

Lecture 6: Hamiltonian simulation and the ground-state problem

Computational aspects of quantum simulation
ENS de Lyon

2026

Abstract

We first solve the MPS exercise from Lecture 5. We then introduce three computational tasks associated with a local Hamiltonian: ground-energy estimation, ground-state preparation, and time simulation. Classical constraint satisfaction gives a first concrete example of the ground-state tasks, while also showing why diagonal quantum evolution should not be confused with classical dynamics. We construct a quantum algorithm for time simulation using the first-order Lie–Trotter product formula, with an explicit error bound and a transverse-field Ising example. Applying the MPS circuit algorithm to the same circuit leads to the classical TEBD method. The lecture ends with the variational characterization of the ground energy.

Learning objectives. After this lecture, students should be able to:

- construct a local circuit approximating e^{-itH} , and bound both its product-formula error and its compilation error;
- recognize TEBD as MPS simulation of a product-formula circuit;
- encode a classical constraint problem as a diagonal local Hamiltonian;
- distinguish ground-energy estimation, ground-state preparation, and time simulation;
- use the variational principle to obtain an upper bound on the ground energy from a trial state.

Main message. Local Hamiltonians place classical optimization and quantum dynamics in a common language. Diagonal Hamiltonians encode classical cost functions, although their unitary evolution is not classical dynamics. General local time evolution can be approximated by a circuit of local gates through a product formula. Ground-state algorithms will be developed in Lecture 7.

Contents

1	Exercise from Lecture 5: an MPS for the W state	2
1.1	Statement	2
1.2	Solution	2
2	Local Hamiltonians and three computational tasks	3
2.1	Three different tasks	4
2.2	A classical example: constraint satisfaction	4
2.3	Hamiltonian time evolution	5
2.4	The simulation task and its output	6

3 Local Hamiltonian time evolution on a quantum computer	6
3.1 The commuting case	6
3.2 First-order product formula	7
3.3 The algorithm and its cost	8
3.4 Worked example: the transverse-field Ising chain	9
4 Ground states	10
5 Assigned exercise: a classical constraint Hamiltonian	10
References	11

1 Exercise from Lecture 5: an MPS for the W state

1.1 Statement

For $n \geq 2$, let

$$|W_n\rangle = \frac{1}{\sqrt{n}} \sum_{j=1}^n |0^{j-1} 1 0^{n-j}\rangle.$$

- (a) Determine the smallest possible MPS bond dimension at each nontrivial consecutive cut, using the Schmidt ranks.
- (b) Put

$$B^0 = I_2, \quad B^1 = \frac{1}{\sqrt{n}} \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \quad \ell^T = (1 \ 0), \quad r = (0 \ 1)^T.$$

Verify the amplitudes $\ell^T B^{i_1} \cdots B^{i_n} r$, and convert them into open-boundary MPS tensors with specified sizes.

- (c) Contract this MPS with $|1\rangle\langle 1|$ at site j to obtain $\Pr[\text{qubit } j \text{ is } 1] = 1/n$. Continue through site n .

1.2 Solution

At the cut after site k , where $1 \leq k < n$,

$$|W_n\rangle = \sqrt{\frac{k}{n}} |W_k\rangle |0^{n-k}\rangle + \sqrt{\frac{n-k}{n}} |0^k\rangle |W_{n-k}\rangle.$$

Here $|W_1\rangle = |1\rangle$. The two vectors on each side are orthonormal, and both coefficients are nonzero. This is a Schmidt decomposition of rank two. By the MPS characterization from Lecture 5, the smallest possible bond dimension is therefore two at every internal bond.


Write $N = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$, so that $N^2 = 0$. If the string contains no 1, the amplitude is $\ell^T r = 0$. If it contains exactly one 1, the product is $N/\sqrt{n}$, and the amplitude is $1/\sqrt{n}$. Two or more 1's give zero because $N^2 = 0$. Thus the tensors are

$$A^{[1],i} = \ell^T B^i \in \mathbb{C}^{1 \times 2}, \quad A^{[k],i} = B^i \in \mathbb{C}^{2 \times 2} \quad (2 \leq k < n), \quad A^{[n],i} = B^i r \in \mathbb{C}^{2 \times 1}.$$

For part (c), we can contract directly using these matrices. The probability of obtaining 1 at site j is

$$p_j = \langle W_n | (|1\rangle\langle 1|)_j | W_n \rangle = \sum_{\substack{i_1, \dots, i_n \in \{0,1\} \\ i_j = 1}} \left| \ell^T B^{i_1} \dots B^{i_n} r \right|^2.$$

Indeed, the projector fixes the j -th bit to 1, while the identity operators at all other sites sum over their possible values. Graphically, this is the following contraction (shown for $1 < j < n$):



Each B node has entries $(B^i)_{\alpha,\beta}$: horizontal edges sum the matrix indices, and vertical edges sum physical indices. The upper row gives an amplitude, the lower row its complex conjugate. The omitted sites have the same vertical connections as sites 1 and n . At site j , the projector forces the physical index in both rows to be 1. All tensors here are real, so complex conjugation changes nothing.

