Quantum Computing in the NISQ* Era and Beyond

John Preskill

Astroinformatics 2019

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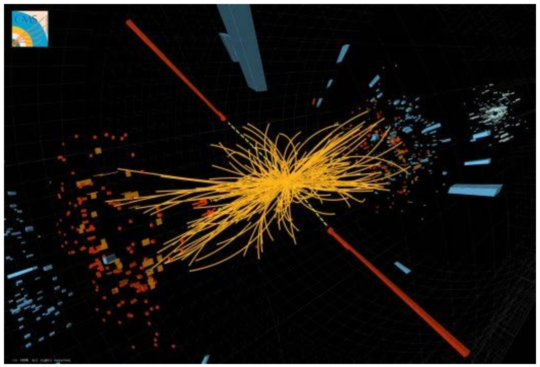		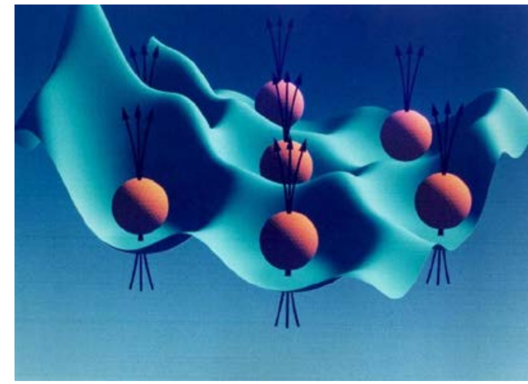
Higgs boson	Large scale structure	“More is different”
Neutrino masses	Cosmic microwave background	Many-body entanglement
Supersymmetry	Dark matter	Phases of quantum matter
Quantum gravity	Dark energy	Quantum computing
String theory	Gravitational waves	Quantum spacetime

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

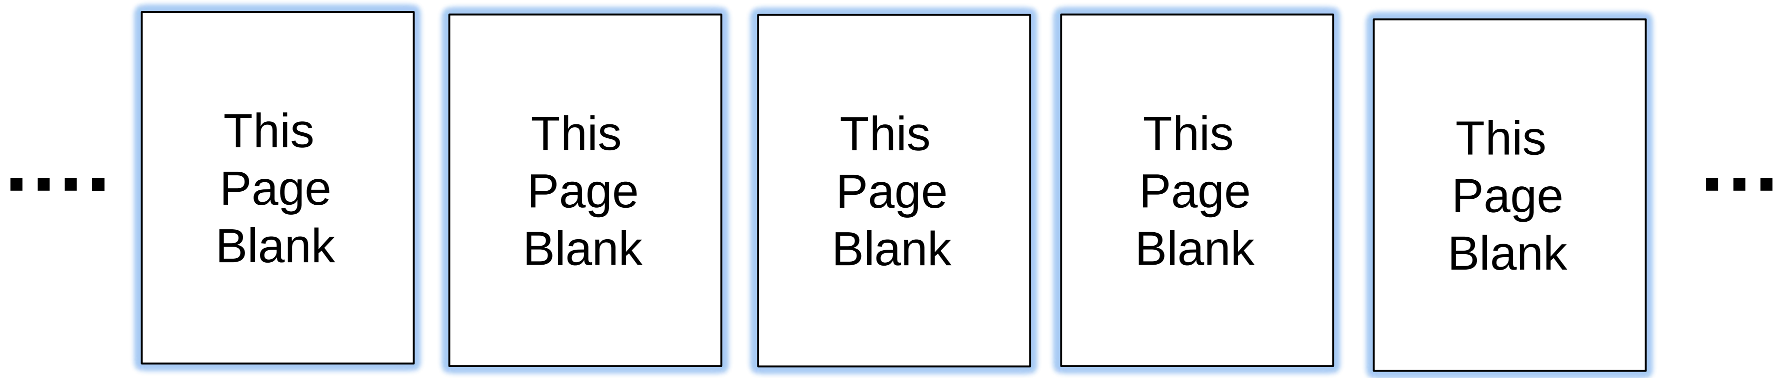

Nearly all the information in a typical entangled "quantum book" is encoded in the correlations among the "pages".

