

Lecture 5: Matrix product states and quantum Hamiltonians

Computational aspects of quantum simulation
ENS de Lyon

2026

Abstract

We complete our study of matrix product states (MPS): minimal bond dimensions, contractions, local gate updates, and exact classical simulation when all intermediate Schmidt ranks remain small. We then introduce quantum Hamiltonians, ground energy, locality, and the transverse-field Ising chain. Classical constraint Hamiltonians and Hamiltonian time simulation will be developed in Lecture 6.

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

- prove that the minimal bond dimensions equal the Schmidt ranks;
- construct simple MPSs and contract them to compute norms, expectations, and computational-basis probabilities;
- update an MPS under one-site and neighboring two-site unitary gates;
- state and prove the exact MPS simulation principle;
- define local Hamiltonians, ground energy, and ground states;
- interpret the terms in the transverse-field Ising chain.

Main message. An MPS gives an efficient classical representation only while the Schmidt ranks across the chain remain small. Quantum Hamiltonians provide a different concise description: they specify the energy of a many-body system as a sum of local interactions. Standard MPS algorithms are discussed in [1].

Contents

1	Exercise from Lecture 4 and recap	2
2	Matrix product states	3
2.1	Two explicit MPS examples	6
3	Classical simulation of low-entanglement circuits	6
3.1	Full proof of the circuit update	7
3.2	What can be computed from an MPS representation	9
4	Quantum Hamiltonians: energy and locality	9
4.1	Hamiltonians and locality	9
5	Optional training exercises	10

6 Assigned exercise: an explicit MPS for the W state	10
References	12

1 Exercise from Lecture 4 and recap

Problem

For $n \geq 2$, define

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

and set $|W_1\rangle = |1\rangle$. This exercise uses only Schmidt decomposition and optimal truncation; no result about MPS is needed.

- (a) For $1 \leq k < n$, prove the decomposition

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

and show that it is a Schmidt decomposition.

- (b) Determine the Schmidt rank and the two Schmidt coefficients across each nontrivial consecutive cut.
- (c) For a fixed cut k , give a best *unnormalized* approximation to $|W_n\rangle$ of Schmidt rank at most one, and compute the minimum squared Euclidean error. Justify optimality using Proposition 4.2 from Lecture 4. What happens when n is even and $k = n/2$?

Solution

Split the defining sum according to whether the unique 1 is among the first k or the last $n - k$ sites. This gives

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

Each $|W_j\rangle$ is normalized and orthogonal to $|0^j\rangle$, since its computational-basis strings have exactly one 1. The two left vectors are therefore orthonormal, as are the two right vectors. Both coefficients are positive for $1 \leq k < n$, so this is a Schmidt decomposition of rank two, with coefficients $\sqrt{k/n}$ and $\sqrt{(n-k)/n}$.

By the optimal truncation proposition of Lecture 4, keep a term with the larger coefficient. One optimal unnormalized approximation is

$$|\widetilde{W}_{n,k}\rangle = \begin{cases} \sqrt{k/n} |W_k\rangle |0^{n-k}\rangle, & k \geq n-k, \\ \sqrt{(n-k)/n} |0^k\rangle |W_{n-k}\rangle, & k < n-k. \end{cases}$$

Its minimum squared error is the discarded squared coefficient,

$$\min_{\text{Schmidt rank}(|\varphi\rangle) \leq 1} \| |W_n\rangle - |\varphi\rangle \|^2 = \frac{\min(k, n-k)}{n}.$$

When n is even and $k = n/2$, either displayed Schmidt term can be kept, giving squared error $1/2$; the optimum is not unique. The retained vector is intentionally not normalized.

The linear-algebra tool we will use

For a state $|\psi\rangle = \sum_{a,b} M_{ab}|a\rangle|b\rangle$, the Schmidt rank is $\text{rank}(M)$. A compact SVD $M = U\Sigma V^\dagger$, keeping exactly the nonzero singular values, gives the Schmidt decomposition. We will repeatedly reshape state or tensor coefficients into a matrix and apply this factorization.