Now evaluate the sum. Every term already contains $B^1 = N/\sqrt{n}$ at site j . If any other bit is 1, the matrix product contains two copies of N , separated only by identity matrices, and is zero because $N^2 = 0$. Thus the *only surviving assignment* is $0^{j-1}10^{n-j}$. Contracting its upper row all the way to the right boundary gives

$$\underbrace{\ell^T B^0 \dots B^0}_{j-1 \text{ sites}} B^1 \underbrace{B^0 \dots B^0 r}_{n-j \text{ sites}} = (1 \ 0) \frac{1}{\sqrt{n}} \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \frac{1}{\sqrt{n}}.$$

The lower row contributes the same factor. Consequently,

$$p_j = \frac{1}{\sqrt{n}} \frac{1}{\sqrt{n}} = \frac{1}{n}.$$

For $j = 1$ or $j = n$, the corresponding empty product is the identity, and the calculation is unchanged. This also explains why the contraction must include every site and the right boundary: the desired scalar is the *complete* matrix product, multiplied by its complex conjugate.

2 Local Hamiltonians and three computational tasks

Recall that a Hamiltonian is a Hermitian operator. A k -local Hamiltonian on n qubits is specified as

$$H = \sum_{a=1}^L h_a,$$

where each h_a acts nontrivially on at most k qubits. We assume $k = O(1)$ and $L = \text{poly}(n)$. Thus the input is a polynomial list of small matrices and their supports, rather than one exponentially large matrix. Identities outside the support of a local term are implicit.

2.1 Three different tasks

Let $E_0 = \lambda_{\min}(H)$ be the ground energy, and let Π_0 be the orthogonal projector onto the corresponding ground space. We will distinguish three computational tasks:

1. *Ground-energy estimation:* output a number approximating E_0 .
2. *Ground-state preparation:* produce a state $|\psi\rangle$ for which $\langle\psi|\Pi_0|\psi\rangle$ is close to one.
3. *Time simulation:* implement the unitary e^{-itH} , which will be defined in Section 2.3.

These tasks use the same local description of H , but ask for different outputs. We first illustrate the ground-state tasks with an entirely classical optimization problem written in Hamiltonian language.

2.2 A classical example: constraint satisfaction

Consider Boolean variables $x_1, \dots, x_n$ and constraints $C_1, \dots, C_L$. Constraint C_a depends on a set S_a of at most k variables. Let

$$v_a(y) = \begin{cases} 0, & \text{if the local assignment } y \text{ satisfies } C_a, \\ 1, & \text{if it violates } C_a. \end{cases}$$

Associate the computational-basis state $|x\rangle$ with the classical assignment x , and define the local penalty

$$h_a = \sum_{y \in \{0,1\}^{S_a}} v_a(y) |y\rangle\langle y|_{S_a} \otimes I_{\overline{S_a}}, \quad H_C = \sum_{a=1}^L h_a.$$

The term h_a has energy zero on assignments satisfying C_a and energy one on assignments violating it. It is a diagonal projector acting on at most k qubits, so H_C is a diagonal k -local Hamiltonian.

Proposition 2.1 (Classical costs are ground energies). *For every $x \in \{0,1\}^n$,*

$$H_C|x\rangle = \left(\sum_a v_a(x_{S_a}) \right) |x\rangle.$$

Consequently,

$$E_0(H_C) = \min_{x \in \{0,1\}^n} \sum_a v_a(x_{S_a}),$$

the minimum number of violated constraints. In particular, the instance is satisfiable if and only if $E_0(H_C) = 0$.