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

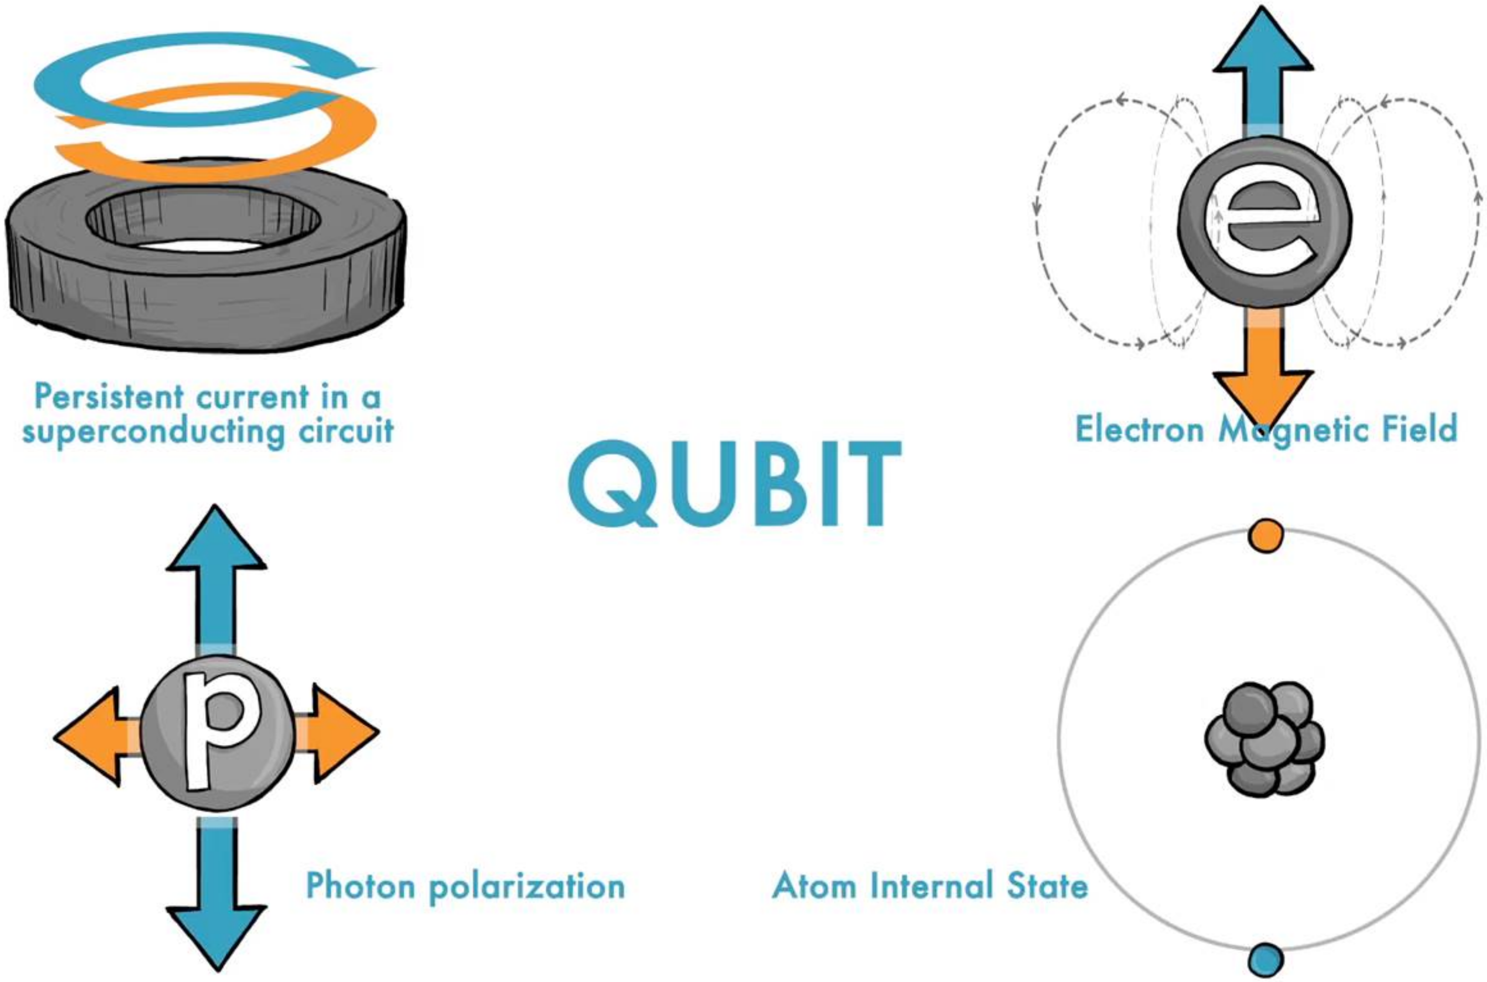
QUBIT
Persistent current in a superconducting circuit
Electron Magnetic Field
Photon polarization
Atom Internal State

