Quantum Computing in the NISQ* Era and Beyond

John Preskill

Bethe Lecture, Cornell

8 April 2019

**Noisy Intermediate-Scale Quantum*

INSTITUTE FOR QUANTUM INFORMATION AND MATTER

CALIFORNIA INSTITUTE OF TECHNOLOGY 1891

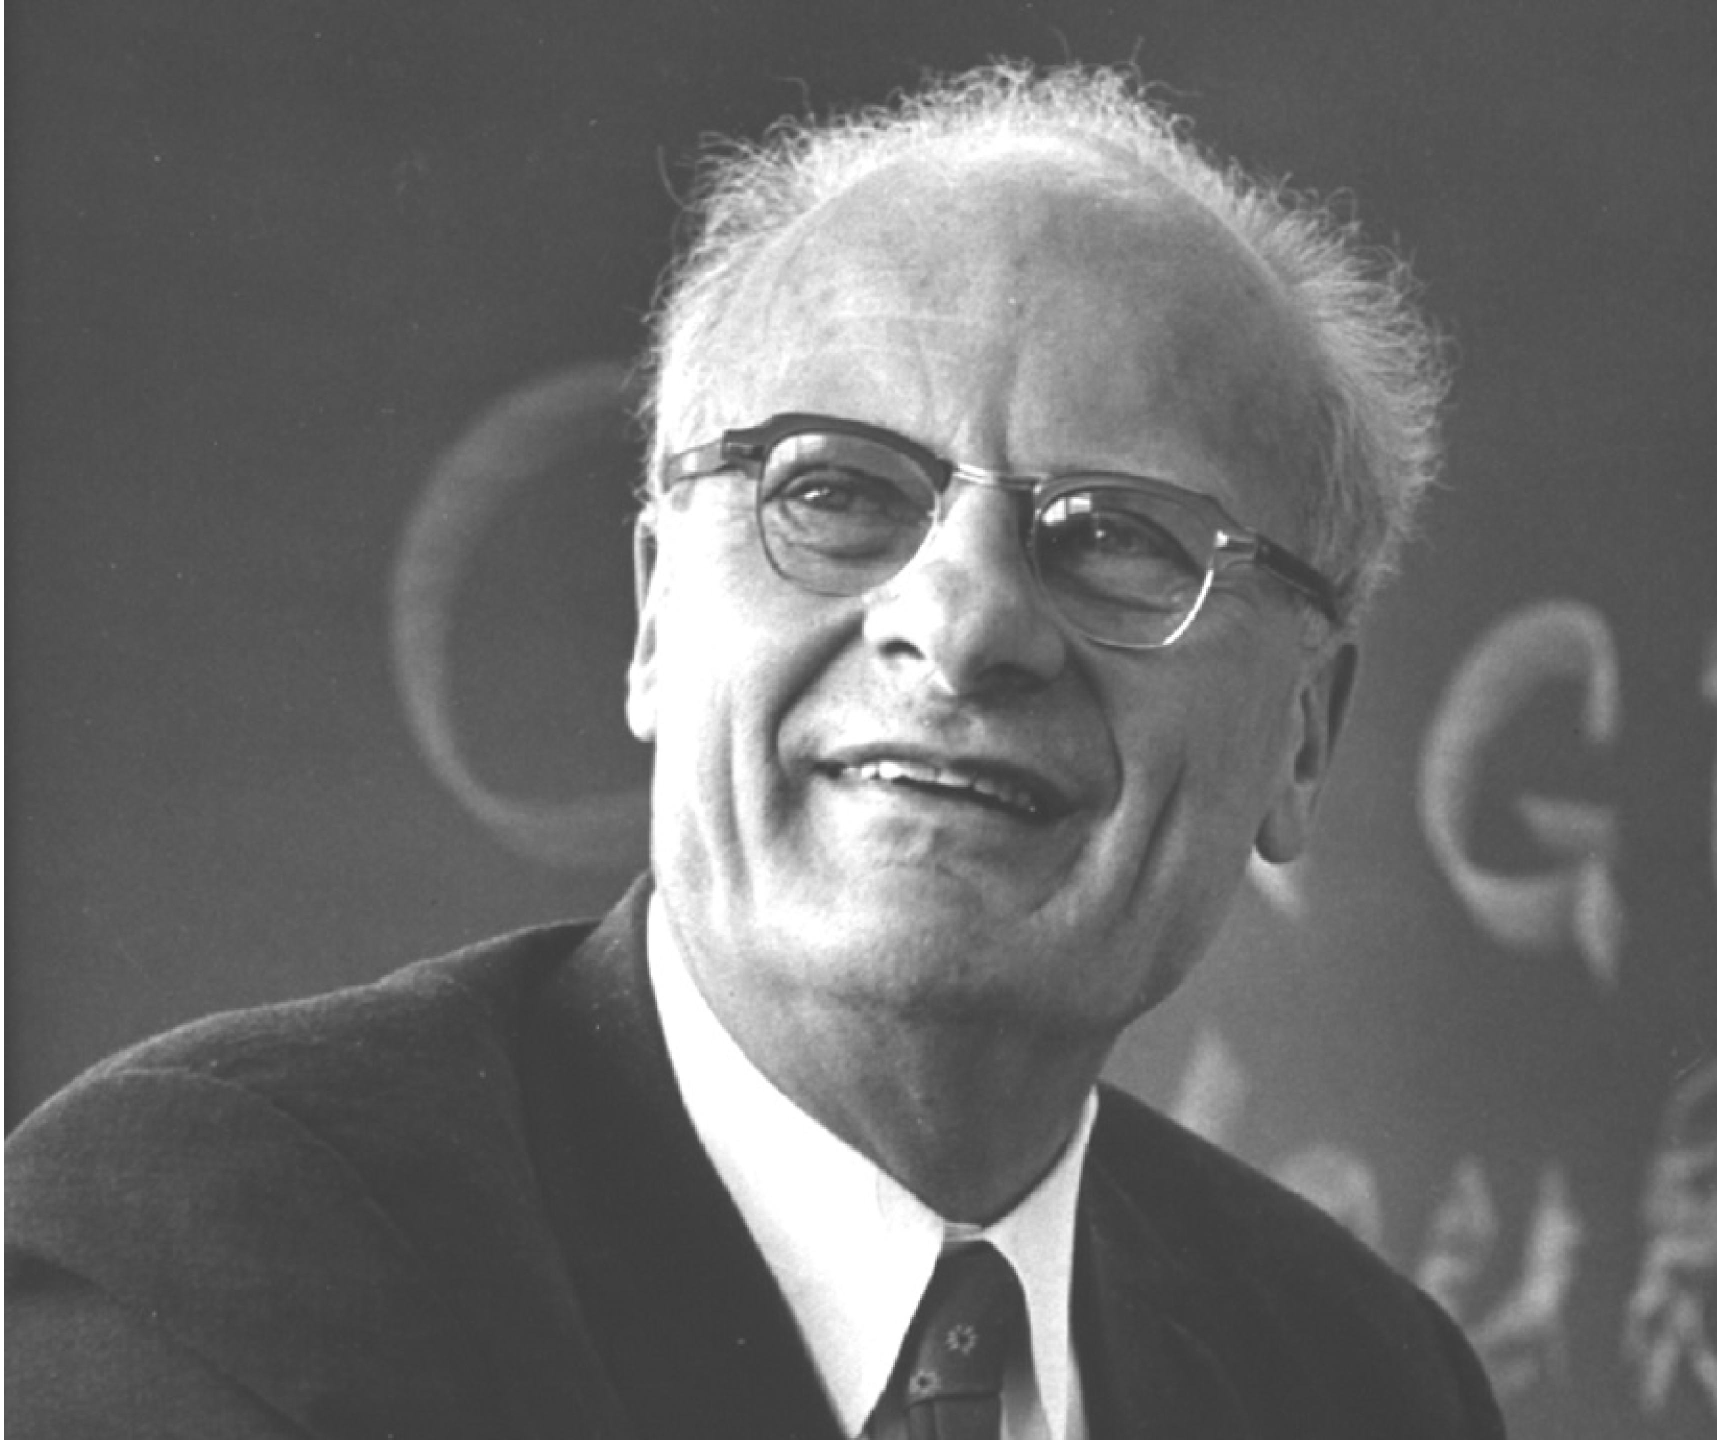

Quantum Information Science

Quantum sensing
Improving sensitivity and spatial resolution.

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.

Frontiers of Physics

short distance	long distance	complexity
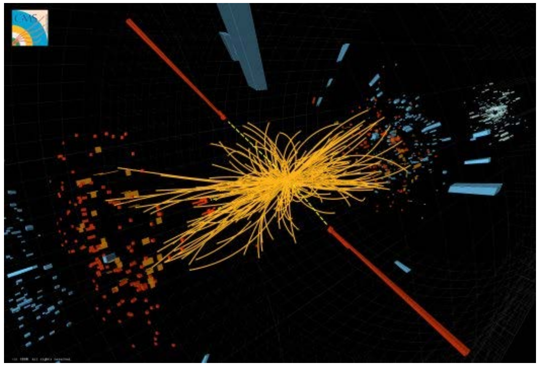		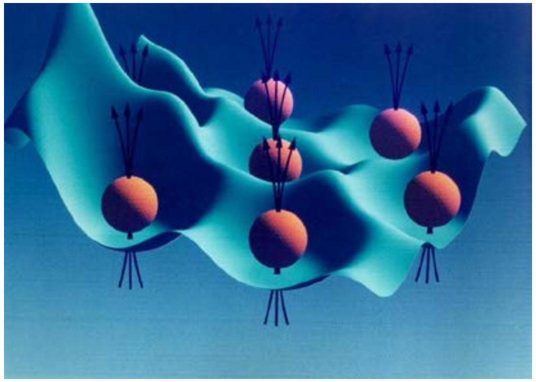
Higgs boson	Large scale structure	“More is different”
Neutrino masses	Cosmic microwave background	Many-body entanglement
Supersymmetry	Dark matter	Phases of quantum matter
Quantum gravity	Dark energy	Quantum computing
String theory	Gravitational waves	Quantum spacetime

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

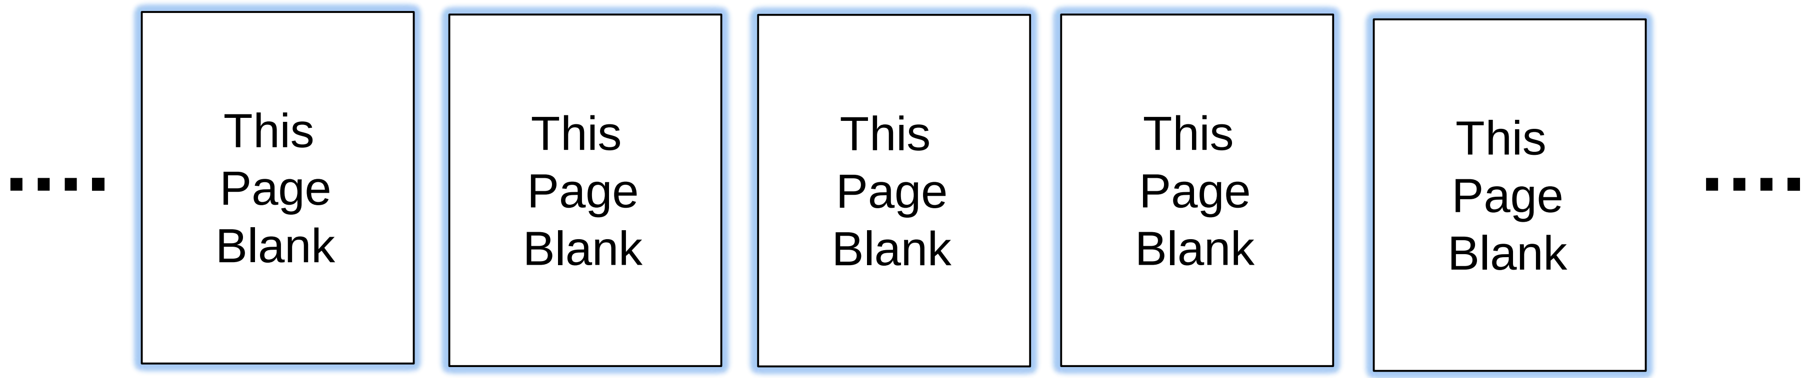

Nearly all the information in a typical entangled "quantum book" is encoded in the correlations among the "pages".

