Our Quantum Future

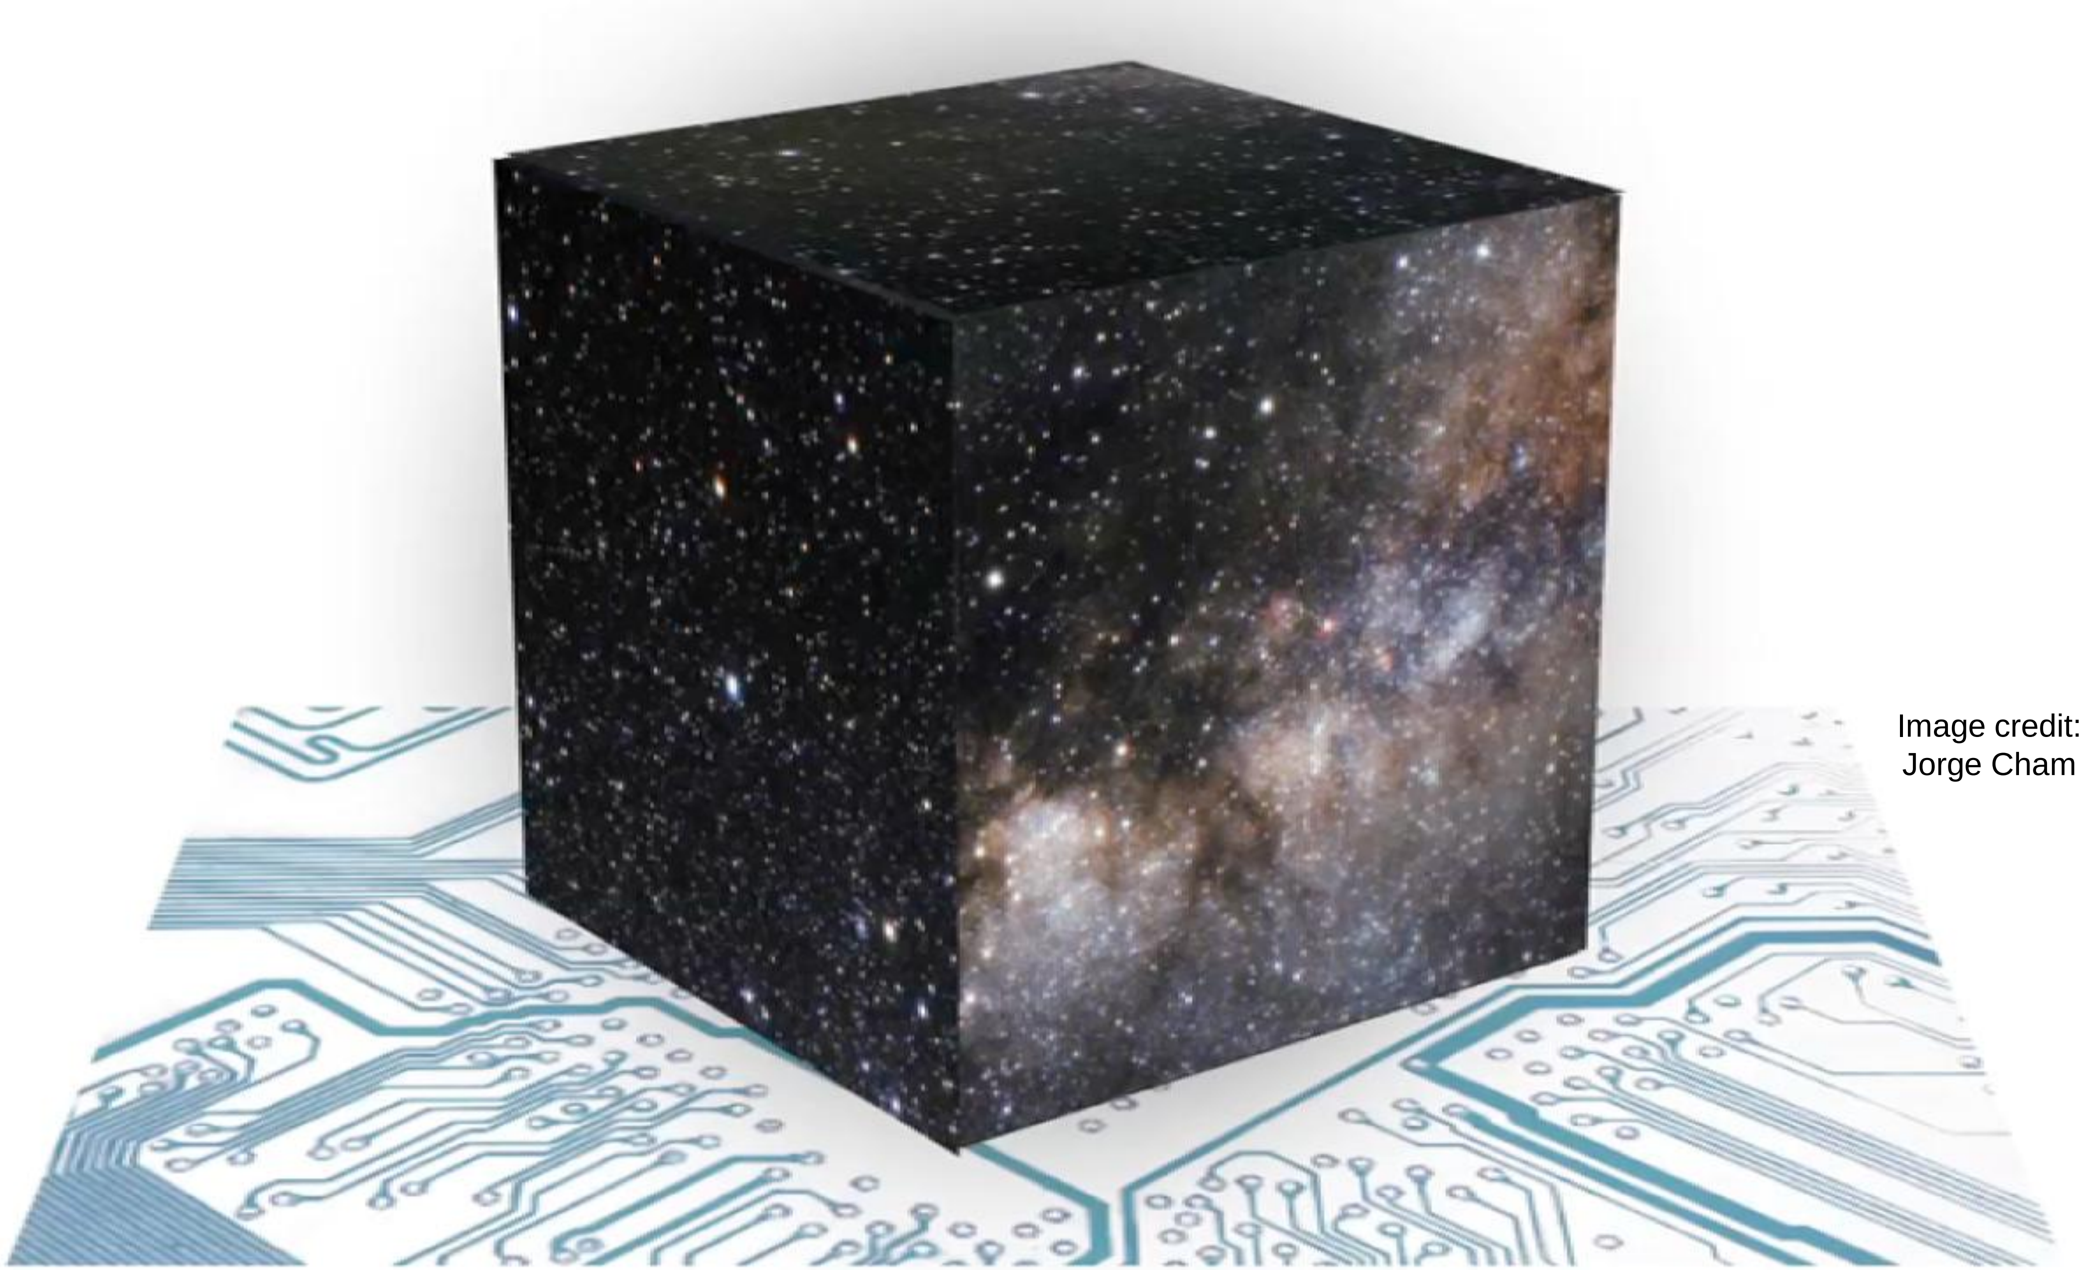

Image credit:
Jorge Cham

INSTITUTE FOR QUANTUM INFORMATION AND MATTER

John Preskill
Caltech X UCLA
4 April 2025

@preskill

Frontiers of Physics

short distance

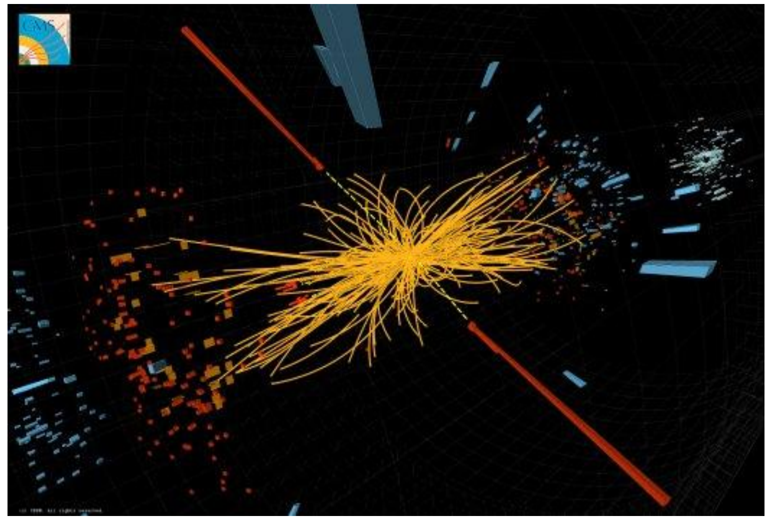

- Higgs boson
- Neutrino masses
- Supersymmetry
- Quantum gravity
- String theory

long distance

- Large scale structure
- Cosmic microwave background
- Dark matter
- Dark energy
- Gravitational waves

complexity

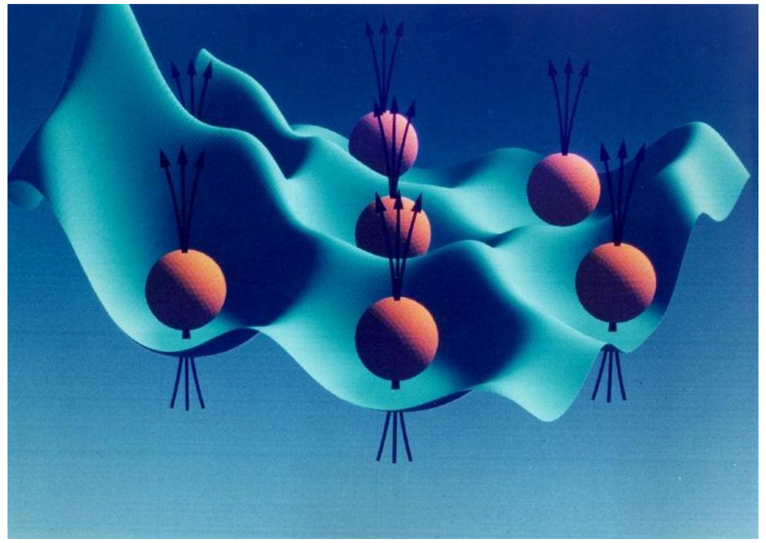

- “More is different”
- Many-body entanglement
- Phases of quantum matter
- Quantum computing
- Quantum spacetime

Open Questions

How will we scale up to quantum computing systems that can solve hard problems?