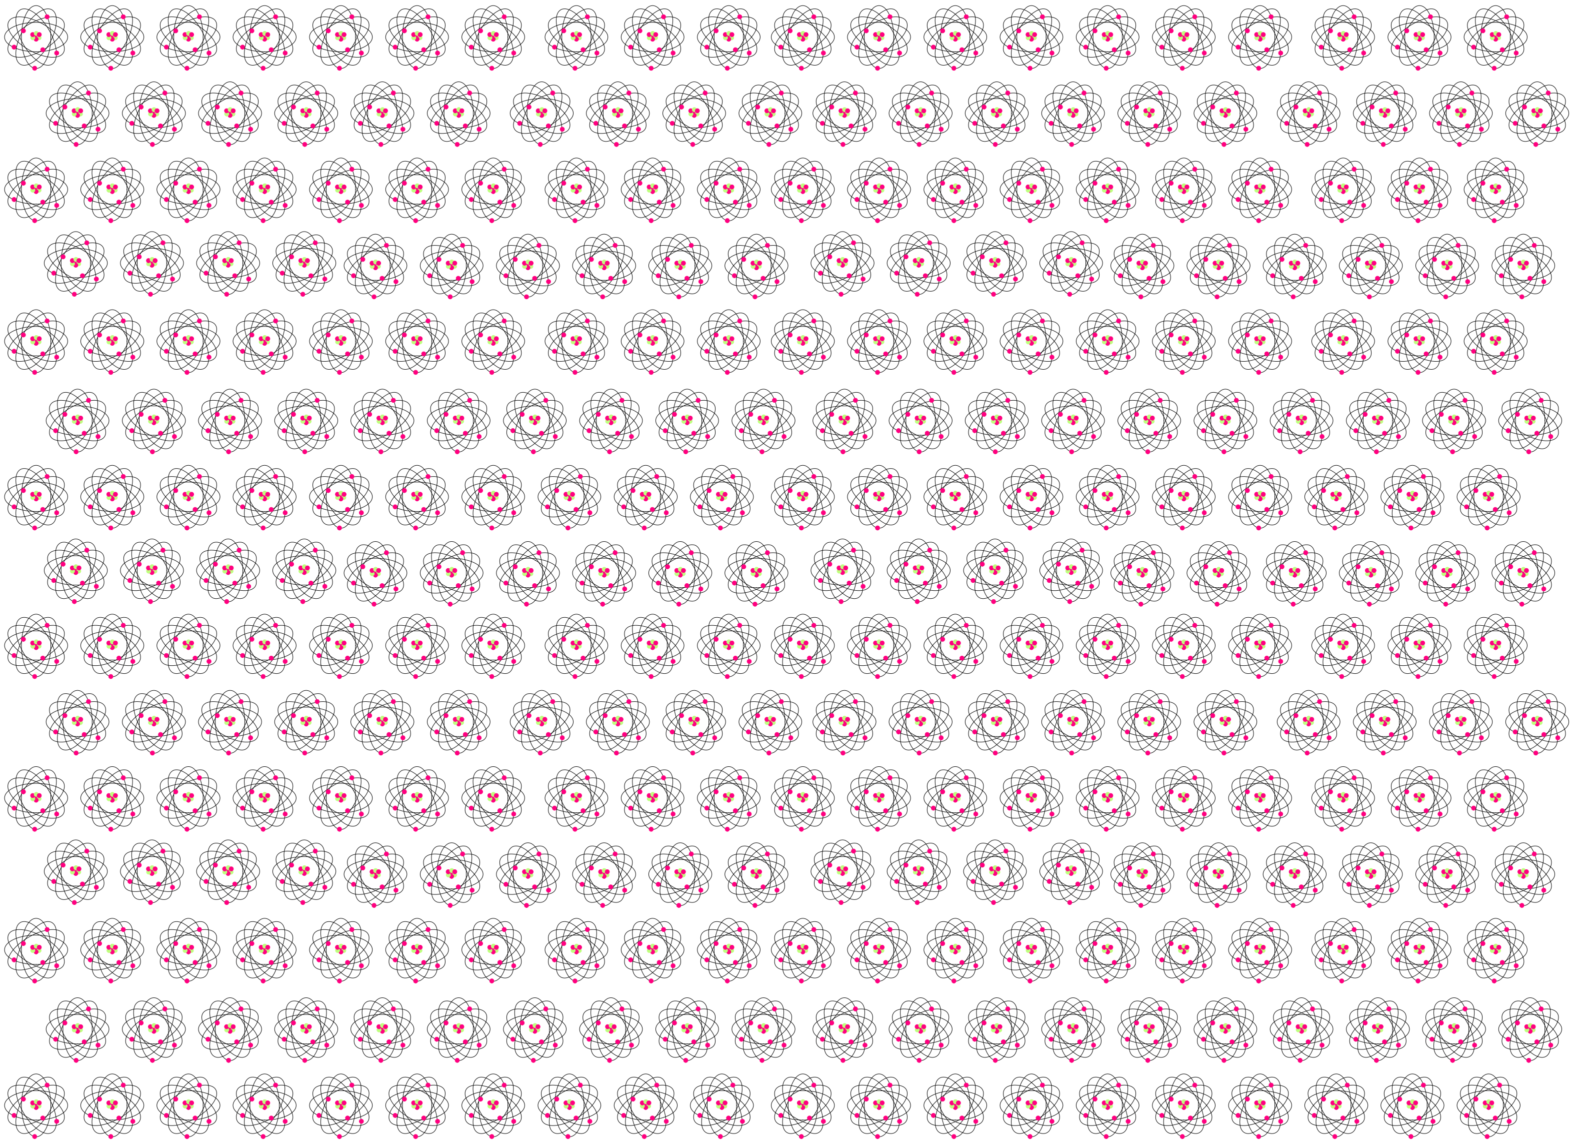

A complete description of a typical quantum state of just 300 qubits requires more bits than the number of atoms in the visible universe.