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

QUBIT

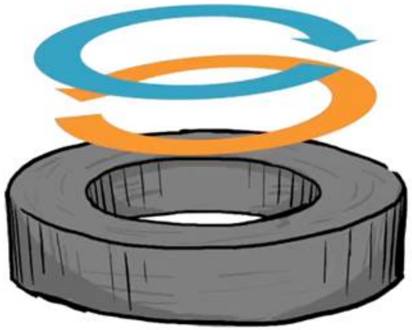
Persistent current in a superconducting circuit

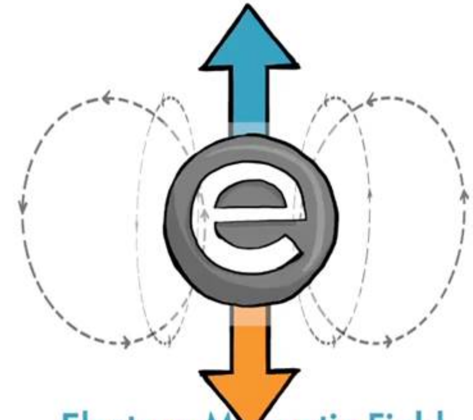
Electron Magnetic Field

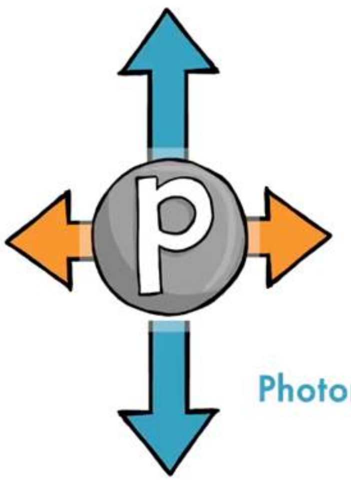
Photon polarization

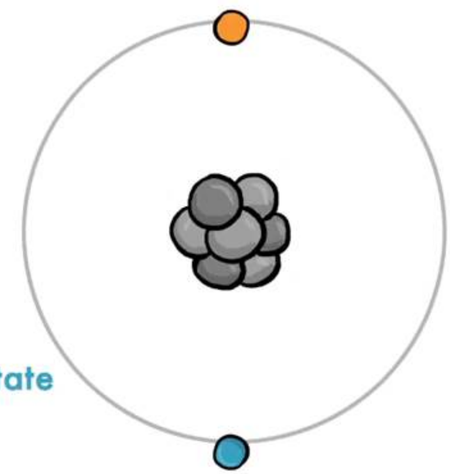
Atom Internal State

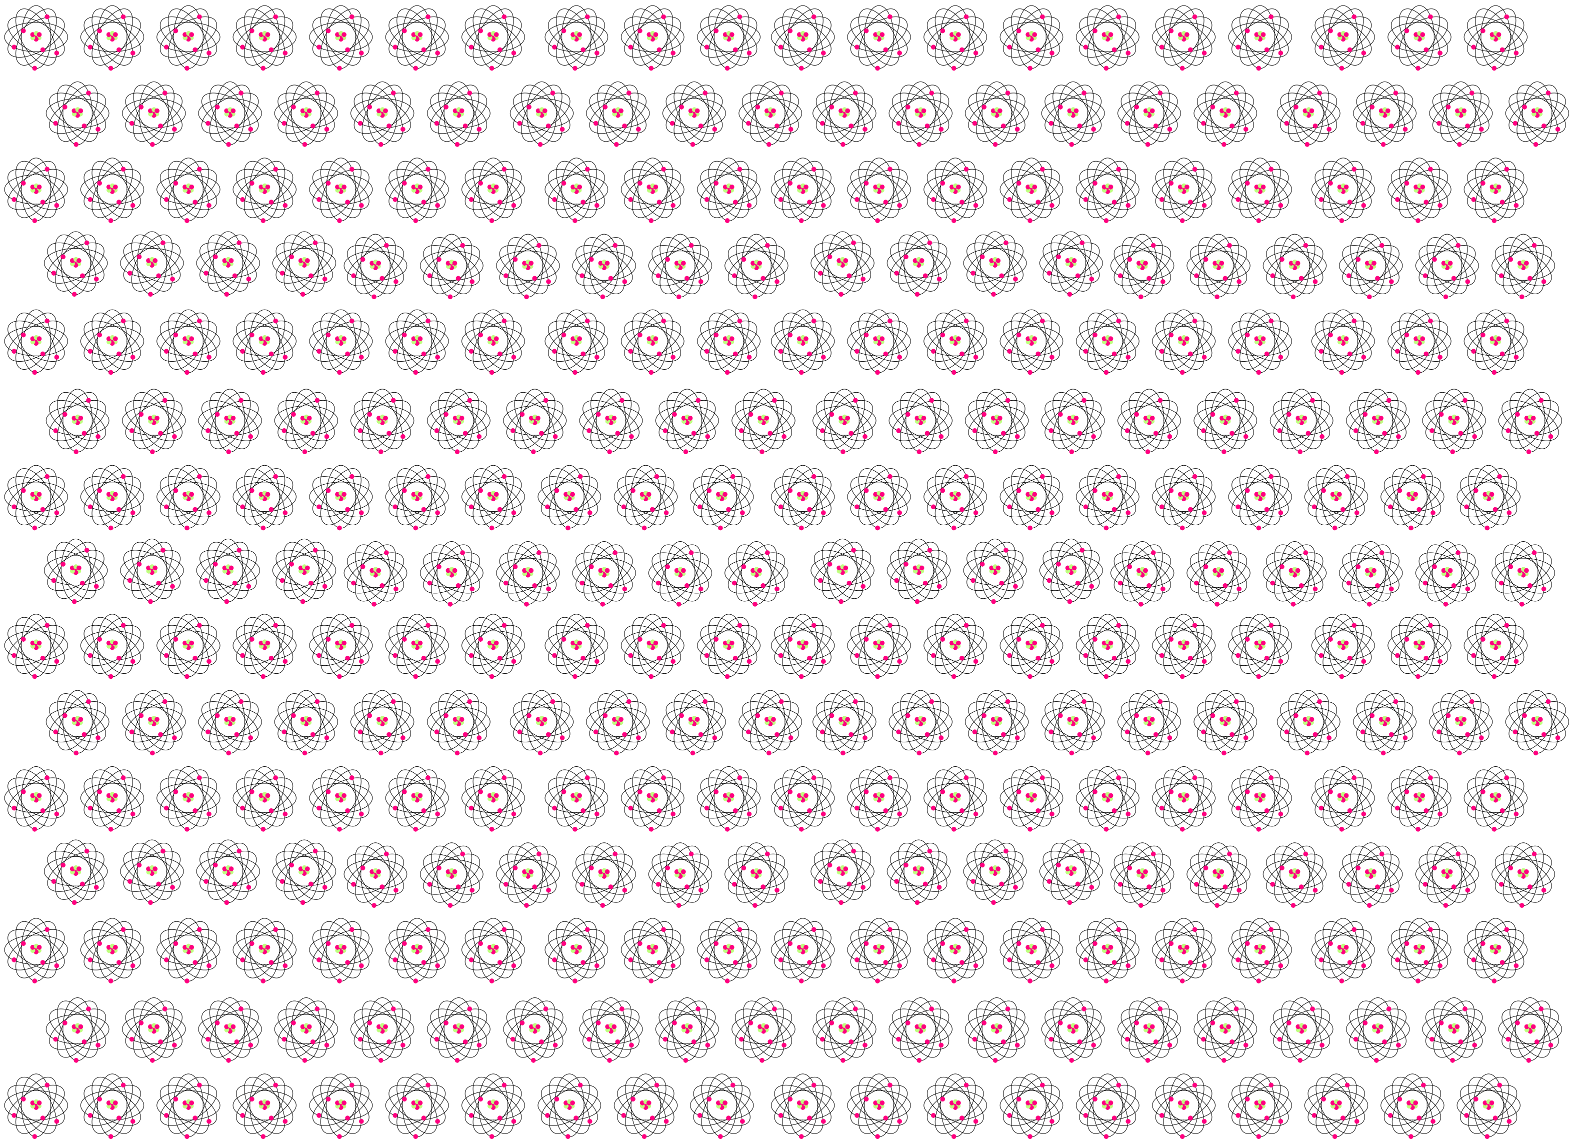

A complete description of a typical quantum state of just 300 qubits requires more bits than the number of atoms in the visible universe.

Why we think quantum computing is powerful

(1) Problems believed to be hard classically, which are easy for quantum computers. Factoring is the best known example.

(2) Complexity theory arguments indicating that quantum computers are hard to simulate classically.

(3) We don’t know how to simulate a quantum computer efficiently using a digital (“classical”) computer. The cost of the best known simulation algorithm rises exponentially with the number of qubits.

