Our Quantum Future

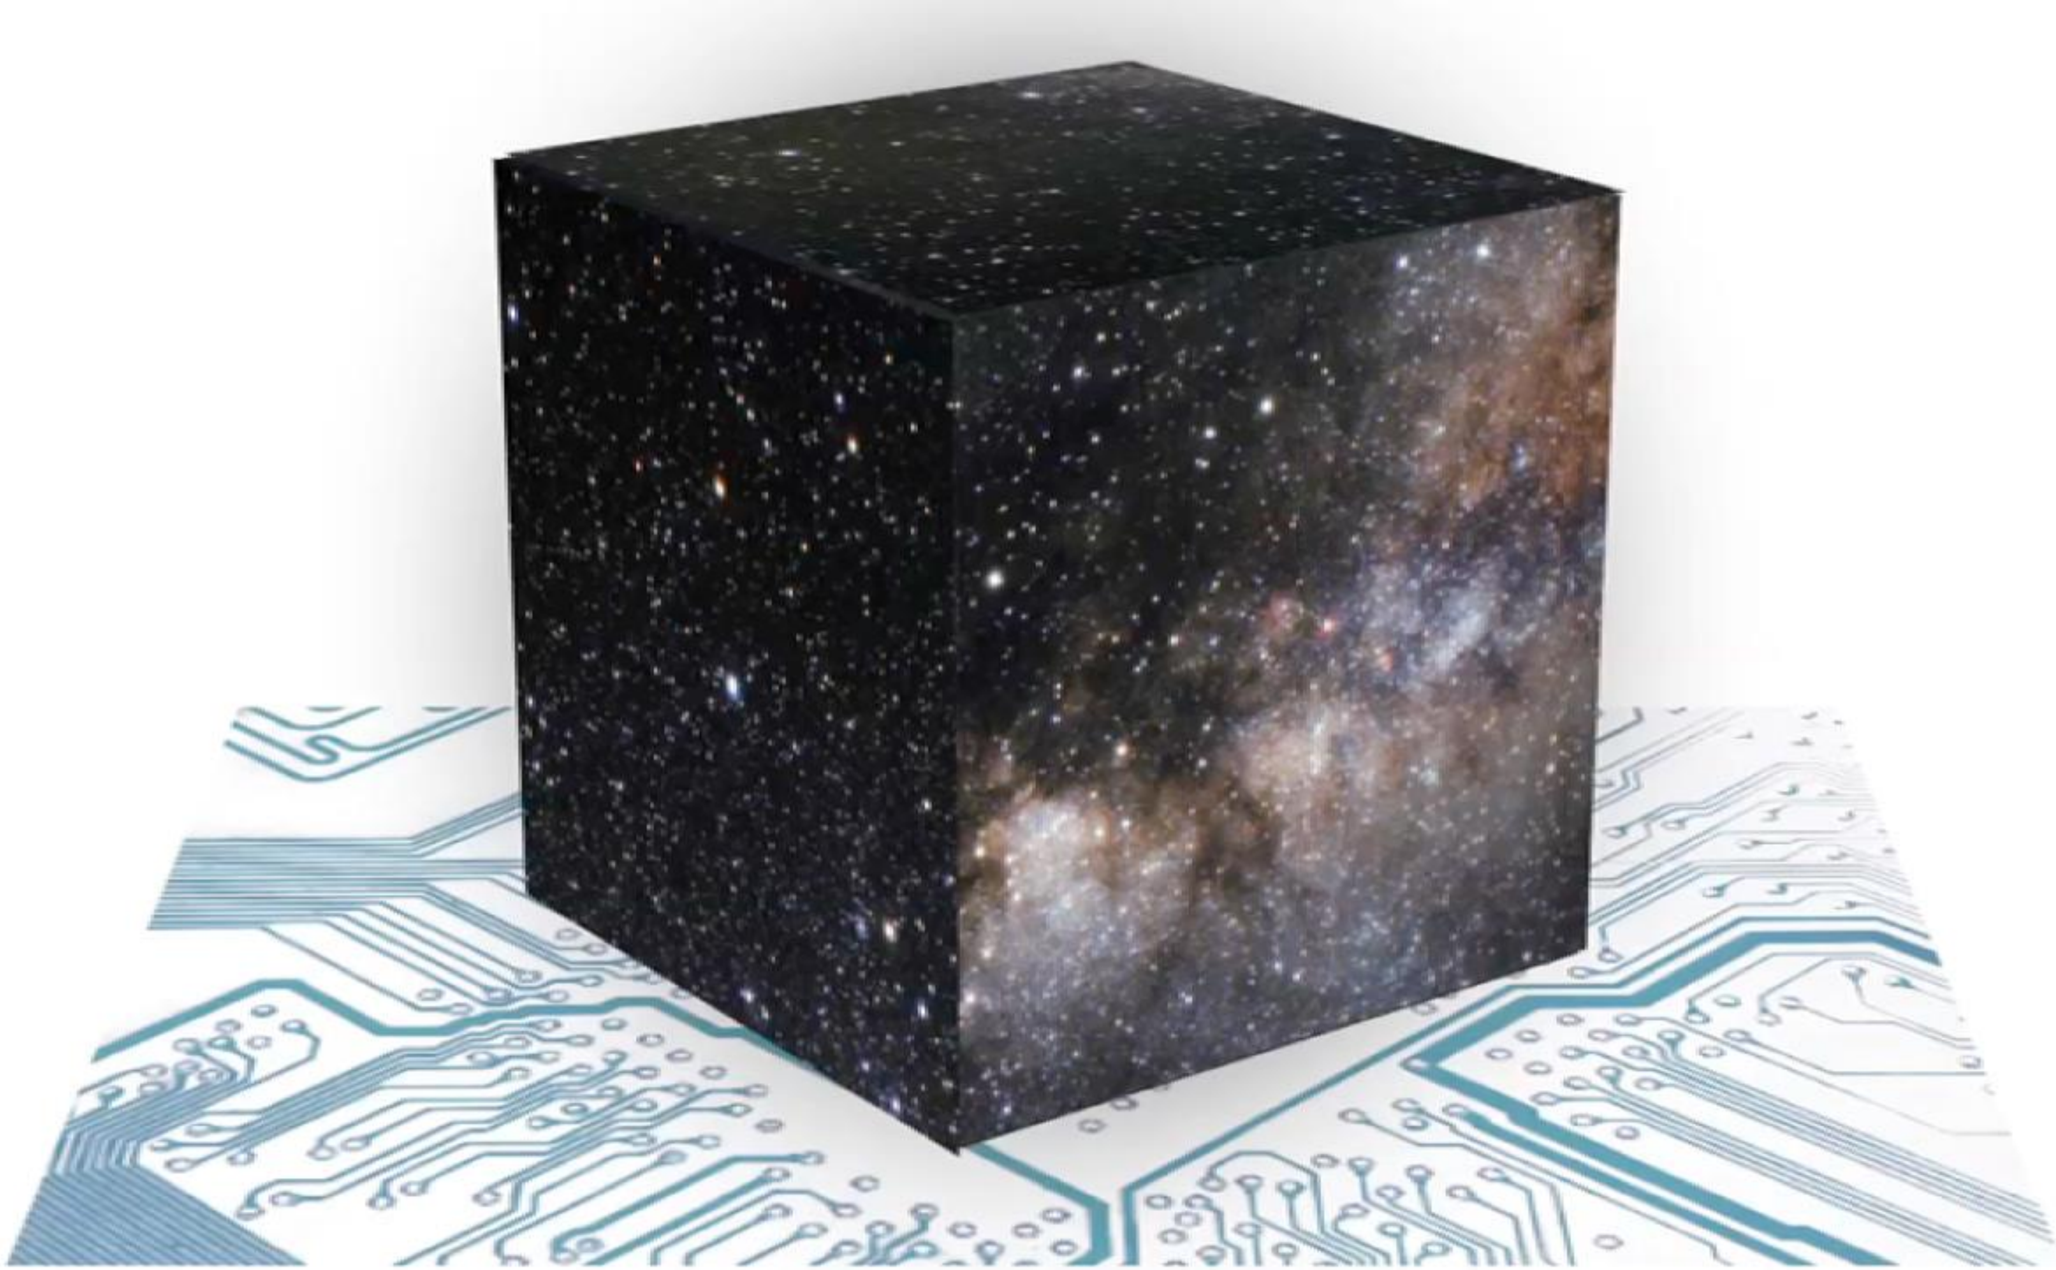

INSTITUTE FOR QUANTUM INFORMATION AND MATTER

John Preskill
US QIS Summer School
Oak Ridge National Laboratory
15 July 2024

@preskill

Quantum Information Science

Quantum sensing
Improving sensitivity, spatial resolution, noninvasiveness.

Quantum cryptography
Privacy founded on fundamental laws of quantum physics.

Quantum networking
Distributing quantumness around the world.

Quantum simulation
Probes of exotic quantum many-body phenomena.

Quantum computing
Speeding up solutions to hard problems.

Hardware challenges cut across all these application areas.

Concepts: entanglement, quantum error correction, computational complexity,

Frontiers of Physics

short distance

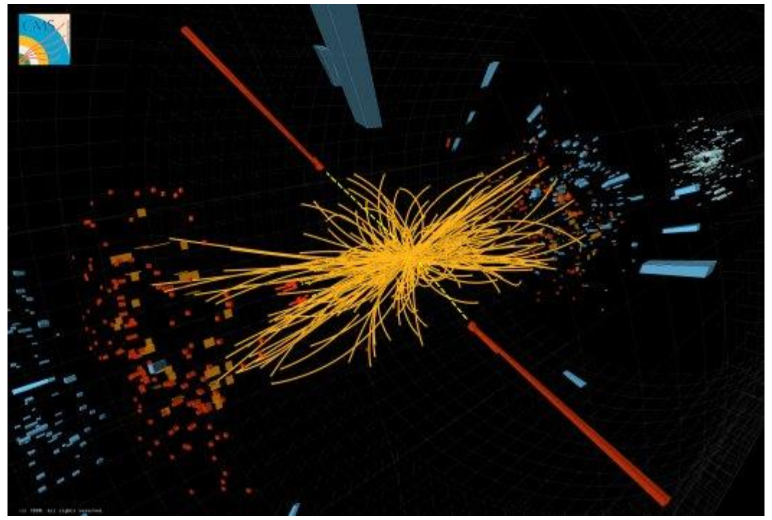

Higgs boson

Neutrino masses

Supersymmetry

Quantum gravity

String theory

long distance

Large scale structure

Cosmic microwave background

Dark matter

Dark energy

Gravitational waves

complexity

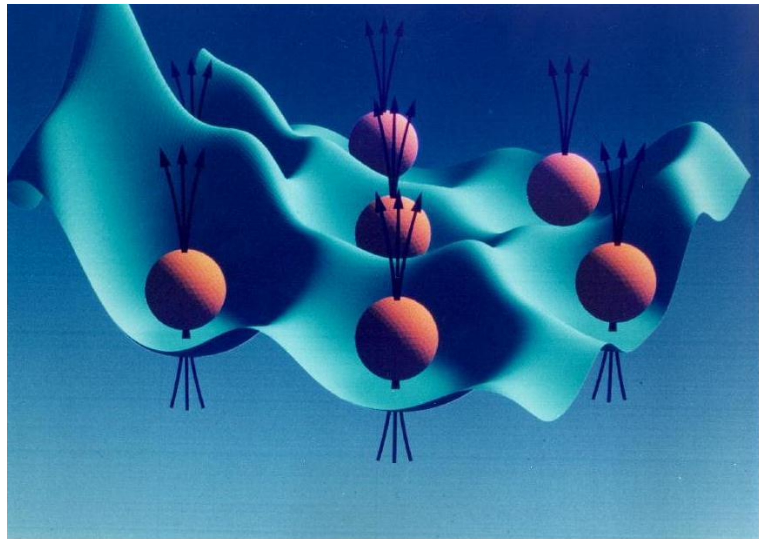

“More is different”

