Q2B:
Quantum Computing for Business

John Preskill

11 December 2018

IQIM

INSTITUTE FOR QUANTUM INFORMATION AND MATTER

CALIFORNIA INSTITUTE OF TECHNOLOGY

1891

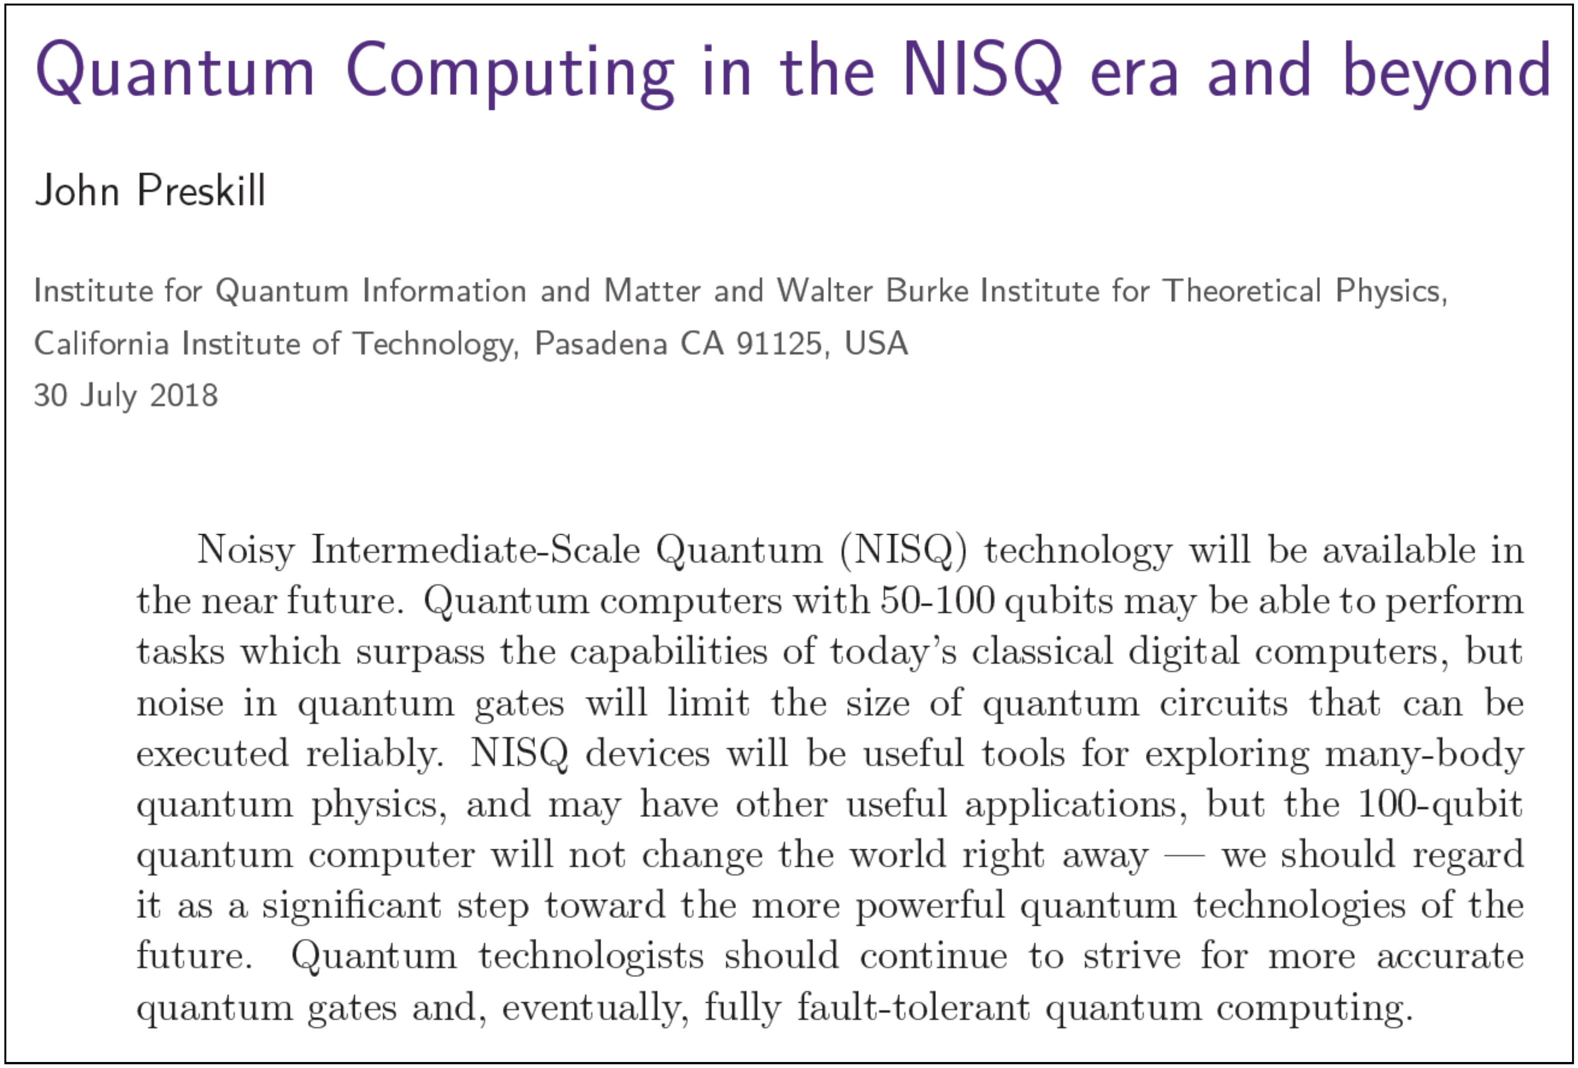

Quantum Computing in the NISQ era and beyond

John Preskill

Institute for Quantum Information and Matter and Walter Burke Institute for Theoretical Physics,
California Institute of Technology, Pasadena CA 91125, USA

30 July 2018

Noisy Intermediate-Scale Quantum (NISQ) technology will be available in the near future. Quantum computers with 50-100 qubits may be able to perform tasks which surpass the capabilities of today's classical digital computers, but noise in quantum gates will limit the size