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

Problems

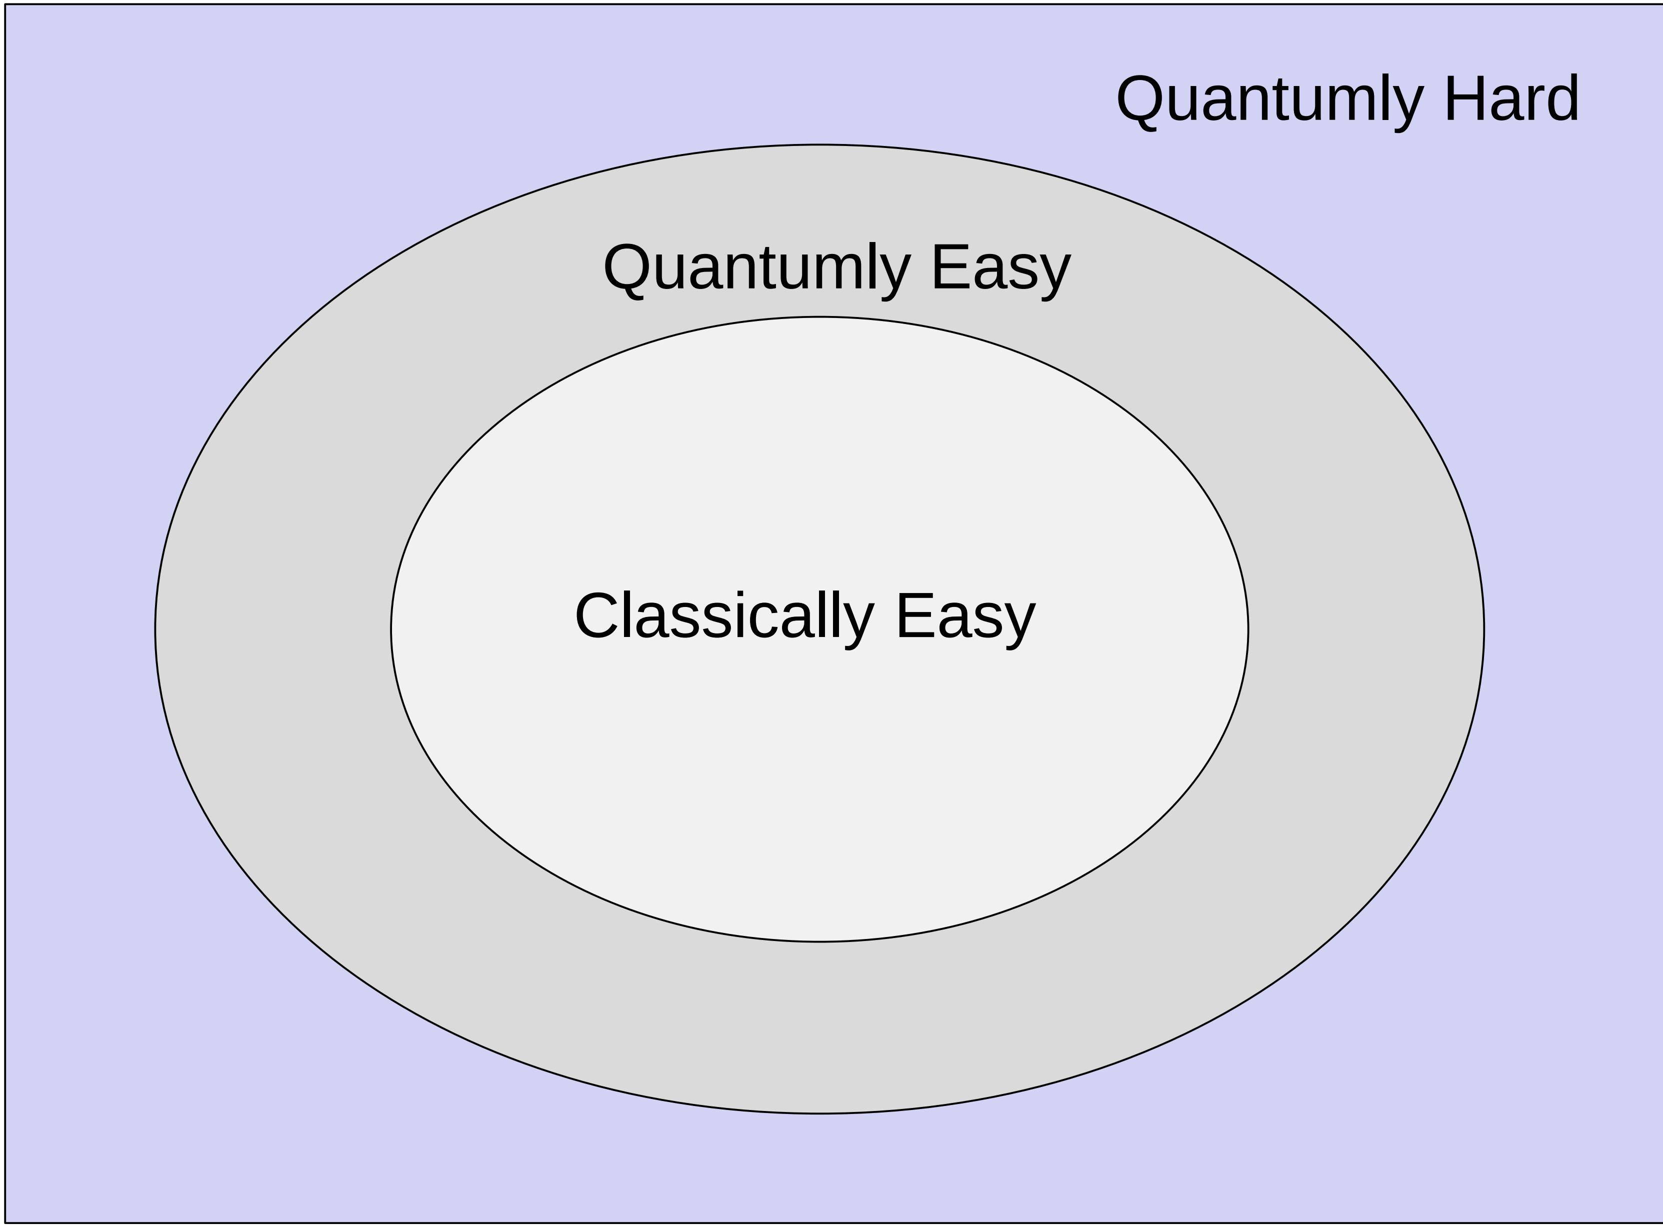

Problems

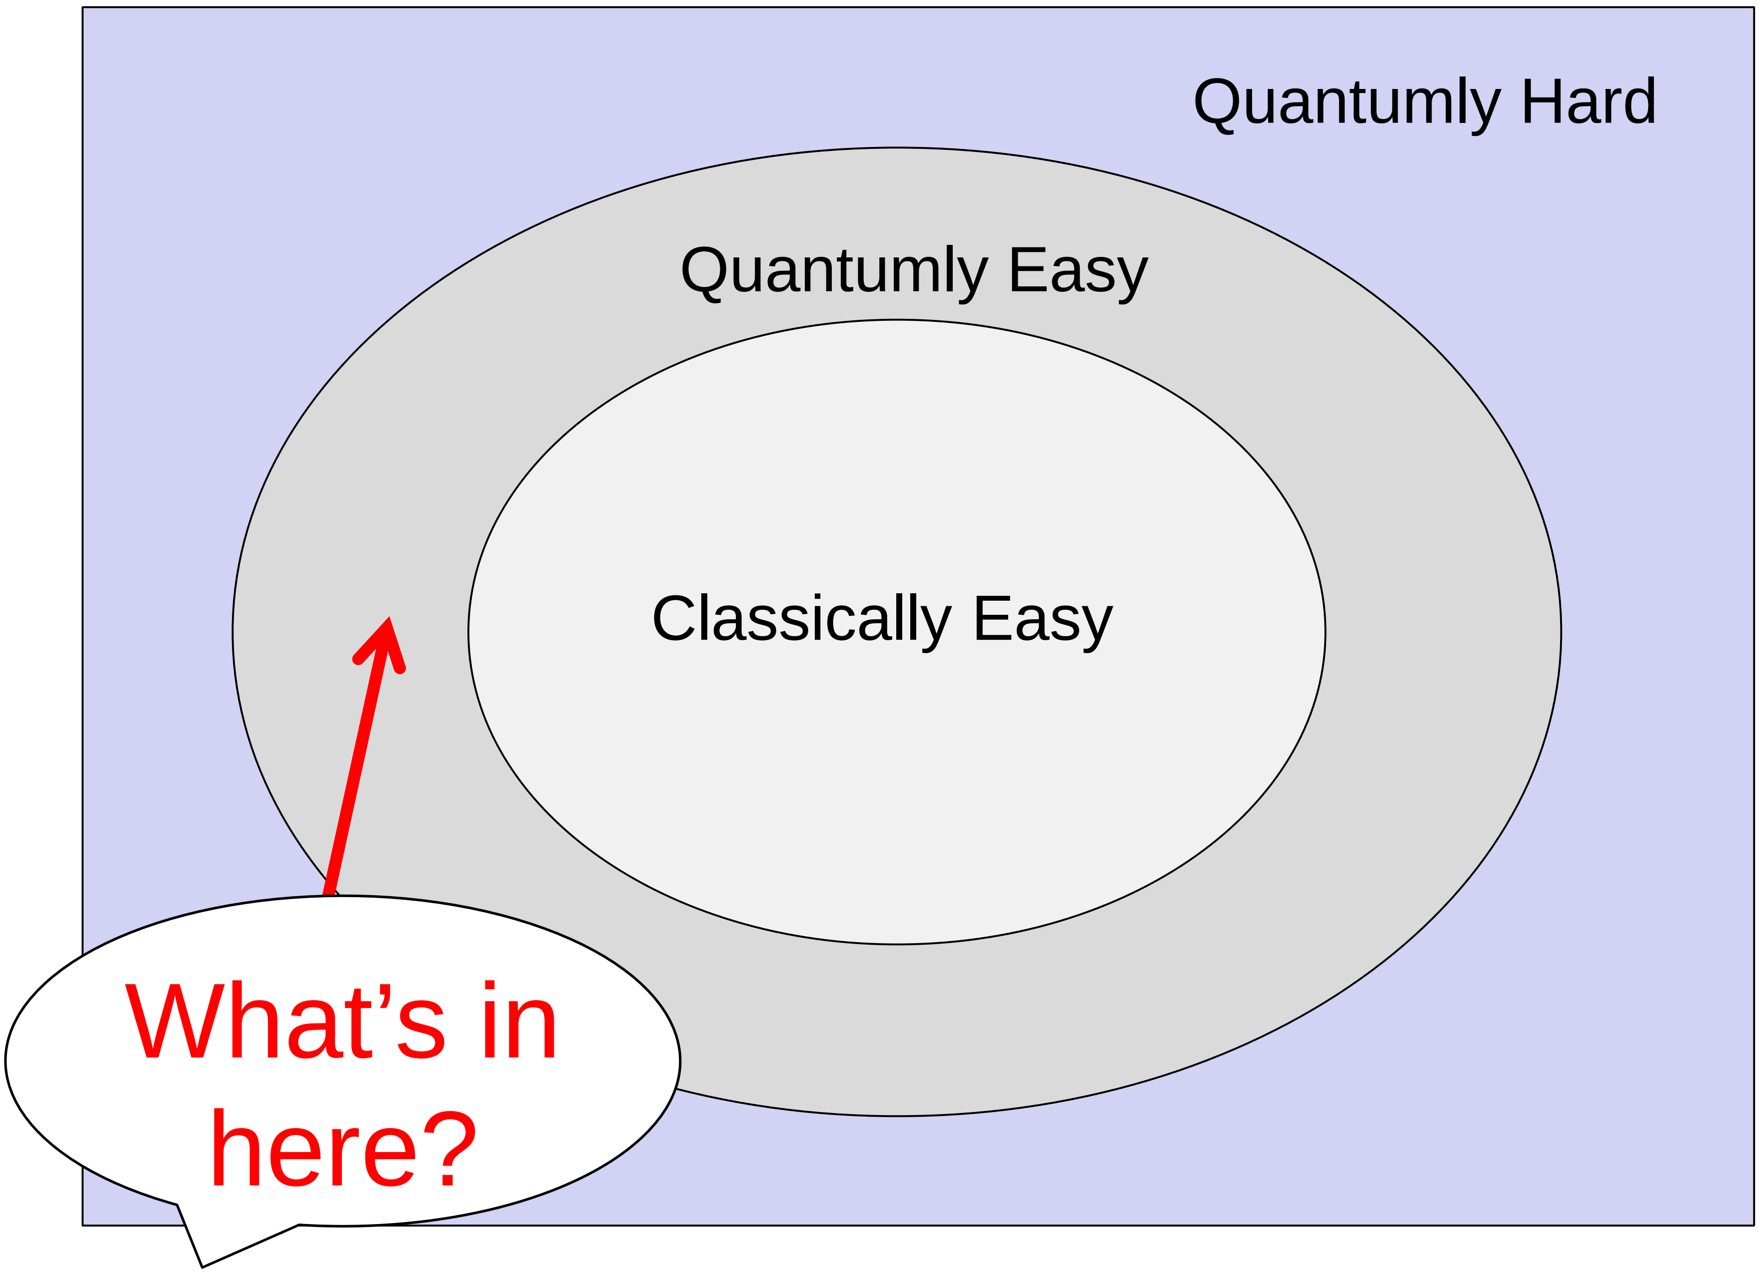
Quantumly Hard
Quantumly Easy
Classically Easy
What’s in here?