Many-body entanglement

Phases of quantum matter

Quantum computing

Quantum spacetime

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

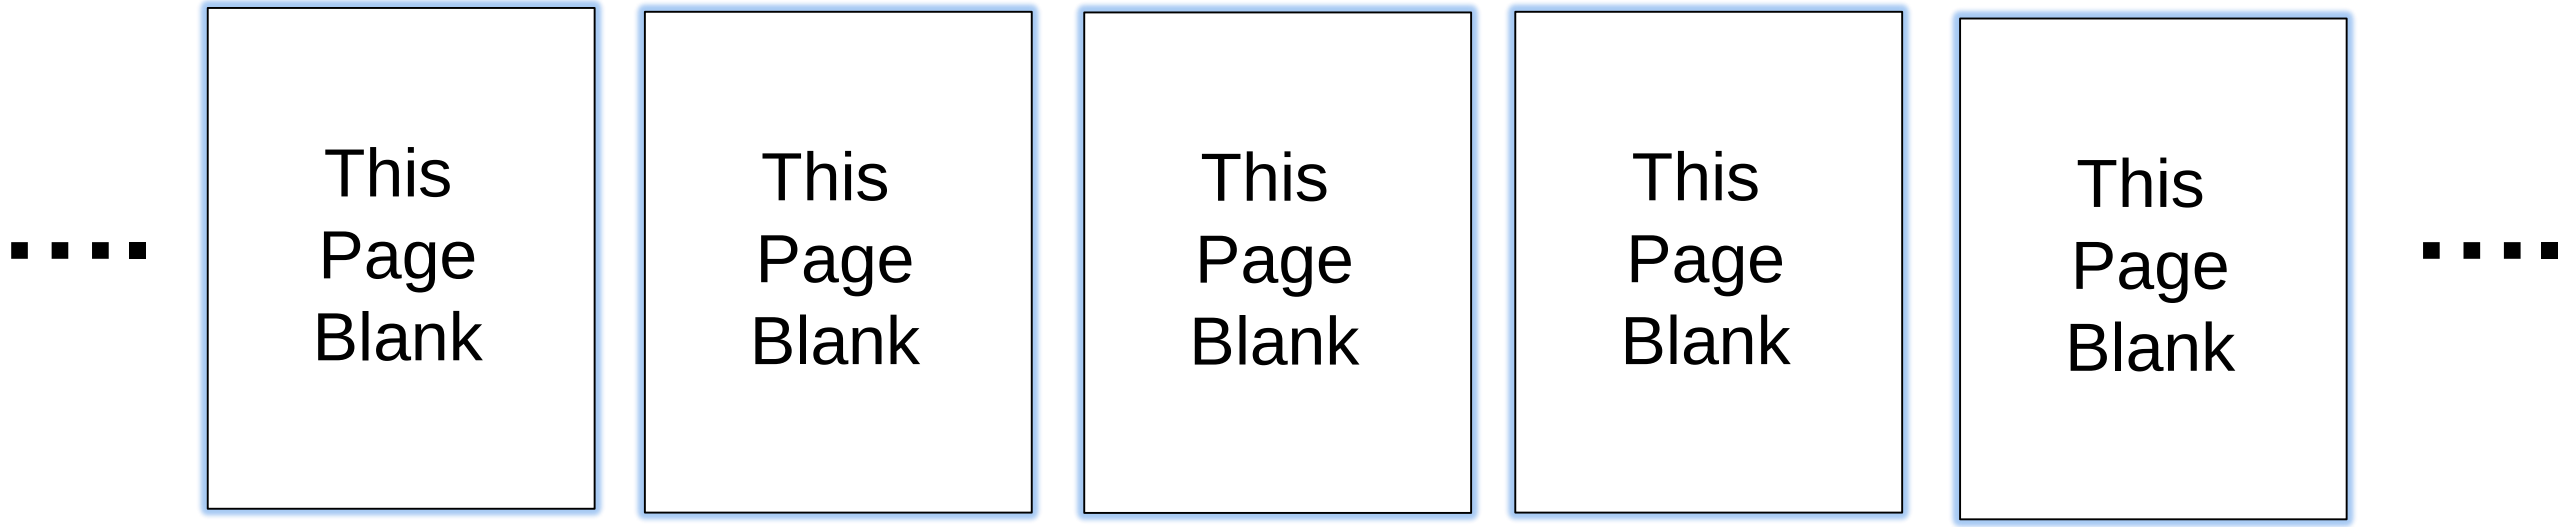

Nearly all the information in a typical entangled “quantum book” is encoded in the correlations among the “pages”.

