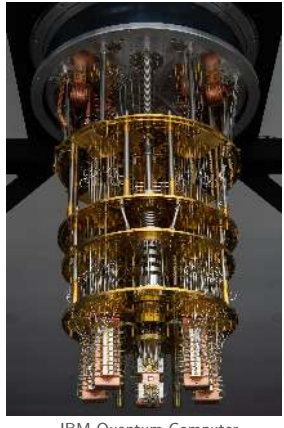
From the quantum many-body problem, to quantum computers and tensor networks

Classical computing as competitor and enabler



IBM Quantum Computer
James Estrin/The New York Times

Antoine Tilloy
April 9th, 2026
56th ORAP Forum
Paris, France

MINES PARIS | PSL | ENS | PSL | Inria | erc

The quantum many-body problem

The thorn in the side of reductionism

Understood by Paul Dirac as early as 1930



Quantum Mechanics of Many-Electron Systems.

By P. A. M. Dirac, St. John's College, Cambridge.

(Communicated by R. H. Fowler, F.R.S.—Received March 12, 1929.)

§ 1. *Introduction.*

The general theory of quantum mechanics is now almost complete, the imperfections that still remain being in connection with the exact fitting in of the theory with relativity ideas. These give rise to difficulties only when high-speed particles are involved, and are therefore of no importance in the consideration of atomic and molecular structure and ordinary chemical reactions, in which it is, indeed, usually sufficiently accurate if one neglects relativity variation of mass with velocity and assumes only Coulomb forces between the various electrons and atomic nuclei. The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble. It therefore becomes desirable that approximate practical methods of applying quantum mechanics should be developed, which can lead to an explanation of the main features of complex atomic systems without too much computation.

The quantum n-body problem

Schrödinger equation

$$\frac{\partial \psi}{\partial t} = -\frac{i}{\hbar} H\psi \quad \text{where} \quad \psi = \begin{bmatrix} \psi_1 \\ \psi_2 \\ \psi_3 \\ \cdots \\ \cdots \\ \psi_N \end{bmatrix}$$

A simple equation (H is a matrix) in a colossal number of variables N!

The quantum n-body problem

Schrödinger equation

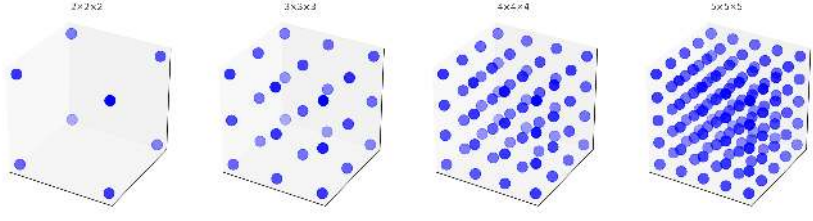
$$\frac{\partial \psi}{\partial t} = -\frac{i}{\hbar} H\psi \quad \text{where} \quad \psi = \begin{bmatrix} \psi_1 \\ \psi_2 \\ \psi_3 \\ \cdots \\ \cdots \\ \psi_N \end{bmatrix}$$

A simple equation (H is a matrix) in a colossal number of variables N!

$N = D^n$ variables (or free parameters)

- $D \geqslant 2$ number of states accessible to an elementary constituent
- n number of constituents

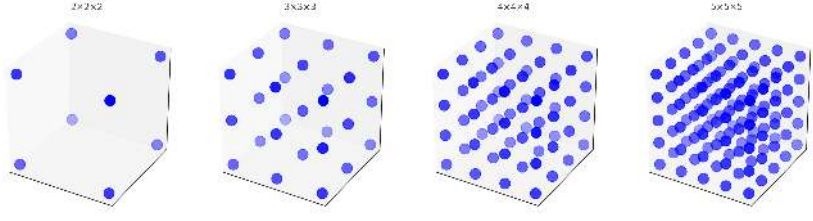
The quantum n-body problem: memory explosion



The simplest case: $D = 2$ possible states

- cube $2 \times 2 \times 2 = 8$ atoms – solvable on a laptop

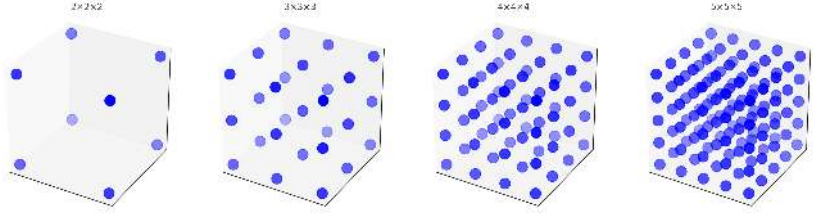
The quantum n-body problem: memory explosion