“The theory of everything?”

“The Theory of Everything is not even remotely a theory of every thing … We know this equation is correct because it has been solved accurately for small numbers of particles (isolated atoms and small molecules) and found to agree in minute detail with experiment. However, it cannot be solved accurately when the number of particles exceeds about 10. No computer existing, or that will ever exist, can break this barrier because it is a catastrophe of dimension … We have succeeded in reducing all of ordinary physical behavior to a simple, correct Theory of Everything only to discover that it has revealed exactly nothing about many things of great importance.”

R. B. Laughlin and D. Pines, PNAS 2000.

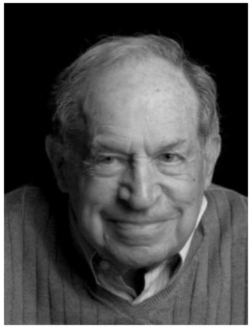

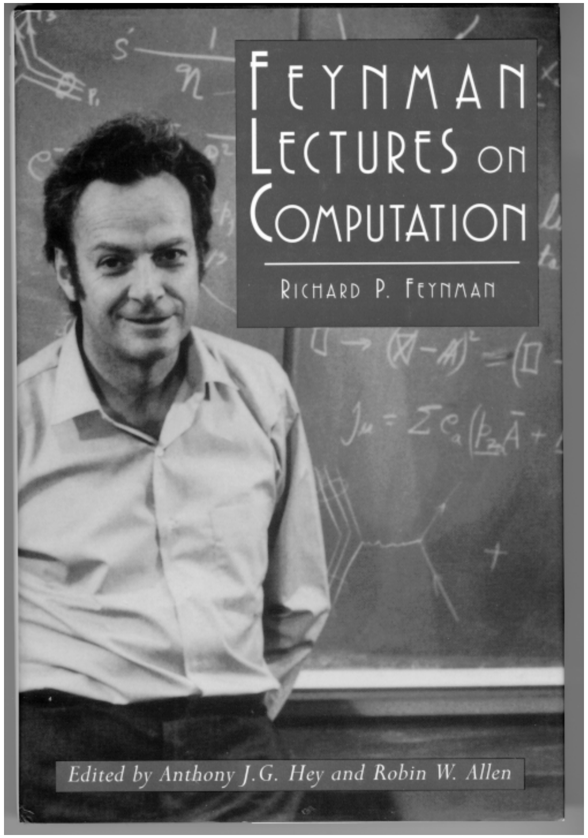

"Nature isn't classical, dammit, and if you want to make a simulation of Nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem because it doesn't look so easy."

R. P. Feynman, 1981

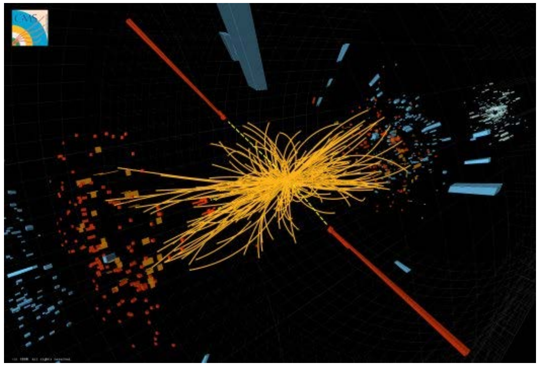

particle collision

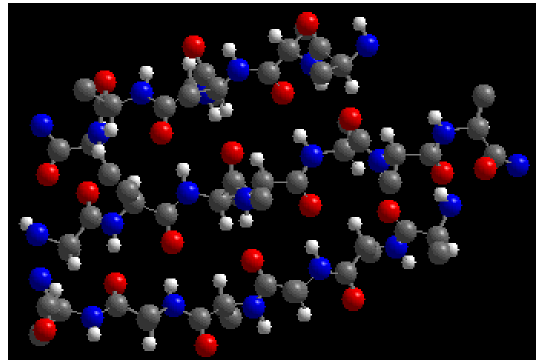

molecular chemistry

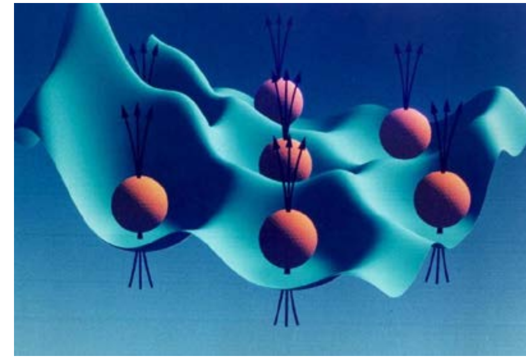

entangled electrons

A quantum computer can simulate efficiently any physical process that occurs in Nature.

(Maybe. We don’t actually know for sure.)

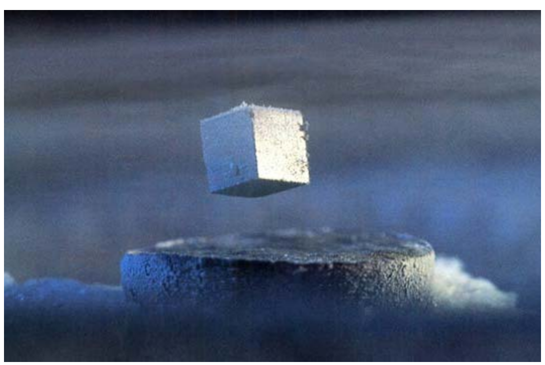

superconductor

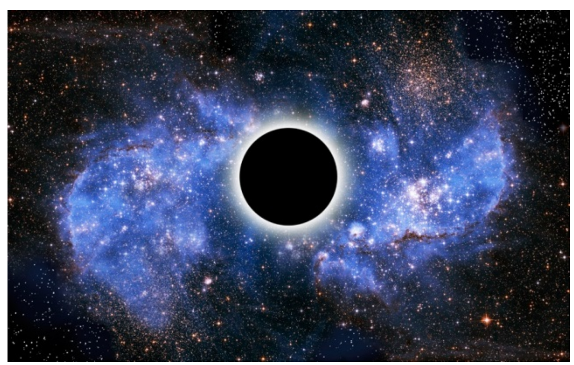

black hole

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.