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

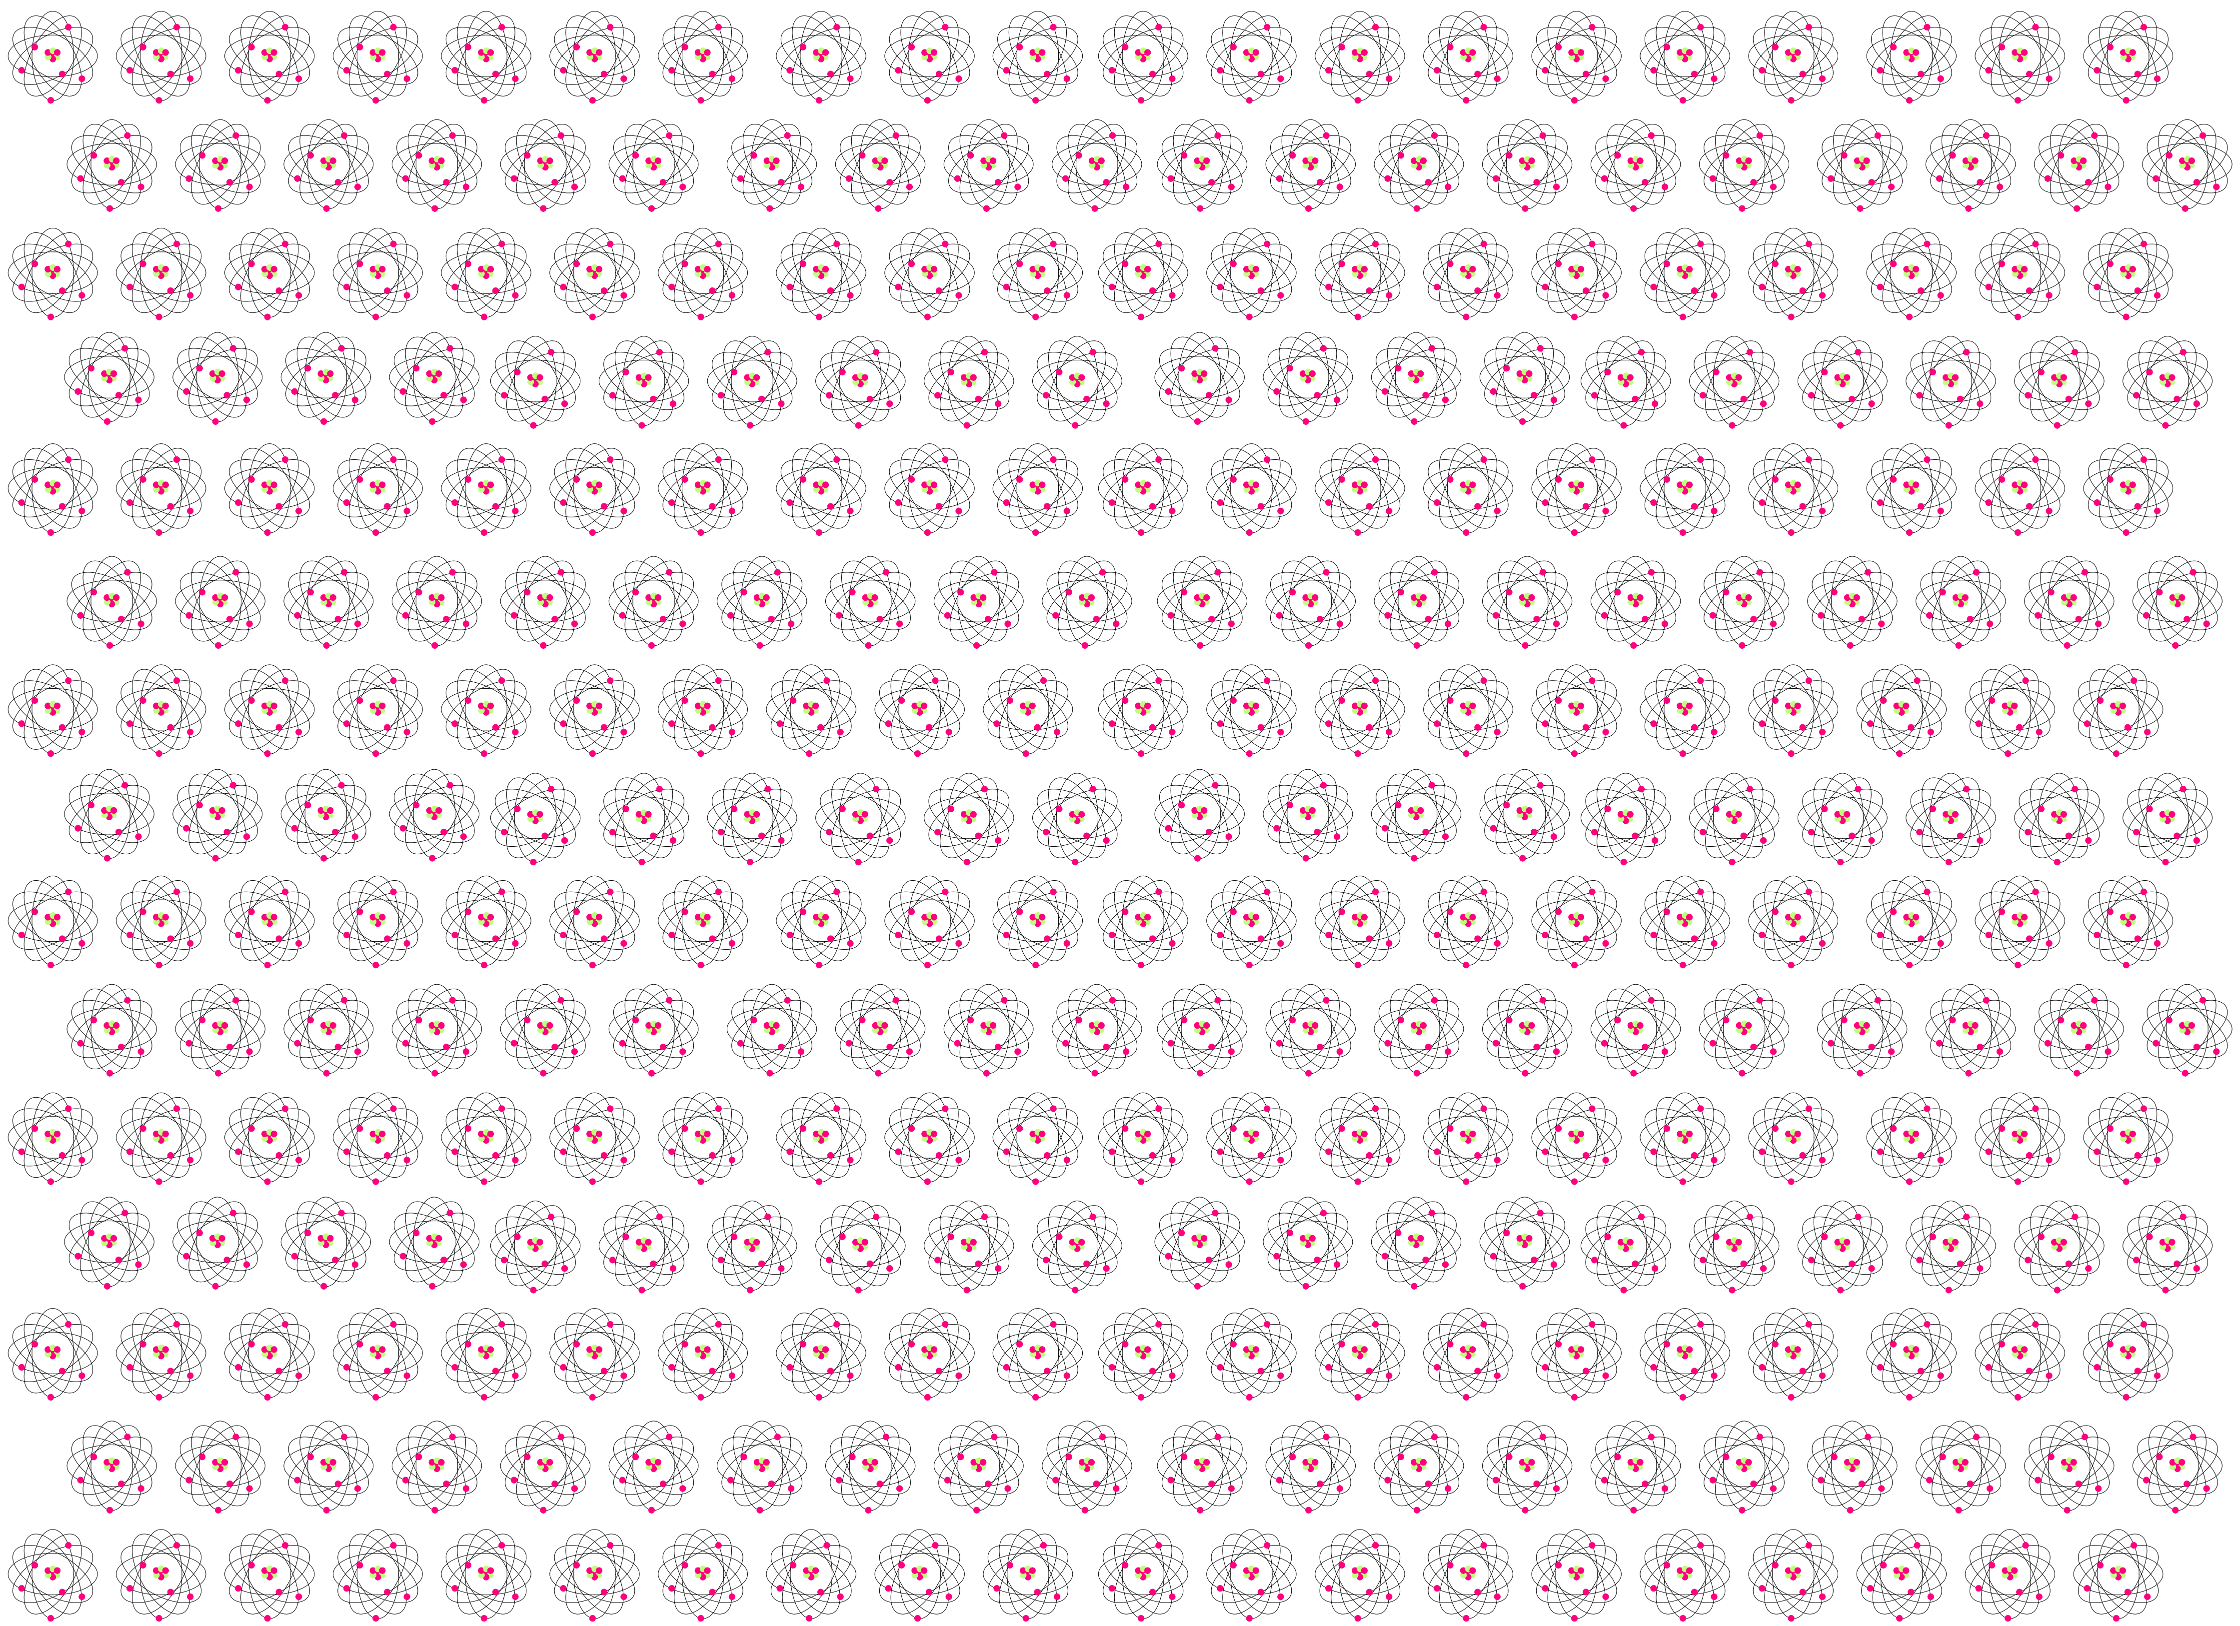

A complete description of a typical quantum state of just 300 qubits requires more bits than the number of atoms in the visible universe.