of quantum circuits that can be executed reliably. NISQ devices will be useful tools for exploring many-body quantum physics, and may have other useful applications, but the 100-qubit quantum computer will not change the world right away — we should regard it as a significant step toward the more powerful quantum technologies of the future. Quantum technologists should continue to strive for more accurate quantum gates and, eventually, fully fault-tolerant quantum computing.

Quantum 2, 79 (2018), arXiv:1801.00862

Based on a Keynote Address delivered at Quantum Computing for Business, 5 December 2017

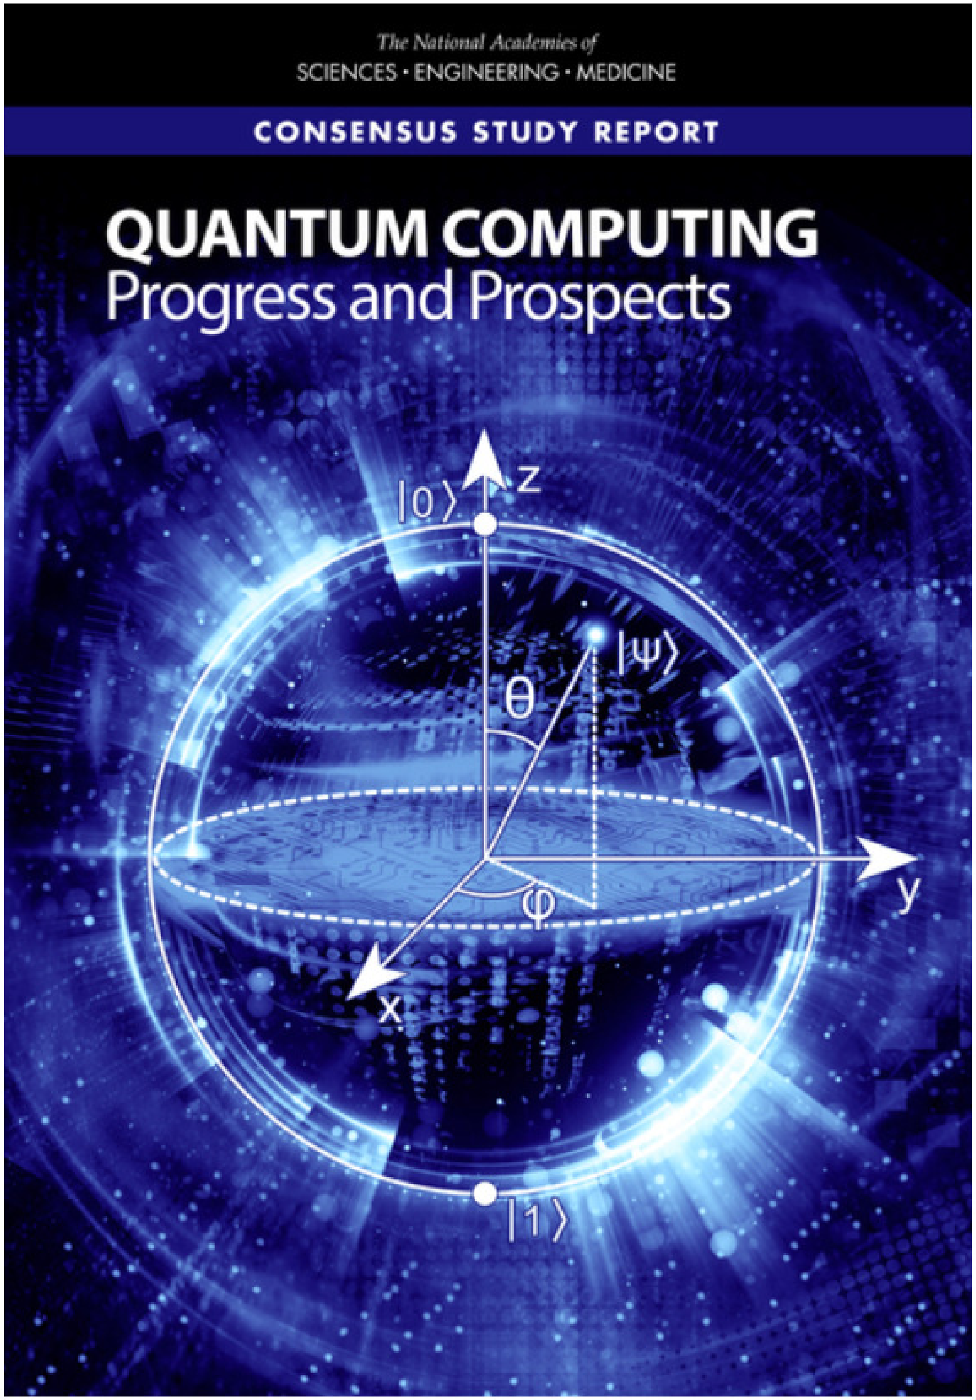
The National Academies of
SCIENCES · ENGINEERING · MEDICINE

