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

Lecture 9: Description Complexity and Low-Energy States

October 2, 2025 *Scribe: Deepu Pradeep*

1 Description Complexity

We can represent an n-qubit quantum state in several ways:

- as a **wavefunction** $|\psi\rangle = \sum_x \psi(x)|x\rangle$;
- by a **quantum circuit** U preparing a state

  $$|\psi\rangle \approx U|0^n\rangle,$$

  where the $\approx$ emphasizes that the prepared state is only approximately the target state due to limited circuit depth or gate set;
- or by a **tensor-network representation**, such as a Matrix Product State (MPS) in 1D or a Projected Entangled Pair State (PEPS) in 2D (refer to Appendix A for a more detailed treatment of this representation).

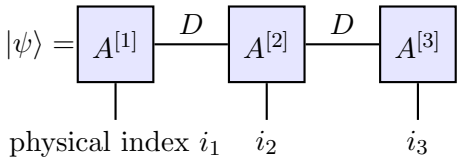

Figure 1: Example of a 1D tensor network (Matrix Product State). Each tensor $A^{[i]}$ has one *physical index* (vertical leg) and two *virtual indices* (horizontal legs) connecting neighboring tensors. Contracting virtual indices encodes local entanglement, and the bond dimension D determines how much entanglement can be represented.

The tensor-network representation explicitly captures the local entanglement structure of quantum states: the geometry of the network constrains how entanglement scales with subsystem size, and efficient tensor networks correspond to states satisfying an area law. This makes them especially well-suited for representing ground states of local Hamiltonians.

Each representation provides a different measure of how "complex" the state is to describe or manipulate.

1.1 Description vs. Operational Complexity

While **description complexity** tells us how many bits are needed to specify a state's representation, it is often not the most *useful* quantity in practice.

Motivation. The notion of description complexity is fine in principle, but limited in practice. If someone provides a state through a quantum circuit, a quantum computer could prepare it and estimate observables. However, if we do not have a quantum computer, can we still compute quantities like expectation values efficiently? To address this, we define a more *operational* notion that captures how difficult it is to extract physical information from a given representation.

1.2 Operational Description Complexity

Definition 1.1 (Operational description complexity)**.** Given a representation of a quantum state (e.g., a circuit, tensor network, or wavefunction) and an integer k, the **operational description complexity** is the classical time required to compute

$$\mathrm{Tr}\,(O\rho)$$

for all observables O of *weight* $\leq k$, where the weight of O denotes the number of non-identity tensor factors in its decomposition. In particular, this includes computing the *energy*

$$\mathrm{Tr}(H\rho) = \sum_x \mathrm{Tr}(h_x\rho)$$

when $H = \sum_x h_x$ is a k-local Hamiltonian.

Example: Local observables in a shallow circuit. Suppose the state is given by a depth-d quantum circuit U acting on n qubits:

$$|\psi\rangle = U|0^n\rangle.$$

Consider computing the expectation value of a k-local observable O:

$$\langle\psi|O|\psi\rangle = \langle 0^n|U^\dagger OU|0^n\rangle.$$

The key insight is that under conjugation by a *local* circuit, the observable O only spreads to a region whose size grows with d. On an all-to-all architecture, the *light-cone* of O may include up to $O(2^d)$ qubits. Thus, evaluating $\langle\psi|O|\psi\rangle$ requires simulating only that portion of the circuit rather than the full n-qubit state.

In particular, if the circuit has depth

$$d = O(\log\log n),$$

the effective support of O remains small, and the computation can be performed efficiently on a classical computer. This reflects that the operational description complexity depends not only on the size of the circuit, but on how entanglement and correlations spread through it.

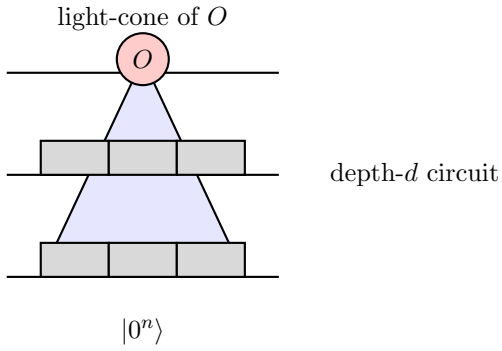

Figure 2: Illustration of the *light-cone structure*: a local observable O conjugated by a finite-depth circuit U spreads only within a bounded region whose size grows with the circuit depth d. This locality limits the classical cost of computing $\mathrm{Tr}\,(O\rho)$.