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

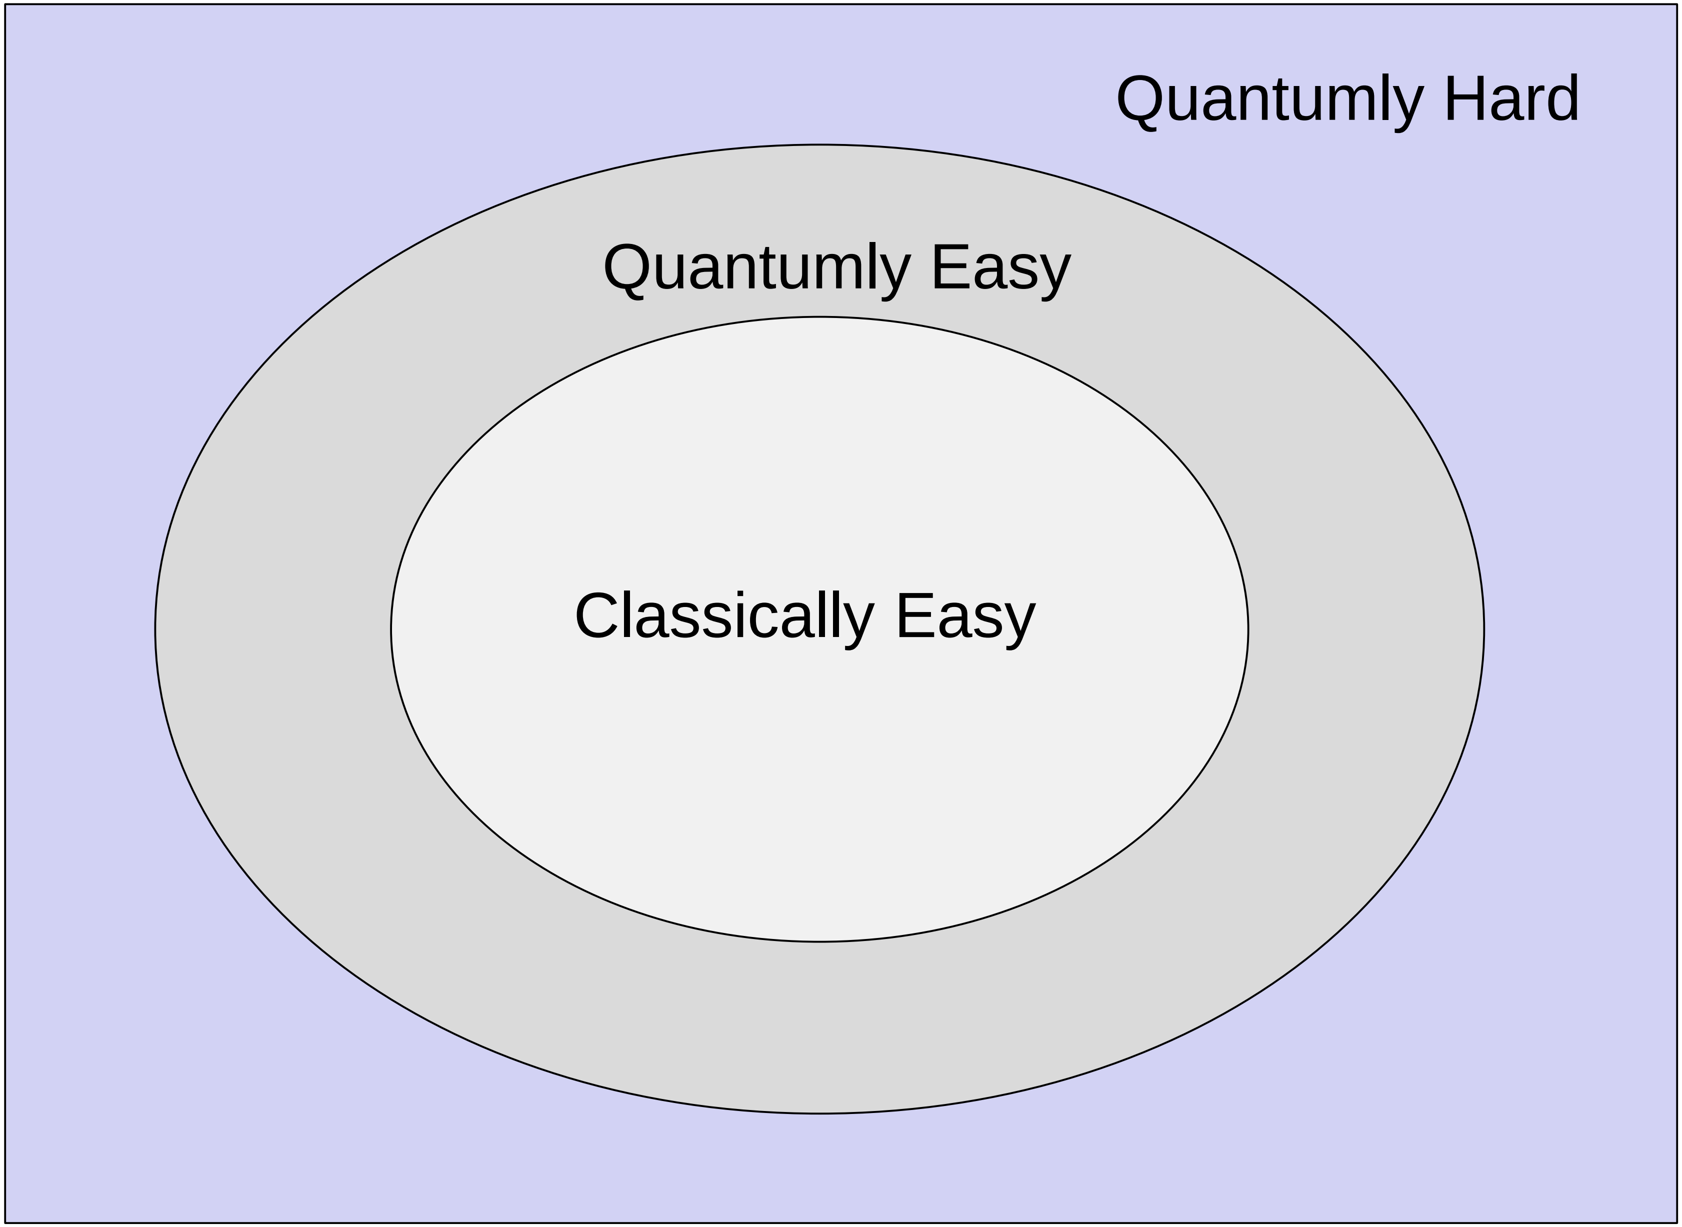

Problems

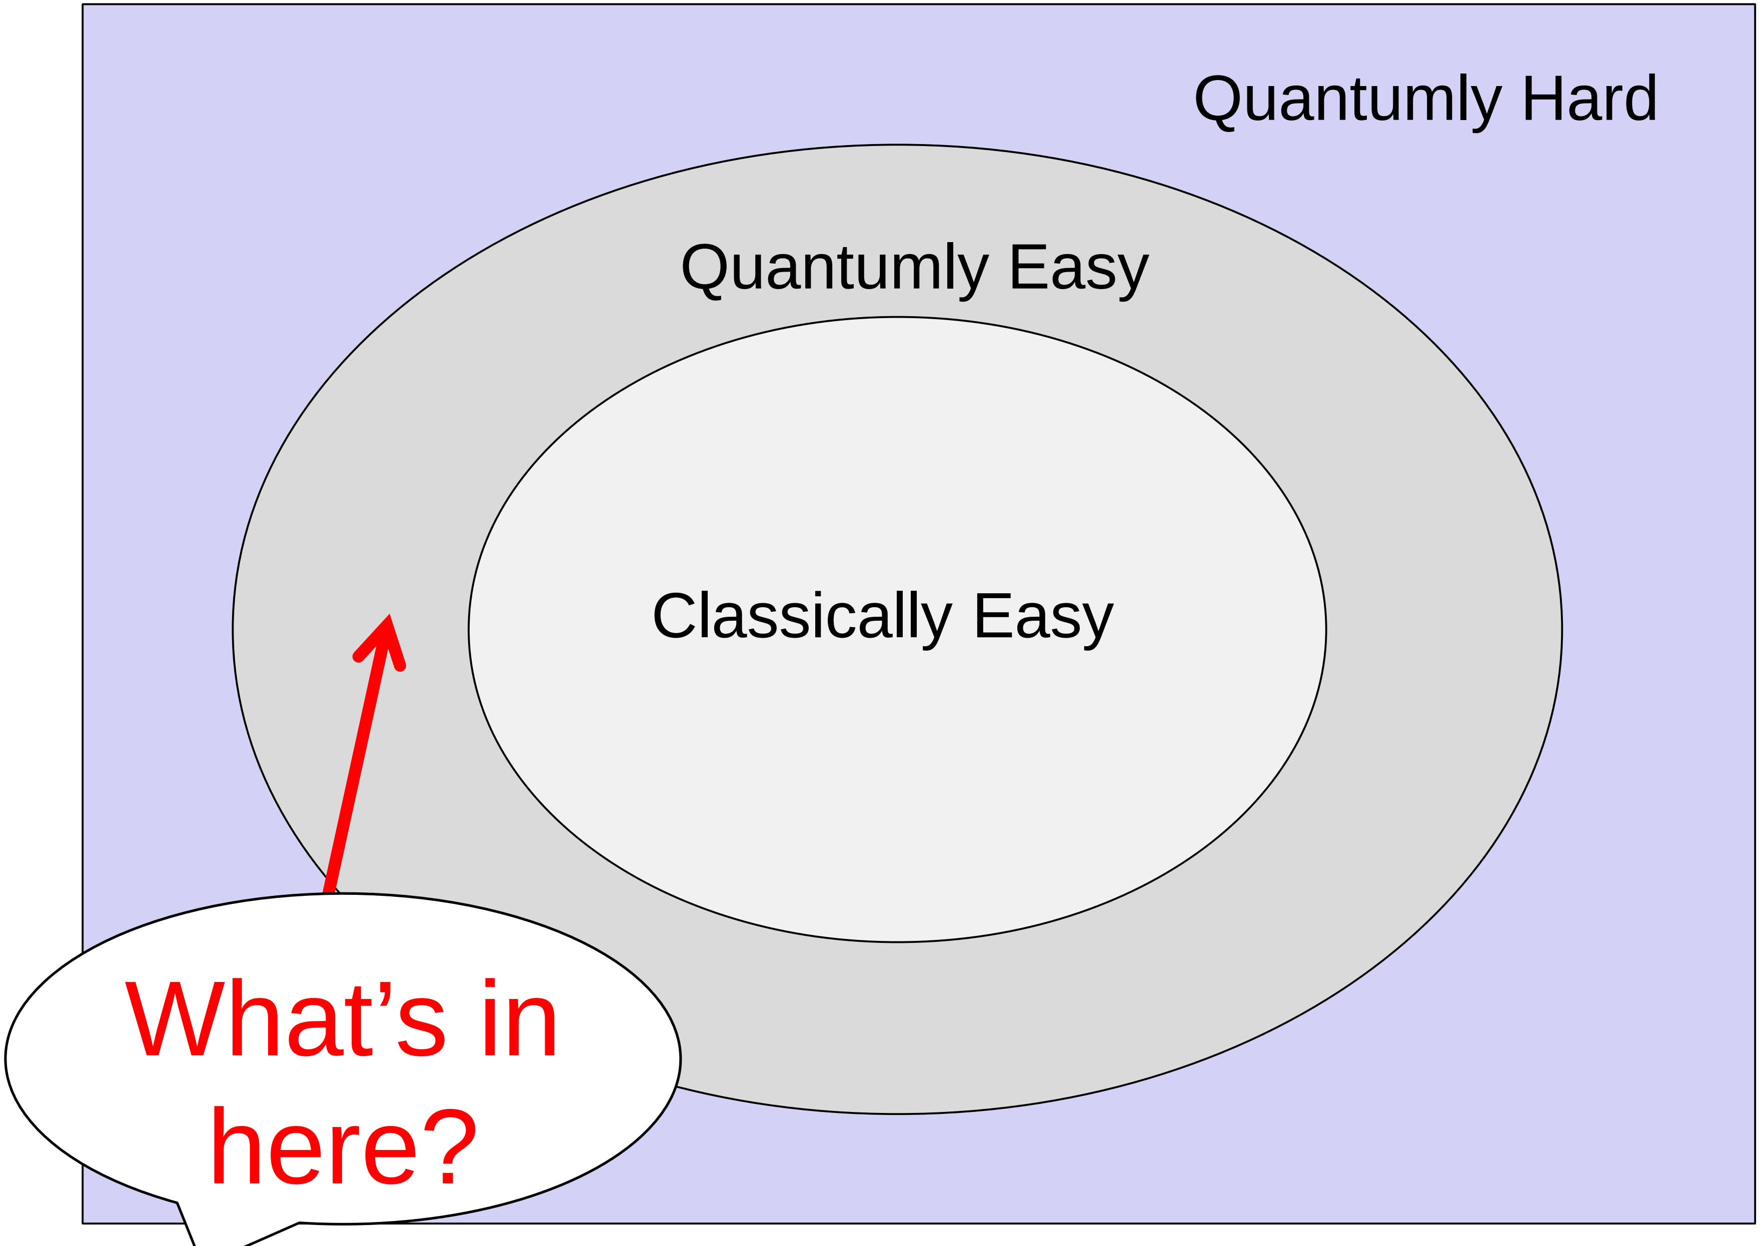
Quantumly Hard
Quantumly Easy
Classically Easy
What’s in here?