The simplest case: $D = 2$ possible states

- cube $2 \times 2 \times 2 = 8$ atoms – solvable on a laptop
- cube $3 \times 3 \times 3 = 27$ atoms – solvable on a laptop

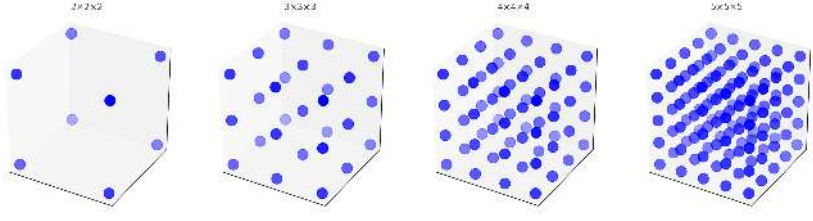
The quantum n-body problem: memory explosion



The simplest case: $D = 2$ possible states

- cube $2 \times 2 \times 2 = 8$ atoms – solvable on a laptop
- cube $3 \times 3 \times 3 = 27$ atoms – solvable on a laptop
- cube $4 \times 4 \times 4 = 64$ atoms – unsolvable even with Frontier

The quantum n-body problem: memory explosion

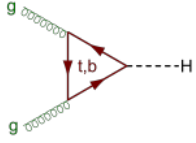


The simplest case: $D = 2$ possible states

- cube $2 \times 2 \times 2 = 8$ atoms – solvable on a laptop
- cube $3 \times 3 \times 3 = 27$ atoms – solvable on a laptop
- cube $4 \times 4 \times 4 = 64$ atoms – unsolvable even with Frontier
- cube $5 \times 5 \times 5 = 125$ atoms – not enough silicon in the universe

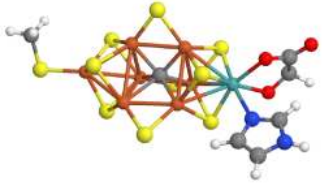
Why it matters: open problems and practical consequences

Nuclear / particle physics



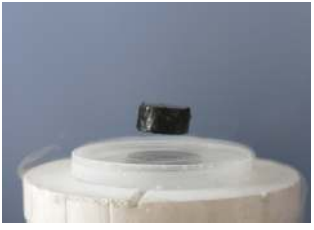
Strong force between quarks
Nuclear properties only measured experimentally