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

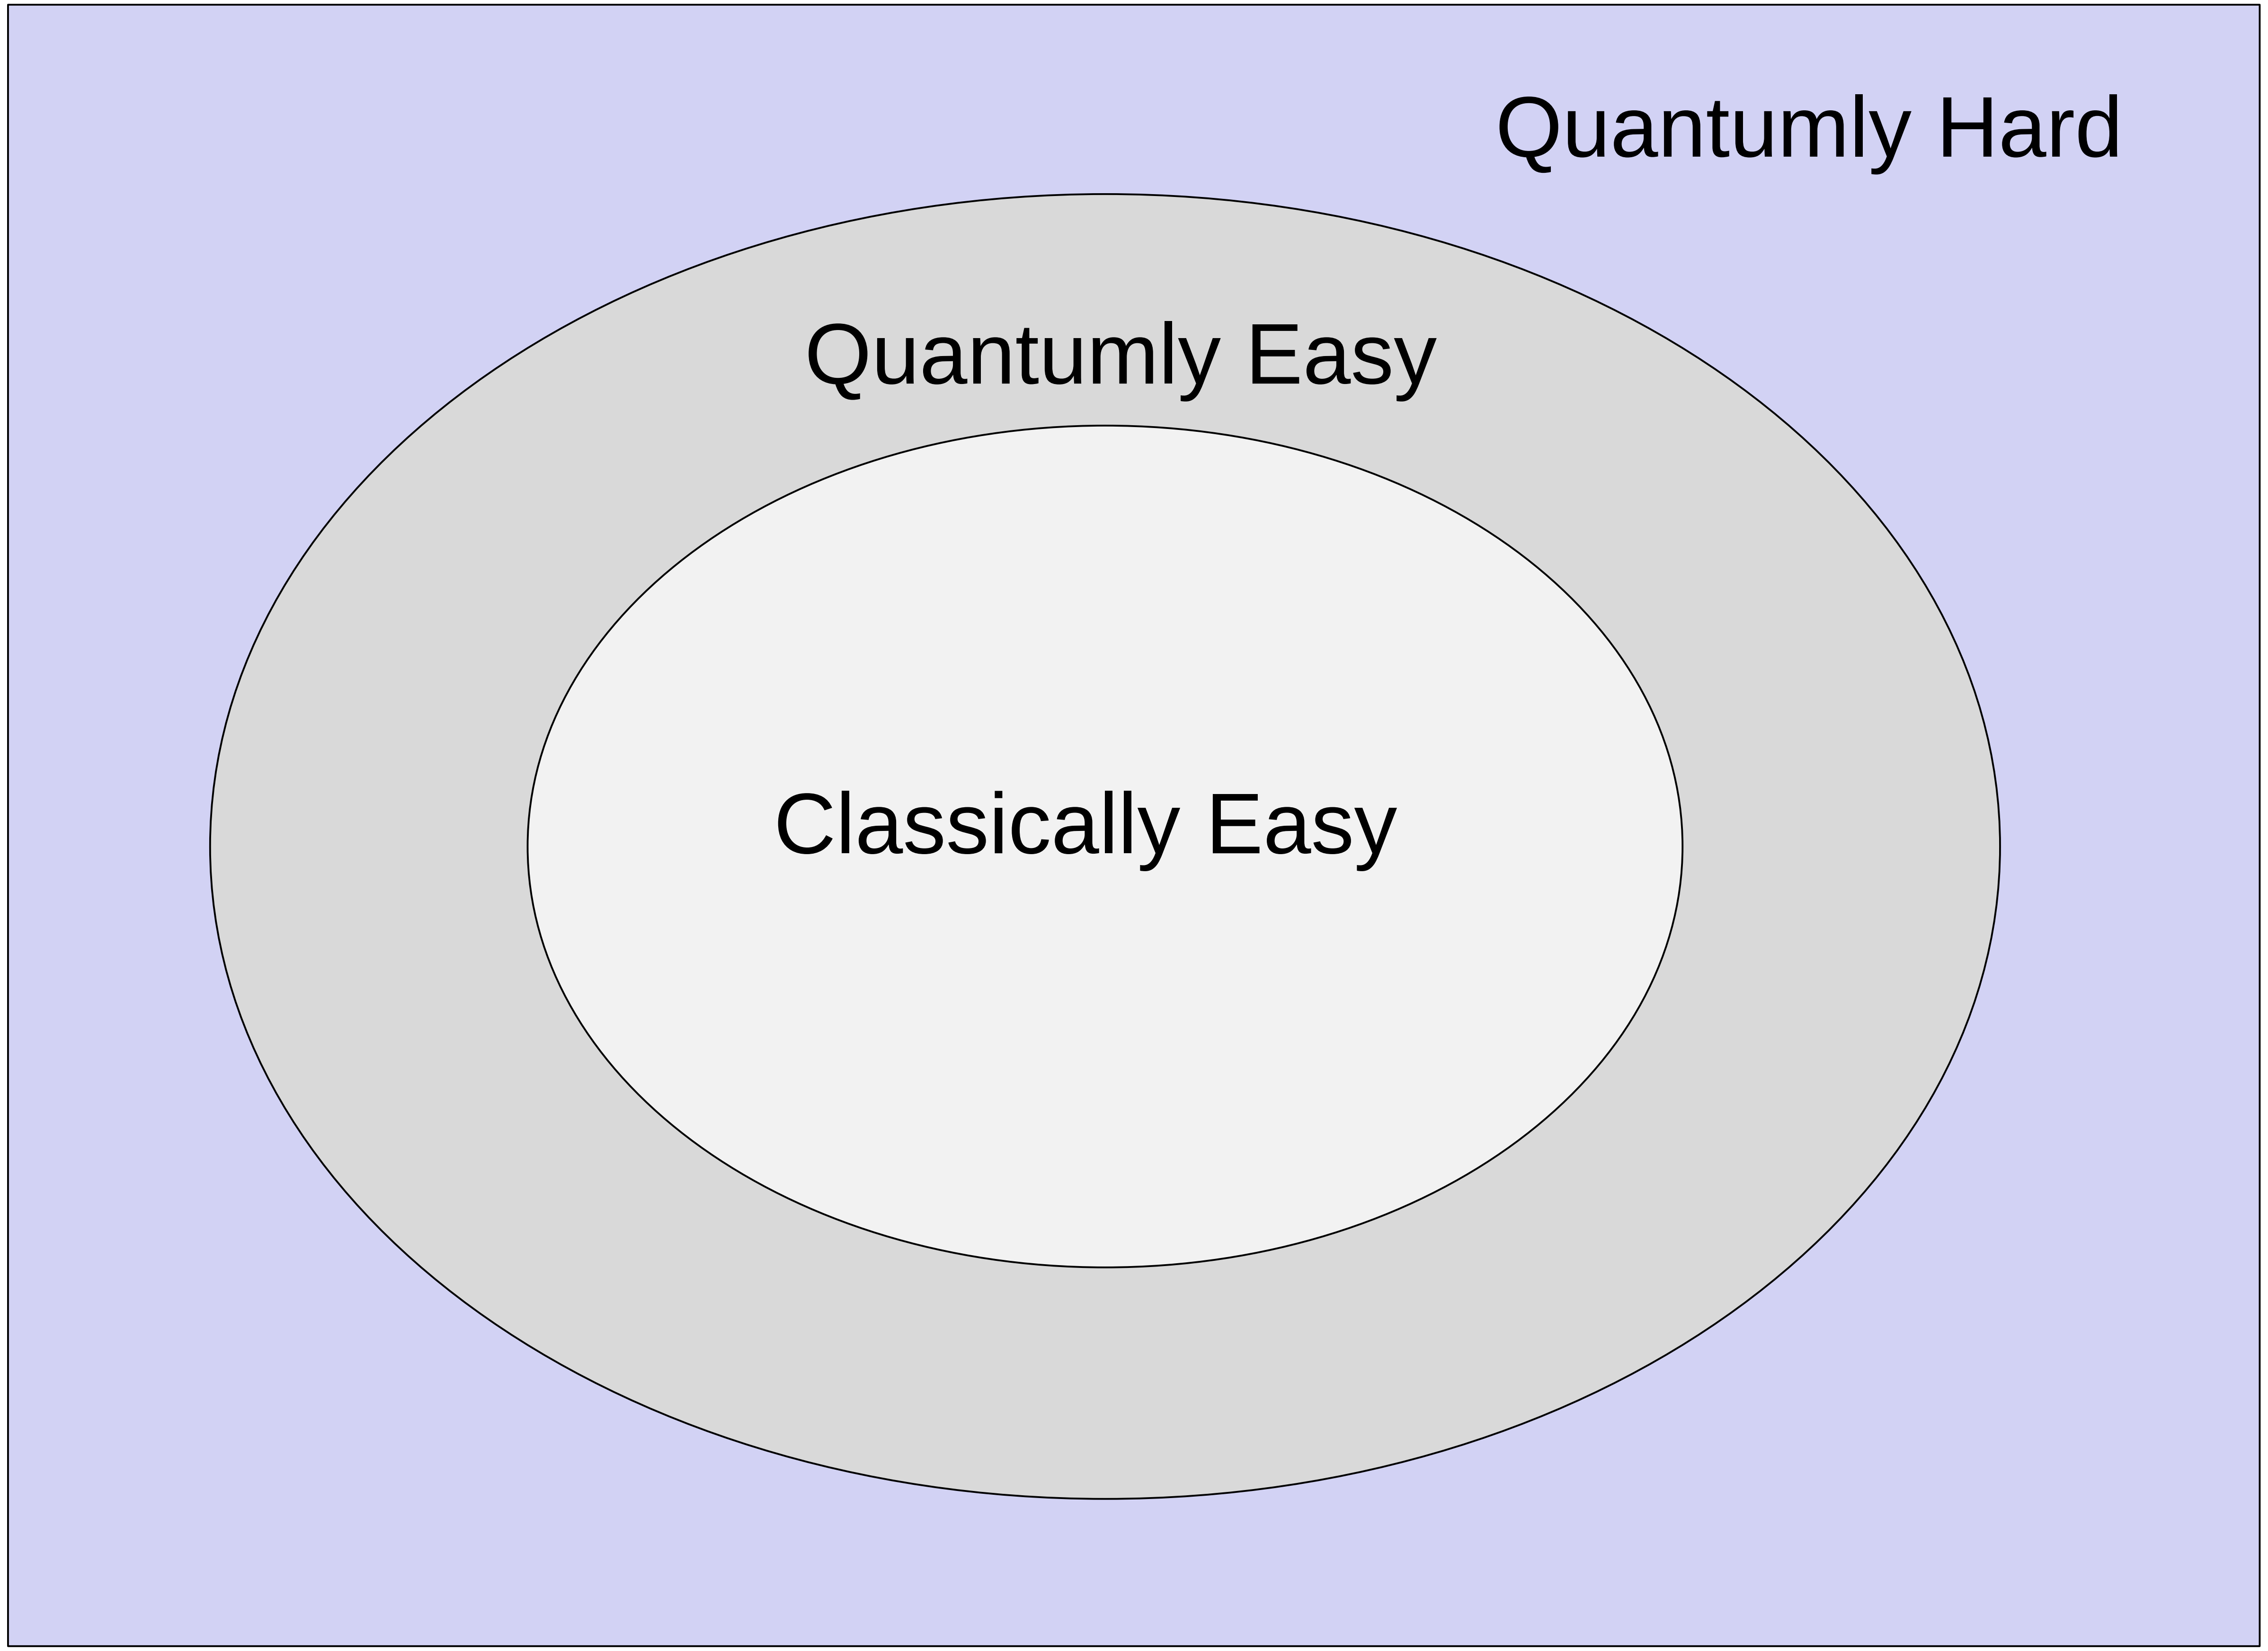

Problems