What are the important applications for science and for industry?

Applications

Quantumly easy.

Classically hard.

Interesting.

Scaling

Devices.

Error correction.

Systems engineering.

Two fundamental ideas

(1) *Quantum complexity*

Why we think quantum computing is powerful.

(2) *Quantum error correction*

Why we think quantum computing is scalable.

Quantum entanglement

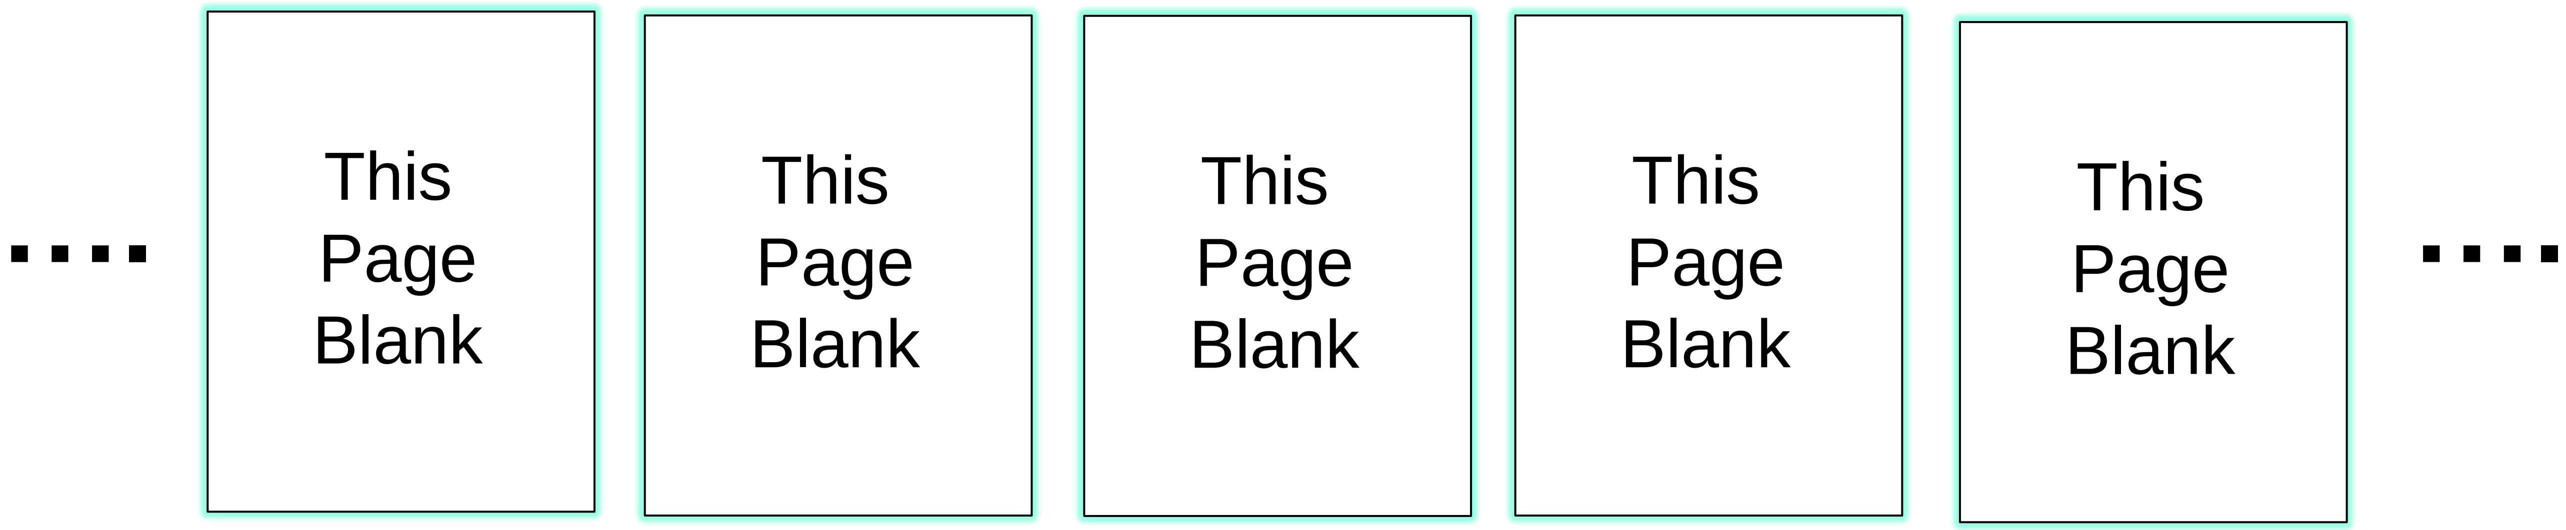

Nearly all the information in a typical entangled “quantum book” is encoded in the correlations among the “pages”.

You can't access the information if you read the book one page at a time.