"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.

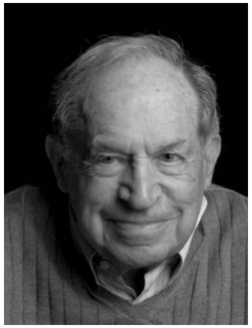

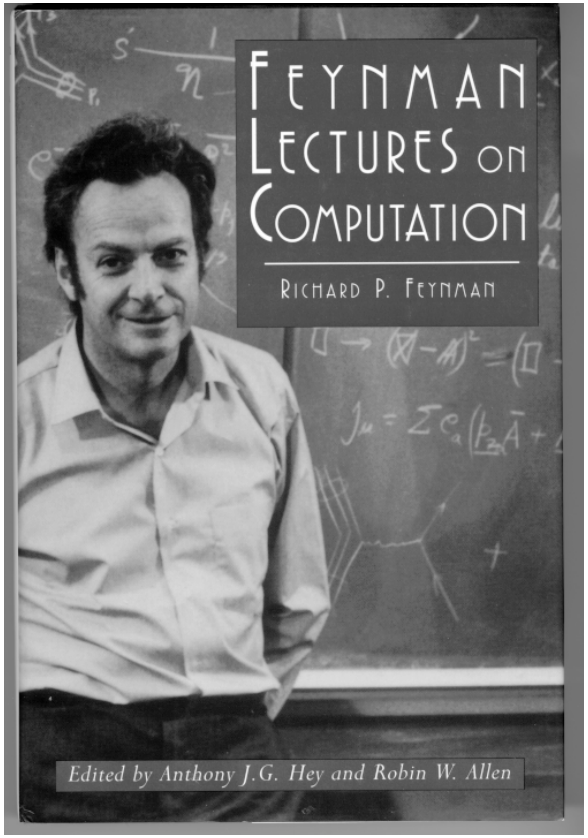

“Nature isn’t classical, dammit, and if you want to make a simulation of Nature, you’d better make it quantum mechanical, and by golly it’s a wonderful problem because it doesn’t look so easy.”