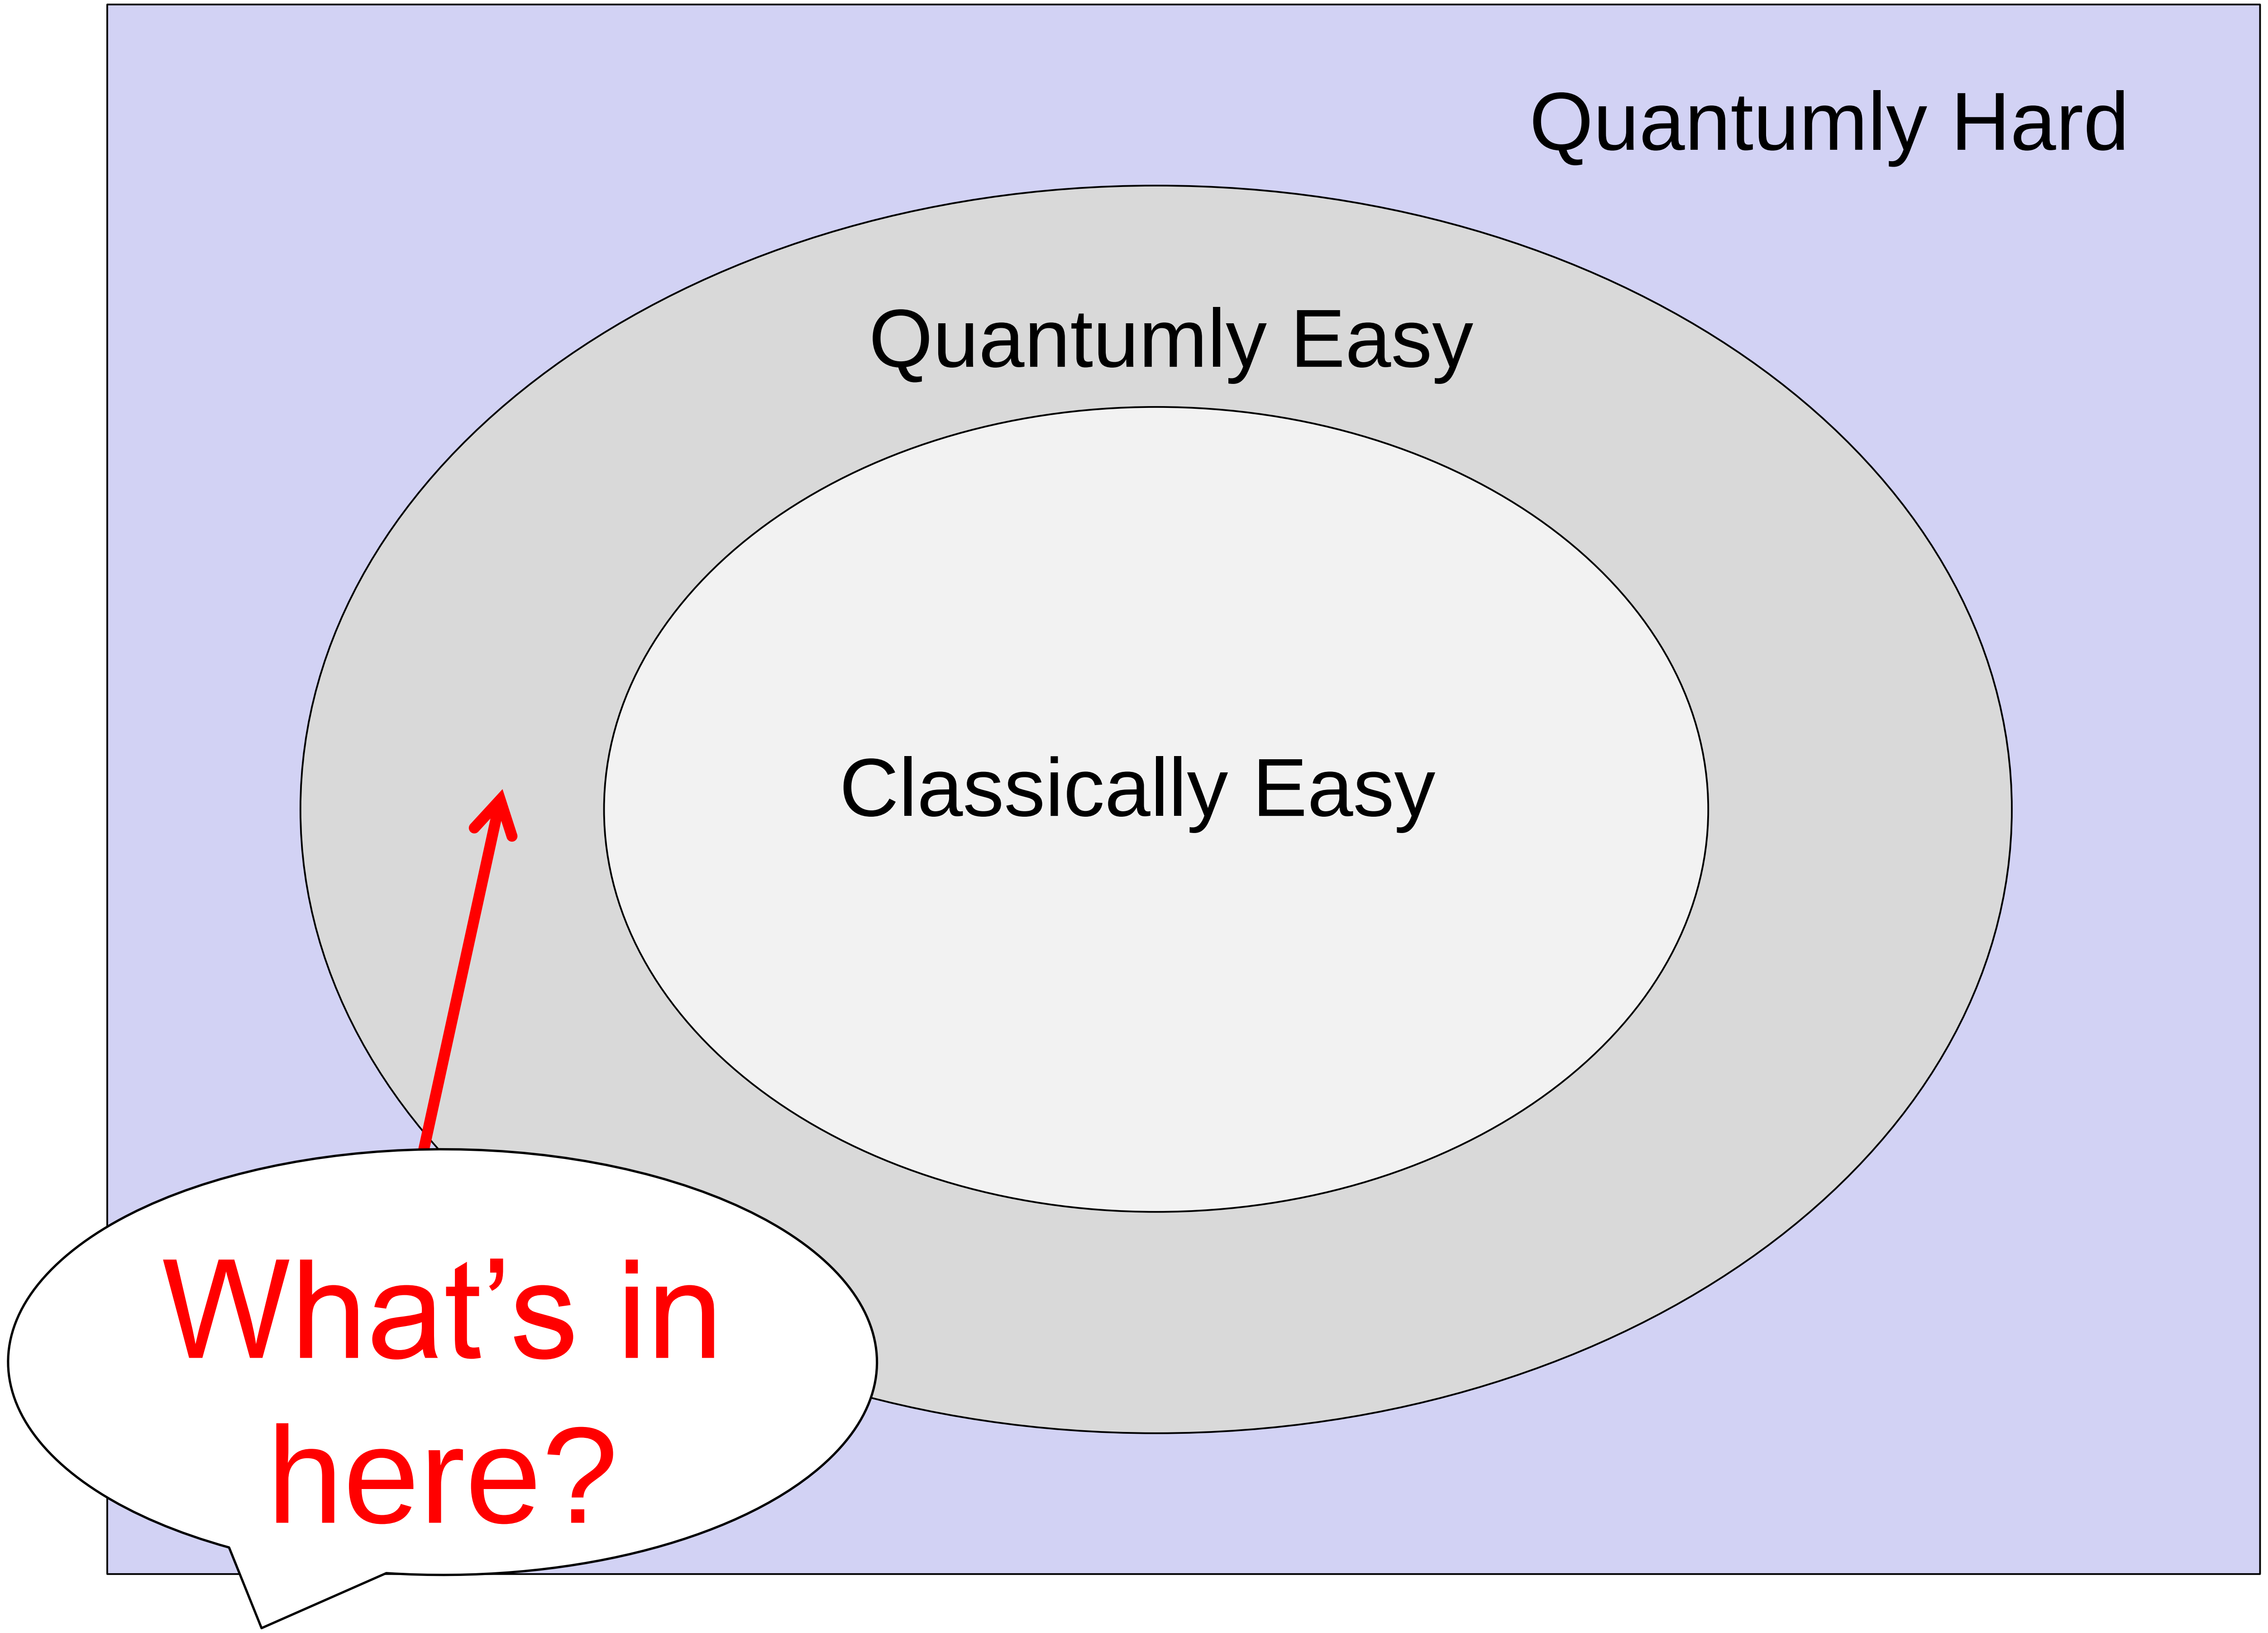
Quantumly Hard
Quantumly Easy
Classically Easy
What's in here?