Decoherence

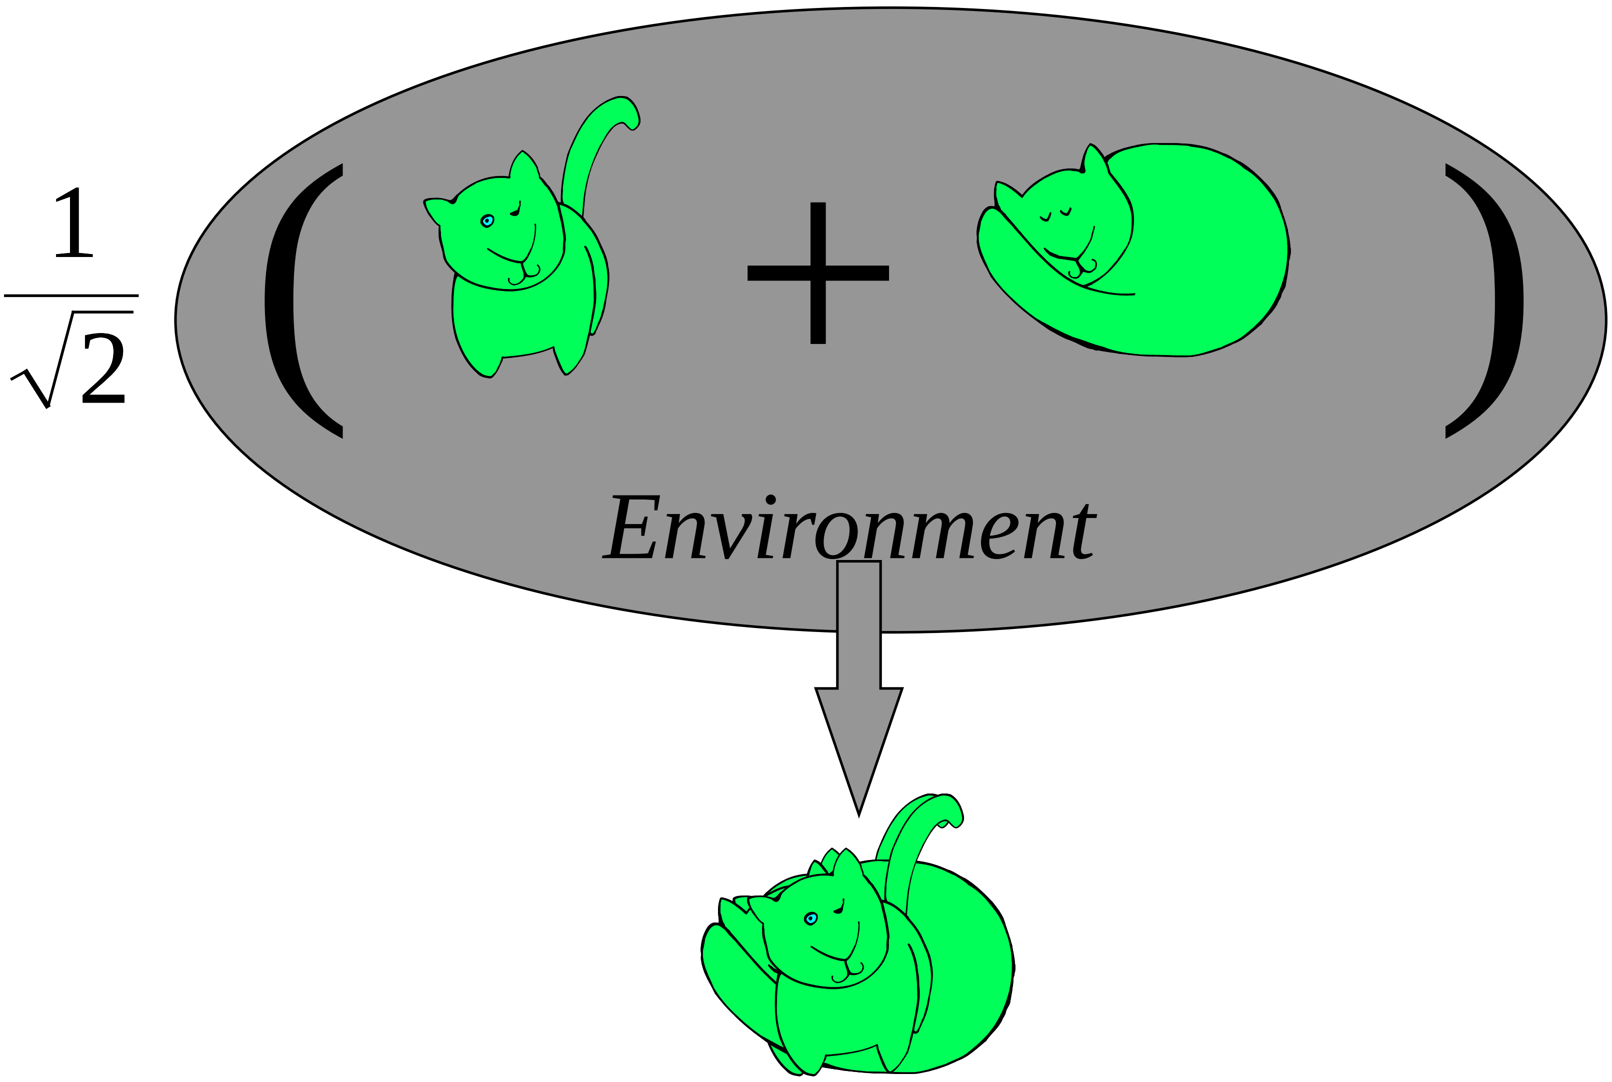

Decoherence

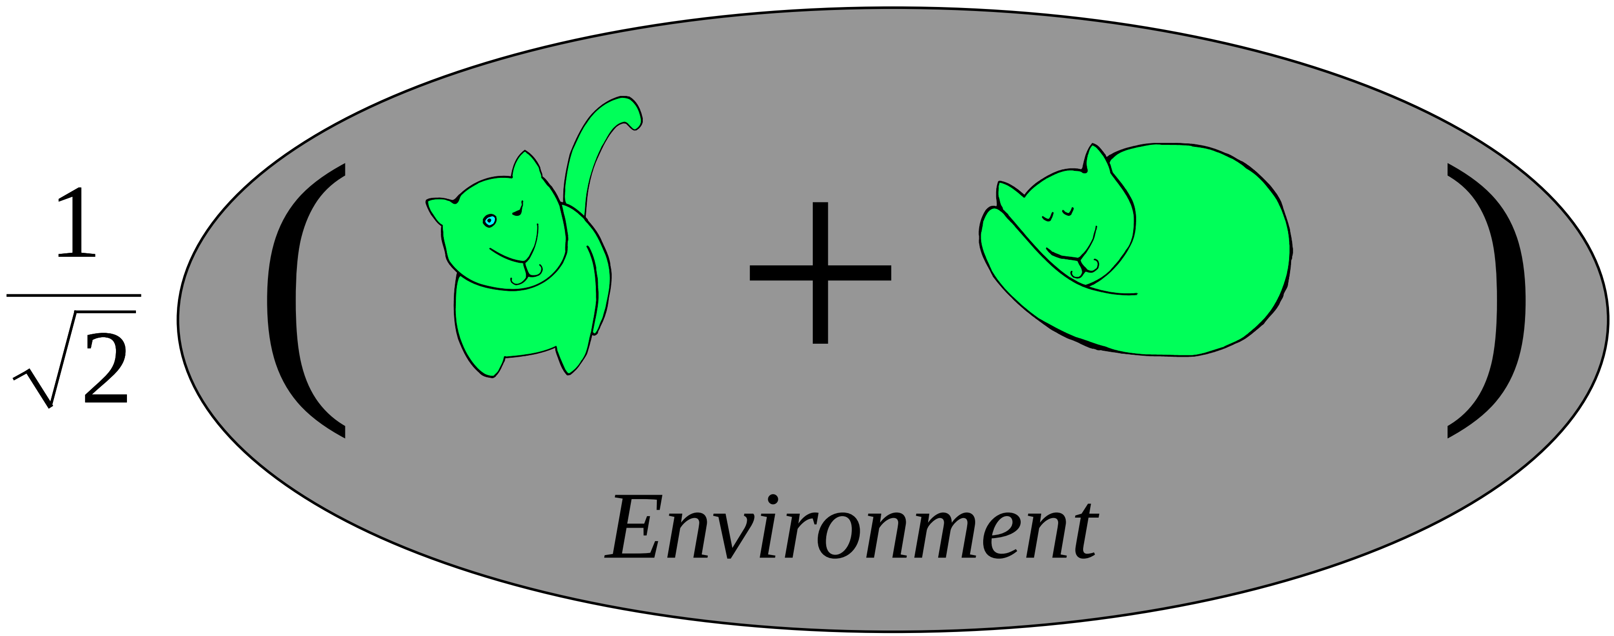

Decoherence explains why quantum phenomena, though observable in the microscopic systems studied in the physics lab, are not manifest in the macroscopic physical systems that we encounter in our ordinary experience.

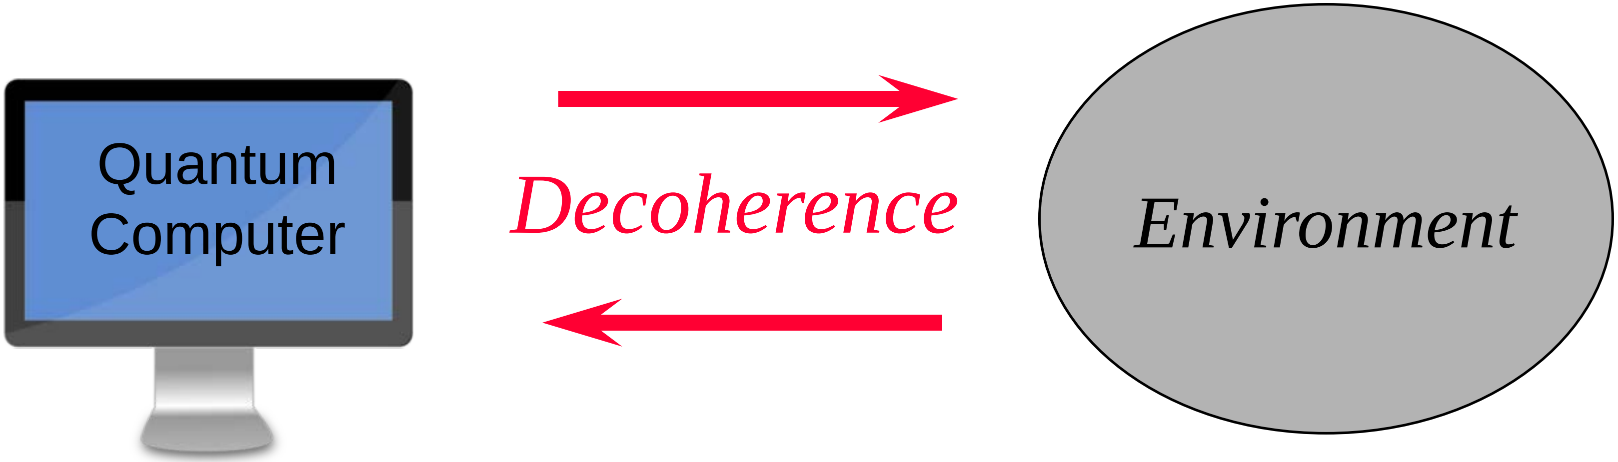

ERROR!

How can we protect a quantum computer from decoherence and other sources of error?

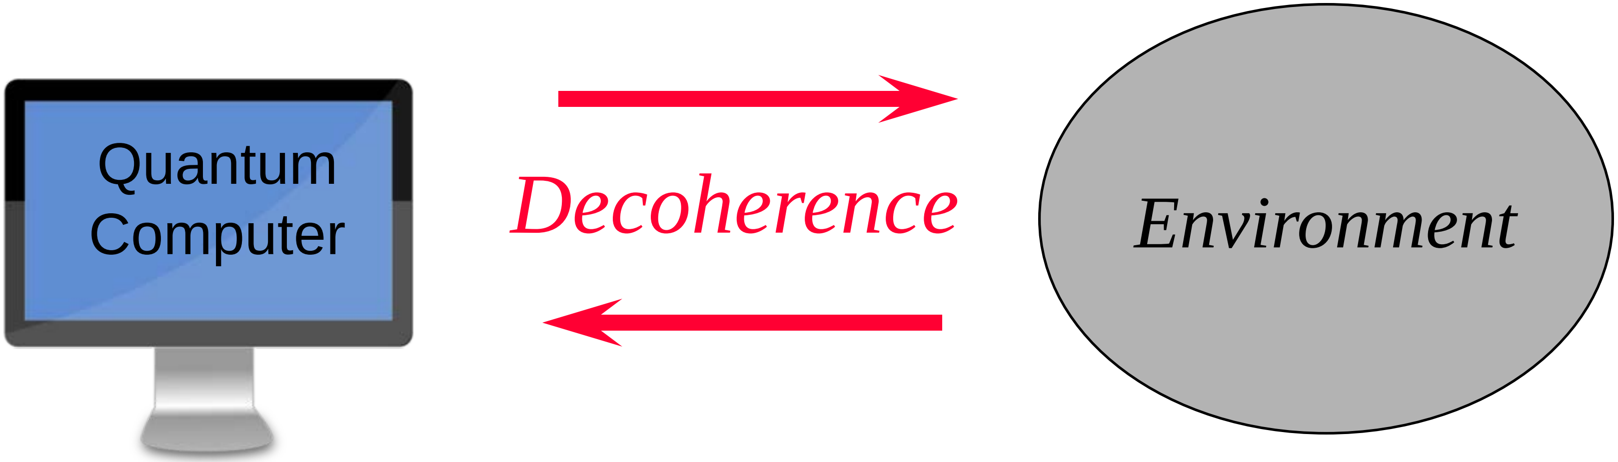

ERROR!