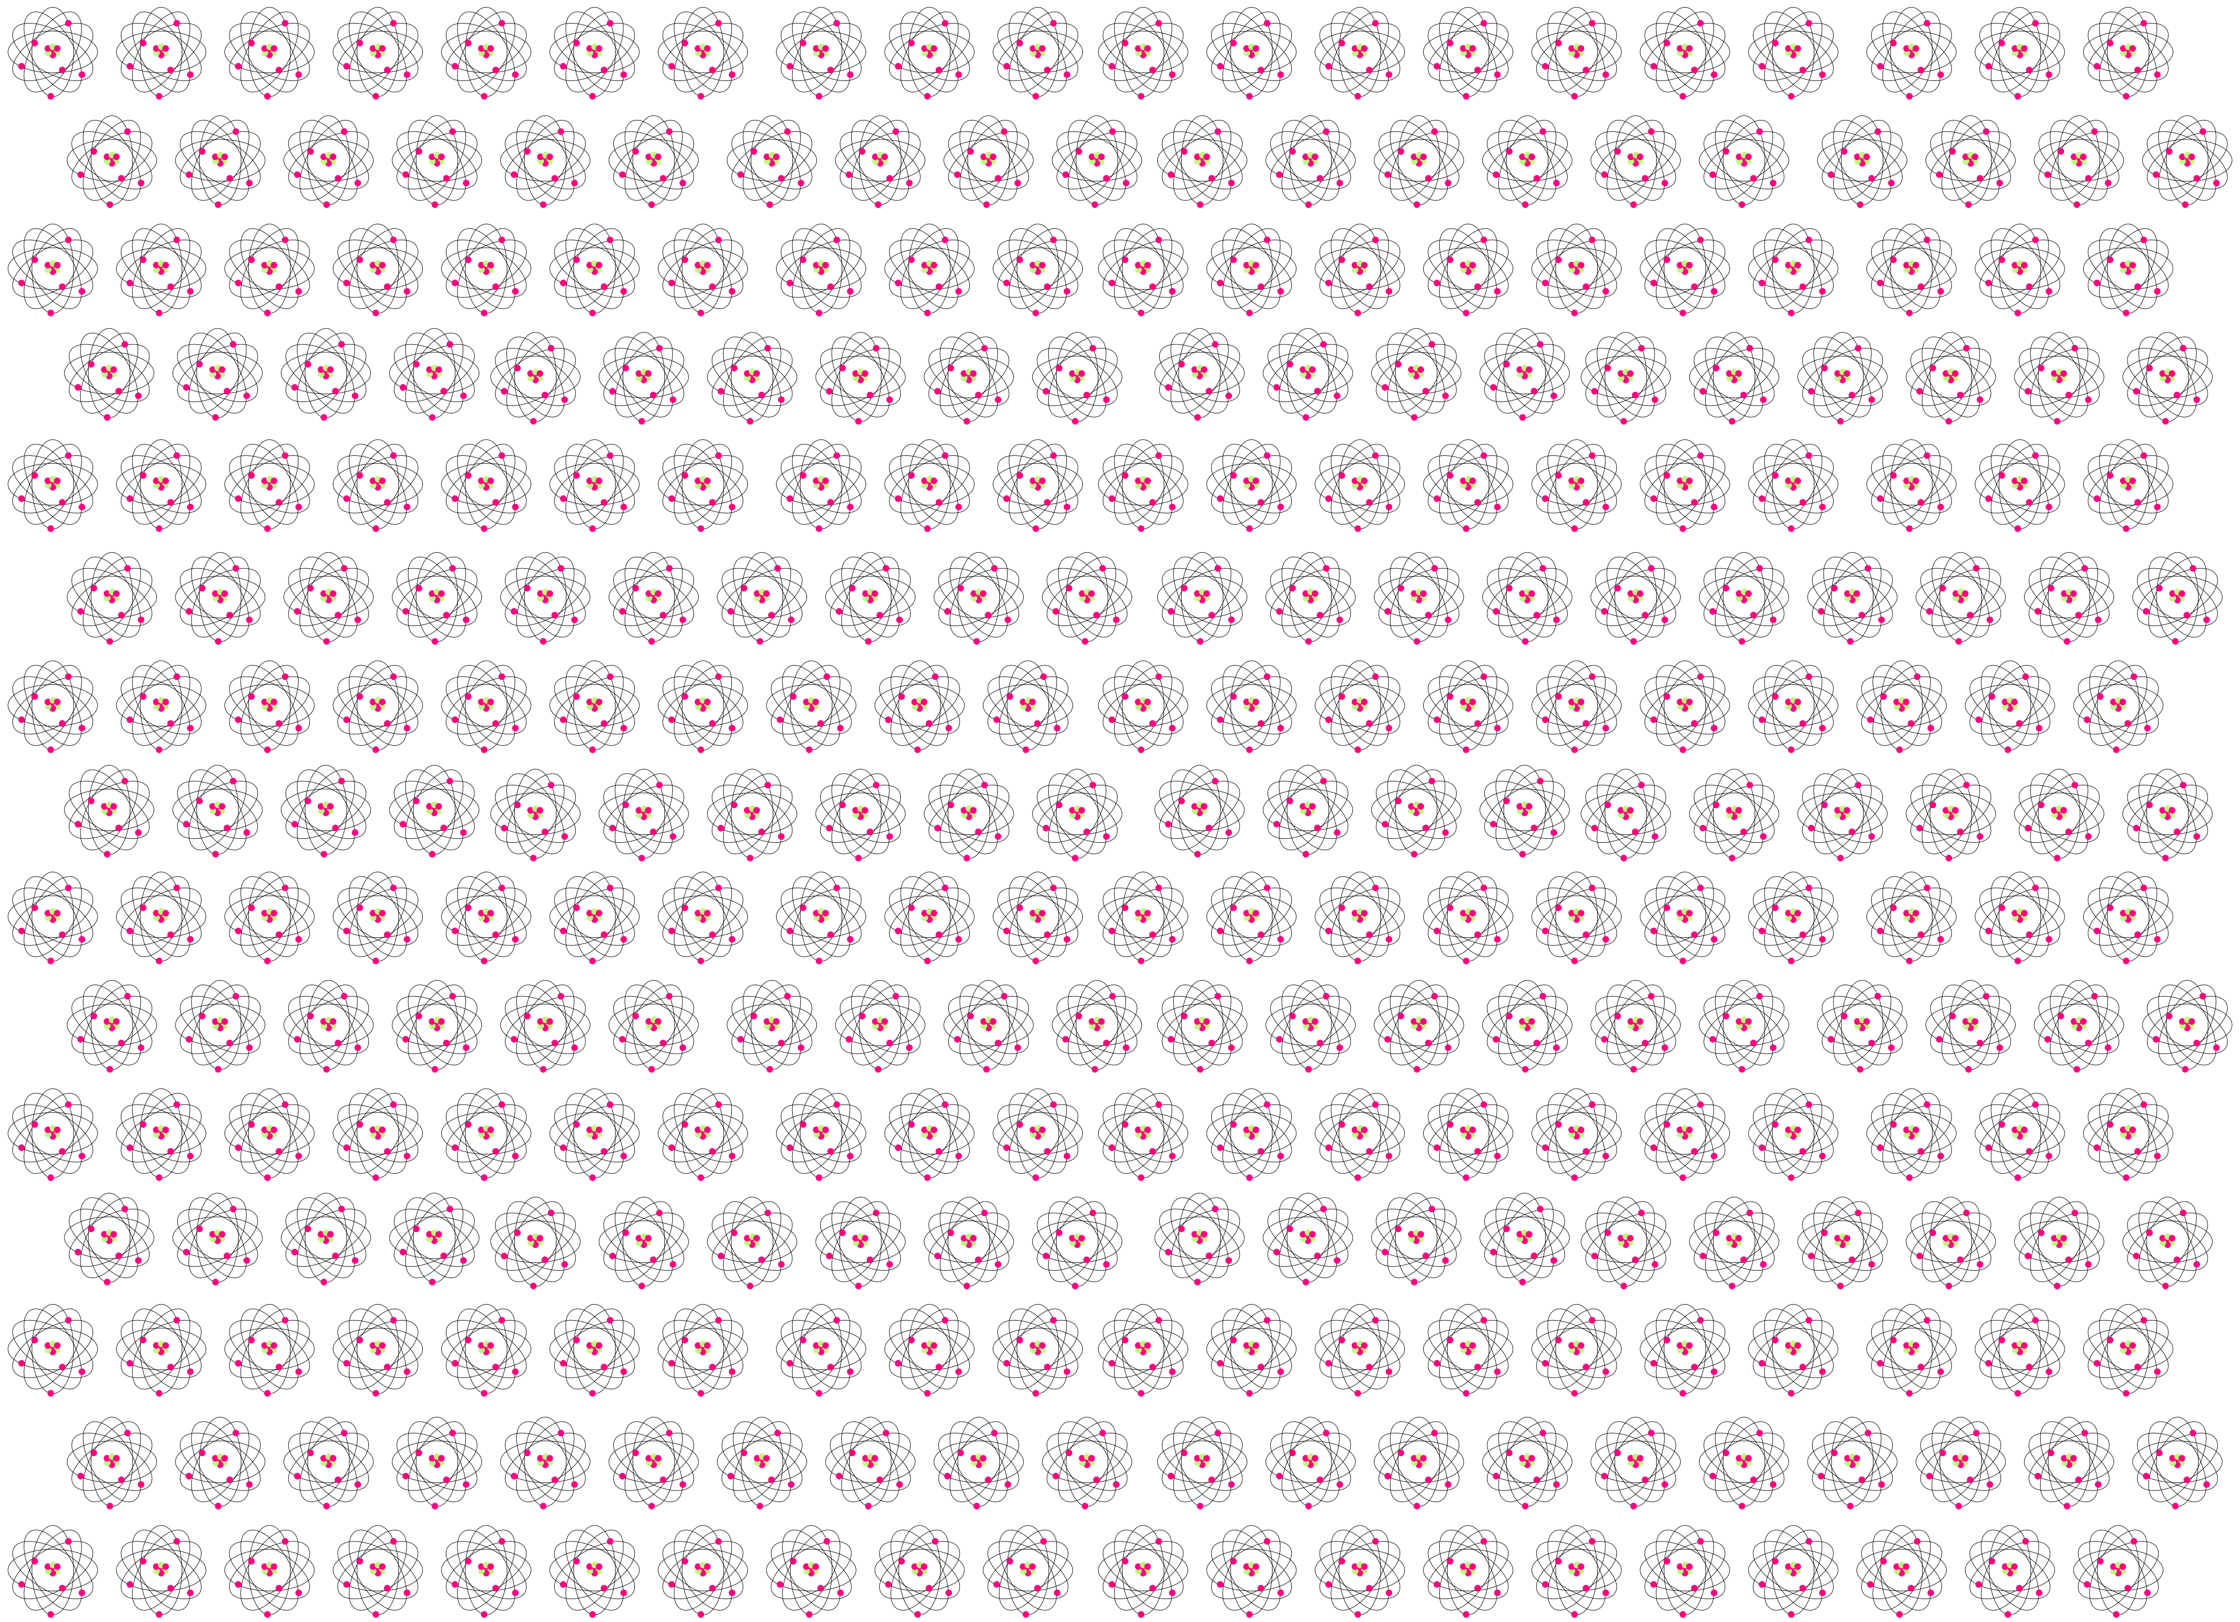

A complete description of a typical quantum state of just 300 qubits requires more bits than the number of atoms in the visible universe.

Why we think quantum computing is powerful

(1) Some problems are believed to be hard for conventional ("classical") computers, yet are easy for quantum computers. Factoring is the best known example.

(2) We don't know how to simulate a quantum computer efficiently using a classical computer.

But … the power of quantum computing is limited. For example, we don't believe that quantum computers can efficiently find exact solutions to worst-case instances of NP-hard optimization problems (e.g., the traveling salesman problem).