2 Matrix product states

Definition 2.1 (Open-boundary matrix product state). *An open-boundary n -qudit matrix product state (MPS) $|\psi\rangle \in (\mathbb{C}^d)^{\otimes n}$ is a state whose amplitudes have the form*

$$\psi_{i_1 \dots i_n} = A^{[1]i_1} A^{[2]i_2} \dots A^{[n]i_n},$$

or equivalently

$$|\psi\rangle = \sum_{i_1, \dots, i_n=1}^d A^{[1]i_1} A^{[2]i_2} \dots A^{[n]i_n} |i_1 \dots i_n\rangle.$$

Here $A^{[k]i_k}$ is a $\chi_{k-1} \times \chi_k$ matrix, with $\chi_0 = \chi_n = 1$. The integers χ_k are the bond dimensions.

Remark 2.2 (MPS as a tensor network). *A matrix product state is a special case of a tensor network: nodes represent tensors, and links joining nodes represent contractions, that is, summation over a shared index. An open link represents an index that is not summed over. For an MPS, the nodes form a chain. The tensor at site k has entries $A_{\alpha_{k-1}, \alpha_k}^{[k]i_k}$; its bond indices connect neighboring sites, while its physical index i_k remains open.*



Horizontal links: contracted bond indices $\alpha_k \in \{1, \dots, \chi_k\}$.

Upward open links: physical indices $i_k \in \{1, \dots, d\}$.

The matrix product in the definition is precisely the contraction of all bond indices:

$$\psi_{i_1 \dots i_n} = \sum_{\alpha_1=1}^{\chi_1} \dots \sum_{\alpha_{n-1}=1}^{\chi_{n-1}} A_{1, \alpha_1}^{[1]i_1} A_{\alpha_1, \alpha_2}^{[2]i_2} \dots A_{\alpha_{n-1}, 1}^{[n]i_n}.$$

The open physical indices label the amplitudes of the resulting state.

If every $\chi_k \leq \chi$, the tensors contain at most

$$\sum_{k=1}^n d \chi_{k-1} \chi_k \leq nd\chi^2$$

complex numbers, rather than the d^n amplitudes of a general state.

Given an n -qudit state $|\psi\rangle$, how small can we make the bond dimensions χ_k ? The answer is determined by the Schmidt ranks.

Theorem 2.3 (Bond dimension equals Schmidt rank). *For a state $|\psi\rangle$ on an ordered chain, let r_k be its Schmidt rank across the cut $1 \dots k \mid k+1 \dots n$. Then every MPS representation of $|\psi\rangle$ with bond dimensions χ_k satisfies $r_k \leq \chi_k$. Moreover, there is an exact MPS representation with $\chi_k = r_k$ simultaneously for every $k = 1, \dots, n-1$.*

Proof. For the first statement, cut an MPS between tensors k and $k + 1$. The contracted left block has one remaining bond index $\alpha_k \in \{1, \dots, \chi_k\}$, and so does the right block. Hence

$$|\psi\rangle = \sum_{\alpha_k=1}^{\chi_k} |L_{\alpha_k}\rangle |R_{\alpha_k}\rangle.$$

The vectors need not be orthogonal. To justify the rank bound, write α for α_k , and choose orthonormal bases $\{|e_a\rangle\}_{a=1}^{d^k}$ on the left and $\{|f_b\rangle\}_{b=1}^{d^{n-k}}$ on the right. Expand each product factor in its respective basis:

$$\begin{aligned} |L_\alpha\rangle &= \sum_{a=1}^{d^k} L_{a\alpha} |e_a\rangle, & L_{a\alpha} &= \langle e_a | L_\alpha \rangle, \\ |R_\alpha\rangle &= \sum_{b=1}^{d^{n-k}} R_{b\alpha} |f_b\rangle, & R_{b\alpha} &= \langle f_b | R_\alpha \rangle. \end{aligned}$$