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

Quantum error correction

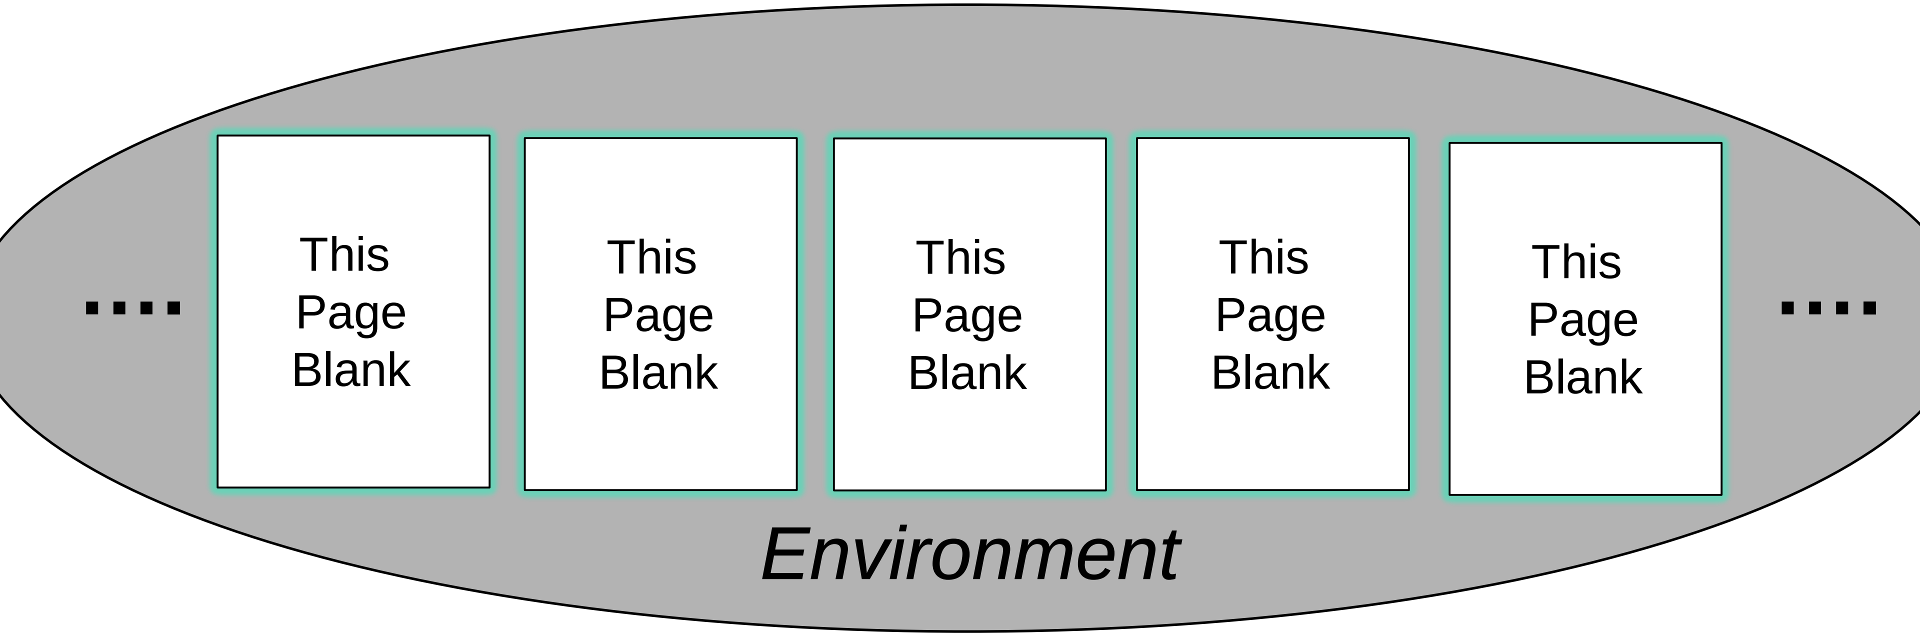

The protected "logical" quantum information is encoded in a highly entangled state of many physical qubits.

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

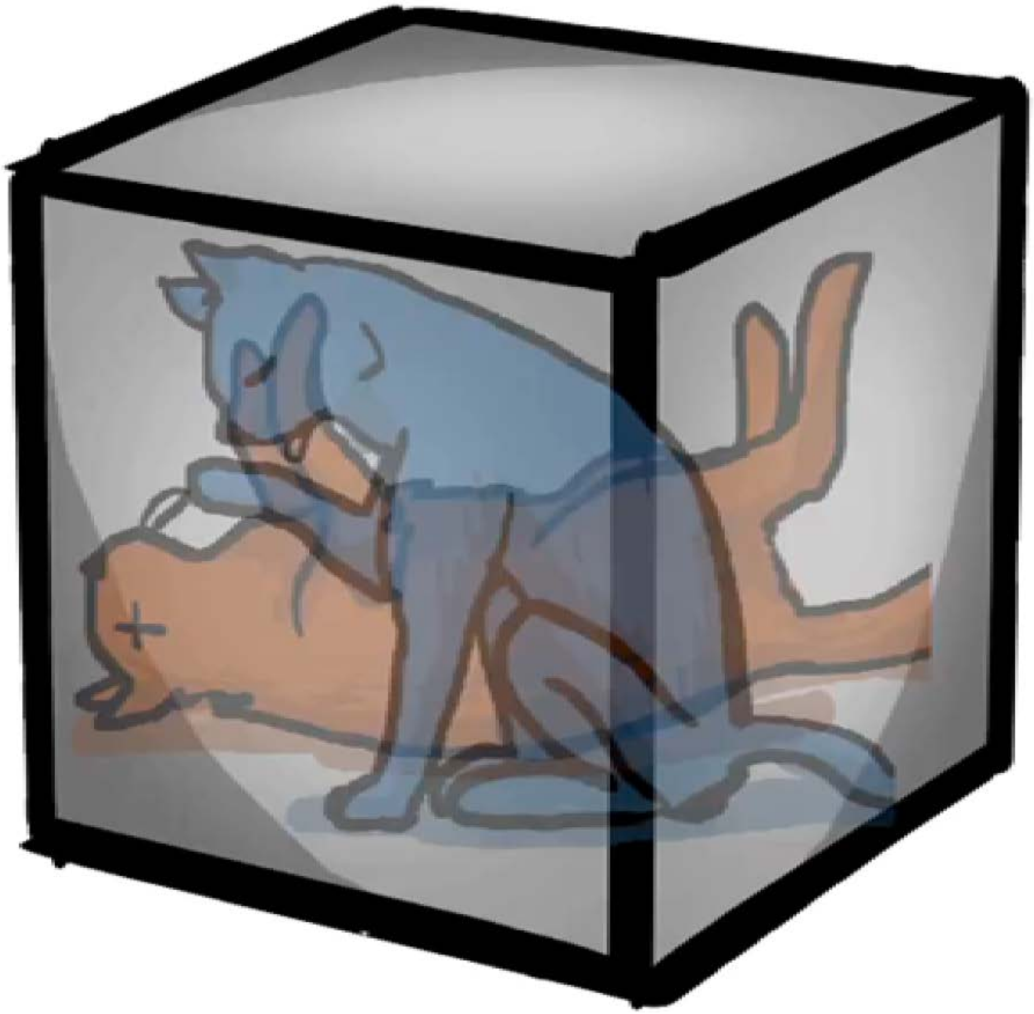

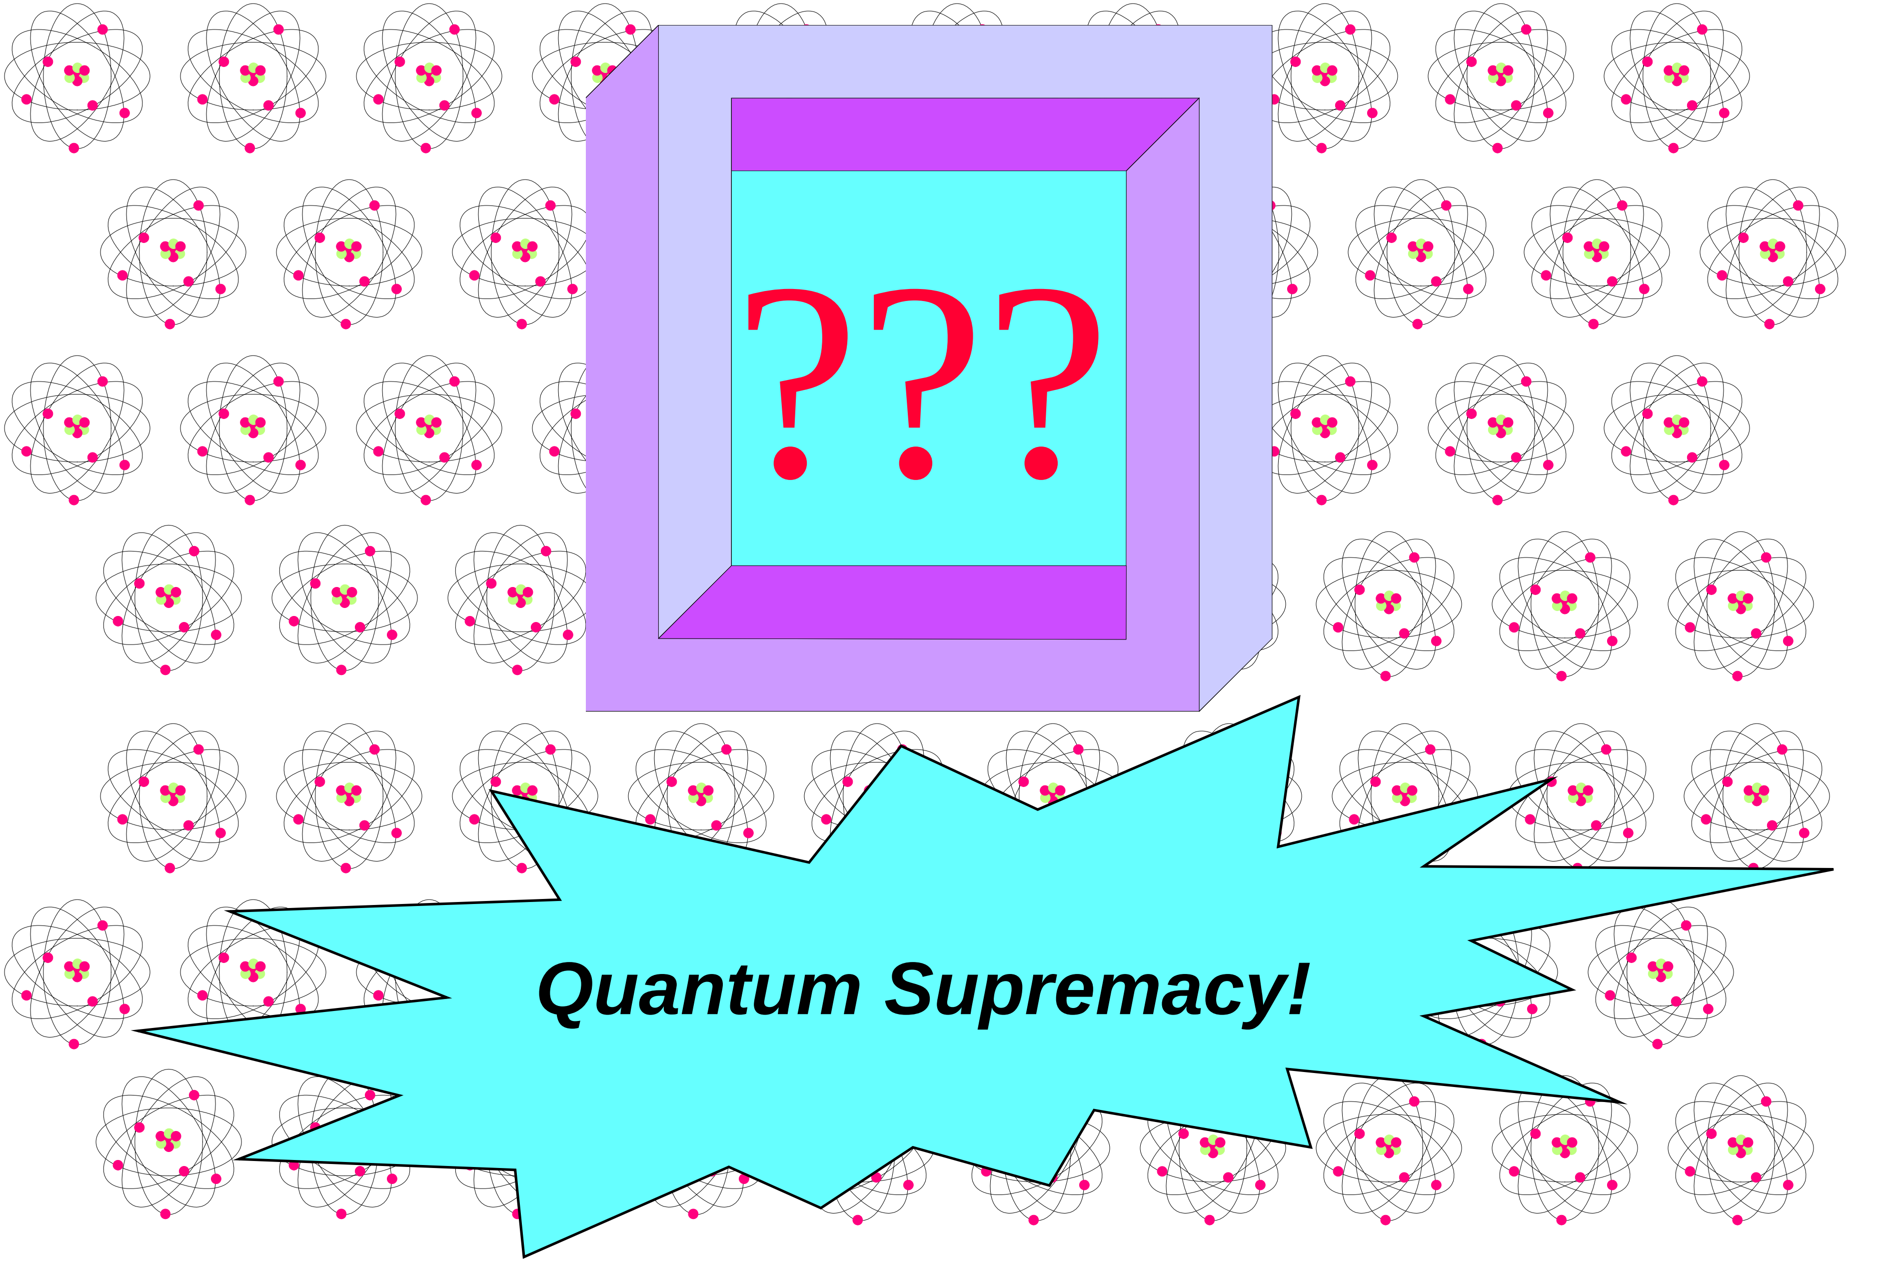
???
Quantum Supremacy!