CONSENSUS STUDY REPORT

QUANTUM COMPUTING
Progress and Prospects

MARK A. HOROWITZ,
Stanford University, *Chair*
ALÁN ASPURU-GUZIK,
University of Toronto
DAVID D. AWSCHALOM,
University of Chicago
BOB BLAKLEY,
Citigroup
DAN BONEH,
Stanford University
SUSAN N. COPPERSMITH,
University of Wisconsin, Madison
JUNGSANG KIM,
Duke University
JOHN M. MARTINIS,
Google, Inc.
MARGARET MARTONOSI,
Princeton University
MICHELE MOSCA,
University of Waterloo
WILLIAM D. OLIVER,
Massachusetts Institute of Technology
KRYSTA SVORE,
Microsoft Research
UMESH V. VAZIRANI,
University of California, Berkeley

What I didn't say much about last year

Quantum cryptography

Privacy founded on fundamental laws of quantum physics.

Quantum networking

Distributing quantumness around the world.

Quantum sensing

Improving sensitivity and spatial resolution.

Instead I focused on quantum computing.

Which uses some of the same hardware as the above applications.

Frontiers of Physics

short distance

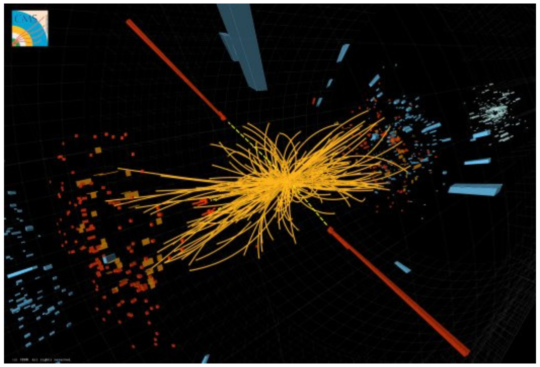

Higgs boson

Neutrino masses

Supersymmetry

Quantum gravity

String theory

long distance

Large scale structure

Cosmic microwave background

Dark matter

Dark energy

Gravitational waves

complexity

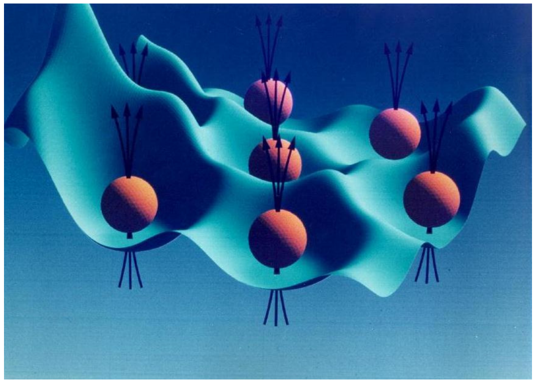

“More is different”

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.