Proof. Every local term acts diagonally as $h_a|x\rangle = v_a(x_{S_a})|x\rangle$. Summing over a proves the first identity. Since the computational-basis states form an eigenbasis, the smallest eigenvalue is the smallest classical cost. Equivalently, for any normalized superposition $|\psi\rangle = \sum_x \alpha_x |x\rangle$,

$$\langle\psi|H_C|\psi\rangle = \sum_x |\alpha_x|^2 \sum_a v_a(x_{S_a})$$

is a weighted average of classical costs and cannot be smaller than their minimum. $\square$

For this example, ground-energy estimation is precisely the classical optimization problem of minimizing the number of violated constraints. Ground-state preparation means finding an optimal assignment x^* and preparing $|x^*\rangle$; allowing a quantum superposition cannot improve the optimal energy.

One clause of a 3-SAT instance. The clause $x_1 \vee \neg x_2 \vee x_3$ is violated only by 010. Its penalty is

$$h_{\text{clause}} = |010\rangle\langle 010| = \frac{1}{8}(I + Z_1)(I - Z_2)(I + Z_3).$$

Adding one such term for every clause turns the number of violated clauses into the energy. Thus a satisfiable formula has ground energy zero, whereas an unsatisfiable formula has ground energy at least one.

2.3 Hamiltonian time evolution

In units where $\hbar = 1$, the Schrödinger equation for a time-independent Hamiltonian is

$$i \frac{d}{dt} |\psi(t)\rangle = H |\psi(t)\rangle.$$

If

$$H = \sum_j E_j |E_j\rangle\langle E_j|$$

is a spectral decomposition, define

$$e^{-itH} = \sum_j e^{-itE_j} |E_j\rangle\langle E_j|.$$

This operator is unitary, and the solution of the Schrödinger equation is

$$|\psi(t)\rangle = e^{-itH} |\psi(0)\rangle.$$

Thus H determines both the energies of its eigenvectors and their time evolution.

This spectral definition gives $e^{-i(s+t)H} = e^{-isH} e^{-itH}$ and $\frac{d}{dt} e^{-itH} = -iH e^{-itH}$. If a Hermitian Pauli P satisfies $P^2 = I$, then

$$e^{-itP} = \cos(t)I - i \sin(t)P.$$

For example, if $H = X$ and the initial state is $|0\rangle$, then

$$|\psi(t)\rangle = \cos(t)|0\rangle - i \sin(t)|1\rangle,$$

so measuring in the computational basis gives outcome 1 with probability $\sin^2(t)$.

Time simulation does not by itself prepare a ground state. Indeed, if $|\phi\rangle = \sum_j c_j |E_j\rangle$, then

$$e^{-itH} |\phi\rangle = \sum_j c_j e^{-itE_j} |E_j\rangle.$$

The phases change, but every weight $|c_j|^2$ remains fixed. In particular, evolution under H preserves the total weight in its ground space.

Remark 2.2 (Returning to the classical example). *For the constraint Hamiltonian,*

$$e^{-itH_C} |x\rangle = e^{-it \sum_a v_a(x_{S_a})} |x\rangle.$$

Thus a basis state remains the same classical assignment, up to a phase. The diagonal Hamiltonian embeds the classical energy function, but its unitary evolution is not a classical update rule or stochastic dynamics.

2.4 The simulation task and its output

Definition 2.3 (Hamiltonian time simulation). *Given a local description of H , a real time t , and a tolerance $\varepsilon > 0$, construct a quantum circuit implementing a unitary $\tilde{U}$ such that*

$$\|\tilde{U} - e^{-itH}\| \leq \varepsilon.$$

The norm is the operator norm, $\|A\| = \max_{\|\psi\|=1} \|A|\psi\rangle\|$. Global phases are irrelevant here: one may choose a representative of the implemented unitary when applying this norm criterion.

This notion of approximation directly implies approximation of expectation values. Indeed, for every normalized state $|\psi\rangle$ and every Hermitian observable O ,

$$\left| \langle \psi | \tilde{U}^\dagger O \tilde{U} | \psi \rangle - \langle \psi | e^{itH} O e^{-itH} | \psi \rangle \right| \leq 2\varepsilon \|O\|.$$

To see this, add and subtract $\langle \psi | e^{itH} O \tilde{U} | \psi \rangle$, and use the triangle inequality, $\|\tilde{U} - e^{-itH}\| \leq \varepsilon$, and the fact that all the unitaries have norm one. Thus approximating the evolution in operator norm is sufficient for approximating the expectation values of observables after the evolution.

3 Local Hamiltonian time evolution on a quantum computer

Hamiltonian time simulation asks for a circuit approximating e^{-itH} . It is one of the tasks most directly matched to a quantum computer: the desired output is itself a unitary evolution on the same n -qubit state space. This does not by itself prove an exponential advantage for every observable or every input, but it avoids writing down the 2^n amplitudes of the evolving state. The product-formula approach below is the elementary construction underlying the quantum simulation result of Lloyd [1].