Quantum computing in the NISQ Era

The (noisy) 50-100 qubit quantum computer is coming soon. (*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 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.

Potentially transformative scalable quantum computers may still be decades away. We're not sure how long it will take.

Qubit "quality"

The *number* of qubits is an important metric, but it is not the only thing that matters.

The *quality* of the qubits, and of the "quantum gates" that process the qubits, is also very important. All quantum gates today are noisy, but some are better than others. Qubit measurements are also noisy.

For today's *best* hardware (superconducting circuits or trapped ions), the probability of error per (two-qubit) gate is about 10^{-3}, and the probability of error per measurement is about 10^{-2} (or better for trapped ions). We don't yet know whether systems with many qubits will perform that well.

Naively, we cannot do many more than 1000 gates (and perhaps not even that many) without being overwhelmed by the noise. Actually, that may be too naïve, but anyway the noise limits the computational power of NISQ technology.

Eventually we'll do much better, either by improving (logical) gate accuracy using quantum error correction (at a hefty overhead cost) or building much more accurate physical gates, or both. But that probably won't happen very soon.

Other important features: The *time* needed to execute a gate (or a measurement). E.g., the two-qubit gate time is about 40 ns for superconducting qubits, 100 μs for trapped ions, a significant difference. Also qubit connectivity, fabrication yield, …

Hybrid quantum/classical optimizers

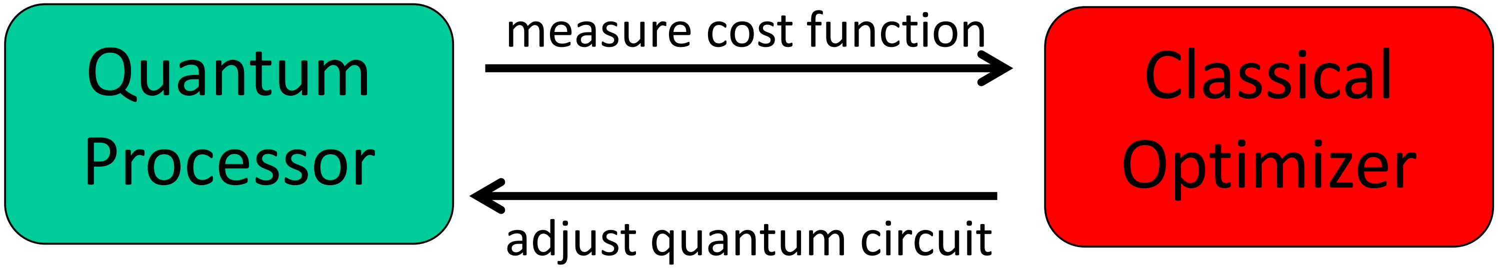

We don't expect a quantum computer to solve worst case instances of NP-hard problems, but it might find better approximate solutions, or find them faster.

Combine quantum evaluation of a cost function with a classical feedback loop for seeking a quantum state with a lower value.

Quantum approximate optimization algorithm (QAOA).
In effect, seek low-energy states of a classical spin glass.

Variational quantum eigensolvers (VQE).
Seek low energy states of a quantum many-body system with a local Hamiltonian.

Classical optimization algorithms (for both classical and quantum problems) are sophisticated and well-honed after decades of hard work. Will NISQ be able to do better? We can try it and see how well it works.

How quantum testbeds might help

Peter Shor: "You don't need them [testbeds] to be big enough to solve useful problems, just big enough to tell whether you can solve useful problems."

Classical examples:
Simplex method for linear programming: experiments showed it works well in practice before theorists could explain why.

Metropolis algorithm: experiments showed it's useful for solving statistical physics problems before theory established criteria for rapid convergence.

Deep learning. Mostly tinkering so far, without much theory input.

Possible quantum examples:
Quantum annealers, approximate optimizers, variational eigensolvers, … playing around may give us new ideas.

But in the NISQ era, imperfect gates will place severe limits on circuit size. In the long run, quantum error correction will be needed for scalability. In the near term, better gates might help a lot!

What can we do with, say, < 100 qubits, depth < 100? We need a dialog between quantum algorithm experts and application users.

Quantum annealing

The D-Wave machine is a (very noisy) 2000-qubit *quantum annealer* (QA), which solves optimization problems. It *might* be useful. But we have no convincing theoretical argument that QAs are useful, nor have QA speedups been demonstrated experimentally.

Theorists are more hopeful that a QA can achieve speedups if the Hamiltonian has a "sign problem" (is "non-stoquastic"). Present day QAs are stoquastic, but non-stoquastic versions are coming soon.

Assessing the performance of QA may already be beyond the reach of classical simulation, and theoretical analysis has not achieved much progress. Further experimentation should clarify whether QAs actually achieve speedups relative to the best classical algorithms.

QAs can also be used for solving quantum simulation problems rather than classical optimization problems.

Noise-resilient quantum circuits

For near-term applications, noise-resilience is a key consideration in quantum circuit design.

For a generic circuit with G gates, a single faulty gate might cause the circuit to fail. If the probability of error per gate is not much larger than 1/G, we have a reasonable chance of getting the right answer.

But, depending on the nature of the algorithm and the circuit that implements it, we might be able to tolerate a much larger gate error rate.

For some physical simulation problems, a constant probability of error per measured qubit can be tolerated, and the number of circuit locations where a fault can cause an error in a particular qubit is relatively small. This could happen because the circuit has low depth, or because an error occurring at an earlier time decays away by a later time.

Circuits with good noise-resilience (based on tensor network constructions like MERA) are among those that might be useful for solving quantum optimization problems using variational quantum eigensolvers (VQE), improving the prospects for outperforming classical methods during the NISQ era (Kim and Swingle 2017).

Quantum machine learning?

Machine learning is transforming technology and having a big impact on the way we do science as well, so it is natural to wonder about the potential of combining deep learning with quantum technology.

Perhaps a quantum deep learning network can be trained more efficiently, e.g. using a smaller training set. We don't know. We'll have to try it to see how well it works.

High-dimensional classical data can be encoded very succinctly in a quantum state. In principle log N qubits suffice to represent a N-dimensional vector. Such "quantum Random Access Memory" (= QRAM) *might* have advantages for machine learning applications.

However, many proposed quantum machine learning applications are hampered by input/output bottlenecks.

Loading classical data into QRAM is slow, nullifying the potential advantage, and the output is a quantum state, and only a limited amount of information can be accessed by measuring the state.

Perhaps it's more natural to consider quantum inputs / outputs; e.g. better ways to characterize or control quantum systems. Quantum networks might have advantages for learning about *quantum* correlations, rather than classical ones.

Quantum linear algebra

QRAM: an N-component vector b can be encoded in a quantum state $|b\rangle$ of log N qubits.

Given a classical N X N input matrix A, which is sparse and well-conditioned, and the quantum input state $|b\rangle$, the HHL (Harrow, Hassidim, Lloyd 2008) algorithm outputs the quantum state $|y\rangle = |A^{-1} b\rangle$, with a small error, in time O(log N). The quantum speedup is exponential in N.

Input vector $|b\rangle$ and output vector $|y\rangle = |A^{-1} b\rangle$ are quantum! We can sample from measurements of $|y\rangle$.

If the input b is classical, we need to load $|b\rangle$ into QRAM in polylog time to get the exponential speedup (which might not be possible). Alternatively the input b may be computed rather than entered from a database.

HHL is BQP-complete: It solves a (classically) hard problem unless BQP=BPP.

Applications typically require pre-conditioning, which can be expensive. The problem becomes easier when the matrix A has low rank.

HHL is not likely to be feasible in the NISQ era.

Quantum simulation

We're confident *strongly correlated* (highly entangled) materials and large molecules are hard to simulate classically (because we have tried hard and have not succeeded).

Quantum computers will be able to do such simulations, though we may need to wait for scalable fault tolerance, and we don't know how long that will take.

Potential (long-term) applications include pharmaceuticals, solar power collection, efficient power transmission, catalysts for nitrogen fixation, carbon capture, etc. These are not likely to be fully realized in the NISQ era.