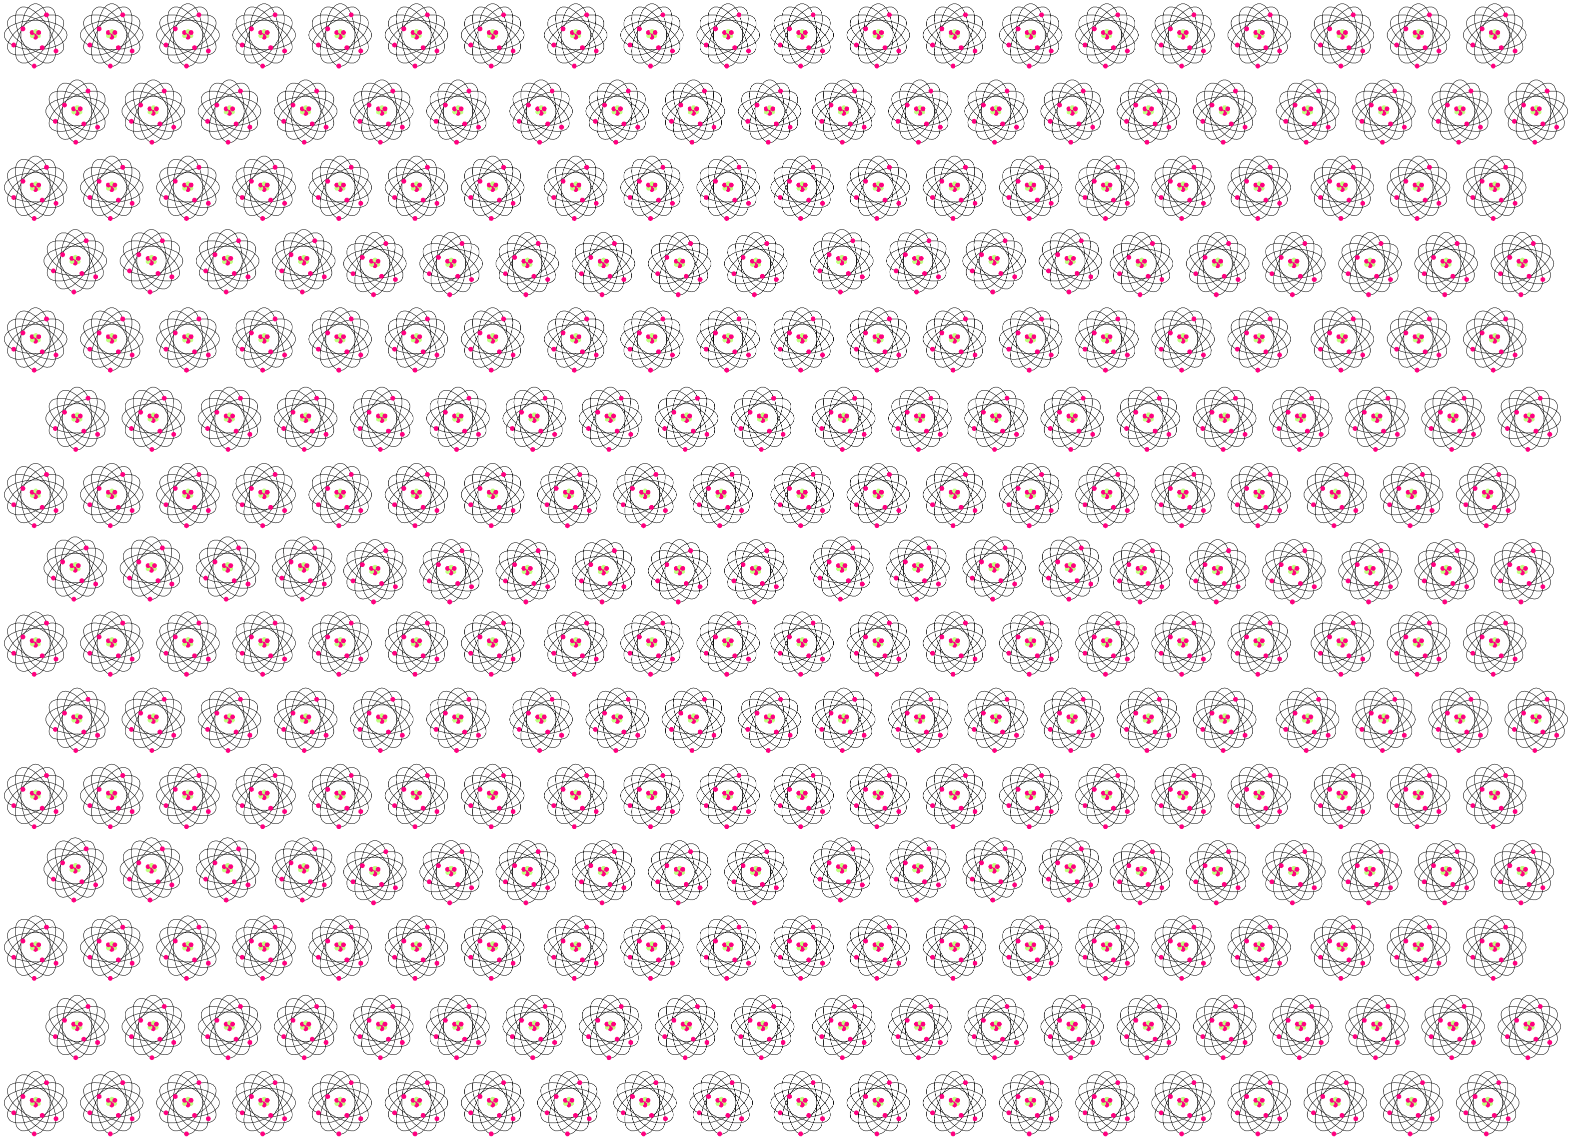

A complete description of a typical quantum state of just 300 qubits requires more bits than the number of atoms in the visible universe.