1.3 Motivation

The study of quantum-state complexity is central to both physics and computer science. States that appear in nature—ground states of gapped Hamiltonians, Gibbs states, etc.—are believed to have polynomial-sized descriptions under some efficient representation, unlike arbitrary states of 2^n amplitudes.

2 Physical States and Description Complexity

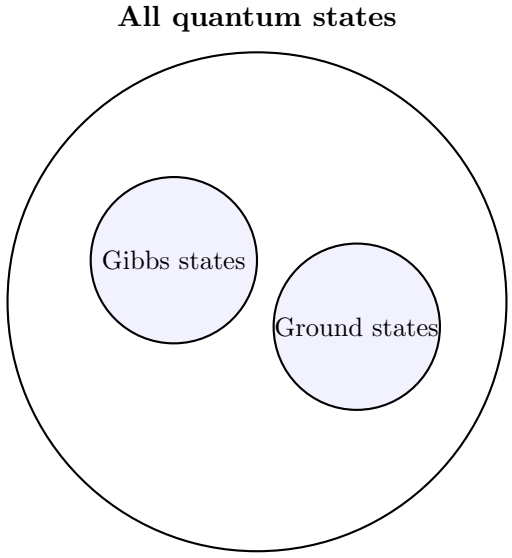

Figure 3: Conceptual picture: the full set of quantum states is exponentially large, but only a tiny subset— such as Gibbs states and ground states of local Hamiltonians—arise in nature and have structured, low-complexity descriptions.

The set of all quantum states on n qubits is enormous, containing exponentially many possi-

bilities in n. However, only a small subset of these states—those that occur naturally in physical systems—tend to have simple descriptions. These *physical states* include:

- **Gibbs states:** thermal equilibrium states of the form

$$\rho_\beta(H) = \frac{e^{-\beta H}}{\mathrm{Tr}(e^{-\beta H})},$$

  which can be described concisely by their Hamiltonian H and temperature β;

- **Ground states:** lowest-energy eigenstates of local Hamiltonians.

A central question in quantum complexity is:

What is the complexity of describing and computing properties of these physically realizable states?

The distinction between these two cases is important. A Gibbs state is uniquely determined by its Hamiltonian and temperature, whereas a Hamiltonian can have many degenerate ground states—different states with the same lowest energy. Thus, "describing the Hamiltonian" does not automatically describe a unique ground state.

2.1 Examples of Representations for Physical States

We now consider which representations provide efficient and meaningful descriptions for physically relevant quantum states such as Gibbs states and ground states of local Hamiltonians.

1. The Hamiltonian itself.

- A Hamiltonian H can serve as a concise representation of a Gibbs state,

$$\rho_\beta(H) = \frac{e^{-\beta H}}{\mathrm{Tr}(e^{-\beta H})}.$$

  This is justified by the **Jaynes principle**: the thermal state maximizes entropy subject to fixed local expectation values. Hence, the local marginals determine H (and vice versa).

- However, the Hamiltonian is *not* a good representation for ground states. It does not capture entanglement structure and cannot distinguish among multiple ground states in a degenerate subspace.

2. Tensor-network representations.

- Tensor networks provide good representations of ground states of gapped 1D Hamiltonians. This follows from the **area law** for entanglement, which implies that these states have limited entanglement across cuts.

- They are also *believed* to be good representations for gapped 2D Hamiltonians.

- In 1D, a matrix product state (MPS) with bond dimension $D = \mathrm{poly}(n)$ serves as an efficient operational description, since expectation values of local observables can be computed in polynomial time in n and D.

3. Quantum circuit representations.

- Quantum circuit descriptions are *not known* to efficiently represent low-energy states of generic local Hamiltonians.
- Formally, suppose for every local Hamiltonian H_c, there exists a circuit U and state $|\psi\rangle$ such that
$$\langle\psi|H_c|\psi\rangle \leq E_0 + \frac{1}{n^3}$$
and the circuit complexity of U is polynomial. Then one could efficiently verify ground-state energies, implying
$$\text{QMA} = \text{QCMA}.$$
- Furthermore, if for every H_c there exists such a $|\psi\rangle$ with
$$\langle\psi|H_c|\psi\rangle \leq E_0 + \frac{1}{n^3}$$
and the circuit has depth $O(\log\log n)$, this would imply QMA = NP. The "if part" is actually false, for instance for Hamiltonians such as the toric code, whose ground states require circuits of depth $\Theta(\log n)$ to prepare.

3 Search for Simple Low-Energy States

A central goal in both quantum complexity theory and condensed-matter physics is to understand *how well simple states can approximate low-energy states* of local Hamiltonians. This is closely related to estimating the ground-state energy—a fundamental task in physics.