Chemistry / catalysis



Complex molecular interactions
Ammonia: $\geqslant 1\%$ of global CO2

Condensed matter



Cuprate superconductivity
No lossless electricity transfer

Turning hardness into power

From quantum simulation to quantum computing

Simulating Physics with Computers

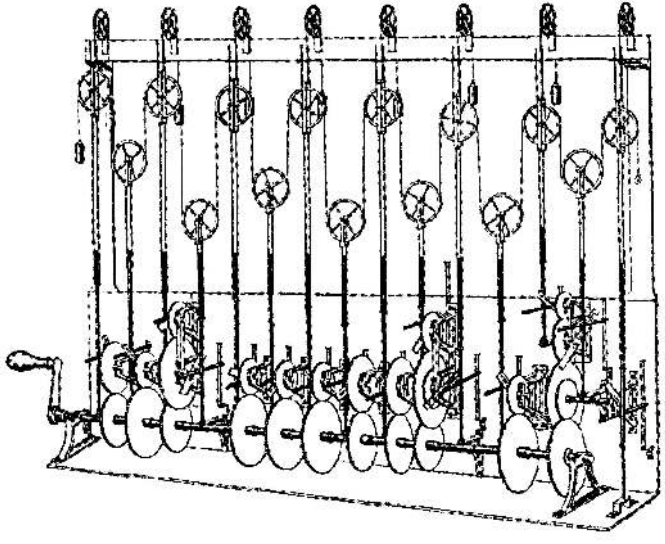
Richard P. Feynman

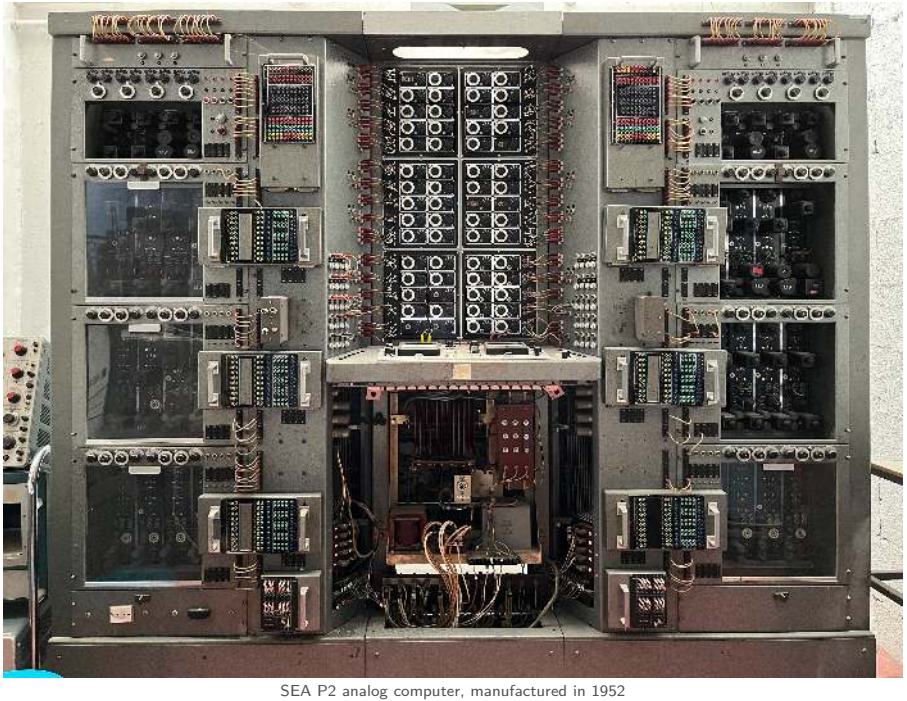
Department of Physics, California Institute of Technology, Pasadena, California 91107

Received May 7, 1981

"Nature is not classical, and if you want to make a simulation of nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy."

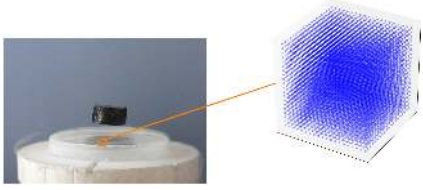
Historical analog computers

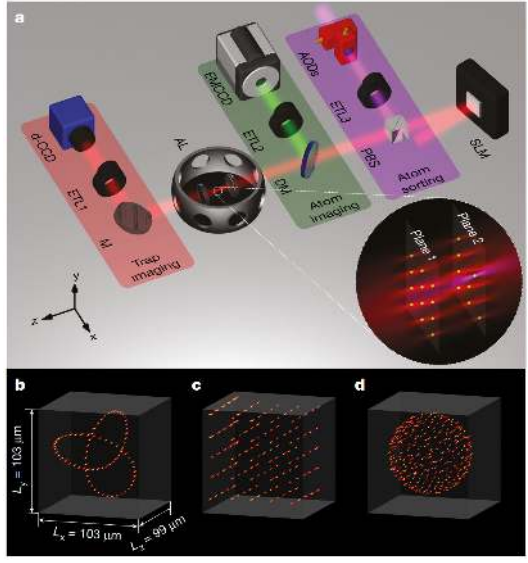




SEA P2 analog computer, manufactured in 1952

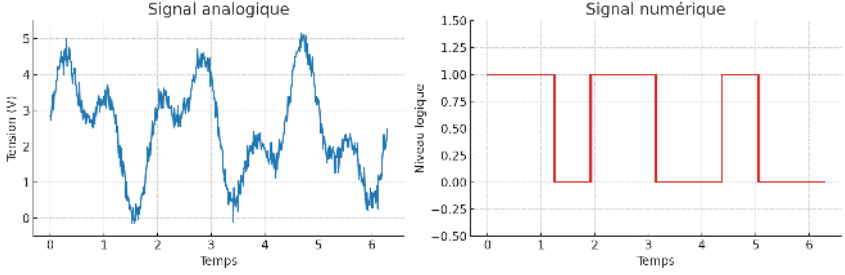
Quantum simulators: atoms in optical tweezers





Barredo, Lienhard, de Léséleuc, Lahaye, Browaeys, *Nature*, 2018

From analog to digital



Restricting the rules of the game:

- Two-level quantum systems $D = 2$ (qubits)
- Only a few discrete allowed operations / dynamics

Two less intuitive surprises

Universality of quantum logic gates

With a **finite** set of simple quantum operations, one can reproduce an arbitrary Hamiltonian dynamics H to arbitrary precision.