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

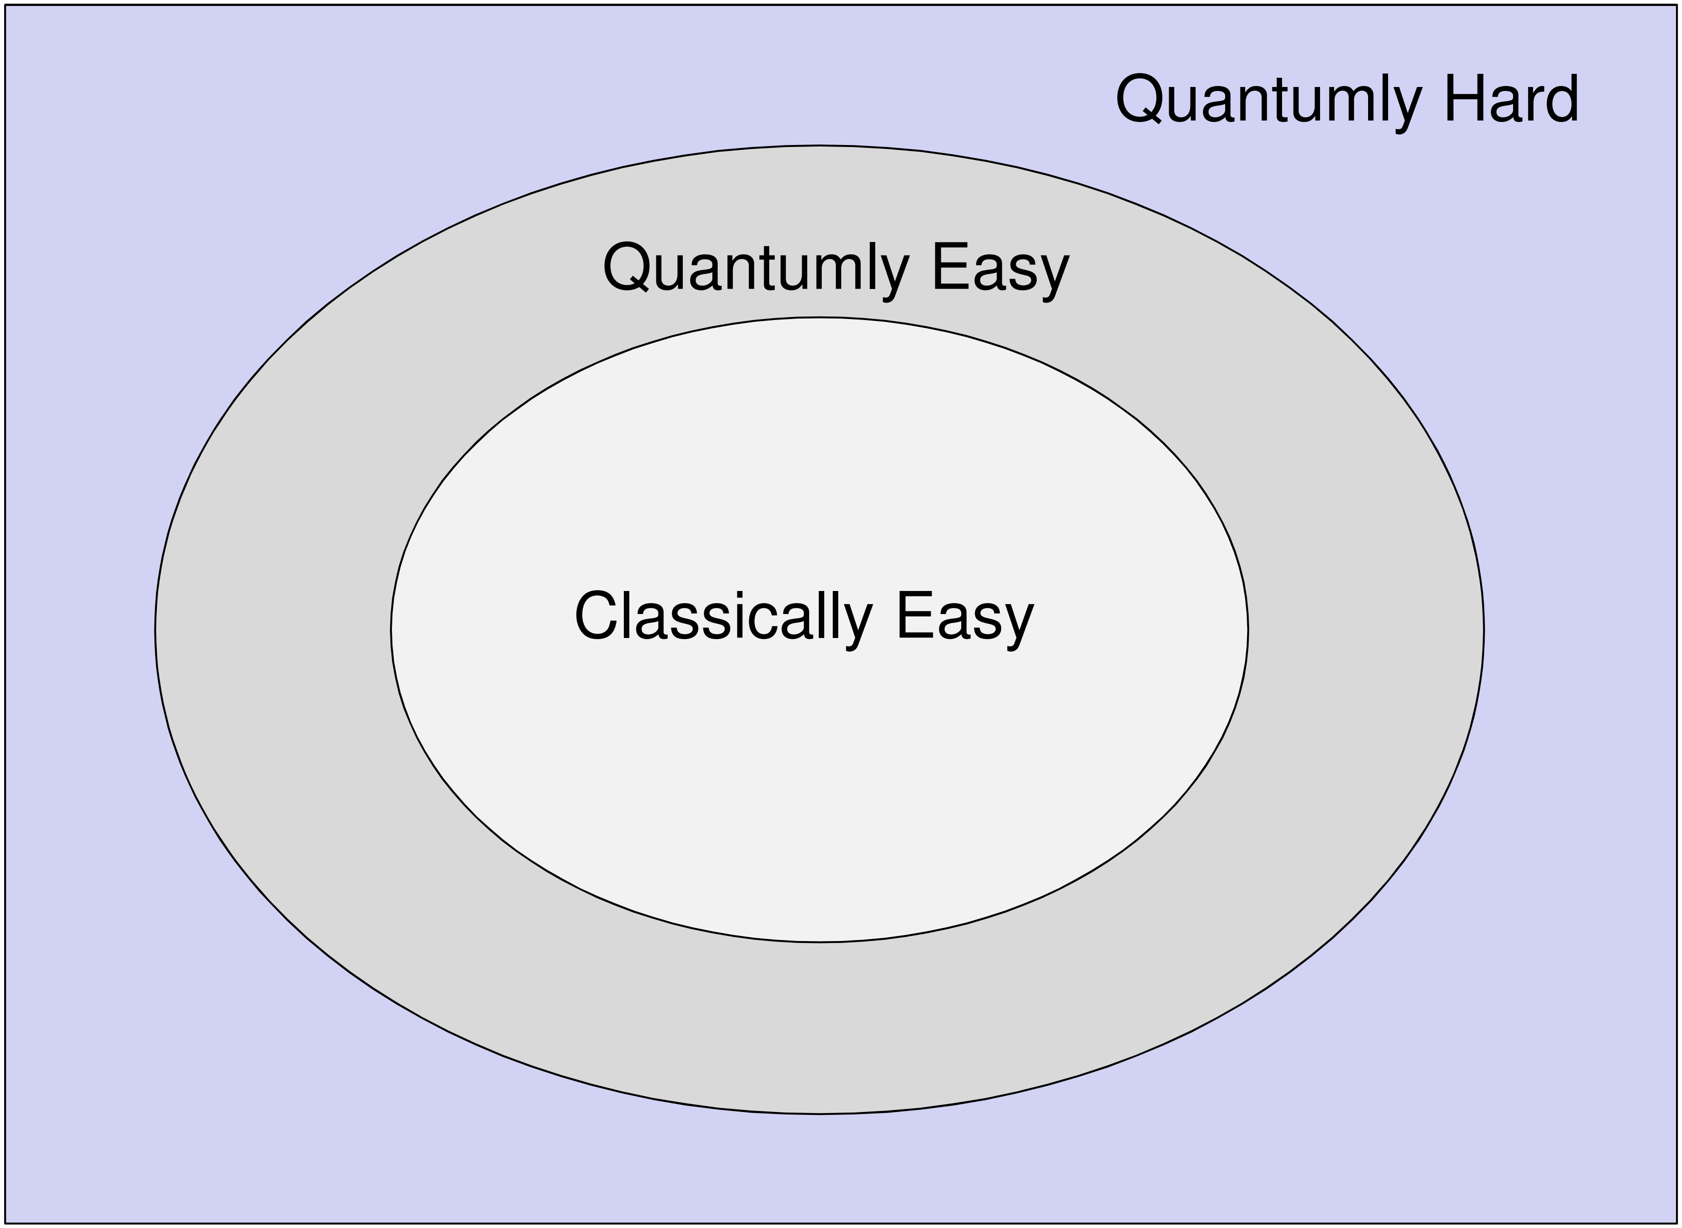