3.1 The commuting case

If the local terms commute, then

$$[h_a, h_b] = 0 \quad \text{for all } a, b \quad \implies \quad e^{-it \sum_a h_a} = \prod_a e^{-it h_a}.$$

Each factor is a constant-width gate when h_a acts on constantly many qubits. Thus the evolution has already been decomposed exactly into a polynomial number of local unitaries. Relative to any fixed universal gate set, each constant-width unitary can then be compiled to the required accuracy. Throughout, a product of operators acts from right to left: in $e^{-ith_1} \dots e^{-ith_L}$, the circuit applies the h_L factor first.

For the Ising interaction with $g = 0$, all $Z_j Z_{j+1}$ terms commute:

$$H_{\text{Ising}} = -J \sum_{j=1}^{n-1} Z_j Z_{j+1}, \quad e^{-itH_{\text{Ising}}} = \prod_{j=1}^{n-1} e^{iJt Z_j Z_{j+1}}.$$

This is a product of $n - 1$ nearest-neighbor two-qubit unitaries. If the circuit model uses a specified universal gate set, these two-qubit unitaries can be decomposed further into that gate set.

3.2 First-order product formula

For noncommuting terms the product is no longer exact, but it becomes accurate over a short time step.

Proposition 3.1 (One-step product-formula error). *Let $H = \sum_{a=1}^L h_a$, with every h_a Hermitian. For every real δ ,*

$$\left\| e^{-i\delta H} - e^{-i\delta h_1} e^{-i\delta h_2} \dots e^{-i\delta h_L} \right\| \leq \frac{\delta^2}{2} \sum_{1 \leq a < b \leq L} \| [h_a, h_b] \|.$$

Here $\| \cdot \|$ is the operator norm.

Proof. We first prove the estimate for two Hermitian operators A and B . It is enough to take $\delta \geq 0$: for negative δ , apply the positive-time result to $-A, -B$, whose commutator is still $[A, B]$. Define

$$F(s) = e^{-i(\delta-s)(A+B)} e^{-isA} e^{-isB}, \quad 0 \leq s \leq \delta.$$

Its endpoints are

$$F(0) = e^{-i\delta(A+B)}, \quad F(\delta) = e^{-i\delta A} e^{-i\delta B}.$$

Differentiating the three factors and cancelling the two terms containing A gives

$$F'(s) = i e^{-i(\delta-s)(A+B)} [B, e^{-isA}] e^{-isB}.$$

The remaining commutator is itself an integral. Indeed, differentiating $e^{-iuA} B e^{-i(s-u)A}$ with respect to u and integrating from 0 to s gives

$$[B, e^{-isA}] = i \int_0^s e^{-iuA} [A, B] e^{-i(s-u)A} du.$$

All exponentials in these expressions are unitary. Therefore

$$\|F'(s)\| \leq \int_0^s \| [A, B] \| du = s \| [A, B] \|.$$

The fundamental theorem of calculus now gives

$$\begin{aligned} \left\| e^{-i\delta(A+B)} - e^{-i\delta A} e^{-i\delta B} \right\| &= \|F(0) - F(\delta)\| \\ &\leq \int_0^\delta \|F'(s)\| ds \\ &\leq \frac{\delta^2}{2} \| [A, B] \|. \end{aligned}$$

We now add the terms of H one at a time. Put

$$H_b = \sum_{a=1}^b h_a, \quad P_b = e^{-i\delta h_1} \dots e^{-i\delta h_b}.$$

Using the two-operator estimate with $A = H_{b-1}$ and $B = h_b$, followed by the triangle inequality, we obtain

$$\begin{aligned} \left\| e^{-i\delta H_b} - P_b \right\| &\leq \frac{\delta^2}{2} \| [H_{b-1}, h_b] \| + \left\| e^{-i\delta H_{b-1}} - P_{b-1} \right\| \\ &\leq \frac{\delta^2}{2} \sum_{a < b} \| [h_a, h_b] \| + \left\| e^{-i\delta H_{b-1}} - P_{b-1} \right\|. \end{aligned}$$

Iterating this inequality from $b = 2$ to L gives the claimed sum over all pairs. $\square$

Set $\delta = t/r$ and repeat the short-time approximation r times:

$$V_r(t) = \left(e^{-i\delta h_1/r} e^{-i\delta h_2/r} \dots e^{-i\delta h_L/r} \right)^r.$$