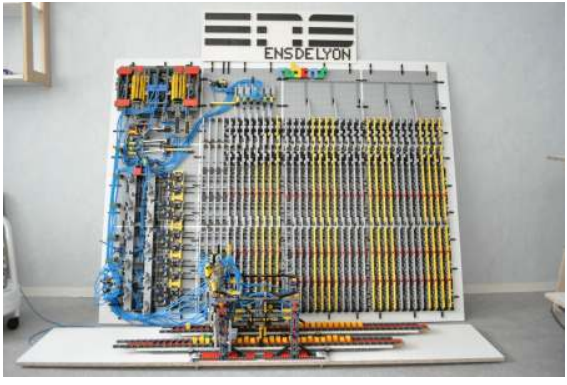
Fault tolerance

Using a clever encoding, one can reach arbitrarily good precision *even if the quantum logic gates are imperfectly implemented.*

Beyond the n-body problem: Shor

Breaking cryptography along the way

Factoring on a classical computer



ENS Lyon – Lego Turing machine $\sim 10^{-2}$flops flops



Oak Ridge – Summit $\sim 10^{17}$

Finding p and q for an N-digit number takes time

$$t_{\text{Lego}} = C_{\text{Lego}} \ \exp(\sqrt[3]{N})$$

$$t_{\text{Summit}} = C_{\text{Summit}} \ \exp(\sqrt[3]{N})$$

Factoring on a quantum computer



- Any classical computer

$$t = C \exp\left(\sqrt[3]{N}\right)$$

- Shor's algorithm on a quantum computer

$$t = C_q N^3$$

The quantum computer is slower per operation: $C_q \gg C_{\text{Summit}}$, but doesn't feel the difficulty increase! $\implies$ breaks strong Church Turing

But quantum computers are insanely slow

A quantum gate takes $\sim 10^{-6}$ s vs $\sim 10^{-10}$ s for a classical transistor
$\Longrightarrow$ overhead factor $C_q / C_{\text{classical}} \sim 10^4 - 10^6$

For Shor (exponential $\rightarrow$ polynomial): this doesn't matter

$$C\, \mathrm{e}^{\sqrt[3]{N}} \quad \text{vs} \quad C_q\, N^3$$

The exponential wins eventually, regardless of the prefactor.

Polynomial speedups?

For Grover (quadratic speedup): the overhead *does* matter

$$C\,N \quad \text{vs} \quad C_q\sqrt{N}$$

Crossing at $N^* = (C_q/C)^2 \sim 10^8 - 10^{12}$

Polynomial speedups?

For Grover (quadratic speedup): the overhead *does* matter

$$C\,N \quad \text{vs} \quad C_q\sqrt{N}$$

Crossing at $N^* = (C_q/C)^2 \sim 10^8 - 10^{12}$

A **galactic algorithm**: asymptotically better, but practically useless

Like Coppersmith-Winograd matrix multiplication $O(N^{2.373})$ vs Strassen $O(N^{2.807})$ vs naive $O(N^3)$ – theoretically better, never used in practice.

$\rightarrow$ Quadratic speedups are almost never enough!

Other applications?

Beyond the N-body problem and fast factoring: **not much!**

- Optimization? $\rightarrow$ unclear
- Artificial intelligence? $\rightarrow$ probably not
- Climate-related problems? $\rightarrow$ perhaps very indirectly
- Testing everything in parallel? $\rightarrow$ **no**

For now: the N-body problem is the most useful application!