R. P. Feynman, 1981

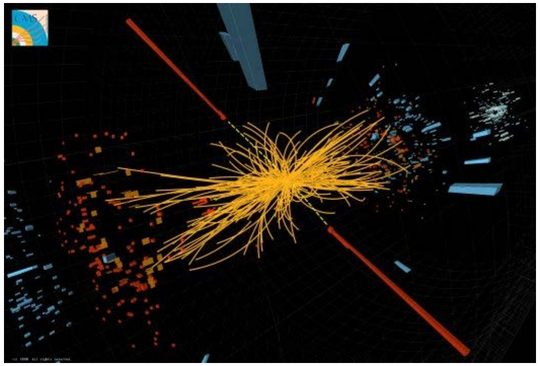

particle collision

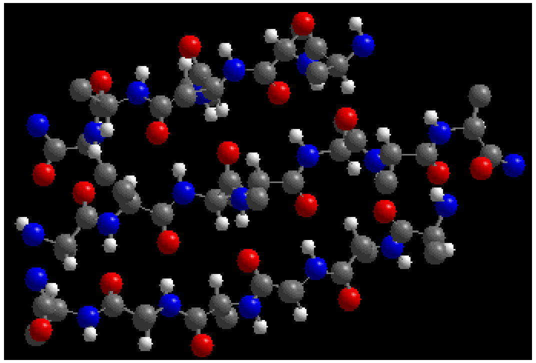

molecular chemistry

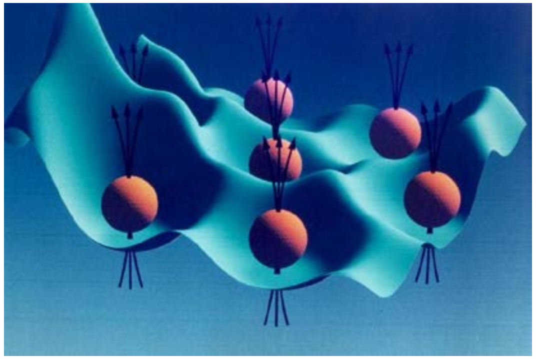

entangled electrons

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

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

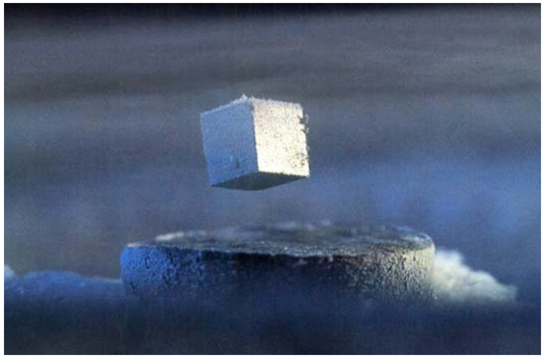

superconductor

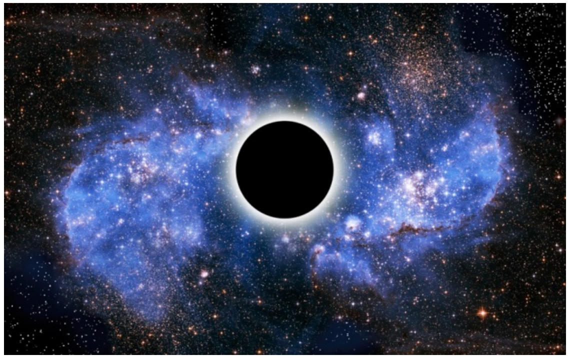

black hole

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.