Question. Given a Hamiltonian H, how well can a simple class of states (e.g., product states or low-depth circuit states) approximate the ground energy of H (or equivalently, the maximum energy of $-H$)?

We define an **approximation ratio** that quantifies this:
$$\text{Approximation ratio} = \frac{E_{\text{simple}}}{E_{\max}},$$
where E_{simple} is the maximum energy achievable within the restricted family of "simple" states and $E_{\max}$ is the true maximum eigenvalue of H. The closer this ratio is to 1, the better the family approximates the optimal energy.

3.1 Product-State Approximation

Let the Hamiltonian be a 2-local operator written as
$$H = \sum_{\alpha} b_\alpha P_\alpha,$$

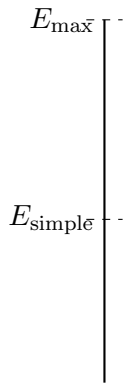

Figure 4: The approximation ratio measures how well a simple state family reaches the true maximum energy.

where each P_α is a Pauli operator with zero identity coefficient. For such Hamiltonians, Lieb, and later Brayvi–Gosset–Temme–Konig proved a universal product-state approximation bound:

$$\frac{E_{\text{product}}}{E_{\max}} \geq \frac{1}{9}.$$

That is, even if the optimal (possibly highly entangled) state achieves energy $E_{\max}$, there always exists a product state whose energy is at least one-ninth of that optimum.

3.1.1 Proof Sketch (via the dephasing channel)

Let $|\psi\rangle$ be the state achieving the maximum energy,

$$E_{\max} = \langle\psi|H|\psi\rangle.$$

Define the quantum channel $\mathcal{M}$ that measures each qubit in the Pauli eigenbasis:

$$\mathcal{M}(\rho) = \frac{1}{3} \sum_{\alpha\in\{X,Y,Z\}} \sum_{b=\pm 1} |\phi_b^\alpha\rangle\langle\phi_b^\alpha|\, \rho\, |\phi_b^\alpha\rangle\langle\phi_b^\alpha|,$$

where $\{|\phi_b^\alpha\rangle\}$ are the eigenstates of the Pauli operators.

Applying $\mathcal{M}^{\otimes n}$ to $|\psi\rangle\langle\psi|$ yields a product (classical) state:

$$\rho_{\text{prod}} = \mathcal{M}^{\otimes n}(|\psi\rangle\langle\psi|).$$

Then,

$$\text{Tr}\left(H\rho_{\text{prod}}\right) = \sum_\alpha b_\alpha \text{Tr}\left(P_\alpha\, \mathcal{M}^{\otimes n}(|\psi\rangle\langle\psi|)\right) = \sum_\alpha b_\alpha \text{Tr}\left(\mathcal{M}^{\otimes 2}(P_\alpha)\, |\psi\rangle\langle\psi|\right).$$

The channel $\mathcal{M}$ acts on the Pauli basis as

$$\mathcal{M}(X) = \tfrac{1}{3}X, \qquad \mathcal{M}(Y) = \tfrac{1}{3}Y, \qquad \mathcal{M}(Z) = \tfrac{1}{3}Z, \qquad \mathcal{M}(I) = I.$$

Therefore, for any 2-local term P_α,

$$\mathcal{M}^{\otimes 2}(P_\alpha) = \tfrac{1}{9}P_\alpha.$$

It follows that

$$\begin{aligned}\operatorname{Tr}\left(H\rho_{\text{prod}}\right) &= \sum_{\alpha} b_\alpha \operatorname{Tr}\left(\left\{P_\alpha\, \mathcal{M}^{\otimes n}(|\psi\rangle\langle\psi|)\right\}\right)\\ &= \sum_{\alpha} b_\alpha \operatorname{Tr}\left(\left\{\mathcal{M}^{\otimes 2}(P_\alpha)\, |\psi\rangle\langle\psi|\right\}\right)\\ &= \frac{1}{9}\sum_{\alpha} b_\alpha \operatorname{Tr}\left(\left\{P_\alpha\, |\psi\rangle\langle\psi|\right\}\right)\\ &= \frac{E_{\max}}{9}.\end{aligned}$$

Hence,

$$\frac{E_{\text{product}}}{E_{\max}} \geq \frac{1}{9}.$$

This shows that product states provide a constant-factor approximation to the maximum energy of any 2-local Hamiltonian.

3.2 Nonnegative Two-Local Hamiltonians

Let

$$H = \sum_{i<j} h_{ij}, \qquad h_{ij} \succeq 0,$$

where each h_{ij} acts nontrivially only on qubits i and j (i.e. 2-local and positive semidefinite).