Classical computers are especially bad at simulating quantum *dynamics* --- predicting how highly entangled quantum states change with time. Quantum computers will have a big advantage in this arena. Physicists hope for noteworthy advances in quantum dynamics during the NISQ era.

For example: Classical *chaos theory* advanced rapidly with onset of numerical simulation of classical dynamical systems in the 1960s and 1970s. Quantum simulation experiments may advance the theory of *quantum* chaos. Simulations with ~ 100 qubits could be revealing, if not too noisy.

Digital vs. Analog quantum simulation

An *analog quantum simulator* is a quantum system of many qubits whose dynamics resembles the dynamics of a model system we wish to study. A *digital quantum simulator* is a gate-based universal quantum computer, which can be used to simulate any physical system of interest when suitably programmed.

Analog quantum simulation has been an active research area for 15 years or more; digital quantum simulation is just getting started now.

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

Although they are becoming more sophisticated and controllable, analog simulators are limited by imperfect control. They are best suited for studying "universal" properties of quantum systems which are hard to access in classical simulations, yet sufficiently robust to be accessible 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.

Surprising dynamics in quantum platforms

How do excited quantum systems converge to thermal equilibrium? Typically, information which is initially accessible locally spreads quickly, hidden by quantum entanglement. The effects of a perturbation become invisible to local probes.

There is a notable exception, called *many-body localization*. Systems that are strongly disordered are less entangled and thermalize very slowly.

Experiments with a 51-atom quantum simulator discovered an unexpected intermediate case. "Type A" quantum states do thermalize quickly, while "Type B" do not --- instead Type B states undergo long lived coherent oscillations due to repulsive interactions (Harvard group 2017).

This seems rather remarkable because Type A and Type B states are otherwise very similar.

The Type B states are the signature of a new class of quantum matter far from equilibrium, exhibiting "quantum many-body scars" --- previously observed for single-particle systems, but not many-body systems (Turner et al. 2018).

Programmable analog quantum simulators

Between digital and analog. Not gate based, but Hamiltonian is rapidly tunable.

Hamiltonian control errors, if reproducible, need not limit power of a variational scheme.

For example, control the native Hamiltonian of an ion trap, with all-to-all coupling.

Recent application by the Innsbruck group (2018): accurate measurement of the low-energy spectrum of a 20-site lattice model (Schwinger model).

Evolve with H_1 for time t_1, H_2 for time t_2, etc. Then measure at the end. Classically optimize over variational parameters to find expectation value of the model Hamiltonian H.

Self verification: Minimize expectation value of $(H-E)^2$, check it's with *zero* when E is an eigenvalue. (Decoherence does not limit accuracy for this system size.)

Should remain feasible with ~ 50 ions.

For quantum advantage: entangling dynamics or higher-dimensional systems.

Quantum imaginary time evolution (QITE)

For hybrid quantum / classical algorithms like QAOA and VQE, the classical parameter optimization is challenging! We should seek alternative ways to explore low energy states of many-body quantum systems with NISQ devices.

Simulating imaginary time evolution $\exp(-\beta H)$ is a powerful classical algorithm for preparing ground states, limited by exponential cost of storing a quantum state.

Nonunitary transformations on a QC require ancilla systems and postselection.

Except maybe not …. (Chan et al. 2019). Instead, it suffices to find the result of applying $\exp(-\varepsilon H)$ to the input state, achieved by some unitary and a renormalization of the state. Find the unitary by state tomography and solving a linear system.

Efficient if the correlation length stays finite, and relatively low depth circuits may suffice.

Not a panacea, but a promising alternative to low-depth VQE and (expensive) phase estimation.

The steep climb to scalability

NISQ-era quantum devices will not be protected by quantum error correction. Noise will limit the scale of computations that can be executed accurately.

Quantum error correction (QEC) will be essential for solving some hard problems. But QEC carries a high overhead cost in number of qubits & gates.

This cost depends on both the hardware quality and algorithm complexity. With today's hardware, solving (say) useful chemistry problems may require hundreds to thousands of physical qubits for each protected logical qubit.

To reach scalability, we must cross the daunting "quantum chasm" from hundreds to millions of physical qubits. This may take a while.

Advances in quantum gate fidelity, systems engineering, algorithm design, and error correction protocols can hasten the arrival of the fully fault-tolerant quantum computer.

Quantum-safe privacy

(1) How long will current systems (e.g. RSA-2048) be safe against quantum attack?
(2) How long will it take to deploy quantum safe alternatives (e.g. lattice based)?
(3) How long should keys be secure?

What's the solution? (Longer keys will not suffice.)

(A) Post-quantum cryptography? Works on conventional hardware, but how safe are the computational assumptions?

(B) Quantum cryptography? New quantum infrastructure needed for global communication. But no computational assumptions.

Some users will prefer (A), others might choose (B).

Further research/development focused on quantum resistance will strengthen (A). Standards will be needed; that takes time.

Satellite-based QKD and quantum repeaters will boost (B).

Cryptographers should be quantum savvy!

Blockchain: Proof of work is hash-based, so pretty safe. RSA/ECC-based digital signature is vulnerable to Shor's algorithm, if broken before transaction is placed on the blockchain.

Quantum networks

(1) End nodes, (2) quantum channels, (3) quantum repeaters, (4) classical channels.

Quantum channel: photons sent through free space or fiber.

Fiber: 17 dB per 100 km. And not much improvement for 20 years. So 100 km is possible, 1000 km is impossible.

Extending the range. Satellite based or ground based (repeaters).

For repeater, quantum memory is needed (cannot measure & resend.) Can "purify" and "swap" entanglement. Easier than fault-tolerant quantum computing. E.g. might use atomic ensembles or rare earth ions in crystals.

End node need not be trusted (in "device independence" protocol).

Might need transducers: e.g. traveling optical photons stored in quantum memory at microwave frequency. These could be optomechanical devices.

Other applications for quantum networking: scalable and secure multiparty quantum computing, global quantum sensors and clocks, etc.

Quantum sensing

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.

What's coming? Quantum enhancements from entanglement, squeezing, error correction. Hybrid quantum technologies for multi-modal function.

Quantum radar (a.k.a. quantum illumination). Entanglement enhances signal to noise. Transduction from microwave to visible.

What quantum states of multiple sensors provide the best sensing enhancements? Exploring this is a potential task for quantum machine learning.

Better sensing might be employed to detect noise in quantum devices, and improve noise mitigation.

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

Quantum sensing

Detecting axions and other (low-mass) dark matter candidates:
-- Superconducting nanowire detectors for hidden photons, axions, etc.
-- Nondemolition measurement of single microwave photons using transmons.
Magnetic tunneling junction arrays for ultrafast magnetic field detection.
-- Entangled-states in ion traps for detection of weak forces at the Heisenberg limit.
-- Quantum sensors based on photon upconversion from RF to microwave.
-- Distinguish recoil from WIMPS and neutrinos, e.g. using NV centers in diamond.

Multidisciplinary effort: HEP theory and experiment, QIS experimental strategies, new materials and platforms

Other HEP goals: Detecting drift of fundamental constants, electric dipole moments.

LIGO: Improved sensitivity by frequency-dependent squeezing of the light. With nonlinear crystals now, with optomechanical devices eventually.

Distantly separated optical telescopes which share entanglement can perform interferometry by teleporting photons (someday). Detect a city on another planet.

Quantum simulation of quantum field theories.

Beyond Euclidean Monte Carlo on classical computers?

-- Improved predictions for QCD backgrounds in collider experiments

-- Equation of state for nuclear matter, quark gluon plasma, early universe

-- Exploration of other strongly-coupled theories, beyond-standard-model physics

-- Stepping stone to quantum gravity, e.g. through holographic duality

-- New insights!

No sign problem!

-- Sample accurately from outgoing states in simulation of scattering event.

-- Real-time correlation functions, including at nonzero temp and chem potential.

-- Transport properties, far from equilibrium phenomena.

Quantum simulation of quantum field theories.

Where are we now?

-- Resource scaling estimates (number of qubits and gates) for scattering simulations in scalar and Yukawa theories.
-- *Classical* tensor-network simulation of massive 1D QED.
Static and dynamic studies of strings and string breaking.
-- Few-site *quantum* simulations of 1D QED with trapped ions and superconducting circuits.
-- Proposals for analog simulation using ultracold atoms, etc.
-- In progress: Classical and quantum simulations of nonabelian gauge symmetry, higher dimensions.

Prospects for quantum advantage (e.g. in one dimension)?

-- Beyond what can be simulated classically using tensor networks?
-- Classical simulation methods fail for highly entangled states.
-- High-energy scattering with multiple particle production.
-- Dynamics after a quench, or many successive scattering events.

Quantum Gravity and Quantum Information

Why quantum gravity?

(1) Erect a complete theory of fundamental interactions.
(2) Resolve deep puzzles about the quantum physics of black holes.
(3) Understand the very early history of the universe.

Anti-de Sitter space

-- We live in de Sitter space, which has no boundary (positive dark energy).
-- AdS space has a boundary; this makes quantum mechanics easier.
-- Eventually we'll need to learn how to do quantum mechanics in dS. It's hard.

Holographic duality

-- Amazingly, quantum gravity in AdS is equivalent to quantum field theory (without gravity) on its boundary.
-- Remarkably, geometry in the bulk spacetime is encoded as quantum entanglement in the boundary theory. ("Emergent geometry")
-- Delightfully, the mapping from bulk to boundary is a quantum error-correcting code!

Challenges abound

-- Understanding the black hole interior.
-- Further elucidation of the AdS/CFT code.
-- De Sitter space!

Exploring quantum gravity with a quantum simulator

Holographic duality opens a path to simulating nonperturbative quantum gravity using quantum computers and quantum simulators.

Probe bulk geometry by measuring boundary entanglement structure.

Probe bulk locality by measuring commutators of nonlocal boundary operators, perhaps by studying linear response.

Study the formation and evaporation of a black hole in the bulk; on the boundary a highly excited state settles down to thermal equilibrium.

Probe fast scrambling behavior with out-of-time-order correlators (NMR, ion traps, atoms in cavities, superconducting circuits). Not just the scrambling time but more fine grained information like the full Lyapunov spectrum.

Traversal of a wormhole in the bulk as coherent teleportation between two boundaries.

Prospects for QIS

Can noisy intermediate-scale quantum computing (NISQ) surpass exascale classical hardware running the best classical algorithms?

Near-term quantum advantage for useful applications is possible, but not guaranteed.

Hybrid quantum/classical algorithms (like QAOA and VQE) can be tested.

Quantum dynamics of highly entangled systems is especially hard to simulate, and is therefore an especially promising arena for quantum advantage.

Truly transformative quantum computing technology may need to be fault tolerant, and so may still be far off. But we don't know for sure how long it will take. Progress toward fault-tolerant QC must continue to be a high priority for quantum technologists.

Quantum sensing, networking, and computing will advance together. Next-generation quantum sensors can provide unprecedented capabilities of potential commercial interest, while also enabling new methods for exploring fundamental physics.

Quantum simulators can (someday) probe aspects of quantum field theory and quantum gravity which are beyond the reach of classical simulators, thus illuminating the nature of emergent spacetime.