if a 0 is never followed by a 1. Therefore we choose

$$H_{clock} = \sum_{t=1}^{T-1} \left(|01\rangle\langle 01|\right)_{t,t+1}$$

For valid encodings, the projection onto $|t\rangle$ only needs to act on qubits numbered t and $t+1$ (for $t \in \{1,2,\ldots,T-1\}$ and only qubit t for $t \in \{0,T\}$:

$$|t\rangle\langle t| = \left(|10\rangle\langle 10|\right)_{t,t+1}$$

The terms that advance or retard the time act on three qubits:

$$\left.\begin{array}{l} |t\rangle\langle t-1| = \left(|110\rangle\langle 100|\right)_{t-1,t,t+1} \\ |t-1\rangle\langle t| = \left(|100\rangle\langle 110|\right)_{t-1,t,t+1} \end{array}\right\} \text{which are adjoint to one another}$$

These terms act on three qubits, so that $U_t \otimes |t\rangle\langle t-1|$ acts on 5 qubits if U_t is a 2-qubit gate. Thus with this clock encoding, the Hamiltonian $H = H_{in} + H_{out} + H_{prop} + H_{clock}$ is 5-local.

This completes the demonstration that any QMA problem is reducible to the problem of estimating the ground state energy (with a $1/\text{poly}(n)$ promise gap) of a 5-local Hamiltonian.

Thus we have shown that the 5-local Hamiltonian problem is a "natural" QMA-complete problem, much as 3SAT is a natural NP-complete problem. But while in the classical case many "practical" problems have been shown to be NP-complete, the family of problems that have been shown to be QMA-complete is still rather small, and the problems seem relatively "artificial".

In any case, its interesting to see that quantum local Hamiltonian problems seem to be harder than classical ones (if QMA $\neq$ NP).

I won't discuss the tricks for reducing the QMA-complete problem to $k=2$ (which involves clever use of perturbation theory) or for making H geometrically local (which involves encoding the clock more cleverly, among other things).

Another interesting direction to pursue using these ideas is to show that any problem in BQP can be solved using adiabatic quantum computing. The idea is to replace

$$H_{prop} \longrightarrow H_{prop}(s) = (1-s)\, H_{\substack{clock\\init}} + s\, H_{prop}$$

where the null space of $H_{\substack{clock\\init}}$ fixes the clock at $|t=0\rangle$ and s varies in $[0,1]$. Then the ground state of $H(s=0)$ is easy to construct, and the ground state of $H(s=1)$ is the valid history state. The eigenvalue gap of $H(s)$ stays $\geq \frac{1}{poly(n)}$ for $s \in [0,1]$, so the history state can be prepared in polynomial time by adiabatically varying s.
Once we have prepared the history state, we can measure the clock, projecting out $|t=T\rangle$ with probability $1/(T+1) \geq 1/poly(n)$. And once we have $|\psi(T)\rangle$ we can measure the output qubit to find out if the circuit accepts.