Then for such Hamiltonians,

$$\frac{E_{\text{prod}}}{E_{\max}} \geq \frac{1}{2},$$

as shown by Gharibian–Kempe (2011).

Let $|\psi\rangle$ be the state achieving maximal energy, so that

$$\langle\psi|H|\psi\rangle = E_{\max}.$$

This bound is optimal, as can be seen from the two-qubit EPR pair example.

Consider the following channel $\mathcal{M}_{1\ldots n}$: measure qubit 1 in its eigenbasis; condition on the measurement outcome, measure qubit 2 in a new eigenbasis, and so on sequentially. This defines a completely positive trace-preserving (CPTP) map that produces a product state from $|\psi\rangle$.

We evaluate

$$\begin{aligned}\operatorname{Tr}\left(\left\{\mathcal{M}_{1\ldots n}(|\psi\rangle\langle\psi|)\, H\right\}\right) &= \sum_{i<j} \operatorname{Tr}\left(\left\{\mathcal{M}_{1\ldots n}(|\psi\rangle\langle\psi|)\, h_{ij}\right\}\right)\\ &= \sum_{i<j} \operatorname{Tr}\left(\left\{|\psi\rangle\langle\psi|\, \mathcal{M}_{1\ldots n}(h_{ij})\right\}\right)\\ &= \sum_{i<j} \operatorname{Tr}\left(\left\{|\psi\rangle\langle\psi|\, \mathcal{M}_{1\ldots j}(h_{ij})\right\}\right),\end{aligned}$$

where we use that terms h_{ij} act only on sites i, j and later measurements act trivially on them.

Proceeding inductively,

$$\begin{aligned}\operatorname{Tr}\left(\left\{\mathcal{M}_{1\ldots n}(|\psi\rangle\langle\psi|)\, H\right\}\right) &= \sum_{i<j} \operatorname{Tr}\left(\left\{\mathcal{M}_{1\ldots(j-1)}(|\psi\rangle\langle\psi|)\, h_{ij}\right\}\right) \\ &\geq \frac{1}{2}\sum_{i<j} \operatorname{Tr}\left(\left\{|\psi\rangle\langle\psi|\, h_{ij}\right\}\right) \\ &= \frac{E_{\max}}{2},\end{aligned}$$

where the inequality uses the following fact for all positive semidefinite h_{ij}:

$$(\langle v|\otimes\mathbb{I})\, h_{ij}\, (|v\rangle\otimes\mathbb{I}) + (\langle v^{\perp}|\otimes\mathbb{I})\, h_{ij}\, (|v^{\perp}\rangle\otimes\mathbb{I}) \succeq \frac{h_{ij}}{2}.$$

Hence, the product state $\mathcal{M}_{1\ldots n}(|\psi\rangle\langle\psi|)$ achieves at least half of the optimal energy.

3.3 Dense Hamiltonians

For dense 2-local Hamiltonians (where each qubit interacts with many others), Brandão–Harrow (2013) showed an even stronger guarantee:

$$\frac{E_{\text{prod}}}{E_{\max}} \geq 1 - \frac{1}{3\sqrt{d}},$$

where d denotes the average degree of the interaction graph.

Intuitively, as the Hamiltonian becomes denser (large d), correlations average out, and product states approximate the optimal energy more closely.

A Tensor Networks and Matrix Product States

Tensor networks provide a structured and efficient way of representing quantum states, particularly those with limited entanglement. They describe high-rank tensors (multidimensional arrays) using diagrams where nodes represent tensors and edges represent contracted indices. Each line, or *leg*, corresponds to one index of the tensor, and contracting two tensors along a leg represents summing over that shared index, analogous to matrix multiplication.

A.1 Motivation: Entanglement as a Resource

The usefulness of tensor networks stems from the observation that entanglement acts as a *resource* in quantum mechanics. Many physical quantum states— such as ground states of local Hamiltonians—exhibit **limited entanglement** across spatial partitions. For such states, the full exponential parameterization $|\psi\rangle = \sum_x \psi(x)|x\rangle$ is redundant. Instead, one can store only the *locally entangled structure* using a tensor network representation that scales polynomially in system size.

A.2 Tensor Network Diagrams

A few basic tensor objects are:

- Vector (rank-1): A_j — drawn as a node with one leg.
- Matrix (rank-2): M_{ij} — node with two legs.
- General tensor (rank-k): $T_{\alpha\beta\gamma\ldots}$ — node with k legs.

Connecting legs represents summing over shared indices (a *contraction*). When all legs are contracted, the network represents a scalar value. The order of contraction determines the computational cost of evaluating the network; minimizing the number of open legs at intermediate steps reduces the memory required for simulation.