Quantumly Hard

Quantumly Easy

Classically Easy

Problems

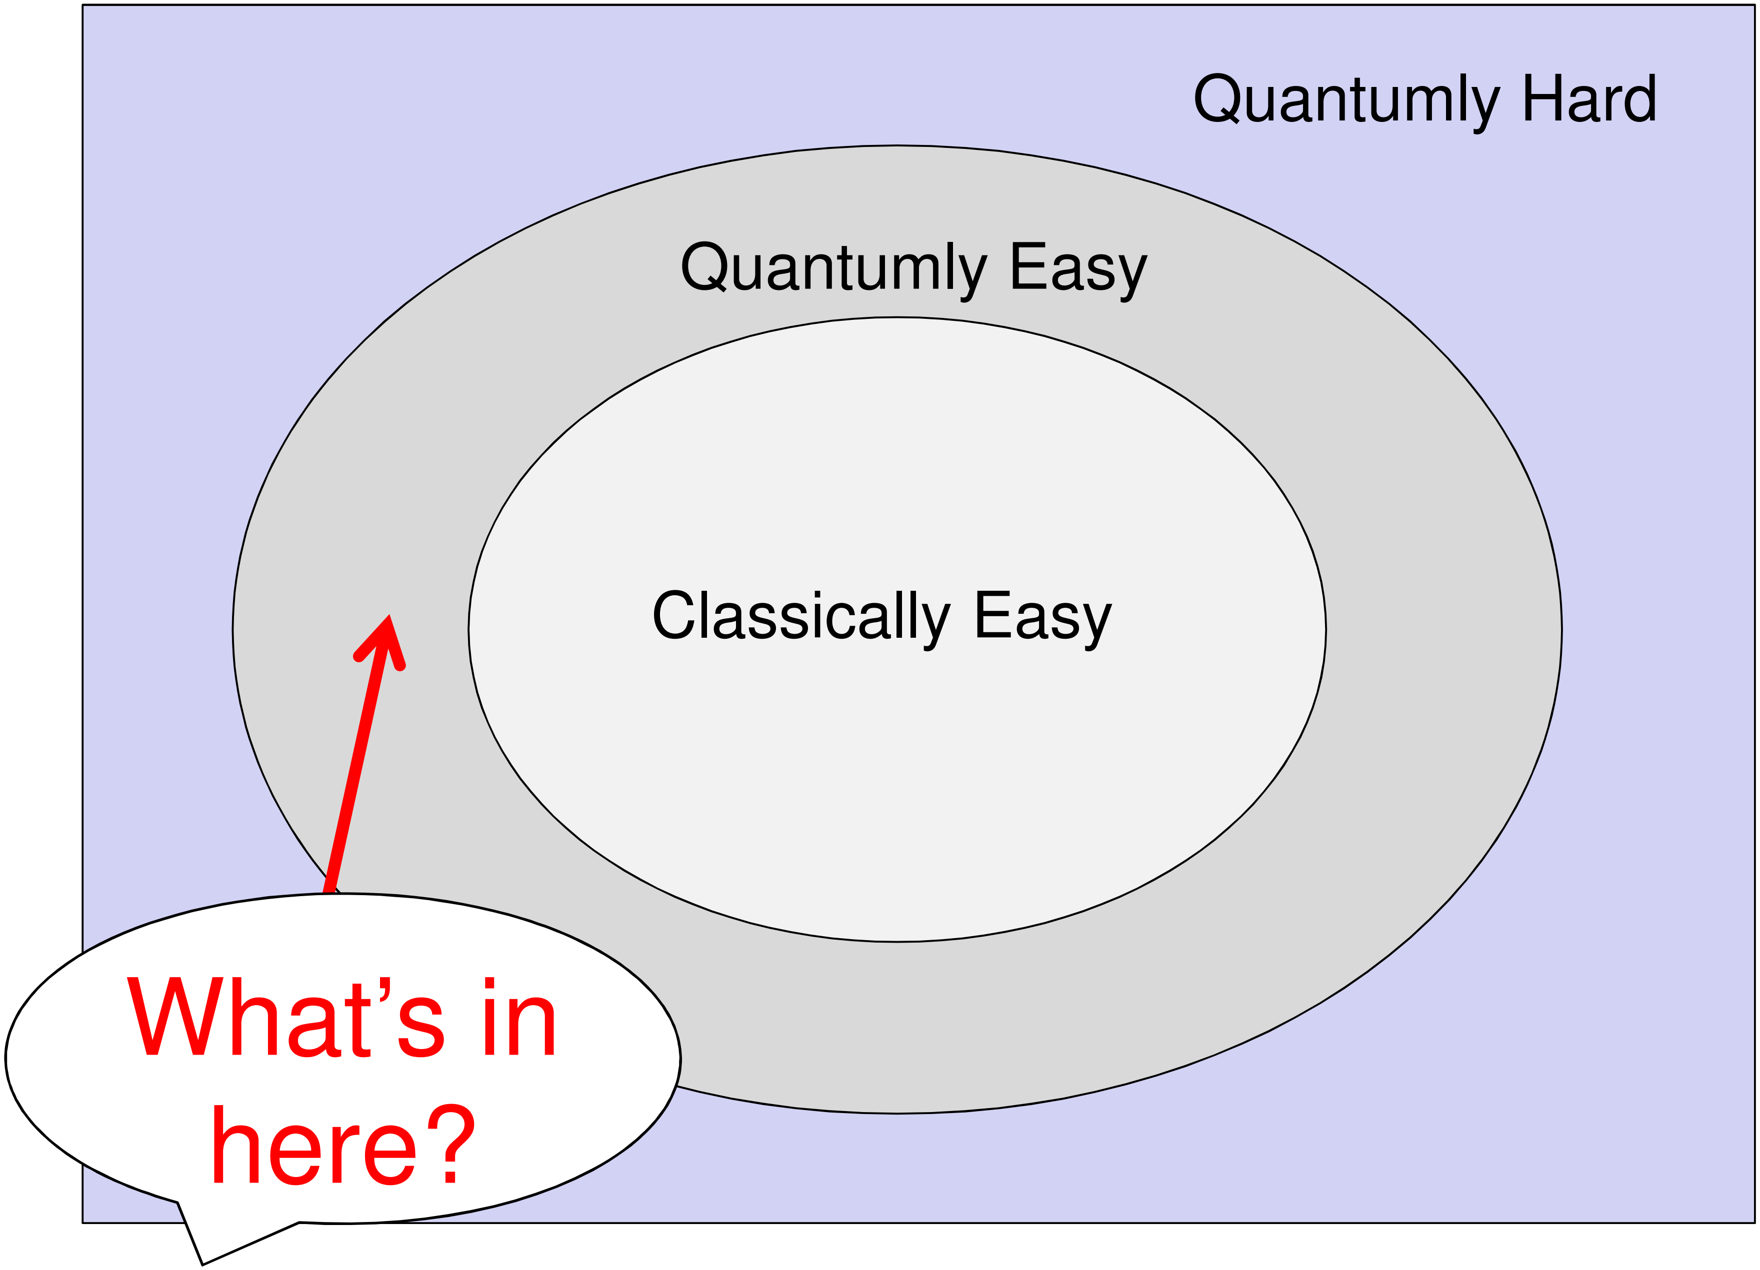

Quantumly Hard

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.