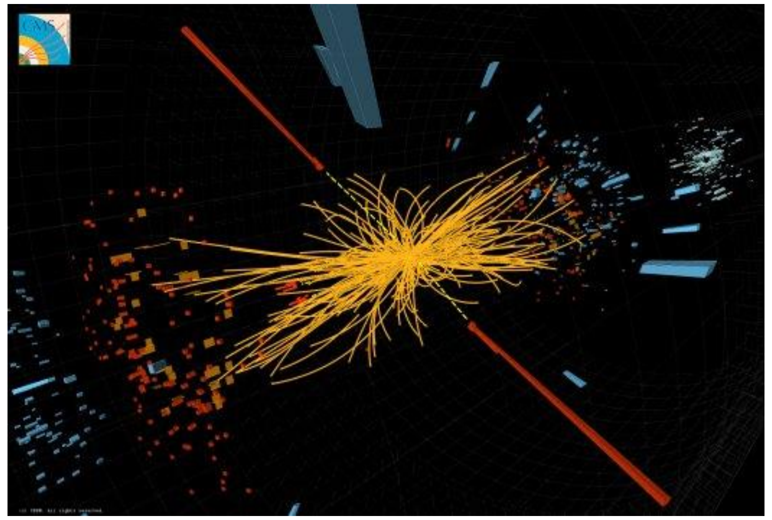
particle collision

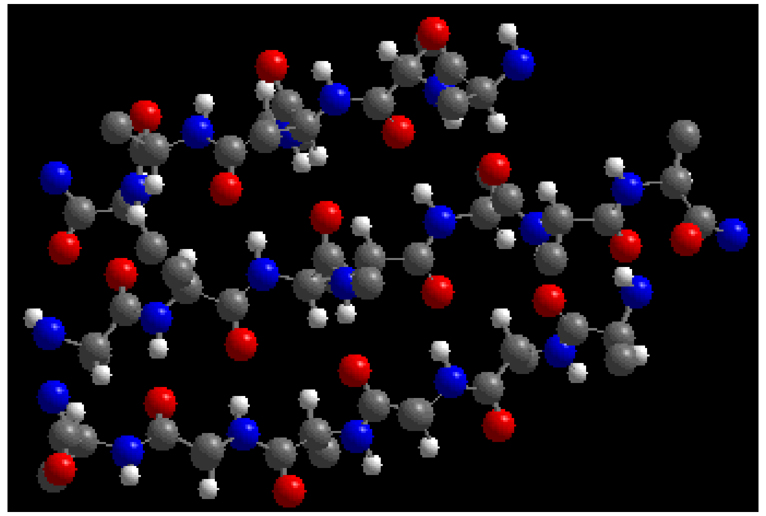
molecular chemistry

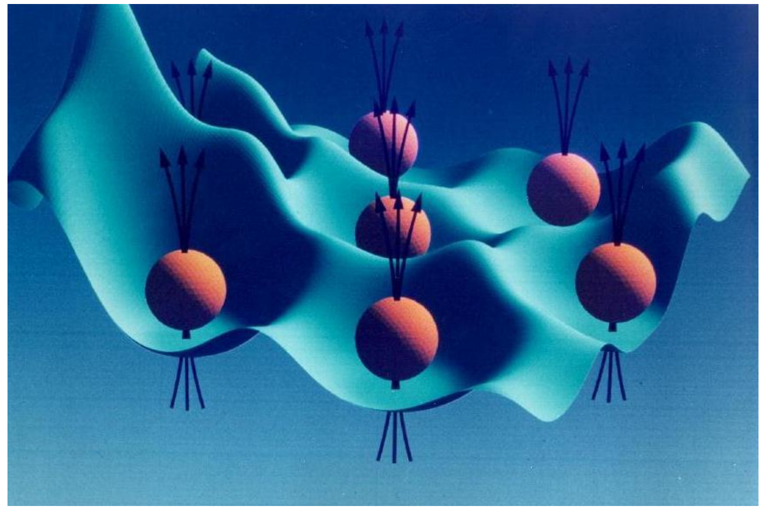
entangled electrons

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

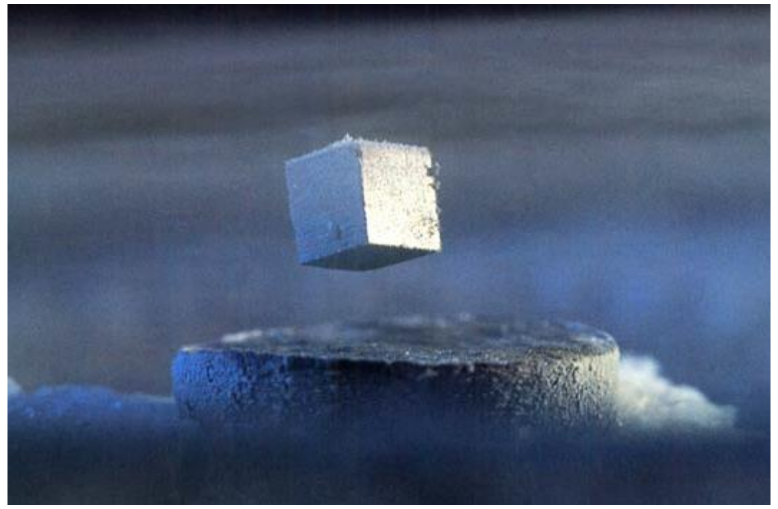
superconductor

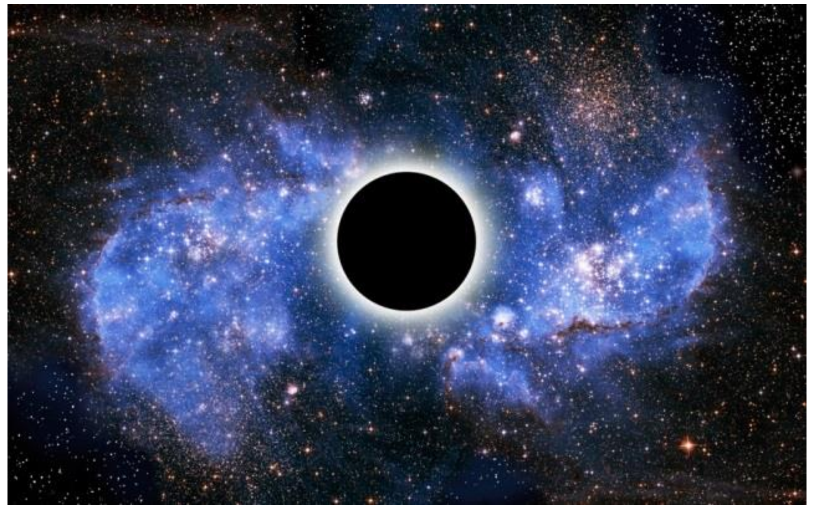
black hole


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

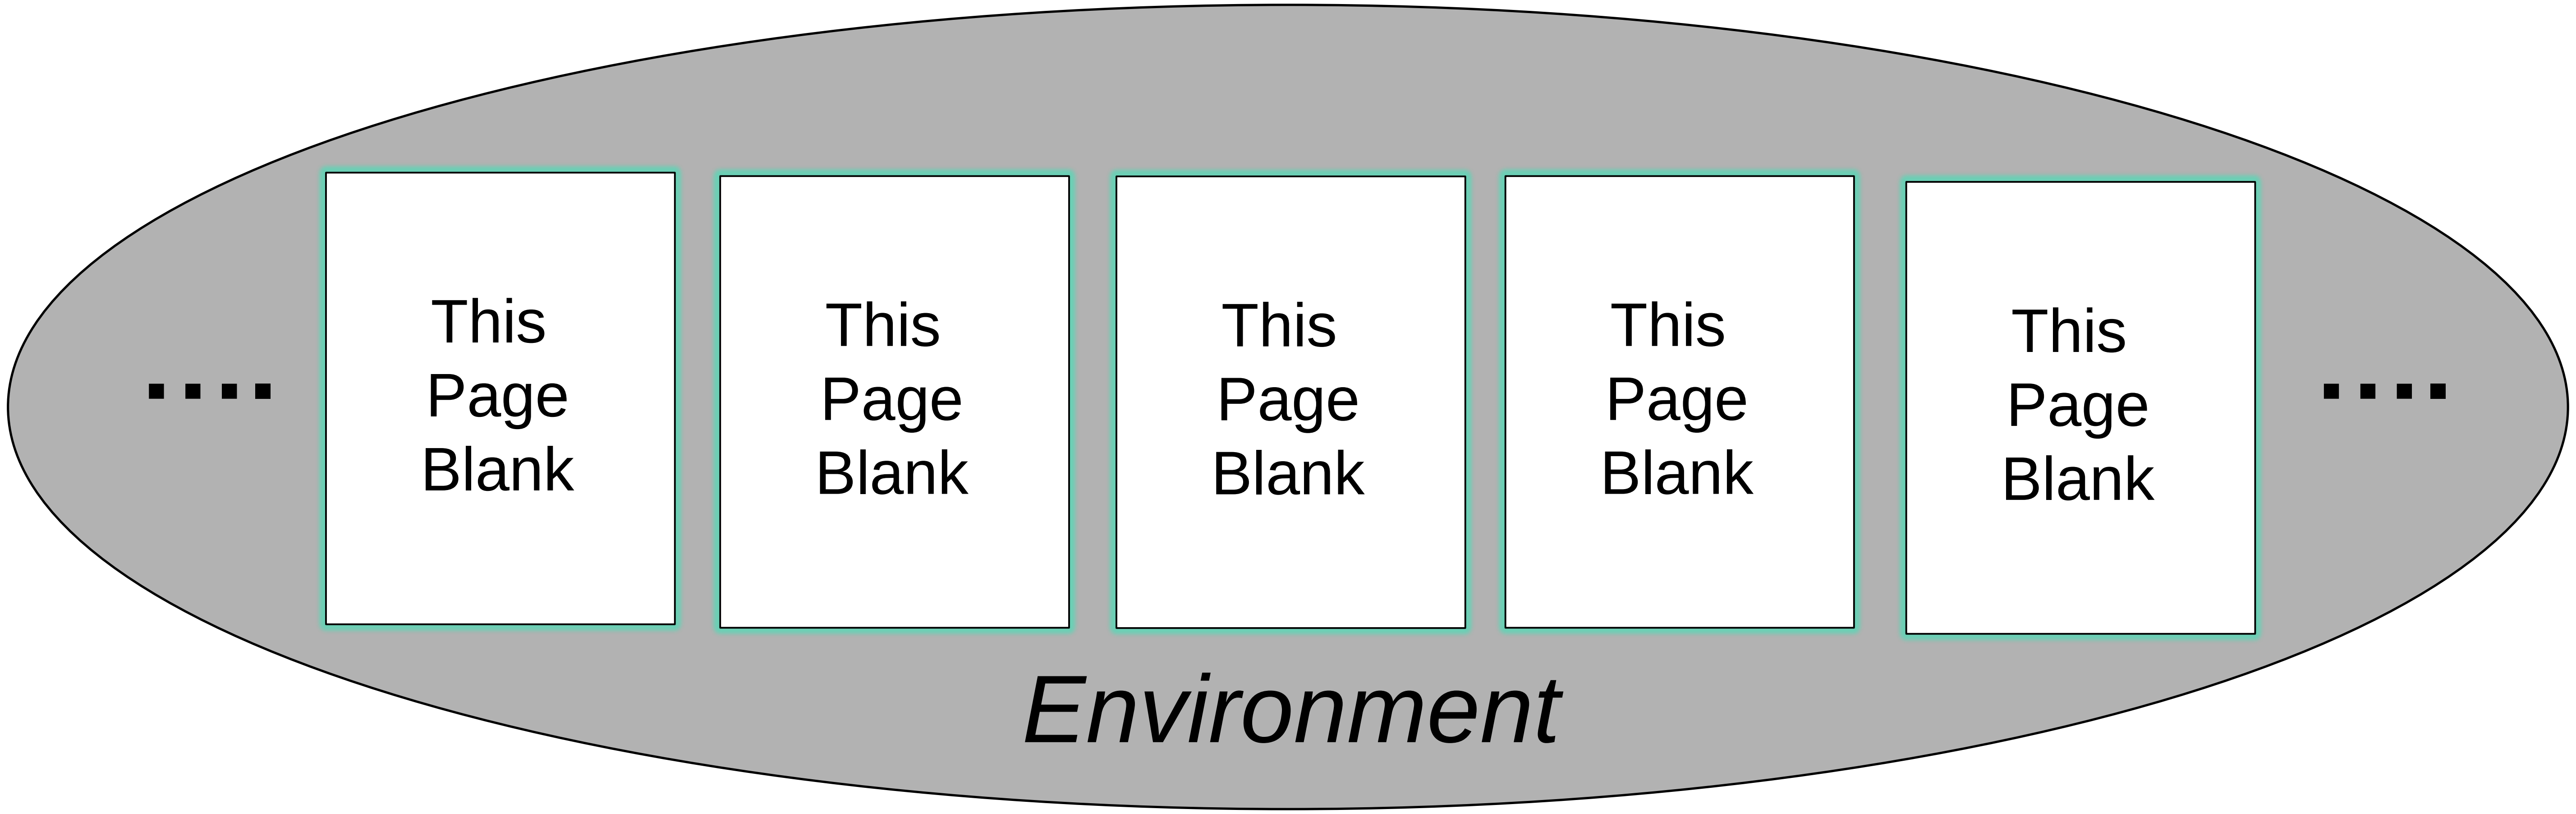

The protected "logical" quantum information is encoded in a highly entangled state of many physical qubits.