A.3 Schmidt Decomposition and Low-Entanglement States

A single bipartition provides the starting point for understanding tensor networks. Any bipartite pure state can be written as

$$|\psi\rangle = \sum_{\mu=1}^{\chi} \sqrt{\lambda_\mu} |\phi_\mu^A\rangle |\phi_\mu^B\rangle,$$

where χ is the Schmidt rank and λ_μ are the Schmidt coefficients. The number χ quantifies the effective entanglement between the two parts: if χ is small, the state can be represented compactly. The entanglement entropy $S = -\sum_\mu \lambda_\mu \log \lambda_\mu$ determines the number of parameters required.

A.4 Matrix Product States (MPS)

For a system of n qubits, we can recursively apply the Schmidt decomposition between successive partitions to express the global state as a *matrix product*:

$$|\psi\rangle = \sum_{i_1,\ldots,i_n} A_{i_1}^{[1]} A_{i_2}^{[2]} \cdots A_{i_n}^{[n]} |i_1 i_2 \cdots i_n\rangle,$$

where each $A_{i_k}^{[k]}$ is a $\chi_{k-1} \times \chi_k$ matrix. The internal dimensions χ_k, known as *bond dimensions*, encode the entanglement across each bipartition. The number of parameters required scales as

$$O(nd\chi^2),$$

where d is the local Hilbert-space dimension (2 for qubits). When the state obeys an *area law* for entanglement, χ grows only polynomially with n, making this representation efficient:contentReference[oaicite:4]index=

A.5 Constructive Algorithm (SVD Approach)

An MPS can be constructed iteratively:

1. Start with the full wavefunction $\psi_{i_1 i_2 \ldots i_n}$.
2. Perform an SVD between the first qubit and the rest: $\psi_{i_1,(i_2\ldots i_n)} = U\Lambda V$.

3. Define $A^{[1]} = U$, absorb Λ into V, and repeat the SVD between the next qubit and the remainder.

After $n-1$ steps, we obtain an MPS with tensors $\{A^{[1]}, A^{[2]}, \ldots, A^{[n]}\}$ and diagonal Schmidt matrices $\{\Lambda^{[1]}, \Lambda^{[2]}, \ldots\}$ between them.

A.6 Canonical Forms

Different conventions exist for absorbing the Λ matrices into adjacent tensors. In the **left-canonical form**, each tensor to the left of a central site forms an orthonormal basis; similarly, the right-canonical form orthogonalizes the right side. The central tensor where both orthogonality conditions meet is called the *orthogonality center*. This canonicalization is numerically stable and ensures efficient computation of observables such as

$$\langle O_j \rangle = \mathrm{Tr}(\rho_j O_j),$$

by moving the orthogonality center to site j before contracting the network.

A.7 Projected Entangled Pair States (PEPS)

An alternative but equivalent picture is the **Projected Entangled Pair State (PEPS)**. Imagine creating maximally entangled pairs between neighboring sites and then applying a local projection operator at each site that maps the virtual indices of those pairs to the physical Hilbert space. The resulting state is a tensor network whose geometry matches that of the lattice. PEPS generalize MPS to higher dimensions.

A.8 The Area Law and Efficiency

For local Hamiltonians with ground states obeying the area law,

$$S_A \propto |\partial A|,$$

the entanglement entropy of a region A scales with the size of its boundary rather than its volume. In one-dimensional gapped systems, this implies constant entanglement entropy, guaranteeing that MPS with small bond dimension provide an exponentially accurate approximation to the ground state. In higher dimensions, PEPS states play a similar role, though the contraction cost grows more rapidly.

A.9 Computational Use and Physical Insight

Tensor networks are not merely efficient data structures—they encode the physical *structure of entanglement*. The geometry of the network mirrors the causal or spatial structure of correlations in the system. The bond dimension D quantifies the amount of entanglement that can be represented, while the network topology reflects the locality of interactions. For example:

- MPS efficiently represent 1D ground states of gapped Hamiltonians.
- PEPS generalize this to 2D lattice systems.
- Tree tensor networks and MERA capture critical or scale-invariant systems.

A.10 Summary

Tensor networks compress quantum states by encoding entanglement structure directly into network geometry. The MPS formalism, derived from successive Schmidt decompositions, provides an exponentially efficient representation for 1D systems obeying an area law. The key insight is that quantum states of physical interest—especially ground states of local Hamiltonians—occupy a tiny, structured corner of the exponentially large Hilbert space, and tensor networks reveal that structure.