Problems

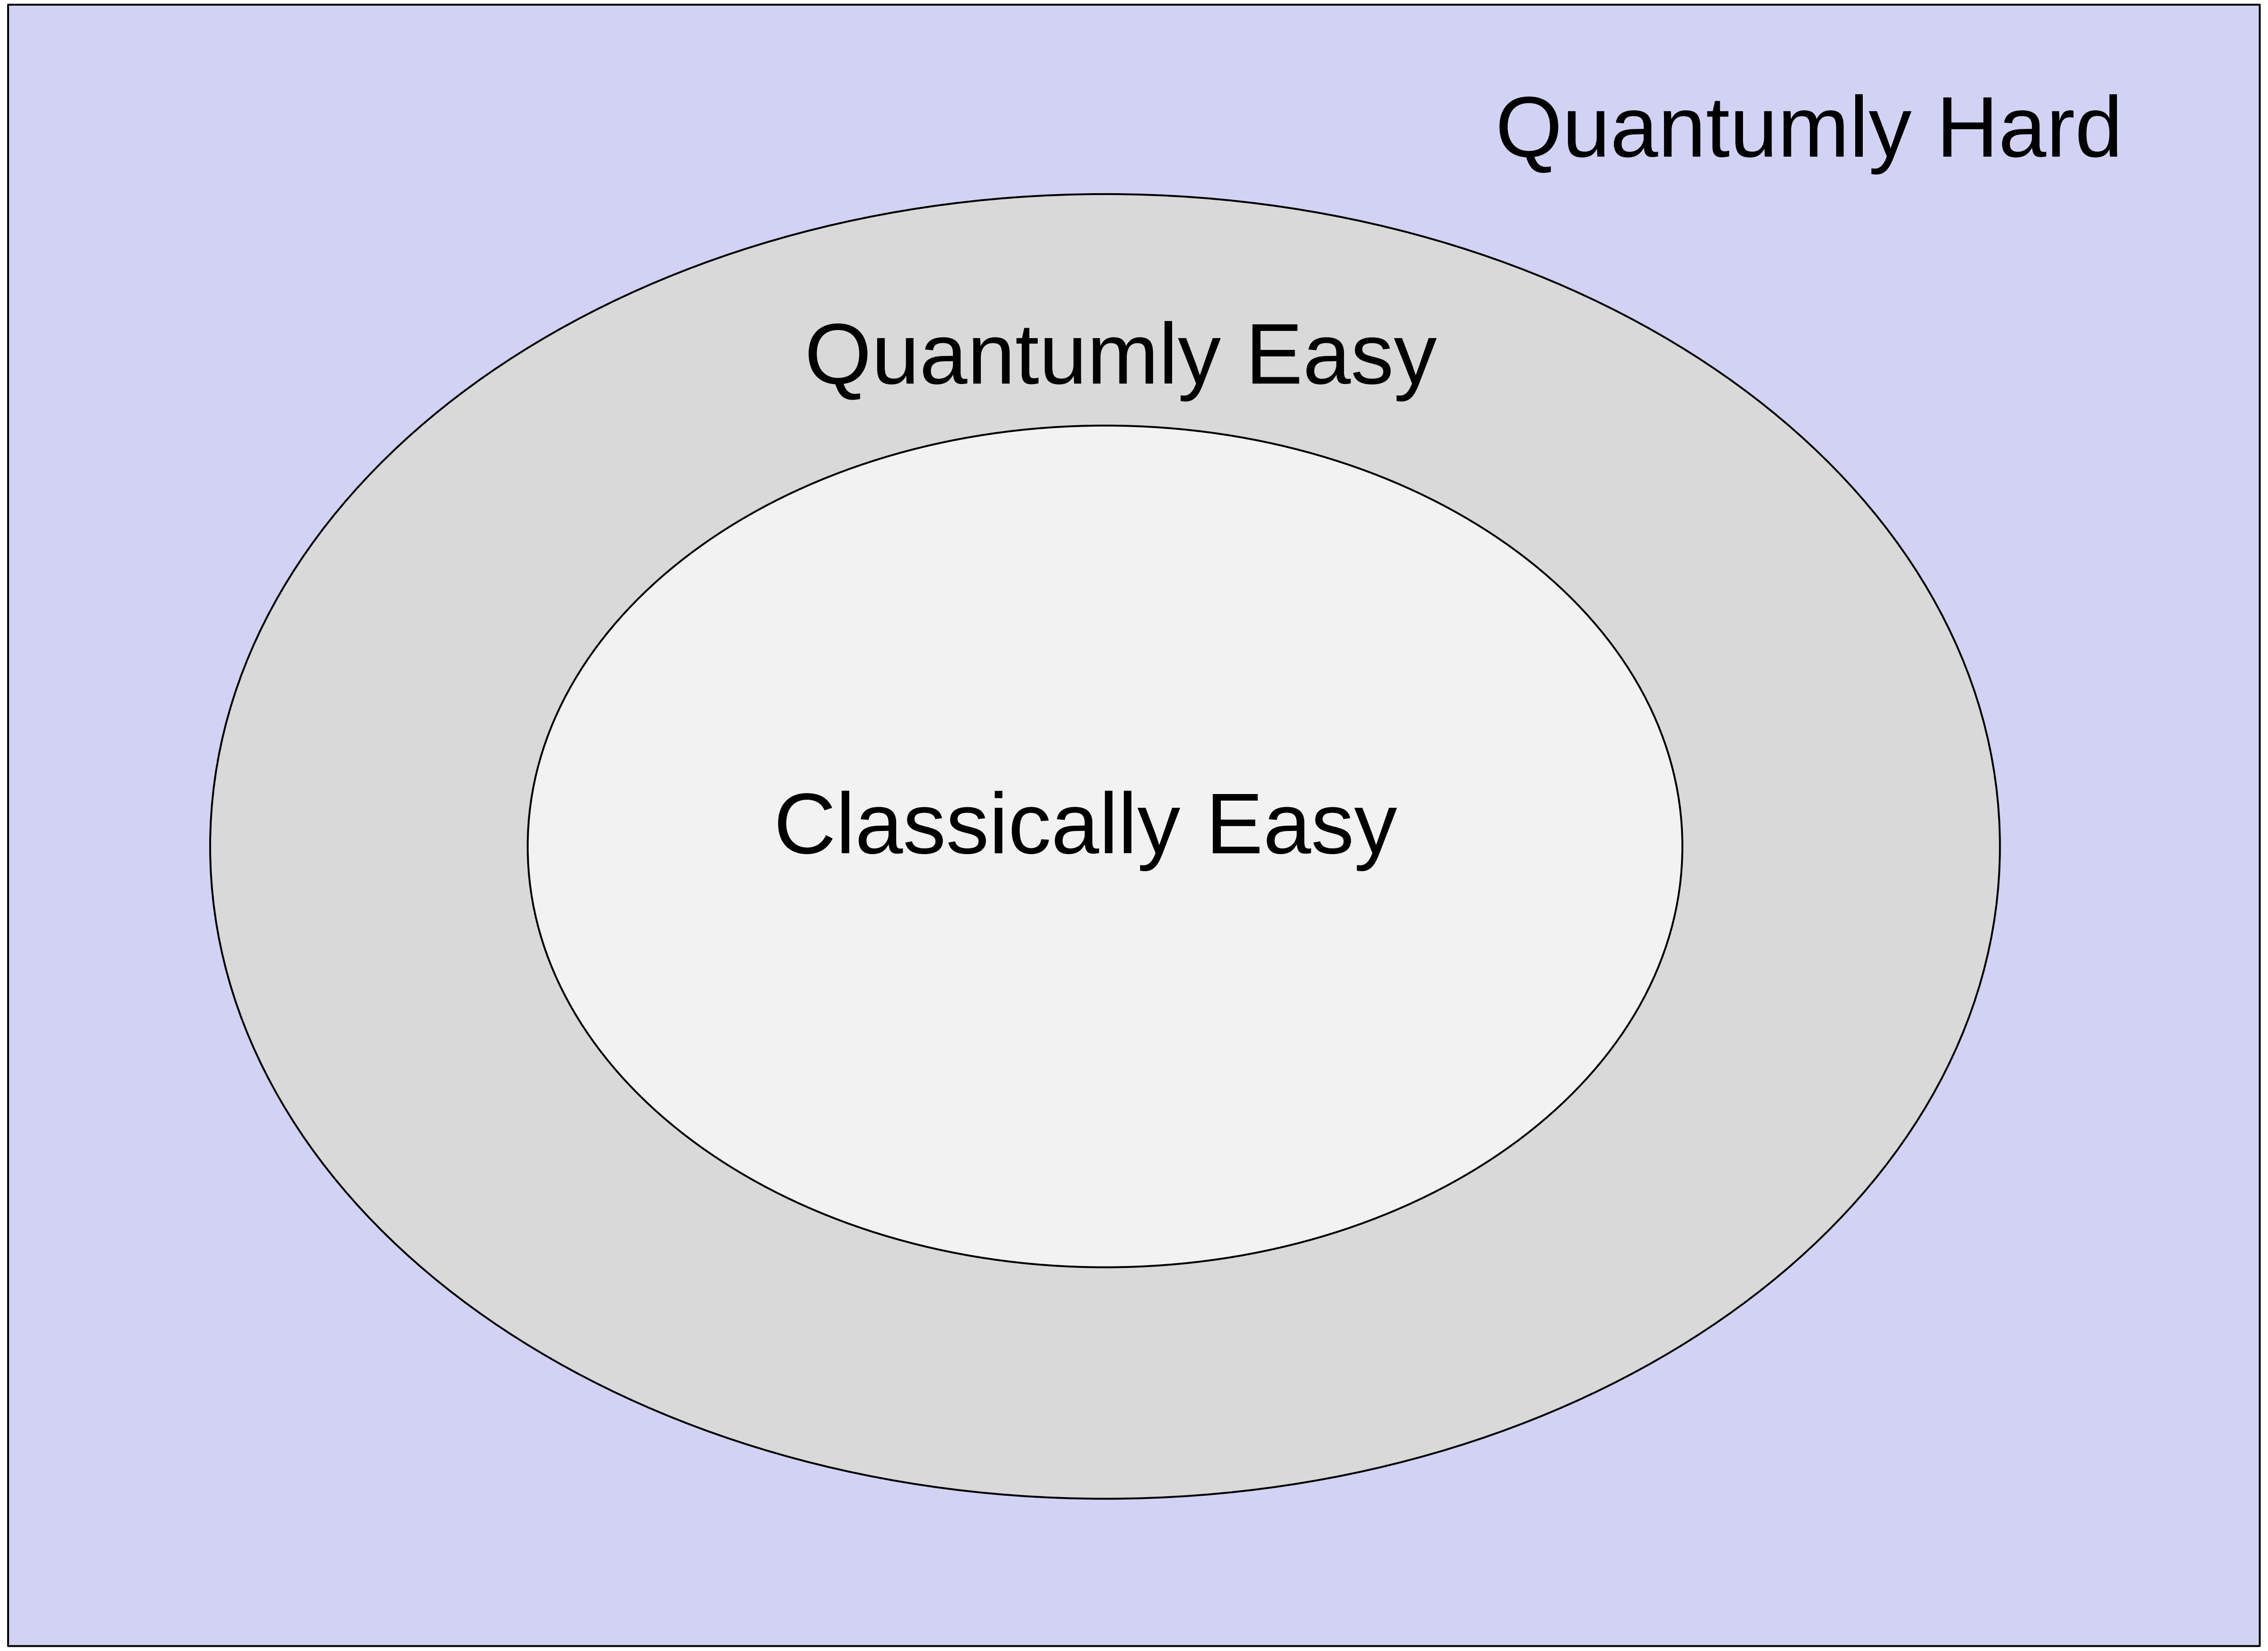

Problems

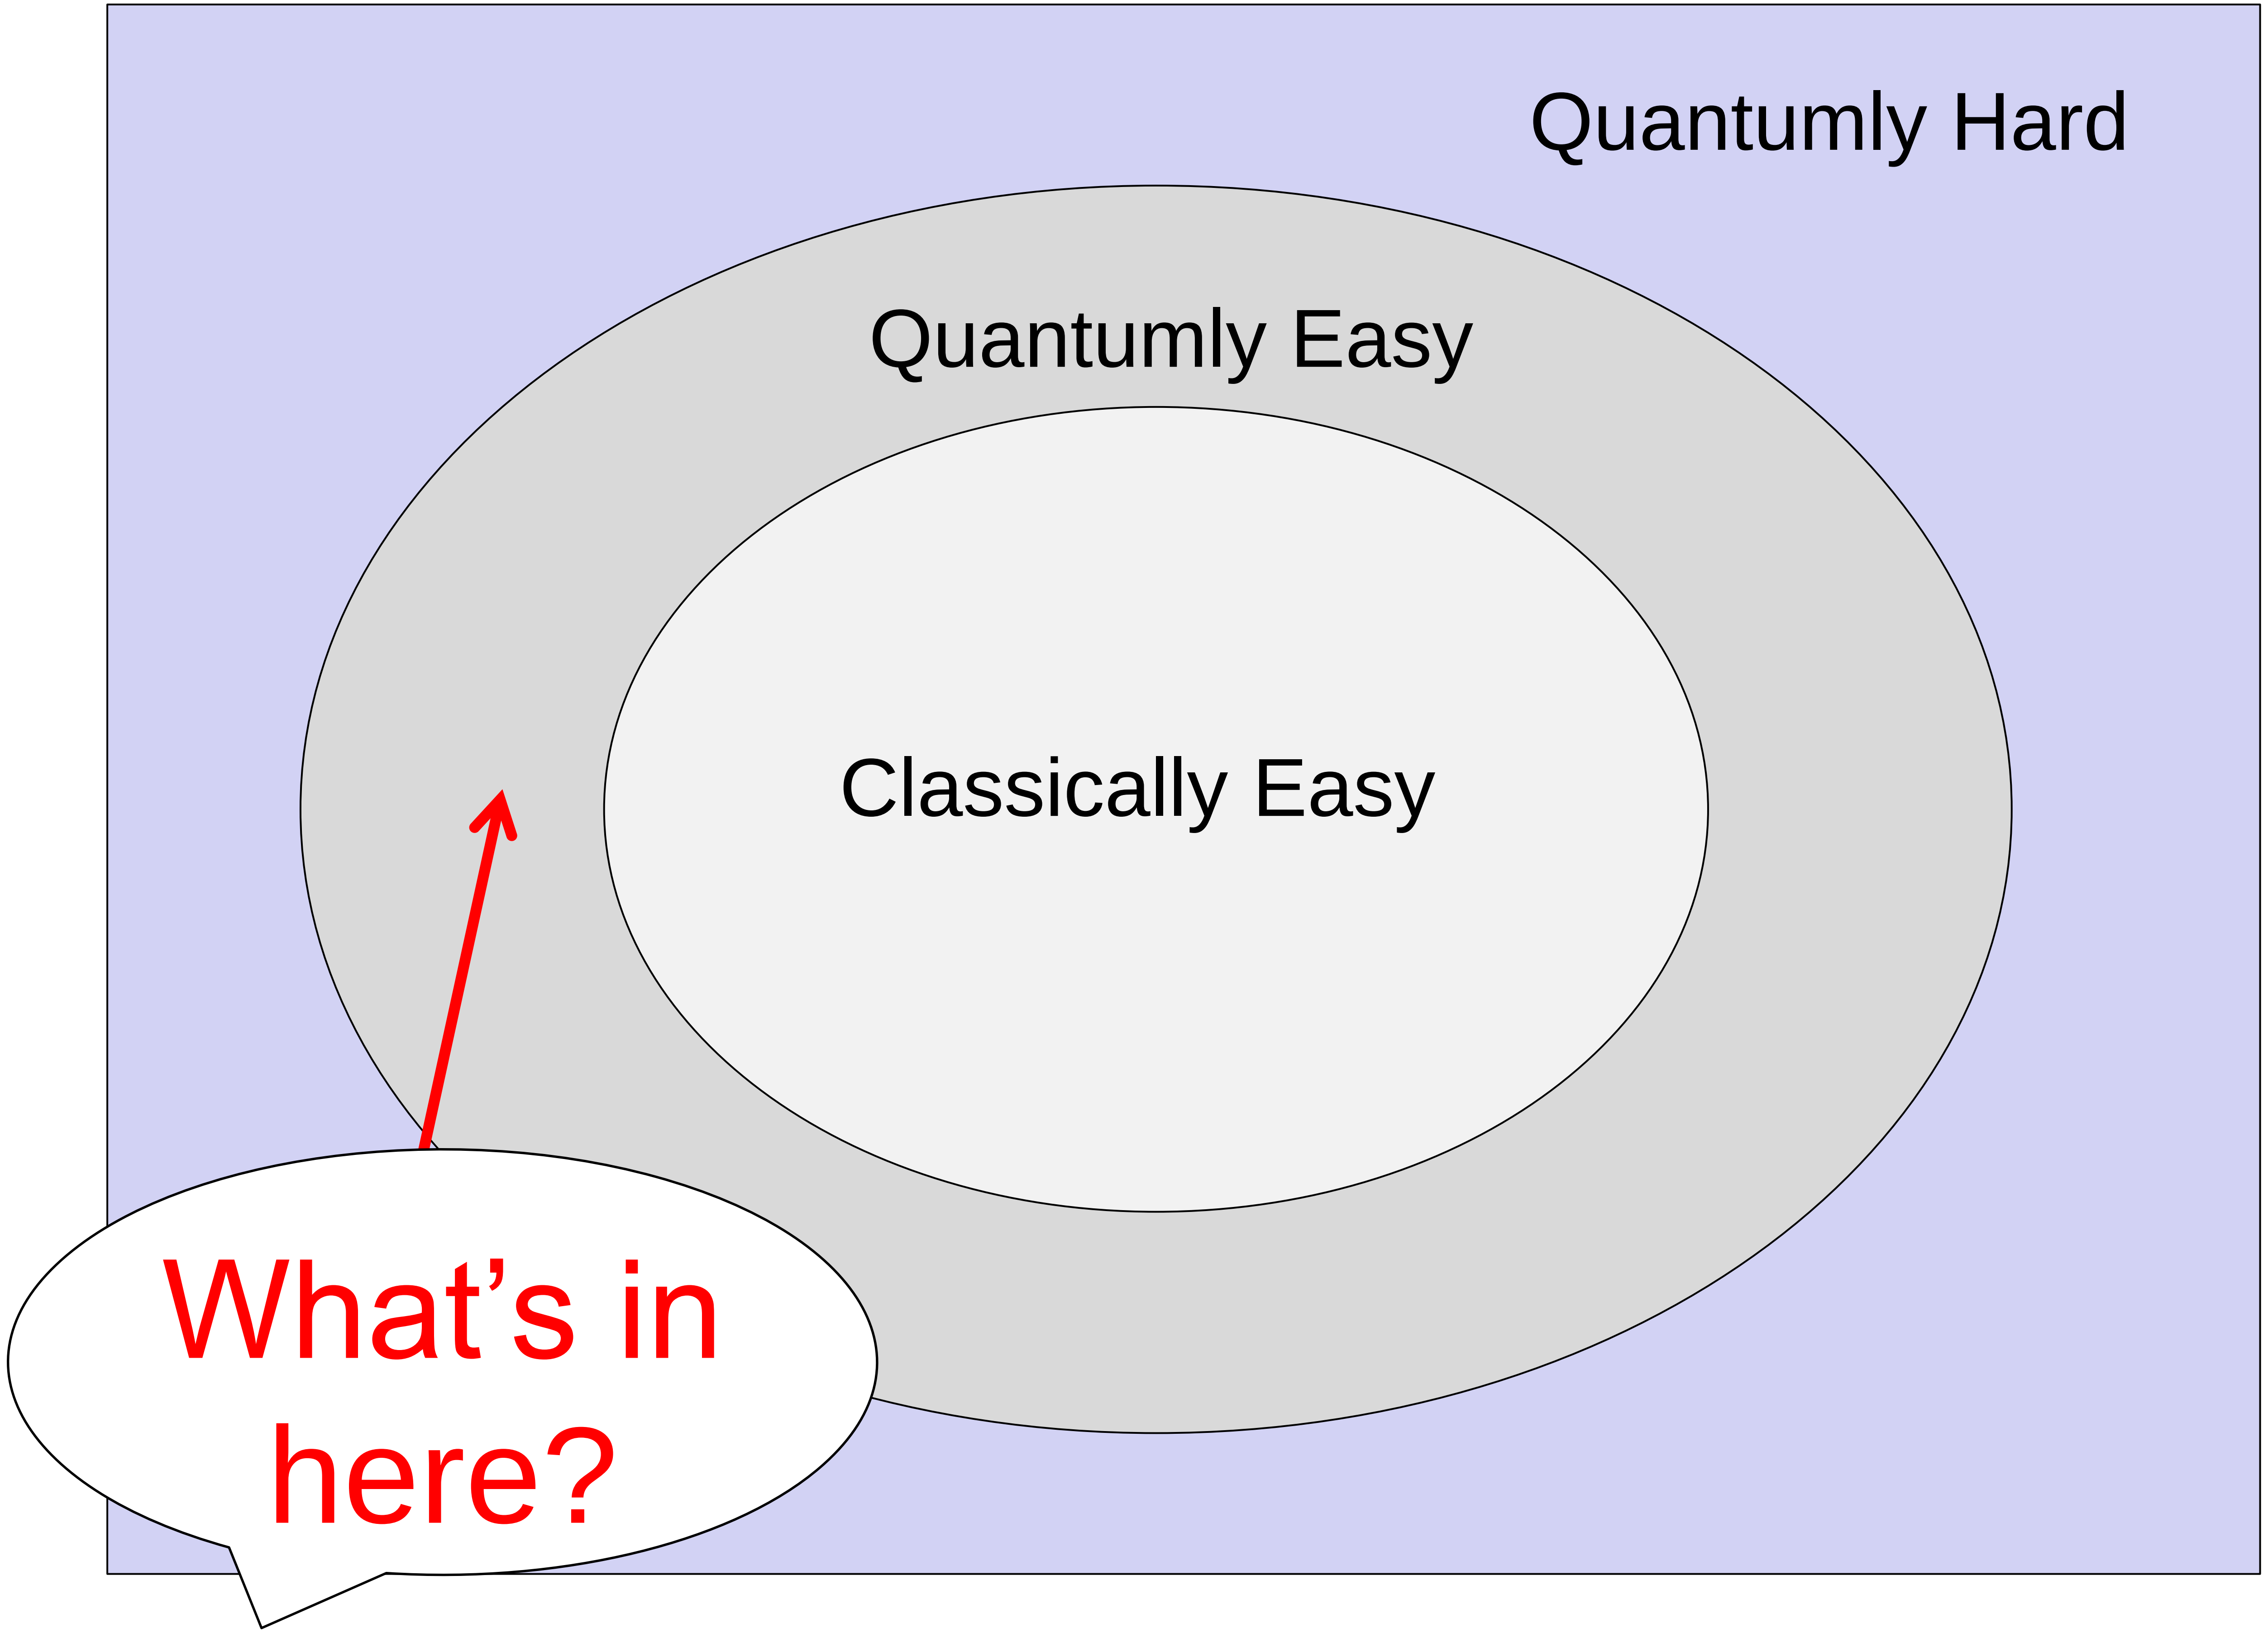
Quantumly Hard
Quantumly Easy
Classically Easy
What’s in here?