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

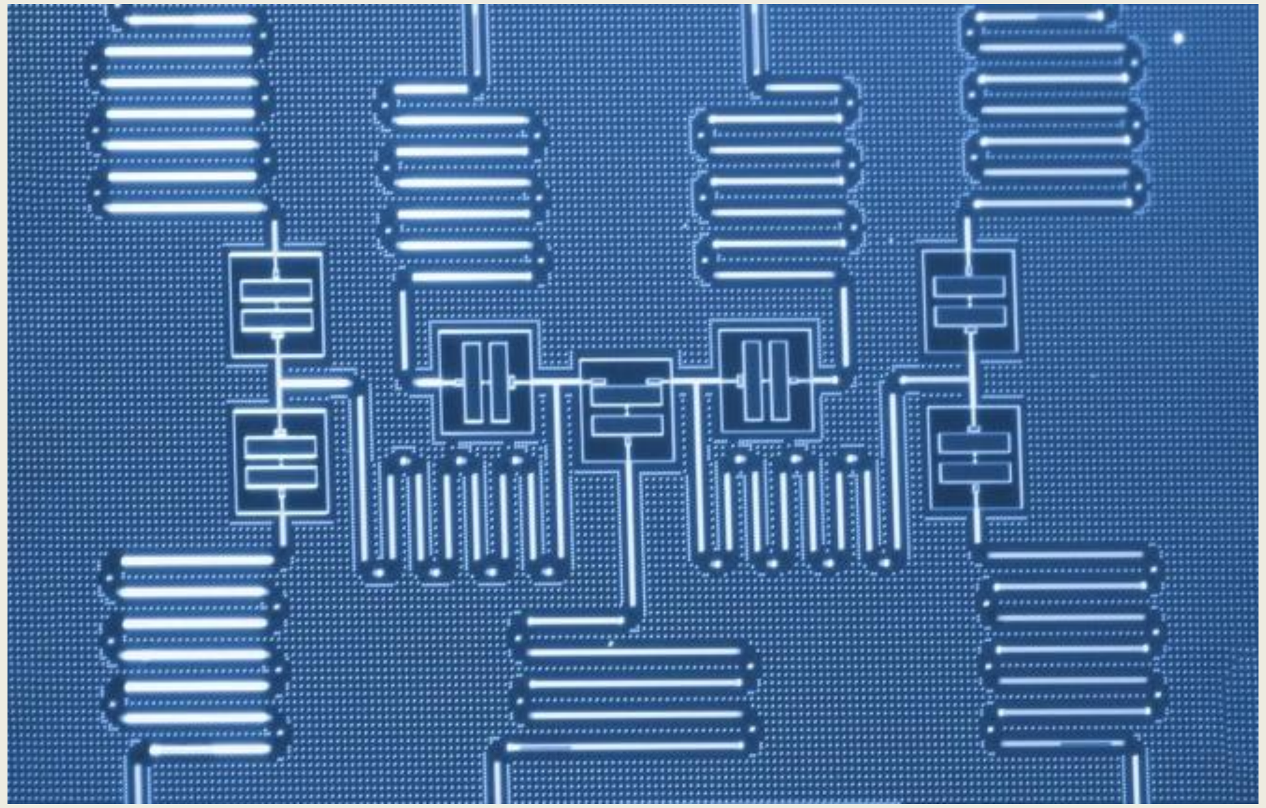

superconducting qubits

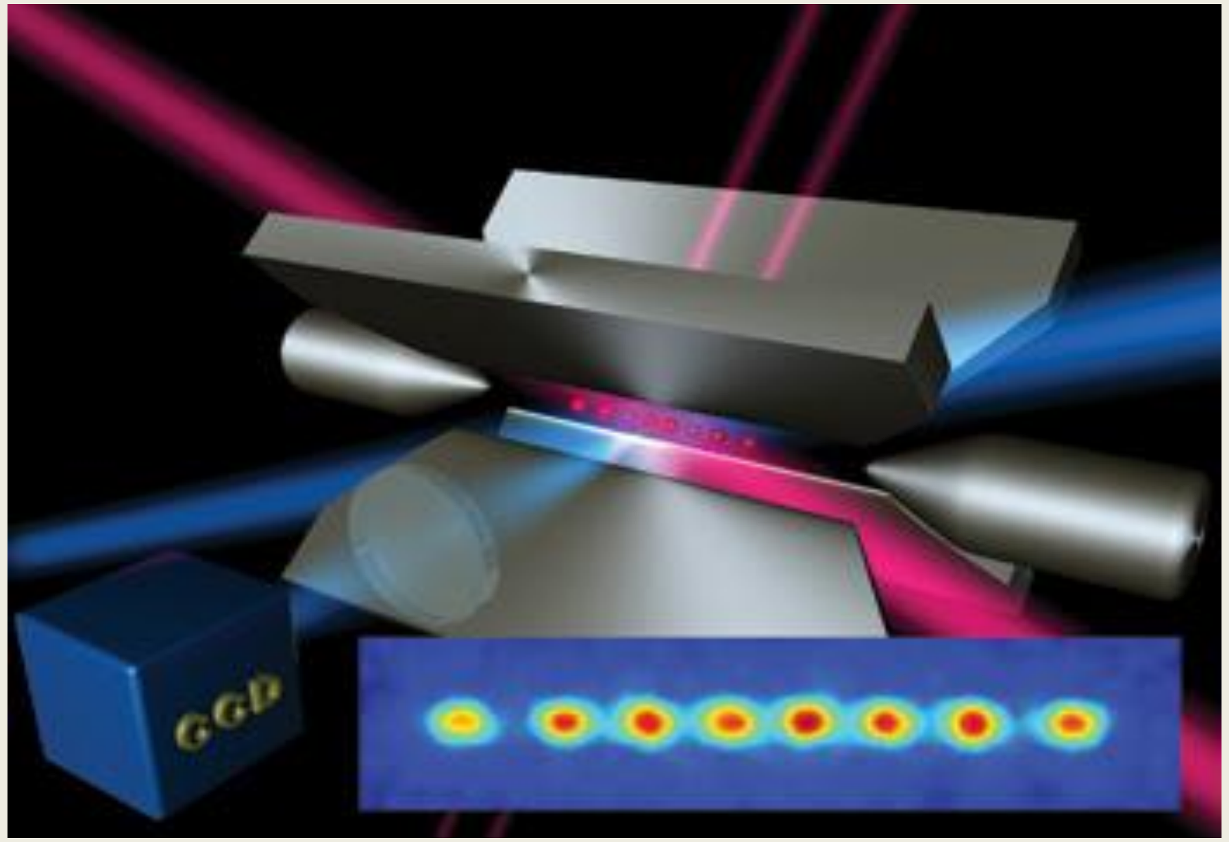

trapped atoms/ions

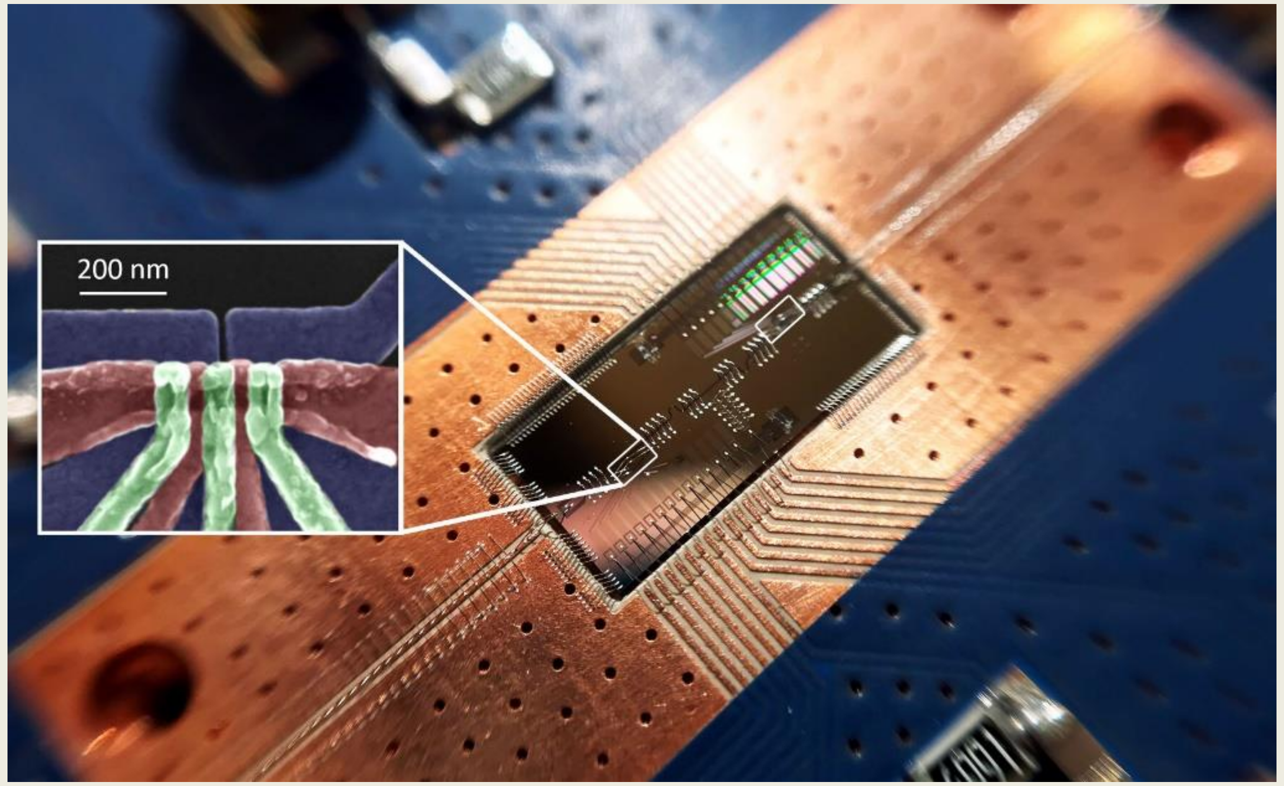

silicon spin qubits

photonics

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.

Superconductors

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

Transmons: artificial atoms, carefully fabricated and 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.

Classical systems cannot simulate quantum systems efficiently (a widely believed but unproven conjecture).

Arguably the most interesting thing we know about the difference between quantum and classical.

Nature **574**, pages 505–510 (2019), 23 October 2019

Article

Quantum supremacy using a programmable superconducting processor

https://doi.org/10.1038/s41586-019-1666-5

Received: 22 July 2019

Accepted: 20 September 2019

Published online: 23 October 2019

Frank Arute[1], Kunal Arya[1], Ryan Babbush[1], Dave Bacon[1], Joseph C. Bardin[1,2], Rami Barends[1], Rupak Biswas[3], Sergio Boixo[1], Fernando G. S. L. Brandao[1,4], David A. Buell[1], Brian Burkett[1], Yu Chen[1], Zijun Chen[1], Ben Chiaro[5], Roberto Collins[1], William Courtney[1], Andrew Dunsworth[1], Edward Farhi[1], Brooks Foxen[1,5], Austin Fowler[1], Craig Gidney[1], Marissa Giustina[1], Rob Graff[1], Keith Guerin[1], Steve Habegger[1], Matthew P. Harrigan[1], Michael J. Hartmann[1,6], Alan Ho[1], Markus Hoffmann[1], Trent Huang[1], Travis S. Humble[7], Sergei V. Isakov[1], Evan Jeffrey[1],

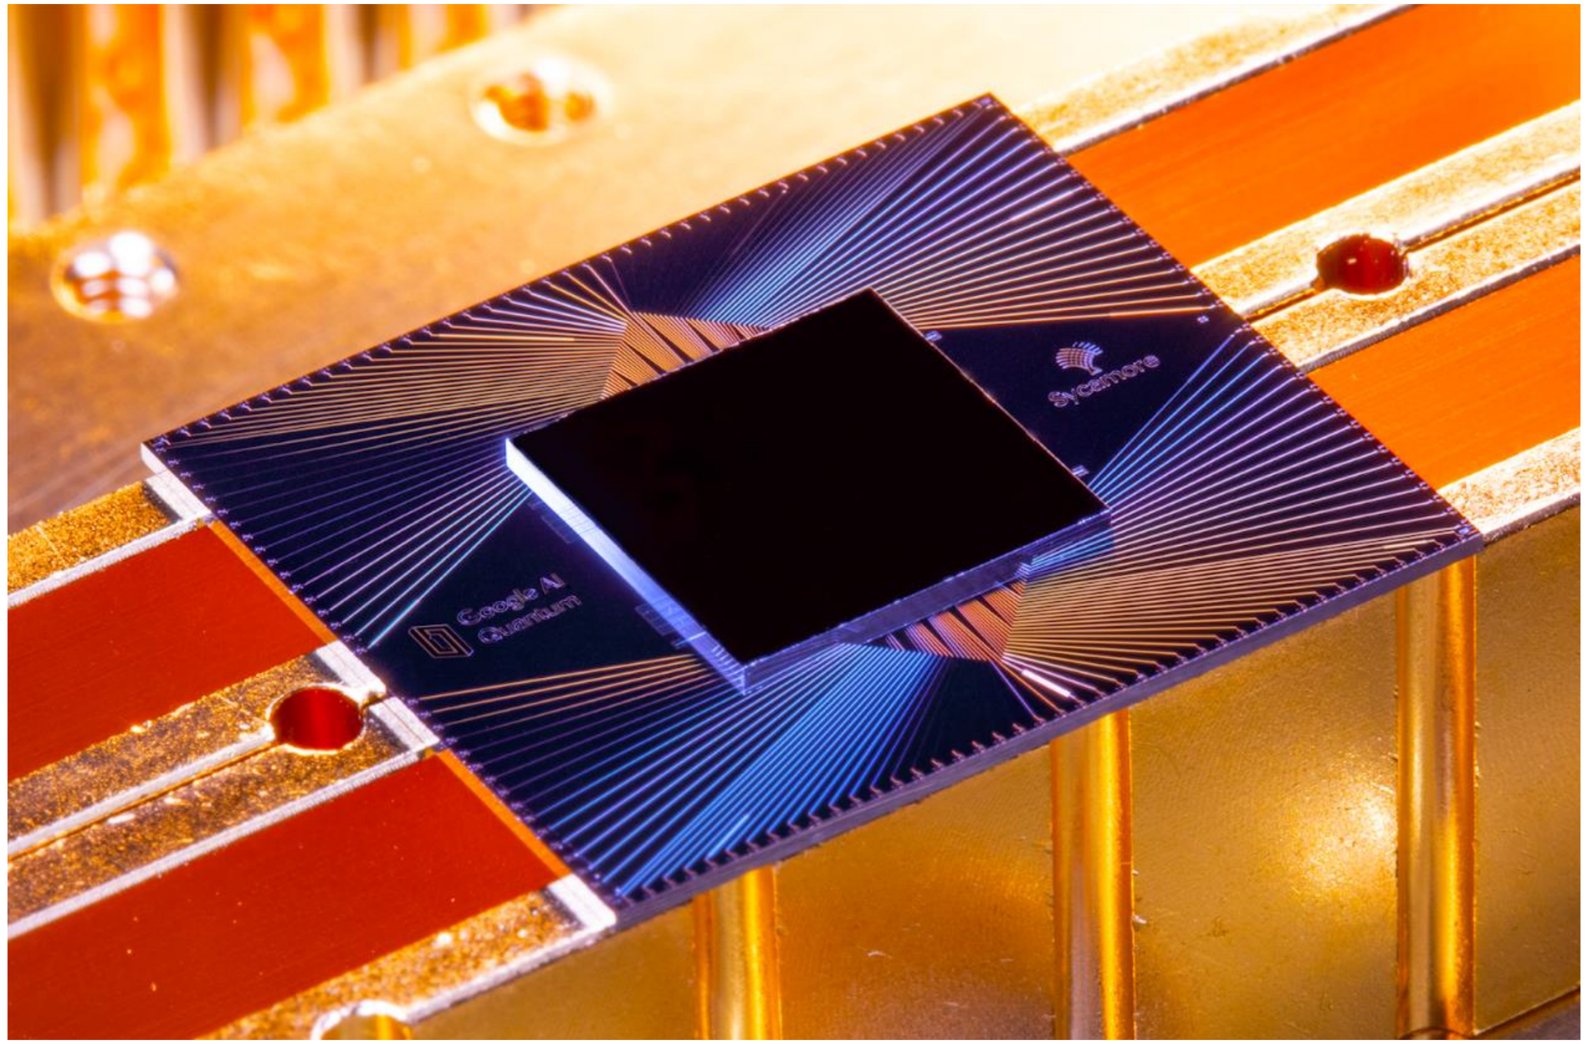

Credit: Erik Lucero/Google

About Sycamore

“Quantum David vs. Classical Goliath”

In 2023: A fully programmable circuit-based quantum computer. *n*=67 working qubits in a two-dimensional array with coupling of nearest neighbors.

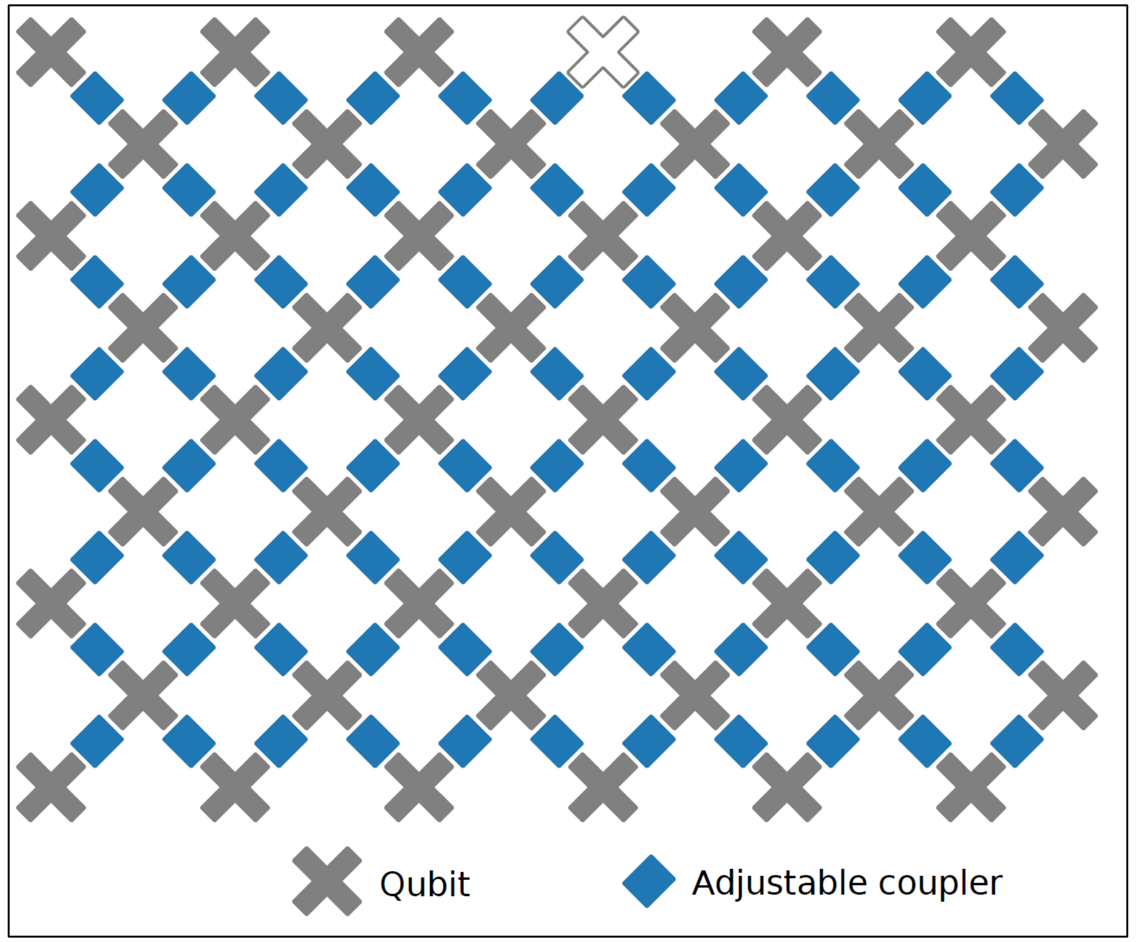

A circuit with 32 layers of 2-qubit gates *can be executed millions of times in a few minutes*, yielding verifiable results.