For unitaries U, V , the telescoping identity

$$U^r - V^r = \sum_{q=0}^{r-1} U^{r-1-q} (U - V) V^q$$

implies $\|U^r - V^r\| \leq r\|U - V\|$. Taking

$$U = e^{-i(t/r)H}, \quad V = e^{-i\delta h_1/r} \dots e^{-i\delta h_L/r},$$

therefore gives

$$\|e^{-itH} - V_r(t)\| \leq \frac{t^2}{2r} \sum_{a < b} \|[h_a, h_b]\|.$$

We record the resulting bound as a theorem.

Theorem 3.2 (First-order Lie–Trotter formula). *For Hermitian $h_1, \dots, h_L$, real t , and integer $r \geq 1$,*

$$\left\| e^{-it \sum_{a=1}^L h_a} - \left(e^{-i\delta h_1/r} \dots e^{-i\delta h_L/r} \right)^r \right\| \leq \frac{t^2}{2r} \sum_{a < b} \|[h_a, h_b]\|.$$

In particular, this error tends to zero as r tends to infinity.

3.3 The algorithm and its cost

Product-formula simulation algorithm. Given $H = \sum_a h_a$, t , and $0 < \varepsilon < 1$:

1. Obtain an upper bound $C \geq \sum_{a < b} \|[h_a, h_b]\|$. If $\|h_a\| \leq 1$, one may always use $C = L(L-1)$.
2. Set $r = \max\{1, \lceil t^2 C / \varepsilon \rceil\}$ and $\delta = t/r$.
3. For each of the r steps, apply the local gates $e^{-i\delta h_L}, \dots, e^{-i\delta h_1}$, in that order. If using a finite gate set, approximate each local gate to error at most $\varepsilon/(2rL)$.

There are rL local gates before compilation. The two error budgets are $\varepsilon/2$ for the product formula and $\varepsilon/2$ for compilation.

Theorem 3.3 (Efficient quantum simulation of local Hamiltonians). *Let $H = \sum_{a=1}^L h_a$ be a k -local Hamiltonian on n qubits, where $k = O(1)$, $L = \text{poly}(n)$, $\|h_a\| \leq 1$, and the local matrices are efficiently computable to the requested precision. For an efficiently specified time t with $|t| \leq \text{poly}(n)$ and a tolerance $0 < \varepsilon < 1$ with $1/\varepsilon \leq \text{poly}(n)$, a classical polynomial-time algorithm constructs a polynomial-size quantum circuit $\tilde{U}$ such that, up to an irrelevant global phase,*

$$\|\tilde{U} - e^{-itH}\| \leq \varepsilon.$$

Proof. Let

$$C = \sum_{a < b} \|[h_a, h_b]\|.$$

Since $\|[h_a, h_b]\| \leq 2 \|h_a\| \|h_b\| \leq 2$, we have $C = O(L^2) = \text{poly}(n)$. Choose

$$r \geq \max \left\{ 1, \frac{t^2 C}{\varepsilon} \right\}.$$

The product-formula error is then at most $\varepsilon/2$. The circuit contains rL local exponentials $e^{-ith_a/r}$. Each acts on at most $k = O(1)$ qubits. We use a standard efficiently compilable universal gate set, such as Clifford plus T : a constant-width unitary can be approximated, up to global phase, to error η with $\text{poly}(\log(1/\eta))$ gates and classical compilation time polynomial in the requested precision. Compile every factor to error at most $\varepsilon/(2rL)$. A second telescoping argument shows that the total compilation error is at most $\varepsilon/2$. Because rL and the required number of bits of precision are polynomial, the resulting uniform circuit has polynomial size. Combining the two errors proves the claim. $\square$

The theorem is often summarized informally by saying that local Hamiltonian simulation is in BQP. More precisely, Hamiltonian simulation is a unitary implementation task rather than a decision problem: the theorem constructs a uniform polynomial-size quantum circuit that can be used as a subroutine in BQP algorithms.