Beating the many-body problem classically

State compression with tensor networks

The compression approach

Variational method for ground state search

1. Guess a manifold $\mathcal{M} \subset \mathscr{H}$ with few parameters ν, *i.e.* $\dim\mathcal{M} \ll \dim\mathscr{H}$
2. Tune ν to minimize energy $\nu = \operatorname{argmin}_{\nu\in\mathcal{M}} \frac{\langle\nu|H|\nu\rangle}{\langle\nu|\nu\rangle}$ and get $|\text{ground state}\rangle \simeq |\nu\rangle$

The compression approach

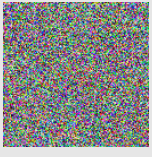
Variational method for ground state search

1. Guess a manifold $\mathcal{M} \subset \mathscr{H}$ with few parameters ν, *i.e.* $\dim\mathcal{M} \ll \dim\mathscr{H}$
2. Tune ν to minimize energy $\nu = \operatorname{argmin}_{\nu\in\mathcal{M}} \frac{\langle\nu|H|\nu\rangle}{\langle\nu|\nu\rangle}$ and get $|\text{ground state}\rangle \simeq |\nu\rangle$

Classical compression



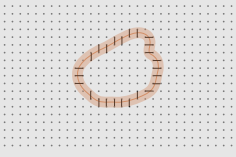
cat image



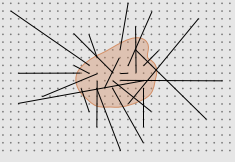
"typical" image

atypical $\Longrightarrow$ compressible

Quantum compression



low energy state

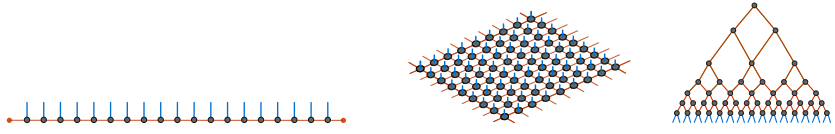


random state

area law = atypical $\Longrightarrow$ compressible

Tensor network states in a nutshell

.zip or **.jpg** for complex quantum states that appear in Nature



- **Exponential reduction:** $2^N \longrightarrow N \times D^{2d}$ parameters
- **Efficient compression:** error $\leqslant e^{-D}$ or $1/\text{superpoly(D)}$
- **Computational cost:** D^3 in $1d$, $D^9 - D^{11}$ in $2d$ $\leftarrow$ dense matrix multiplication

Hard to reach the scaling regime in $3d$ (analogous to conv nets before GPUs?)

Usually **memory reasonable, computationally bottlenecked**

Where does classical compression fail?

Tensor networks efficiently solve "low-entanglement" problems

They struggle with:

- Real-time dynamics (entanglement grows linearly)
- Systems in $d \geqslant 2$ space dimensions

Why this matters

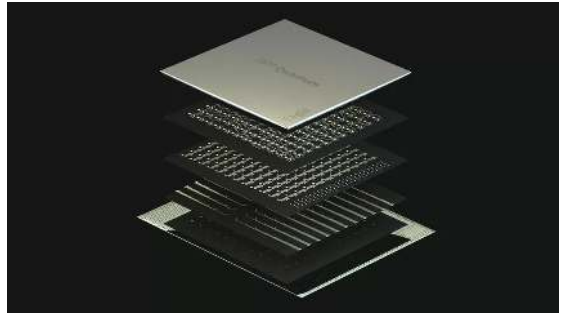
The problems that **resist** classical compression are exactly where quantum computers could provide an advantage.

Fairly recent example

IBM – June 14, 2023 in *Nature*

Evidence for the utility of quantum computing before fault tolerance

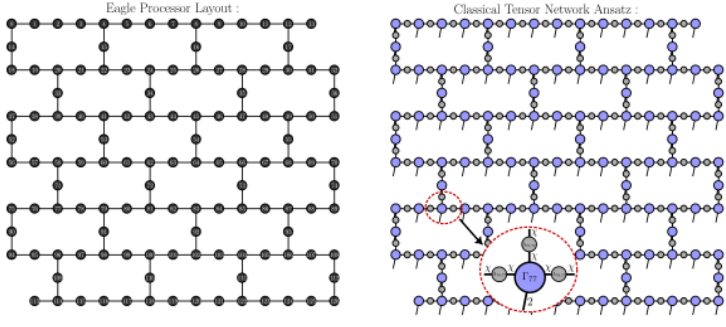




- 127 qubits based on superconducting loops
- First example of useful quantum supremacy?

Not yet

Flatiron – June 26, 2023 on `ArXiv`
Efficient tensor network simulation of IBM's Eagle [...]