Let $L = (L_{a\alpha})$ and $R = (R_{b\alpha})$ be the matrices whose α -th columns contain these expansion coefficients. These columns are simply the coordinate representations of $|L_\alpha\rangle$ and $|R_\alpha\rangle$, not new state vectors. Substituting the expansions, the coefficient matrix of $|\psi\rangle$ in the product basis $\{|e_a\rangle|f_b\rangle\}$ is

$$M_{ab} = \sum_{\alpha=1}^{\chi_k} L_{a\alpha} R_{b\alpha}, \quad M = LR^T.$$

The ordinary transpose appears here, rather than a conjugate transpose, because tensor-product amplitudes multiply without complex conjugation. Every column of M is a linear combination of the χ_k columns of L , so $\text{rank}(M) \leq \chi_k$. Since the Schmidt decomposition comes from the SVD of M , its Schmidt rank equals $\text{rank}(M)$, giving $r_k \leq \chi_k$.

For the converse, we construct the tensors from left to right using exact SVDs. At each step the left singular vectors form the next MPS tensor; the singular values and right singular vectors form the residual tensor to be processed at the next step. If $n = 1$, simply take $A^{[1]i_1} = \psi_{i_1}$. Assume below that $n \geq 2$.

First site. Reshape the amplitudes into the $d \times d^{n-1}$ matrix

$$M_{i_1, (i_2, \dots, i_n)}^{[1]} = \psi_{i_1 \dots i_n}.$$

Its rank is r_1 , the Schmidt rank across the first cut. Take a compact SVD $M^{[1]} = U^{[1]} \Sigma^{[1]} V^{[1]\dagger}$, retaining all and only the r_1 nonzero singular values. Thus $U^{[1]}$ has r_1 orthonormal columns. Define

$$A_{1, \alpha_1}^{[1]i_1} = U_{i_1, \alpha_1}^{[1]}, \quad R_{\alpha_1, i_2, \dots, i_n}^{[1]} = (\Sigma^{[1]} V^{[1]\dagger})_{\alpha_1, (i_2, \dots, i_n)}.$$

We then have the exact factorization

$$\psi_{i_1 \dots i_n} = \sum_{\alpha_1=1}^{r_1} A_{1, \alpha_1}^{[1]i_1} R_{\alpha_1, i_2, \dots, i_n}^{[1]}.$$

In particular, the first bond dimension is r_1 .

Next sites and the rank at each step. Suppose that sites $1, \dots, k-1$ have been processed, where $2 \leq k \leq n-1$. Let $W^{[k-1]}$ be the matrix obtained by contracting those tensors, with row index $(i_1, \dots, i_{k-1})$ and column index α_{k-1} . Explicitly,

$$W_{(i_1, \dots, i_{k-1}), \alpha_{k-1}}^{[k-1]} = \left(A^{[1]i_1} \dots A^{[k-1]i_{k-1}} \right)_{1, \alpha_{k-1}}.$$

The current factorization is

$$\psi_{i_1 \dots i_n} = \sum_{\alpha_{k-1}=1}^{r_{k-1}} W_{(i_1, \dots, i_{k-1}), \alpha_{k-1}}^{[k-1]} R_{\alpha_{k-1}, i_k, \dots, i_n}^{[k-1]}. \quad (1)$$

We maintain that $W^{[k-1]}$ has r_{k-1} orthonormal columns, that is, $W^{[k-1]\dagger} W^{[k-1]} = I_{r_{k-1}}$. This holds initially because $W^{[1]} = U^{[1]}$.

Reshape the residual into the $(r_{k-1}d) \times d^{n-k}$ matrix

$$M_{(\alpha_{k-1}, i_k), (i_{k+1}, \dots, i_n)}^{[k]} = R_{\alpha_{k-1}, i_k, \dots, i_n}^{[k-1]}.$$

We claim that $\text{rank}(M^{[k]}) = r_k$. Write $\Psi^{[k]}$ for the coefficient matrix of $|\psi\rangle$ across cut k :

$$\Psi_{(i_1, \dots, i_k), (i_{k+1}, \dots, i_n)}^{[k]} = \psi_{i_1 \dots i_n},$$

Equation (1) gives the first identity below. The second follows by multiplying the first on the left by $W^{[k-1]\dagger} \otimes I_d$ and using $W^{[k-1]\dagger} W^{[k-1]} = I_{r_{k-1}}$:

$$\Psi^{[k]} = (W^{[k-1]} \otimes I_d) M^{[k]}, \quad M^{[k]} = (W^{[k-1]\dagger} \otimes I_d) \Psi^{[k]}.$$

Left multiplication cannot increase matrix rank. Applying this fact to the two identities in turn gives

$$\text{rank}(\Psi^{[k]}) \leq \text{rank}(M^{[k]}) \leq \text{rank}(\Psi^{[k]}).$$

The two ranks are therefore equal. By definition, $\text{rank}(\Psi^{[k]}) = r_k$, proving the claim.