Simulating this quantum circuit using a classical computer is challenging.

Furthermore, the cost of the classical simulation grows exponentially with the number of qubits.

Conclusion: the hardware is working well enough to produce meaningful results in a regime where classical simulation is very difficult.

Quantum computing in the NISQ Era

The (noisy) 100-qubit quantum computer has arrived.
(*NISQ* = noisy intermediate-scale quantum.)

NISQ devices cannot be simulated by brute force using the most powerful currently existing supercomputers.

Noise limits the computational power of NISQ-era technology.

NISQ will be an interesting tool for exploring physics. It *might* also have other useful applications. But we're not sure about that.

NISQ will not change the world by itself. Rather it is a step toward more powerful quantum technologies of the future.

Hybrid quantum/classical optimizers

Eddie Farhi: "Try it and see if it works!"

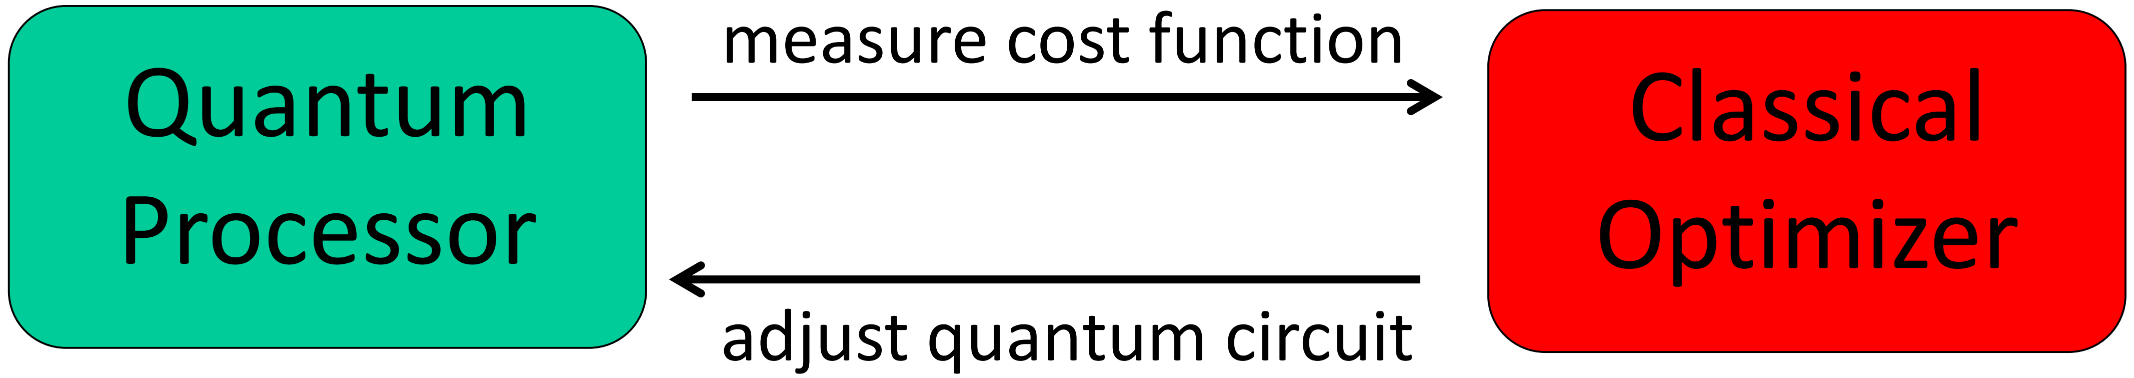

We don't expect a quantum computer to find exact solutions to worst-case instances of NP-hard problems efficiently, but it *might* find better approximate solutions, or find them faster.

Classical optimization algorithms (for both classical and quantum problems) are sophisticated and well-honed after decades of hard work.

We don't know whether NISQ devices can do better, but we can try it and see how well it works.

Status of NISQ applications

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.

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 applications

Dirac (1929): “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.”

Feynman (1981): “You can simulate this with a quantum system, with quantum computer elements. It’s not a Turing machine, but a machine of a different kind.”

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

Ground states in chemistry and materials

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

Yet, heuristic classical algorithms have been very successful.

We are targeting the relatively small "strongly correlated" corner of chemistry and materials science, where such methods falter.

Can quantum computers efficiently solve for ground states in cases where classical methods fail?

Quantum computers cannot find ground states for QMA-hard cases, but that's okay. Nature does not find these states either.

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

Ground states in chemistry and materials

We are seeking problems that are (1) quantumly easy, (2) classically hard, (3) physically relevant.

A patchwork of heuristic classical methods including: HF, DFT, CC, QMC, DMRG, TN, NN, … These lack performance guarantees, but often work. Cost need not scale exponentially with problem size.

Quantum algorithms are heuristic, too. We need an initial state that has sufficient overlap with the ground state.

Strong correlations can result in competing phases, first-order quantum phase transitions, … Adiabatic state preparation may fail.

Garnet Chan et al., Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry (2023)

Ground states in chemistry and materials

We are seeking problems that are (1) quantumly easy, (2) classically hard, (3) physically relevant.

Perhaps exponential quantum advantage should not be expected.

But a significant polynomial advantage is a reasonable expectation and could be quite impactful.

These applications require deep quantum circuits. Fault-tolerant quantum computation will be needed, at a high cost in physical qubits and gates.

Garnet Chan et al., Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry (2023)

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, …)

Digital vs. Analog quantum simulation

Near-term quantum simulators can be either *digital* (circuit based) or *analog* (tunable Hamiltonians).

Analog quantum simulation has been an active research area for 20 years or more; digital quantum simulation is just getting started in the past few years.

Digital provides more flexible Hamiltonian and initial state preparation. But gate based simulations of time evolution are expensive.

Analog platforms include: ultracold (neutral) atoms and molecules, trapped ions, superconducting circuits, etc. These same platforms can also be used for circuit-based computation.

Although they are becoming more sophisticated and controllable, analog simulators are limited by imperfect control. They are best suited for studying “robust” properties that can be accessed using noisy quantum systems.

Eventually, digital (circuit-based) quantum simulators will surpass analog quantum simulators for studies of quantum dynamics, but perhaps not until fault tolerance is feasible.

Experience with near-term digital simulators will lay foundations for fault-tolerant simulations in the future (applies to NISQ more broadly).

Prototypical quantum dynamics simulation task

(1) State preparation. E.g., incoming scattering state.

(2) Hamiltonian evolution. E.g. Trotter approximation.

(3) Measure an observable. E.g., a simulated detector.

Goal: sample accurately from probability distribution of outcomes.

Determine how computational resources scale with: error, system size, particle number, total energy of process, energy gap, …

Resources include: number of qubits, number of gates, …

Hope for polynomial scaling! Or even better: polylog scaling.

Need an efficient preparation of initial state.

Approximating a continuous system incurs discretization cost (smaller lattice spacing improves accuracy).

Is circuit based dynamical simulation too expensive for the NISQ era?

Quantum simulation of quantum field theory

50 years since Ken Wilson proposed lattice gauge theory!

Real-time evolution (collider physics), nonzero chemical potential (early universe, neutron stars), …

What's classically hard? Processes that produce highly entangled states, e.g., multiparticle production, quench, …

Laying the foundations for more revealing future work.

Stepping stone to quantum gravity.

New concepts and insights?

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

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

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

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 $P_{\text{physical}} = 10^{-4}$.)

Progress toward QEC

Erasure conversion. Dominant errors occur at known locations, hence easier to correct.

Biased noise. Physical suppression of bit flips, error-correcting codes for the phase flips.

More efficient codes. But geometrically nonlocal syndrome measurements required.

Co-design. Adapt the coding to the hardware, adapt the hardware to the code.

Erasure conversion

Dominant errors are heralded, occur at known circuit locations, hence easier to correct.

By design, dominant errors exit the computational space of the qubit, and can be detected without disturbing the coherence of undamaged qubits.

Alkaline earth Rydberg atoms [*Princeton, Caltech*].
$|1\rangle \rightarrow |g\rangle$, not $|1\rangle \rightarrow |0\rangle$.

Dual-rail superconducting qubit [*Yale, AWS*]. $|01\rangle, |10\rangle \rightarrow |00\rangle$
Encode using two transmons or two resonators.

Biased noise

Physically suppress the bit flips, use coding to suppress the phase flips. Gates must preserve bias.

Outer code: Repetition code or asymmetric surface code.

Example: the repetition cat code [*Yale, Alice & Bob, AWS*].

Code states $|0\rangle$, $|1\rangle$ are coherent states, well separated in phase space. Bit flips suppressed exponentially as mean photon number n increases.

Photon loss induces phase errors, at rate increasing linearly with n.

More efficient codes

Constant-rate qLDPC (quantum low-density parity-check) codes exist, including "good" codes with constant relative distance.

High accuracy thresholds, efficient decoders, schemes for executing fault-tolerant gates.

But syndrome extraction requires geometrically nonlocal operations, e.g. movable qubits or long-range coupling.

Example [*IBM*]: [[144 physical qubits, 12 logical qubits, distance 12]]

Co-design

Adapt the coding to the hardware.

Adapt the hardware to the code.

An exciting time for Rydberg atom arrays!

May lead the progress in quantum error correction for the next few years, if two-qubit gate fidelities continue to improve.

Thousands of qubits, and movement of atoms enables geometrically nonlocal operations and syndrome measurements [*Harvard/MIT/QuEra*].

Further improvement from erasure conversion.

Repeated syndrome measurement yet to be demonstrated.

Continuous loading of fresh atoms will be needed.

Atomic movement and readout are relatively slow.

Movable qubits

Schemes involving moveable atomic qubits have advantages in the short run.

But in the long run, movement imposes serious limitations on clock speed, unless much faster movement can be achieved.

Fast readout and reset are also important.

Quantum sensing

A *quantum sensor* is (typically) a few-level quantum system that senses something.

Goals: improved sensitivity, better spatial resolution, less invasive, …

High resolution scanning probes of living cells and advanced materials. E.g., NV center = Nitrogen vacancy color center in diamond.

Accelerometers, gyrometers, gravitometers, gravity gradiometers for navigation and surveying. E.g., atom interferometry.

Detection of axions and other dark matter candidates with e.g. superconducting devices. Probing fundamental symmetry violation with ultracold molecules.

Wanted: Better materials, more precise coherent control, longer coherence times, more efficient readout, compact devices, … and new ideas.

“Next-generation” quantum sensing

Higher sensitivity by exploiting squeezing and entanglement. But there is a tradeoff … what enhances sensitivity may also reduce coherence time.

Laser Interferometer Gravitational-Wave Observatory (LIGO): Enhanced sensitivity from injecting squeezed light into the dark port of an interferometer.

Optical-lattice and optical tweezer atomic clocks: improved precision through better control of a quantum many-body system

What entangled quantum states of multi-qubit sensors provide the best sensing enhancements?

Entangled sensor arrays for geodesy and geophysics: Improved predictions of earthquakes and volcanoes.

Maybe someday: Seeing a city on another planet using a long-baseline network of telescopes performing interferometry using shared quantum entanglement.

Quantum networks

Three regimes of distance scale:
(1) data center, (2) km-scale communication, (3) long-range

Quantum channel: photons sent through free space or fiber.
Loss in fiber: 17 dB per 100 km. So 100 km is possible, 1000 km is impossible.

Quantum repeater (cannot measure & resend). Quantum processor to purify and swap entanglement. Easier than fault-tolerant quantum computing.

Quantum transduction. From processor to flying qubit and back.

Nodes need not be trusted (in “device independence” protocol).

Applications for quantum networking: quantum key distribution, scalable and secure multiparty quantum computing, global quantum sensors and clocks, etc.

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?

Prospects for the next 5 years

Encouraging progress toward scalable fault-tolerant quantum computing.

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

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

Outlook

Fulfilling the potential of quantum information science is a grand challenge for 21st century science and technology.

Extraordinary advances in engineering and basic research will be needed to meet these challenges.

This will require sustained, inspired, effort and investment over decades.

It won't be fast or easy, but it's going to be fun!

We’ve only just begun.