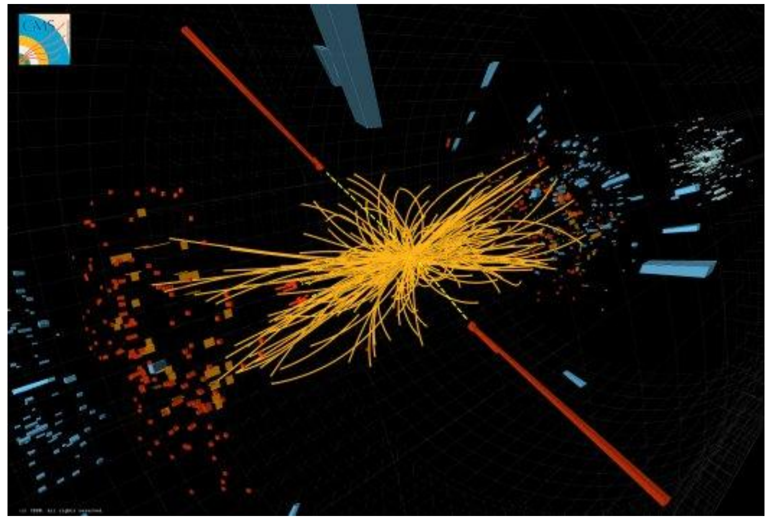
particle collision

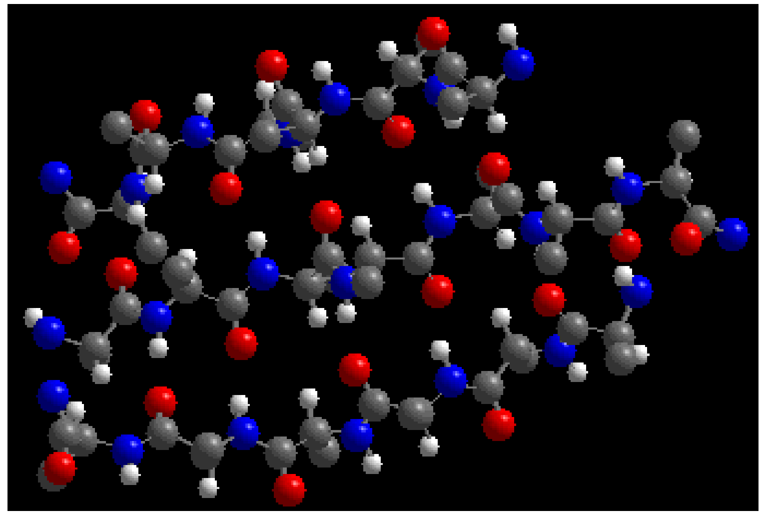
molecular chemistry

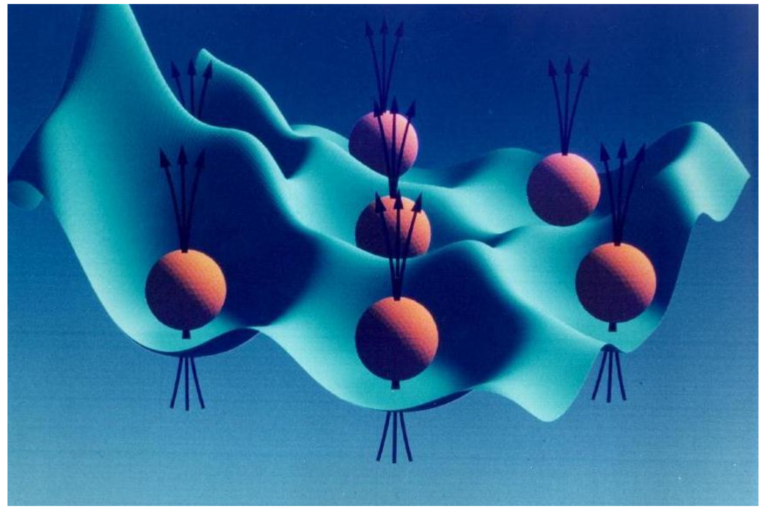
entangled electrons

(We expect that) a quantum computer can simulate efficiently any physical process that occurs in Nature.

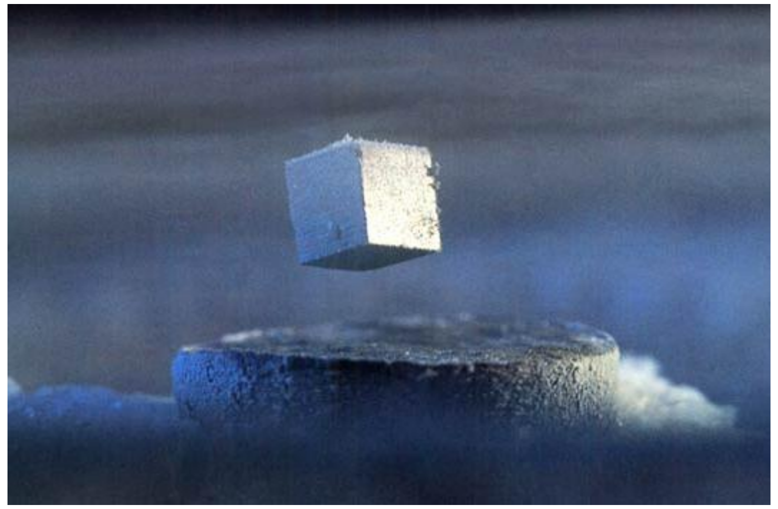
superconductor

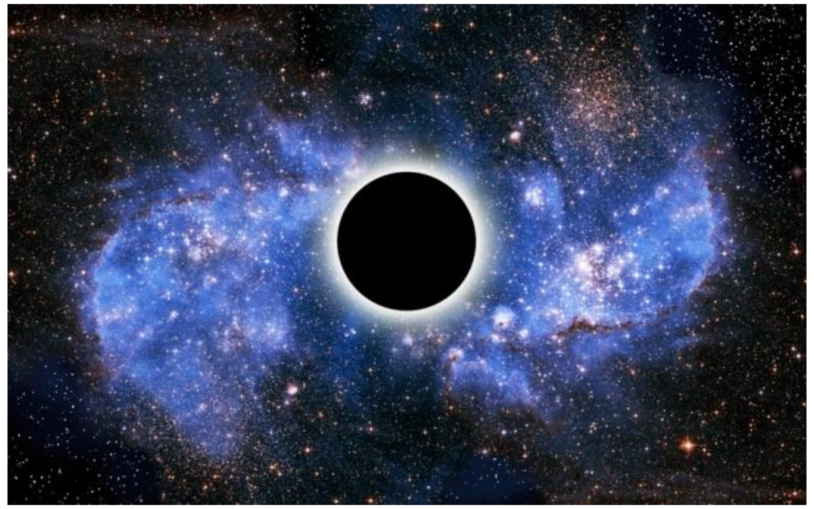
black hole


early universe

Why quantum computing is hard

We want qubits to interact strongly with one another.

We don't want qubits to interact with the environment.

Except when we control or measure them.

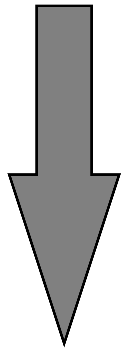

ERROR!

To resist decoherence, we must prevent the environment from "learning" about the state of the quantum computer during the computation.

Quantum error correction

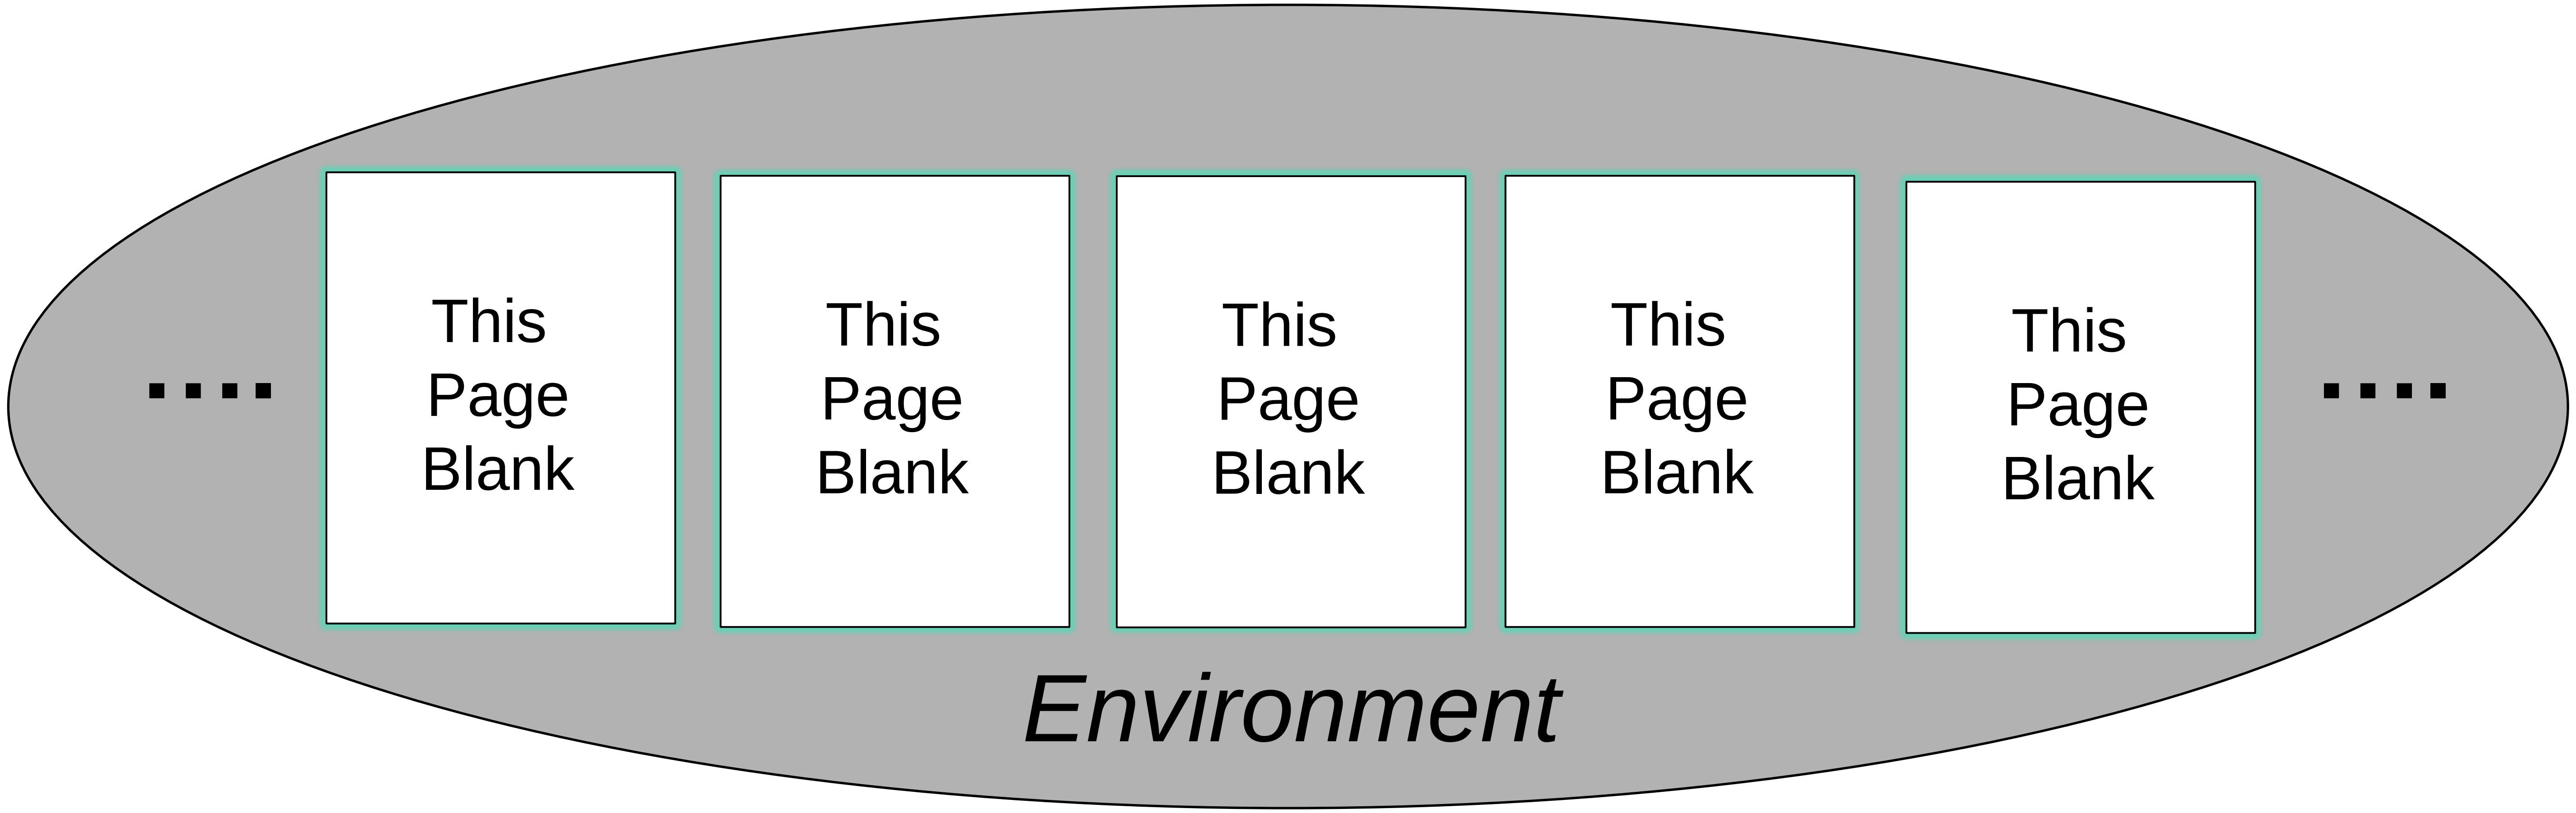

The protected "logical" quantum information is encoded in a highly entangled state of many physical qubits.