Now take a compact SVD $M^{[k]} = U^{[k]} \Sigma^{[k]} V^{[k]\dagger}$ with exactly r_k nonzero singular values, and set

$$\begin{aligned} A_{\alpha_{k-1}, \alpha_k}^{[k] i_k} &= U_{(\alpha_{k-1}, i_k), \alpha_k}^{[k]}, \\ R_{\alpha_k, i_{k+1}, \dots, i_n}^{[k]} &= (\Sigma^{[k]} V^{[k]\dagger})_{\alpha_k, (i_{k+1}, \dots, i_n)}. \end{aligned}$$

Thus $A^{[k] i_k}$ is an $r_{k-1} \times r_k$ matrix, and the residual is factored exactly as

$$R_{\alpha_{k-1}, i_k, \dots, i_n}^{[k-1]} = \sum_{\alpha_k=1}^{r_k} A_{\alpha_{k-1}, \alpha_k}^{[k] i_k} R_{\alpha_k, i_{k+1}, \dots, i_n}^{[k]}.$$

The new contracted left block is $W^{[k]} = (W^{[k-1]} \otimes I_d) U^{[k]}$. Its columns are again orthonormal: the matrices $W^{[k-1]} \otimes I_d$ and $U^{[k]}$ both have orthonormal columns, and explicitly

$$W^{[k]\dagger} W^{[k]} = U^{[k]\dagger} (W^{[k-1]\dagger} W^{[k-1]} \otimes I_d) U^{[k]} = I_{r_k}.$$

This establishes the induction hypothesis for the next site.

Last site. After processing site $n-1$, the residual has only two indices. Define the $r_{n-1} \times 1$ matrices

$$A_{\alpha_{n-1}, 1}^{[n] i_n} = R_{\alpha_{n-1}, i_n}^{[n-1]}.$$

Substituting the successive factorizations gives $\psi_{i_1 \dots i_n} = A^{[1] i_1} \dots A^{[n] i_n}$. No singular values have been discarded: all their weights were retained in the residual tensors and ultimately in the last tensor. The construction therefore gives an exact MPS with $\chi_k = r_k$ simultaneously at every cut. $\square$

2.1 Two explicit MPS examples

A product state $|\phi_1\rangle \otimes \cdots \otimes |\phi_n\rangle$ has bond dimension 1: all tensors are scalars

$$A^{[k]i} = \langle i | \phi_k \rangle.$$

For $|\text{GHZ}_n\rangle$, take

$$A^{[1]0} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 \end{pmatrix}, \quad A^{[1]1} = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 1 \end{pmatrix},$$


for $1 < k < n$,

$$A^{[k]0} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad A^{[k]1} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix},$$

and

$$A^{[n]0} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad A^{[n]1} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

The contraction can be pictured as two possible paths for the virtual bond index:



The first physical index selects channel 1 for the symbol 0 and channel 2 for the symbol 1. Every interior tensor is diagonal: on channel 1, only $A^{[k]0}$ acts nontrivially, while on channel 2, only $A^{[k]1}$ does. Thus the virtual index cannot switch channels. Only the strings 0^n and 1^n have nonzero contractions, and each has amplitude $1/\sqrt{2}$. Hence $\chi = 2$, in agreement with Theorem 2.3.

3 Classical simulation of low-entanglement circuits

The main computational consequence of the MPS representation is the following circuit-simulation theorem. From this point onward we specialize to qubits, so every physical index belongs to $\{0, 1\}$.