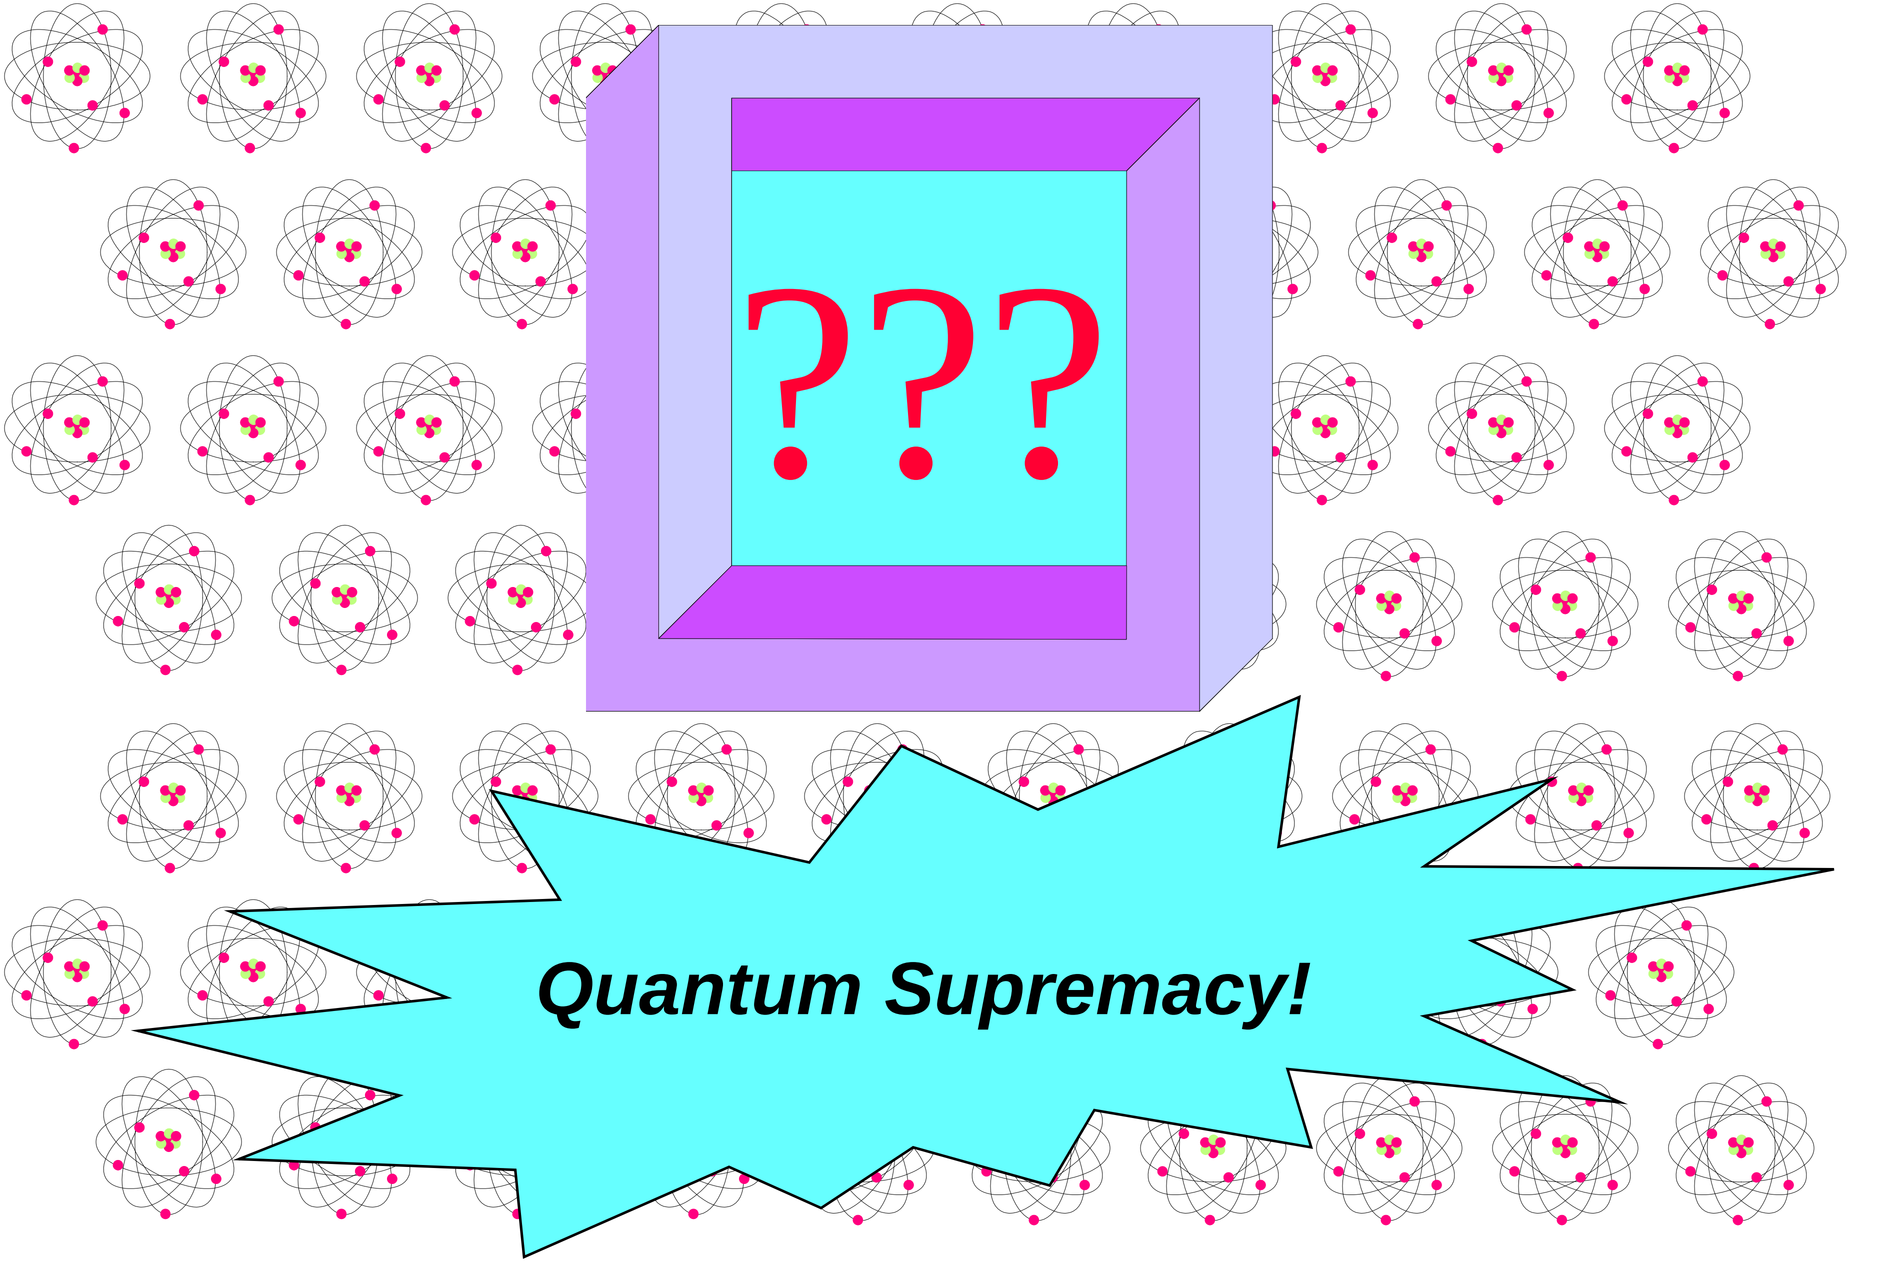
???
Quantum Supremacy!

Quantum computing in the NISQ Era

The (noisy) 50-100 qubit quantum computer is (almost) here. (*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.

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

Quantum 2, 79 (2018), arXiv:1801.00862

Quantum hardware: state of the art

IBM Quantum Experience in the cloud: now 16 qubits (superconducting circuit). 50-qubit device "built and measured."

Google 22-qubit device (superconducting circuit), 72 qubits built.

ionQ: 32-qubit processor planned (trapped ions), with all-to-all connectivity.

Rigetti: 128-qubit processor planned (superconducting circuit).

Harvard 51-qubit quantum simulator (Rydberg atoms in optical tweezers). Dynamical phase transition in Ising-like systems; puzzles in defect (domain wall) density.

UMd 53-qubit quantum simulator (trapped ions). Dynamical phase transition in Ising-like systems; high efficiency single-shot readout of many-body correlators.

And many other interesting platforms … spin qubits, defects in diamond (and other materials), photonic systems, …

There are other important metrics besides number of qubits; in particular, the two-qubit gate error rate (currently > 10^{-3}) determines how large a quantum circuit can be executed with reasonable signal-to-noise.

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 more than a few thousand 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.

Neven's Law

(1) Gate error rates for two-qubit quantum gates are improving exponentially with time. (Debatable, and can't go on for long. But a conservative estimate is that the error rate is decreasing by a factor of two every two to three years.)

(2) Therefore, the volume of a quantum circuit that can be executed with fixed circuit fidelity is increasing exponentially with time. (Not exactly, but close enough to make a point.)

(3) Furthermore, the classical cost of simulating the quantum circuit increases exponentially with the circuit volume. (Maybe not exactly, but definitely superpolynomial.)

(4) Therefore (Neven): for the largest quantum circuit that can be executed with fixed fidelity, the classical cost of the simulation is increasing doubly exponentially with time.

(5) That's really fast. (Even if you don't believe the details.)