The environment can't access this information if it interacts locally with the protected system.

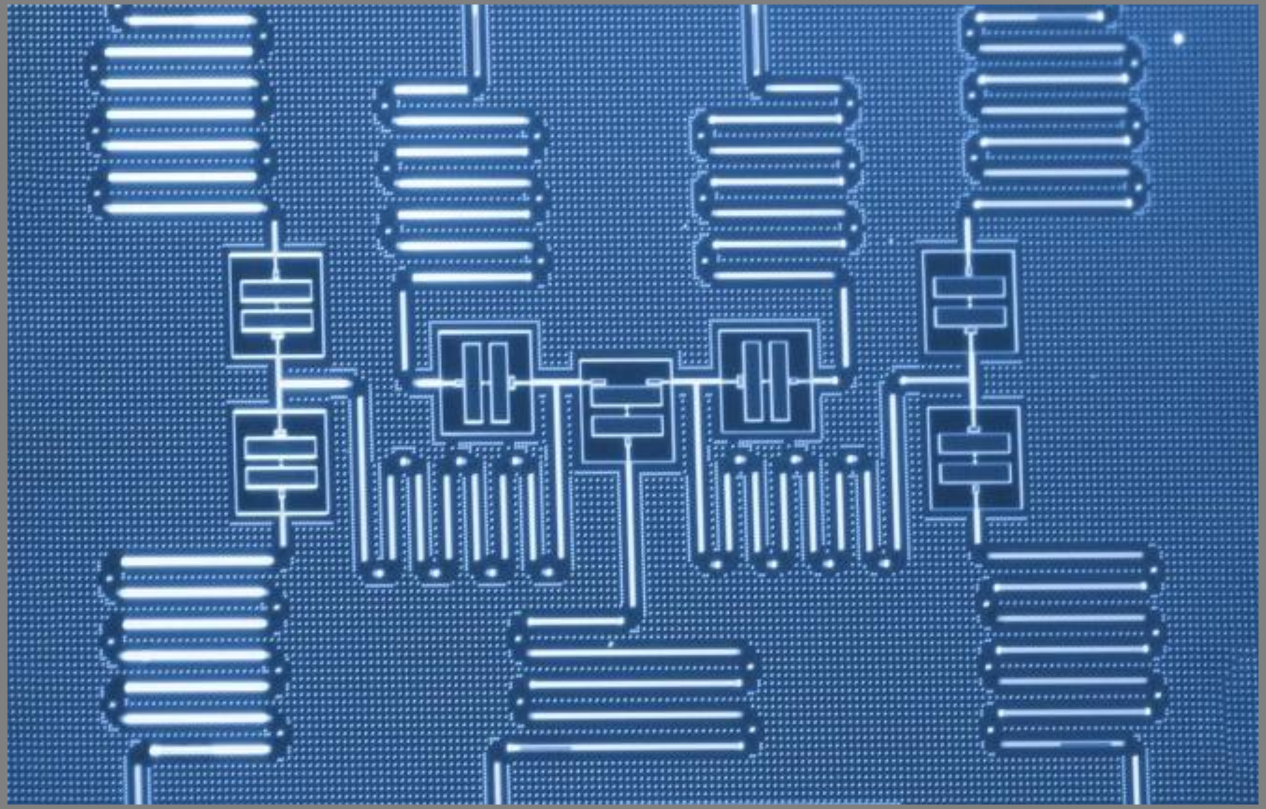

superconducting qubits

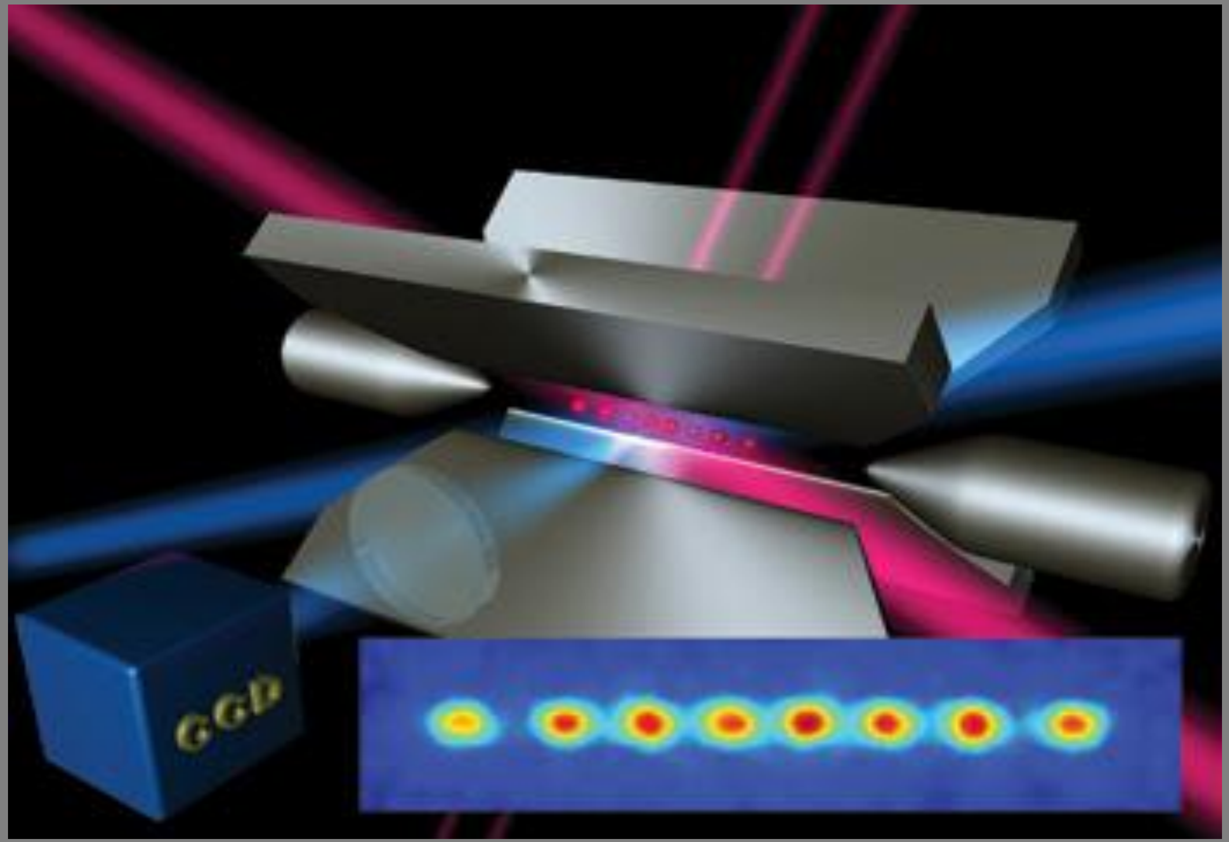

trapped atoms/ions

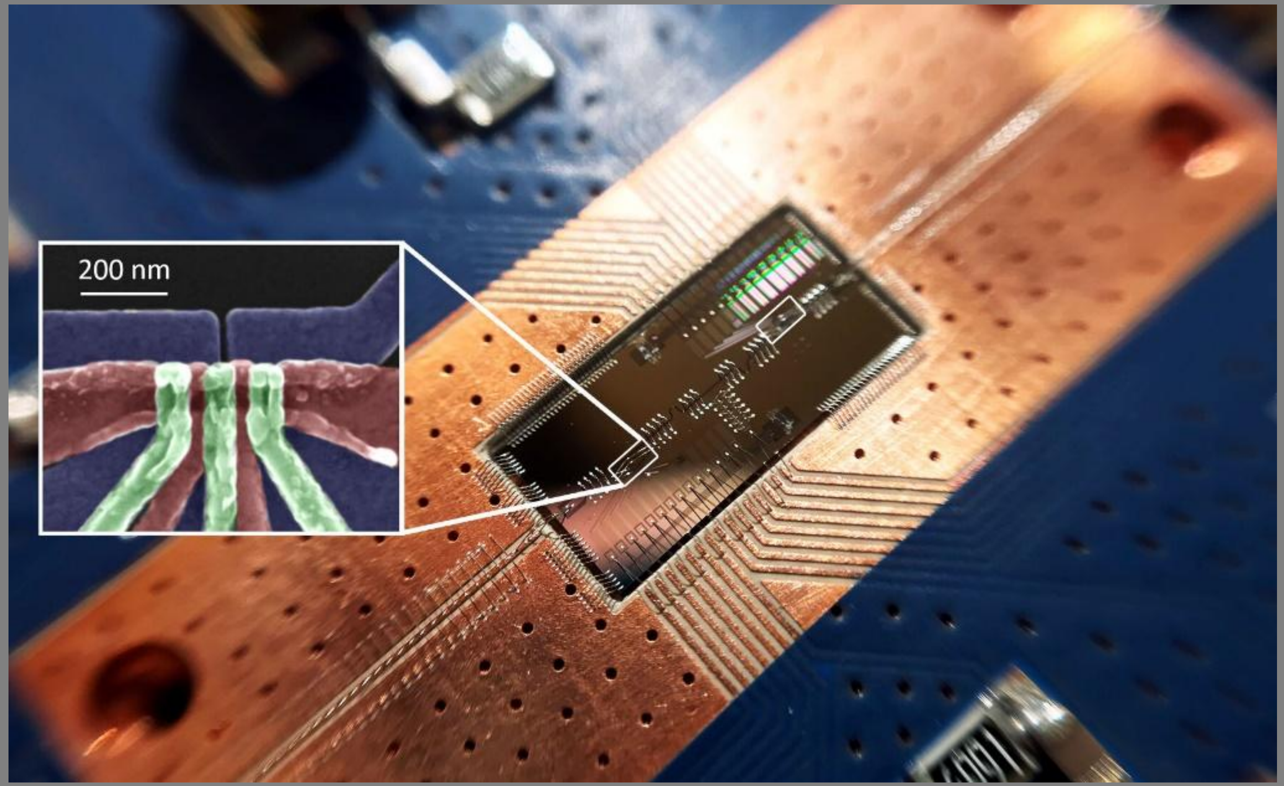

silicon spin qubits

photonics

Atomic Ions

Tens of qubits in a (linear) trap.