Theorem 3.1 (Exact MPS simulation principle). *Consider a circuit on an ordered chain of n qubits, starting from a product state and containing m gates, each of which is either a one-qubit gate or a two-qubit gate on neighboring qubits. Thus n is the number of qubits and m is the number of gates. Suppose that, after every gate, the Schmidt rank across every consecutive cut is at most $\chi = \text{poly}(n, m)$. Then there is an exact classical algorithm that produces an MPS representation of the circuit's output state. Throughout the computation, the maintained MPS has bond dimension at most χ . The gate updates take $O(m\chi^3)$ arithmetic operations, and the MPS uses $O(n\chi^2)$ complex numbers. Local expectation values and computational-basis samples can also be obtained in polynomial time.*

The output of the algorithm is therefore a concise representation of the final quantum state, rather than its 2^n amplitudes written explicitly. Section 3.2 explains what can be computed directly from this MPS representation.

The algorithm

Start with the bond-dimension-one MPS for the product input, and process the gates in circuit order.

1. If a one-qubit gate V acts on qubit k , update only that tensor:

$$\tilde{A}^{[k]i} = \sum_{j \in \{0,1\}} \langle i|V|j \rangle A^{[k]j}.$$

2. If a two-qubit gate G acts on neighboring qubits $k, k+1$, first contract the two tensors and apply G :

$$\Theta_{\alpha\beta}^{ij} = \sum_{p,q \in \{0,1\}} \sum_{\gamma=1}^{\chi_k} \langle i,j|G|p,q \rangle A_{\alpha\gamma}^{[k]p} A_{\gamma\beta}^{[k+1]q}.$$

Regard Θ as the matrix

$$\mathcal{M}_{(\alpha,i),(j,\beta)} = \Theta_{\alpha\beta}^{ij},$$

take its compact SVD $\mathcal{M} = U\Sigma V^\dagger$, and split it back into two neighboring tensors:

$$\tilde{A}_{\alpha\gamma}^{[k]i} = U_{(\alpha,i),\gamma}, \quad \tilde{A}_{\gamma\beta}^{[k+1]j} = (\Sigma V^\dagger)_{\gamma,(j,\beta)}.$$

Only the bond between sites k and $k+1$ is changed.

Why the algorithm works: a proof sketch. A one-qubit gate only changes one physical index and does not change any Schmidt rank. For a two-qubit gate, the SVD is an exact factorization of the updated two-qubit tensor. The number of retained singular values is exactly the Schmidt rank across the cut between the two sites; all other cuts have a unitary acting entirely on one side and hence keep the same Schmidt rank. Thus the bond dimensions never exceed χ . The matrix being decomposed has dimensions at most $2\chi \times 2\chi$, so each two-qubit update costs $O(\chi^3)$. This proves the theorem at the level needed in class.

3.1 Full proof of the circuit update

We give the details for reference. The invariant is that every current bond dimension equals the Schmidt rank at that cut:

$$\chi_\ell = r_\ell, \quad \ell = 1, \dots, n-1.$$

It holds for the product input, for which all these numbers equal one.

One-qubit gates. Substituting the one-qubit update into the MPS amplitude gives

$$A^{[1]i_1} \dots \tilde{A}^{[k]i_k} \dots A^{[n]i_n} = \sum_{j \in \{0,1\}} \langle i_k|V|j \rangle A^{[1]i_1} \dots A^{[k]j} \dots A^{[n]i_n},$$

which is exactly the amplitude of $V_k|\psi\rangle$. The bond dimensions are unchanged. Across any consecutive cut, V_k acts entirely on one side; since it is invertible, it does not change the rank of the coefficient matrix. The invariant is therefore preserved.

Two-qubit gates: exactness of the update. Let $\rho = \text{rank}(\mathcal{M})$. The compact SVD in the algorithm retains its ρ nonzero singular values, so

$$\sum_{\gamma=1}^{\rho} \tilde{A}_{\alpha\gamma}^{[k]i} \tilde{A}_{\gamma\beta}^{[k+1]j} = \Theta_{\alpha\beta}^{ij}.$$

Replacing the two old tensors by the two new ones therefore represents $G_{k,k+1}|\psi\rangle$ exactly. The new central bond dimension is $\tilde{\chi}_k = \rho$, while all other bond dimensions are unchanged.