Hybrid quantum/classical optimizers

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

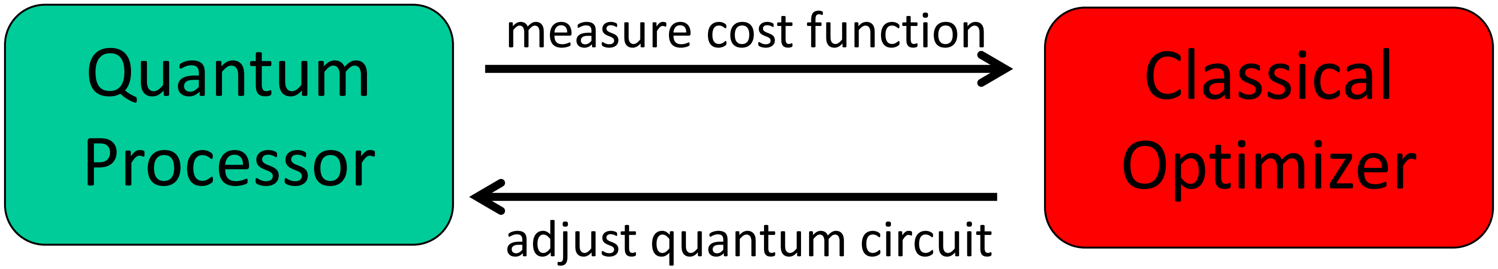

We don't expect a quantum computer to solve worst case instances of NP-hard problems, but it might find better approximate solutions, or find them faster.

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?

Hybrid quantum/classical optimizers

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

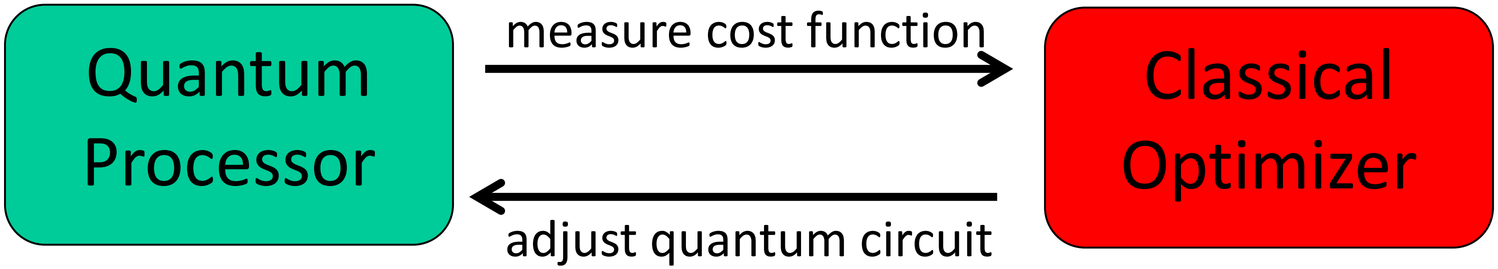

Can we streamline the classical loop? Naively, the number of variational parameters needed scales with the size of the instance. That's a problem.

Solving small instances may provide a good starting point for larger instances (Brandão, Broughton, Farhi, Gutmann, Neven 2018). "If we fix parameters such that the objective function has a high value at some small number of qubits then those same parameters will produce a high value at a larger number of qubits."

Symmetries might help reduce the classical load (e.g., translation invariant quantum optimization).

A concern (Hastings 2019): Why should quantum bounded-depth approximations be better than classical ones?

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.

How to find more applications?

Scott Aaronson: "Instead of thinking of a hard problem and asking how to speed it up, ask what quantum computers are good at and build an application from that."

For example, certified randomness.

Can we have both poly time classical verification and NISQ implementation?
What natural complexity assumptions suffice to ensure security?
More certified randomness by running the same circuit over and over again?

Simulation of quantum dynamics is another application in a similar spirit.

And what else?

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 as well as classical optimization problems.

Quantum machine learning

Jordan Kerenidis: "Overhyped but underestimated"

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 machine learning for quantum states and processes

Quantum machine learning (as for quantum algorithms more generally) is more likely to have an advantage for solving quantum problems.

Find task-specific quantum sensors variationally: What entangled states have advantages? Optimize e.g. the Fisher information by adjusting e.g. a squeezing parameter. For example optimize spin squeezing in an array of many atoms.

Recognizing phases of quantum systems. "It's harder to recognize phases than faces" (if no local order parameter, because of exponentially large Hilbert space).

Quantum convolutional neural networks, combining renormalization group flow with hierarchical error correction to classify quantum correlations (*Cong, Choi, Lukin 2018*).

Quantum codes and decoders for physical noise, e.g. for noise correlations.

Supervised quantum machine learning for pharmaceuticals, catalysts, materials.

Machine learning (quantum and classical) seen through the lens of RG + EC.

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.

Digital vs. Analog quantum simulation in the NISQ era

What hard problems can we solve with (noisy) analog simulators?

What is the potential advantage of digital in the NISQ era?

Simulating time evolution is expensive (Trotter and and other methods).

Digital provides more flexible Hamiltonian and initial state preparation. We can use hybrid quantum/classical methods, e.g. finding a succinct tensor network description of the initial state, which can then be compiled as a small quantum circuit.

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

Today's analog simulators: For example, snapshots of fluctuating string order in the doped Hubbard model, classified using machine learning (*Greiner group*). Spectral response for a strongly coupled Fermi gas (*Zwierlein group*).

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 (MBL)*. 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 (*Lukin 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" --- atypical slightly entangled nonthermal eigenstates in a nonintegrable system. Does this require fine tuning?

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 *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.

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. 20 million physical qubits to break RSA 2048 (*Gidney, Ekerå 2019*), for gate error rate 10^{-3}.

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

Mainstream users may need to be patient.

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

Near-term noise mitigation

No quantum error correction in the near term, and full blown quantum fault tolerance in the long term. But what comes in between?

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 significantly 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 (*Kim 2017*)

We might improve signal-to-noise by extrapolating to the zero-noise limit, or by resampling from a quasi-probability distribution (*Temme, Bravyi, Gambetta 2017*). Or other resampling ideas. (No qubit overhead.)

Quantum speedups in the NISQ era and beyond

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.

Near-term algorithms should be designed with noise resilience in mind.

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

Experimentation with quantum testbeds may hasten progress and inspire new algorithms.

NISQ will not change the world by itself. Realistically, the goal for near-term quantum platforms should be to pave the way for bigger payoffs using future devices.

Lower quantum gate error rates will lower the overhead cost of quantum error correction, and also extend the reach of quantum algorithms which do not use error correction.

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.