Stable laser → state preparation, single-qubit gates, readout.

Manipulate normal modes of vibration → two-qubit gates, all-to-all coupling (tens of microseconds).

Scaling: modular traps with optical interconnects or ion shuttling.

Superconducting Circuits

~ 100 qubits in a two-dimensional array with nearest-neighbor coupling.

Transmons: artificial atoms, carefully fabricated and frequently calibrated.

Microwave resonator for readout, microwave pulses for single-qubit gates.

Two-qubit gates via tunable frequency, tunable couplers, or cross-resonance drive (tens of nanoseconds).

Scaling: modular devices, microwave control lines, materials, fabrication, alternative qubit designs.

Neutral Atoms

100s to 1000s of atoms in optical tweezer arrays. Highly excited Rydberg states for entangling operations.

Atoms are movable, hence no geometrical constraints.

Global control, move to processing zones for local gates.

Atomic movement and readout are relatively slow.

Continuous loading of fresh atoms under development.

Noisy Intermediate Scale Quantum (NISQ) Era

What we have now.

NISQ is valuable for scientific exploration. But there is no proposed application of NISQ computing with *commercial* value for which quantum advantage has been demonstrated *when compared to the best classical hardware running the best algorithms for solving the same problems*.

What we can reasonably foresee.

Nor are there persuasive theoretical arguments indicating that commercially viable applications will be found that do *not* use quantum error-correcting codes and fault-tolerant quantum computing.

Fault-tolerant Application Scale Quantum (FASQ) Era

What we want to have.

-- Quantum computers running a wide variety of useful computations.

-- Machines that can execute of order 10^{12} quantum operations ("teraquop machines").

-- This requires improving quantum gate error rates by about 9 orders of magnitude beyond the current state of the art.

-- Quantum error correction and will be essential for crossing the chasm from NISQ to FASQ. We may need devices with millions of physical qubits.

When will we have it?

No one knows. It might take decades.

Quantum algorithms:

A survey of applications and end-to-end complexities

Alexander M. Dalzell*[1], Sam McArdle*[1], Mario Berta[1,2,3], Przemyslaw Bienias[1],
Chi-Fang Chen[1,4], András Gilyén[5], Connor T. Hann[1], Michael J. Kastoryano[1,6],
Emil T. Khabiboulline[1,7], Aleksander Kubica[1], Grant Salton[1,4,8], Samson Wang[1,3]
and Fernando G. S. L. Brandão[1,4]

[1] *AWS Center for Quantum Computing, Pasadena, CA, USA*
[2] *Institute for Quantum Information, RWTH Aachen University, Aachen, Germany*
[3] *Imperial College London, London, UK*
[4] *Institute for Quantum Information and Matter, Caltech, Pasadena, CA, USA*
[5] *Alfréd Rényi Institute of Mathematics, Budapest, Hungary*
[6] *IT University of Copenhagen, Copenhagen, Denmark*
[7] *Department of Physics, Harvard University, Cambridge, MA, USA*
[8] *Amazon Quantum Solutions Lab, Seattle, WA, USA*

arXiv:2310.03011

Applications: Looking ahead

Optimization, finance, and machine learning. Typical quantum speedups are at best quadratic. Quantum advantage kicks in for very large problem instances and deep circuits.

Quantum many-body physics: Chemistry and materials. Hundreds of logical qubits, hundreds of millions of logical gates or more.

Quantum fault tolerance needed to run these applications. High cost in physical qubits and gates.

Logical gate speed is also important. Run time on the wall clock.

Quantum computing for chemistry and materials

Dirac (1929): "… equations much too complicated to be soluble."

Yet, heuristic classical algorithms are often very successful, and these methods are continually improving.

Quantum computing targets the relatively small "strongly correlated" corner of chemistry and materials science, where such methods falter.

How useful are quantum computers in physically relevant situations that are beyond the reach of classical methods?

Artificial intelligence may drive future progress in (strongly correlated) chemistry and materials science. Eventually, quantum computers can accelerate progress by providing abundant training data.

Quantum algorithms that simulate cooling processes

Polynomial-time quantum algorithms that simulate dissipative dynamics.

Find local minima of the energy efficiently starting with arbitrary states.

Rapid preparation of ground states for some families of Hamiltonians. [arXiv:2503.15827]

In some cases, finding a local minimum is classically hard and quantumly easy. [arXiv:2309.16596]

Simulating quantum dynamics

Classical computers are especially bad at simulating quantum *dynamics.* Quantum computers will have a big advantage.

But …

Many-body localized (MBL) systems, which equilibrate slowly, are only slightly entangled, and *might* therefore be easy to simulate classically.

Systems with strong quantum chaos become highly entangled and are therefore hard to simulate classically. But they might be boring – perhaps they quickly converge to thermal equilibrium and after that "nothing interesting" happens.

If we ask the right questions, scientifically informative surprises should be expected (quantum many-body scars, diabatic evolution in quantum spin liquids, …)

Quantum + Machine Learning

Succinct representations of big data. Severe input/output bottlenecks.

Quantum speedups in matrix inversion. Matrices should be sparse, well conditioned, and not low rank.

Variational quantum ML. Hampered by barren plateaus.

Quantum data + classical ML. Generalize to new quantum systems.

Quantum processing of quantum data. Exponential advantage in some cases.

Exponential speed-ups for optimization?

Goal is finding efficiently a good *approximate* solution that would take exponential time to find classically. (Finding exact solutions is too hard even for quantum.)

Decoded Quantum Interferometry (DQI) uses the quantum Fourier transform to map a problem which seems to be hard to solve classically to a problem that is easy to solve classically.

Example: Optimal Polynomial Intersection (OPI), a variant of polynomial interpolation. Practical uses are unclear. Requires fault tolerance.

A surprising application of the quantum Fourier transform.

How well can the classical team solve OPI?

[arXiv:2408.08292]

Overcoming noise in quantum devices

Quantum error mitigation. Used effectively in current processors. Asymptotic overhead cost scales exponentially.

Quantum error correction. Asymptotic overhead cost scales polylogarithmically. Not yet effective in current processors.

What we need. Better two-qubit gate fidelities, *many* more physical qubits, and the ability to control them. Also fast gates, mid-circuit readout, feed-forward, reset.

Overhead cost of fault tolerance

$$P_{\text{logical}} \approx C\left(P_{\text{physical}} / P_{\text{threshold}}\right)^{(d+1)/2}$$

$$d = \sqrt{n}, \quad C \approx 0.1, \quad P_{\text{threshold}} \approx .01$$

Surface code

Suppose $P_{\text{physical}} = .001, P_{\text{logical}} = 10^{-11}$
$\Rightarrow d = 19, n = 361$ physical qubits per logical qubit
(plus a comparable number of ancilla qubits for syndrome measurement). (Improves to $d = 9$ for $P_{\text{physical}} = 10^{-4}$.)

Quantum error correction: What we want

Repeated rounds of accurate error syndrome measurement.

Quantum memory times that improve sharply as codes increase in size.

Accurate *real-time* decoding of error syndromes, scalable control.

Scalable efficient logical universal gates with (much) higher fidelity than physical gates.

Logical gate fidelities that improve sharply as codes increase in size.

Acceptable overhead cost in space and time.

Speed matters! (Superconducting ~1000 time faster than ions.)

Quantum error correction below the surface-code threshold

[*Google 2024*]

105 qubit Willow processor. Improved transmon lifetime, measurement error, leakage correction.

Millions of rounds of surface-code error syndrome measurement, each lasting ~1 microseconds (600 nanosecond measurement time).

Logical error rate for quantum memory improves by $\Lambda \approx 2$ when code distance increases by 2 (from 3 to 5 to 7).

Accurate *real-time* decoding of error syndromes for distance 3 and 5.

Repetition code: $\Lambda \approx 8.4$ up to d = 29.

Looking ahead: Better Λ, larger codes, high-fidelity logical two-qubit gates.

Toward better superconducting qubits

Transmons: improved materials and fabrication (*IBM*).

Cat qubits with highly biased noise (*Yale, Alice & Bob, AWS*).

Dual-rail qubits for erasure conversion (*Yale, AWS*).

GKP qubits in resonators, well protected against loss, providing soft information (*Yale, ETH*).

Fluxonium for strong anharmonicity and high fidelity: 2Q F > .999 (*MIT, Atlantic*).

Other error correction progress

Movable qubits for all-to-all coupling.

Harvard + MIT + QuEra: Circuit sampling with 48 logical qubits on a 280-qubit device. CCZ gates within an [[8,3,2]] code block.

Atom Computing + Microsoft: Bernstein-Vazirani algorithm with 28 logical qubits in a 256-qubit device. [[4,1,2]] subsystem code.

Quantinuum + Microsoft. Preparation of a cat state with 12 logical qubits on a 56-qubit device. CNOT within a [[16,4,4]] block and transversal between blocks.

Caveats: Few rounds of syndrome measurement and unscalable postselection.

Can movement be much faster?

Error correction and fault tolerance

Surface code. High error threshold, 2D layout, good enough decoders (?), high overhead cost.

High-rate quantum low-density parity-check (LDPC) codes. Geometrically nonlocal, better decoders needed, complex logical operations.

System design. Trade time for space to reduce local control requirements.

General principles. Space and time should be treated in a unified framework, logical operations performed via code deformations.

Megaquop Machine

Logical gate error rate ~ 10^{-6}. Not achievable without QEC.

Error mitigation will continue to be useful in the Megaquop era and beyond.

Beyond classical, NISQ, or analog. E.g., depth 10K and 100 (logical) qubits.

Tens of thousands of high-quality physical qubits.

When will we have it? Less than 5 years? What modality? Rydberg atoms?

What will we do with it? Quantum dynamics?

Commercial as well as scientific applications? [arXiv:2502.17368]

Co-design

Adapt the application and the error correction protocol to the hardware.

Adapt the hardware to the application and the error-correcting code.

Prospects for the next 5 years

Encouraging progress toward scalable fault-tolerant quantum computing.

Scientific insights enabled by programmable quantum simulators and circuit-based quantum computers.

Advances in quantum metrology from improved control of quantum many-body systems.

Prospects for the next 100 years

Past 100 years:

The relatively simple quantum behavior of weakly correlated particles like electrons, photons, etc.

Next 100 years:

The extraordinarily complex quantum behavior of many profoundly entangled particles.