Why ρ is the new Schmidt rank. Let L and R be the coefficient matrices obtained by contracting the tensors strictly to the left and strictly to the right of the updated pair:

$$\begin{aligned} L_{(i_1, \dots, i_{k-1}), \alpha} &= (A^{[1]i_1} \dots A^{[k-1]i_{k-1}})_{1, \alpha}, \\ R_{(i_{k+2}, \dots, i_n), \beta} &= (A^{[k+2]i_{k+2}} \dots A^{[n]i_n})_{\beta, 1}. \end{aligned}$$

At an endpoint the corresponding empty block is the scalar matrix (1).

The columns of L and R are linearly independent. For example, the coefficient matrix across cut $k-1$ has L as a left factor, so

$$r_{k-1} \leq \text{rank}(L) \leq \chi_{k-1}.$$

The invariant $r_{k-1} = \chi_{k-1}$ forces equality. The same argument on the right gives $\text{rank}(R) = \chi_{k+1}$.

Set $F = L \otimes I_2$ and $K = I_2 \otimes R$. The coefficient matrix $\tilde{\Psi}^{[k]}$ of the updated state across the central cut is

$$\tilde{\Psi}^{[k]} = F \mathcal{M} K^T.$$

Because F and K have linearly independent columns, they have left inverses F_L and K_L . Hence

$$\mathcal{M} = F_L \tilde{\Psi}^{[k]} K_L^T.$$

Multiplication cannot increase matrix rank, so the two identities imply

$$\text{rank}(\tilde{\Psi}^{[k]}) \leq \text{rank}(\mathcal{M}) \leq \text{rank}(\tilde{\Psi}^{[k]}).$$


Thus $\rho = \text{rank}(\mathcal{M})$ is exactly the new Schmidt rank $\tilde{r}_k$. At every other cut the gate acts entirely on one side, so the Schmidt rank does not change. The invariant is preserved after every gate.

Cost. The reshaped matrix in a two-qubit update has at most 2χ rows and 2χ columns. Its contraction and compact SVD cost $O(\chi^3)$; a one-qubit update costs $O(\chi^2)$. The n MPS tensors use $O(n\chi^2)$ entries. This proves the simulation and storage bounds. These are exact-arithmetic operation counts; finite-precision and bit-complexity bounds require a separate analysis.

Remark 3.2. *The rank hypothesis is sufficient and concrete, but restrictive. Practical MPS algorithms truncate small Schmidt coefficients after a two-qubit update. The optimal one-cut truncation result from Lecture 4 explains that step, but controlling the accumulated global error requires additional analysis.*

3.2 What can be computed from an MPS representation

Once the output state is stored as an MPS, a local expectation value is obtained by placing the MPS for $|\psi\rangle$ above its complex conjugate, inserting the observable on the desired physical leg, and contracting the resulting network from left to right:



At each stage only a $\chi_k \times \chi_k$ boundary matrix remains. For qubits, the full contraction costs $O(n\chi^3)$ both for the norm and for a one-qubit observable. The same picture with a constant-size operator inserted on several neighboring legs computes a local energy term $\langle\psi|h_a|\psi\rangle$.

Computational-basis probabilities are obtained by inserting basis projectors instead. Computing the conditional probabilities of successive outcomes and sampling them from left to right gives a weak simulation in polynomial time.

4 Quantum Hamiltonians: energy and locality

We now turn from a representation of quantum states to the operators that describe a physical system. A Hamiltonian is the observable representing the energy; in Lecture 6 we will also see how it generates time evolution.

4.1 Hamiltonians and locality

Definition 4.1 (Hamiltonian, ground energy, and locality). *A Hamiltonian on n qubits is a Hermitian operator $H \in \mathbb{C}^{2^n \times 2^n}$. Its smallest eigenvalue*

$$E_0 = \lambda_{\min}(H)$$

is the ground energy, and the corresponding eigenspace is the ground space. A normalized vector in this eigenspace is a ground state.