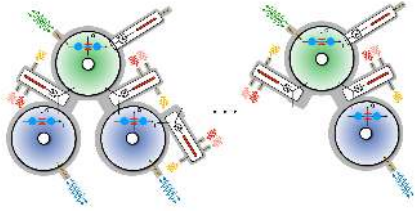
- Classical simulation of the same problem with a tensor network
- Better precision than the quantum computer!

Simulating small parts of realistic quantum computers

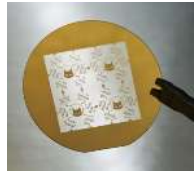
The few-large-body problem

Quantum computing in practice

Building quantum computers requires **simulating their components** classically – crucial for device **calibration**.



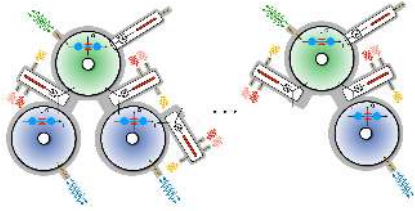
from Guillaud and Mirrahimi, PRX 2019



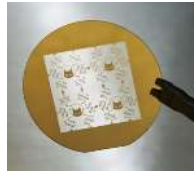
Alice & Bob cat-qubit chip

Quantum computing in practice

Building quantum computers requires **simulating their components** classically – crucial for device **calibration**.



from Guillaud and Mirrahimi, PRX 2019



Alice & Bob cat-qubit chip

Simulation challenges in exact few-body simulations

Already for elementary building blocks!
Memory cost for n "open" qudits $= d^{2n}$

- $d \simeq 100$ for cat-qubits
- $d \simeq 1000$ for GKP qubits

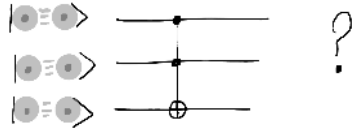
Memory bound, computationally cheap (generators are sparse).

The problem: large local Hilbert space

The vector space of n "cat qubits" with cutoff $N_c \sim 100$:

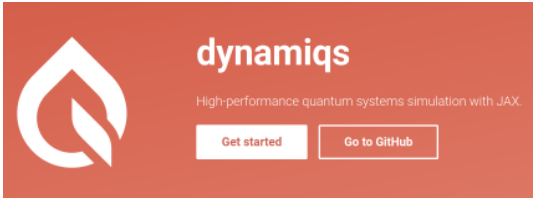
$$|\mathscr{H}| = N_c^n \text{ instead of } 2^n \text{ for "exact" qubits}$$

Simulating basic operations (e.g. Toffoli gate) is already hard:



State of the art in bruteforce: Dynamiqs

Smart use of hardware + autodiff + batching → `dynamiqs`
[started at Alice & Bob, now open source]



Close to the numerical optimum for the exact (uncompressed) problem.

Recent developments:

- GPU acceleration on H100 – significant speedup for Lindblad solves
- Batching + autodiff enable parameter sweeps and gradient-based optimization
- Used routinely for chip design/calibration at Inria/Alice&Bob

Compression: an active frontier

Small exact simulations: memory cost N_c^{2n}

- Mostly solved – well implemented in Dynamiqs
- **At the limit**: 8 x H100 required for some problems

Compression: an active frontier

Small exact simulations: memory cost N_c^{2n}

- Mostly solved – well implemented in Dynamiqs
- **At the limit**: 8 x H100 required for some problems

Larger compressed simulations: remove memory bound

- Local compression: reduce N_c (dynamically shifted Fock, quantics TT)
- Global compression: reduce n (tensor networks: MPS, MPO, MPDO for inter-qubit entanglement)
- Active research – nothing production-ready yet

This is where HPC resources are **mostly used** today for quantum computing R&D.

Conclusion

- The quantum many-body problem is fundamentally **memory-bound**
- Quantum computers turn this hardness into power (Shor, simulation)
  - But they are extremely slow – only exponential speedups matter
- Classical compression (tensor networks) fights back
  - **Compute-bound**: needs HPC, GPUs, dense linear algebra
  - Not for everything: maps the frontier of quantum advantage
- Simulating bricks of quantum hardware is itself an HPC problem
  - GPU-accelerated "exact" Lindblad solvers - `dynamiqs`
  - Compression techniques just starting
- Quantum computing competes with classical compression (for direct many-body solve) and needs it (for few-body design)