3.4 Worked example: the transverse-field Ising chain

A Trotter step for the transverse-field Ising chain. Write

$$H_Z = -J \sum_{j=1}^{n-1} Z_j Z_{j+1}, \quad H_X = -g \sum_{j=1}^n X_j.$$

The terms internal to H_Z commute, as do those internal to H_X . For $\delta = t/r$, one product-formula step is

$$e^{-i\delta H_Z} e^{-i\delta H_X} = \left(\prod_{j=1}^{n-1} e^{iJ\delta Z_j Z_{j+1}} \right) \left(\prod_{j=1}^n e^{ig\delta X_j} \right).$$

In chronological order, apply all one-qubit X -rotations first, and then the ZZ gates. The latter split into two parallel layers: the odd bonds $(1, 2), (3, 4), \dots$, and the even bonds $(2, 3), (4, 5), \dots$. Each step therefore uses $2n - 1$ local gates and at most three layers. Repeating the step r times approximates the complete evolution.

To bound the error, use the two-term formula for $H_Z + H_X$. A bond $Z_j Z_{j+1}$ commutes with X_k unless $k = j$ or $k = j + 1$. In either exceptional case the commutator has norm two: for example, $[Z_j Z_{j+1}, X_j] = 2iY_j Z_{j+1}$. Thus

$$\|[H_Z, H_X]\| \leq 4|Jg|(n-1),$$

and hence

$$\left\| e^{-it(H_Z+H_X)} - \left(e^{-itH_Z/r} e^{-itH_X/r} \right)^r \right\| \leq \frac{2|Jg|(n-1)t^2}{r}.$$

For product-formula error at most ε , it suffices to take $r \geq \max\{1, \lceil 2|Jg|(n-1)t^2/\varepsilon \rceil\}$. For bounded J, g , the number of local gates is $O(n + n^2 t^2/\varepsilon)$. Compilation into a finite gate set adds its own accuracy budget, as above. If $J = 0$ or $g = 0$, one step is exact.

Remark 3.4 (MPS simulation and TEBD). *Apply the MPS update algorithm of Lecture 5 to the same one- and two-qubit gates, optionally truncating the Schmidt decomposition after each two-qubit update. For a one-dimensional nearest-neighbor Hamiltonian, this is the basic time-evolving block decimation (TEBD) method [2]. “Time-evolving” refers to applying the local evolution gates; “block decimation” refers to keeping a smaller Schmidt subspace for the two blocks separated by a bond. It is a standard practical method for classical time simulation using the MPS gate-update rule derived in Lecture 5.*

Its accuracy depends on both the product-formula error and the discarded Schmidt coefficients. Its efficiency depends on the bond dimensions needed to control that truncation error. The quantum circuit has no corresponding Schmidt-rank restriction. We do not repeat the MPS update derivation here.

4 Ground states

We now focus on ground energy.

Proposition 4.1 (Variational principle). *For every Hermitian H ,*

$$E_0 = \min_{\|\psi\|=1} \langle \psi | H | \psi \rangle.$$

Equality holds for a normalized state precisely when it belongs to the ground space.

Proof. In an orthonormal eigenbasis of H ,

$$\langle \psi | H | \psi \rangle - E_0 = \sum_j |c_j|^2 (E_j - E_0) \geq 0.$$

Every summand is nonnegative, so equality means that all coefficients with $E_j > E_0$ vanish. Conversely, any ground state attains equality. $\square$

Every normalized trial state $|\phi\rangle$ therefore gives the rigorous upper bound

$$E_0 \leq \langle \phi | H | \phi \rangle.$$

This simple observation is the starting point for the variational methods of Lecture 7. There we will also ask how to obtain lower bounds on E_0 , which must be valid simultaneously for every normalized state.

5 Assigned exercise: a classical constraint Hamiltonian

Let $G = (V, E)$ be a finite undirected graph, with one Boolean variable x_i and one qubit at each vertex. Each edge $\{i, j\}$ imposes the constraint $x_i \neq x_j$. Define

$$h_{ij} = \frac{I + Z_i Z_j}{2}, \quad H_G = \sum_{\{i,j\} \in E} h_{ij}.$$

Show that h_{ij} is a projector and that its value on $|x\rangle$ is one exactly when the constraint is violated. Prove that

$$E_0(H_G) = |E| - \text{MaxCut}(G),$$

where $\text{MaxCut}(G)$ is the largest number of edges crossing between a subset of vertices and its complement. Explain explicitly why a quantum superposition cannot have energy below the best classical assignment.

References

- [1] Seth Lloyd, *Universal Quantum Simulators*, Science 273, 1073–1078 (1996).
<https://doi.org/10.1126/science.273.5278.1073>
- [2] Guifré Vidal, *Efficient Simulation of One-Dimensional Quantum Many-Body Systems*, Physical Review Letters 93, 040502 (2004).
<https://arxiv.org/abs/quant-ph/0310089>