The Hamiltonian is k -local if it can be written

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

where every term h_a acts nontrivially on at most k qubits. We usually take $k = O(1)$, $L = \text{poly}(n)$, and assume that every local matrix h_a is efficiently specified. Locality refers to the number of qubits in each interaction. Geometric locality additionally requires these qubits to be near each other in a specified geometry.

The Hilbert space has dimension 2^n , so an arbitrary Hamiltonian has exponentially many matrix entries. Physical models are normally specified as sums of simple interactions.

Why is this natural? A term acting on one site describes, for example, the coupling of a spin to an external field. A two-site term describes a direct interaction between two spins. Summing such terms builds a many-body model from local pieces even though its eigenstates and dynamics can be highly entangled.

The transverse-field Ising chain. This is an idealized model of a one-dimensional quantum magnet. Each qubit represents a localized magnetic moment, or spin. The states $|0\rangle$ and $|1\rangle$ describe the two orientations along a preferred z -axis. The nearest-neighbor interaction favors alignment along that axis, while a magnetic field applied in the perpendicular x -direction tends to rotate the spins and creates quantum fluctuations.

For an open chain of n spins, consider

$$H_{\text{Ising}} = -J \sum_{j=1}^{n-1} Z_j Z_{j+1} - g \sum_{j=1}^n X_j, \quad J, g \geq 0.$$

In the computational basis, $Z_j Z_{j+1}$ has eigenvalue $+1$ when the neighboring bits agree and -1 when they disagree. Thus the term $-J Z_j Z_{j+1}$ energetically favors aligned neighboring spins. The operator X_j flips the j -th bit, so the transverse field mixes different classical spin configurations. The two parts generally do not commute: the model expresses a competition between classical alignment and quantum fluctuations.

At $g = 0$, the ground space is spanned by $|0^n\rangle$ and $|1^n\rangle$. At $J = 0$ and $g > 0$, the unique ground state is $|+\rangle^{\otimes n}$. Away from these limits, the ground state need not be a computational-basis state or a product state.

For any normalized state $|\psi\rangle = \sum_j c_j |E_j\rangle$,

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

We therefore have the variational characterization

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

For a local Hamiltonian and an MPS, each term $\langle\psi|h_a|\psi\rangle$ can be evaluated by the contraction method of Section 3.2, with the local observable inserted on the support of h_a . Thus low-bond-dimension MPSs form a concrete variational family for approximating ground states.

5 Optional training exercises

1. Verify directly that the displayed bond-dimension-two tensors represent $|\text{GHZ}_n\rangle$.
2. Derive the left-environment update for $\langle\psi|\psi\rangle$ and check the contraction cost when all internal bond dimensions equal χ .
3. Starting from $|00\rangle$, apply H_1 and $\text{CNOT}_{1\rightarrow 2}$. Use the update equations to obtain an MPS for the Bell state, and explain why its internal bond dimension changes from one to two.

6 Assigned exercise: an explicit MPS for the W state

Use the state $|W_n\rangle$ from the beginning of the lecture, for $n \geq 2$.

- (a) Use Theorem 2.3 and the Schmidt ranks computed above to determine the smallest possible bond dimension at each nontrivial consecutive cut.

- (b) Let

$$B^0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad B^1 = \frac{1}{\sqrt{n}} \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix},$$

with $\ell^T = (1 \ 0)$ and $r = (0 \ 1)^T$. Verify that $\ell^T B^{i_1} \cdots B^{i_n} r$ is $1/\sqrt{n}$ when $i_1 \cdots i_n$ contains exactly one 1, and is zero otherwise. Convert this expression into open-boundary MPS tensors and specify all their sizes.

- (c) Use the contraction procedure of Section 3.2 with the projector $|1\rangle\langle 1|$ at site j to recover $\Pr[\text{qubit } j \text{ is } 1] = 1/n$. Continue the contraction through site n , not only to site j .

References

- [1] Ulrich Schollwöck, *The density-matrix renormalization group in the age of matrix product states*, Annals of Physics 326, 96–192 (2011).
<https://arxiv.org/abs/1008.3477>