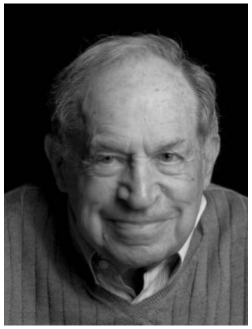

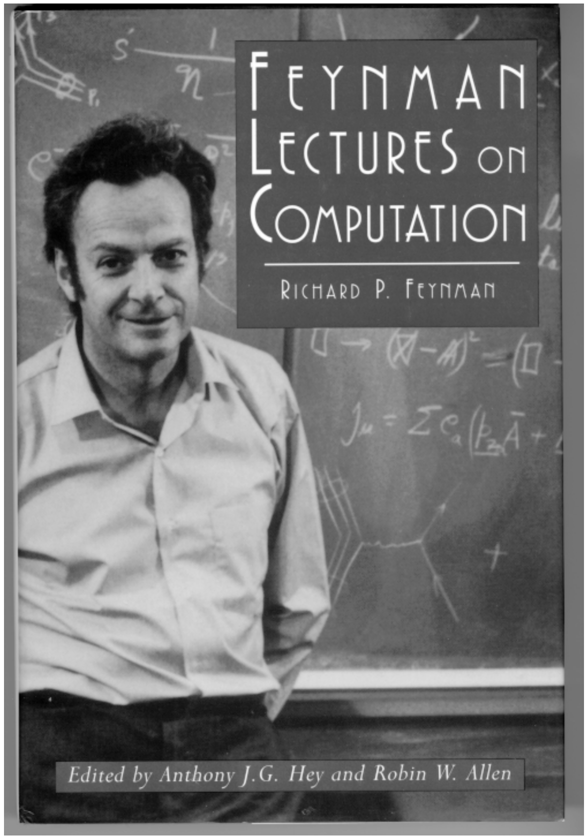

“Nature isn’t classical, dammit, and if you want to make a simulation of Nature, you’d better make it quantum mechanical, and by golly it’s a wonderful problem because it doesn’t look so easy.”

R. P. Feynman, 1981

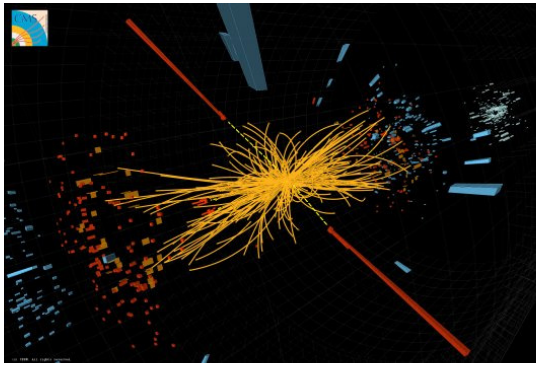
particle collision

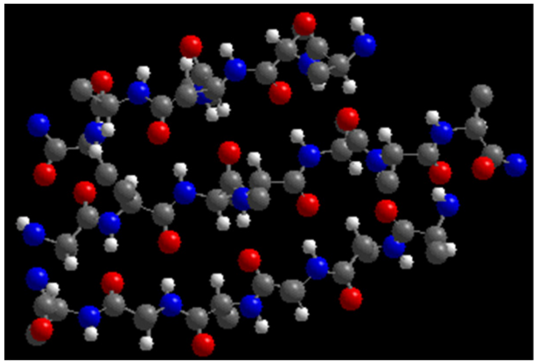
molecular chemistry

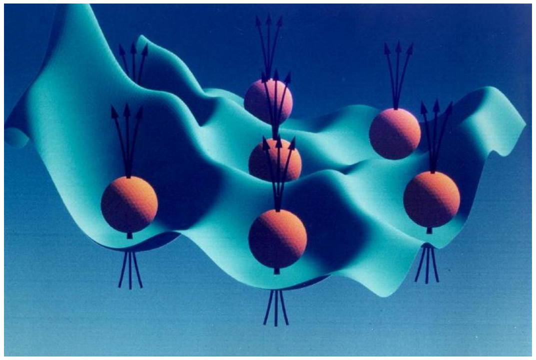
entangled electrons

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

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

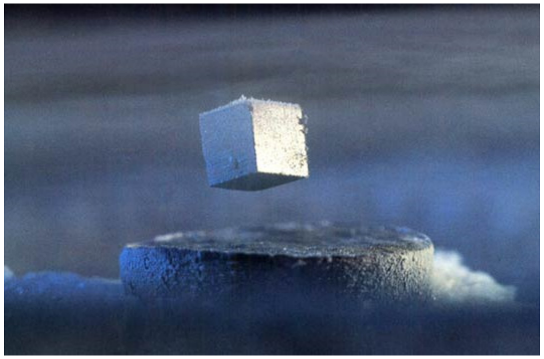
superconductor

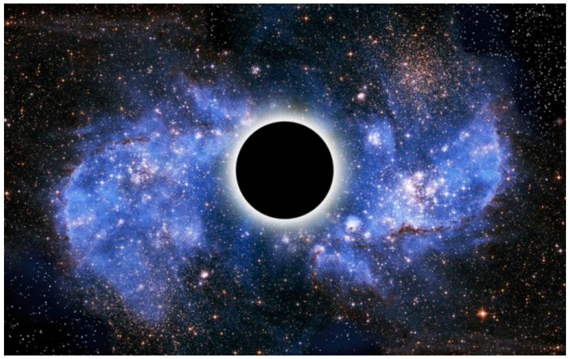
black hole


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.

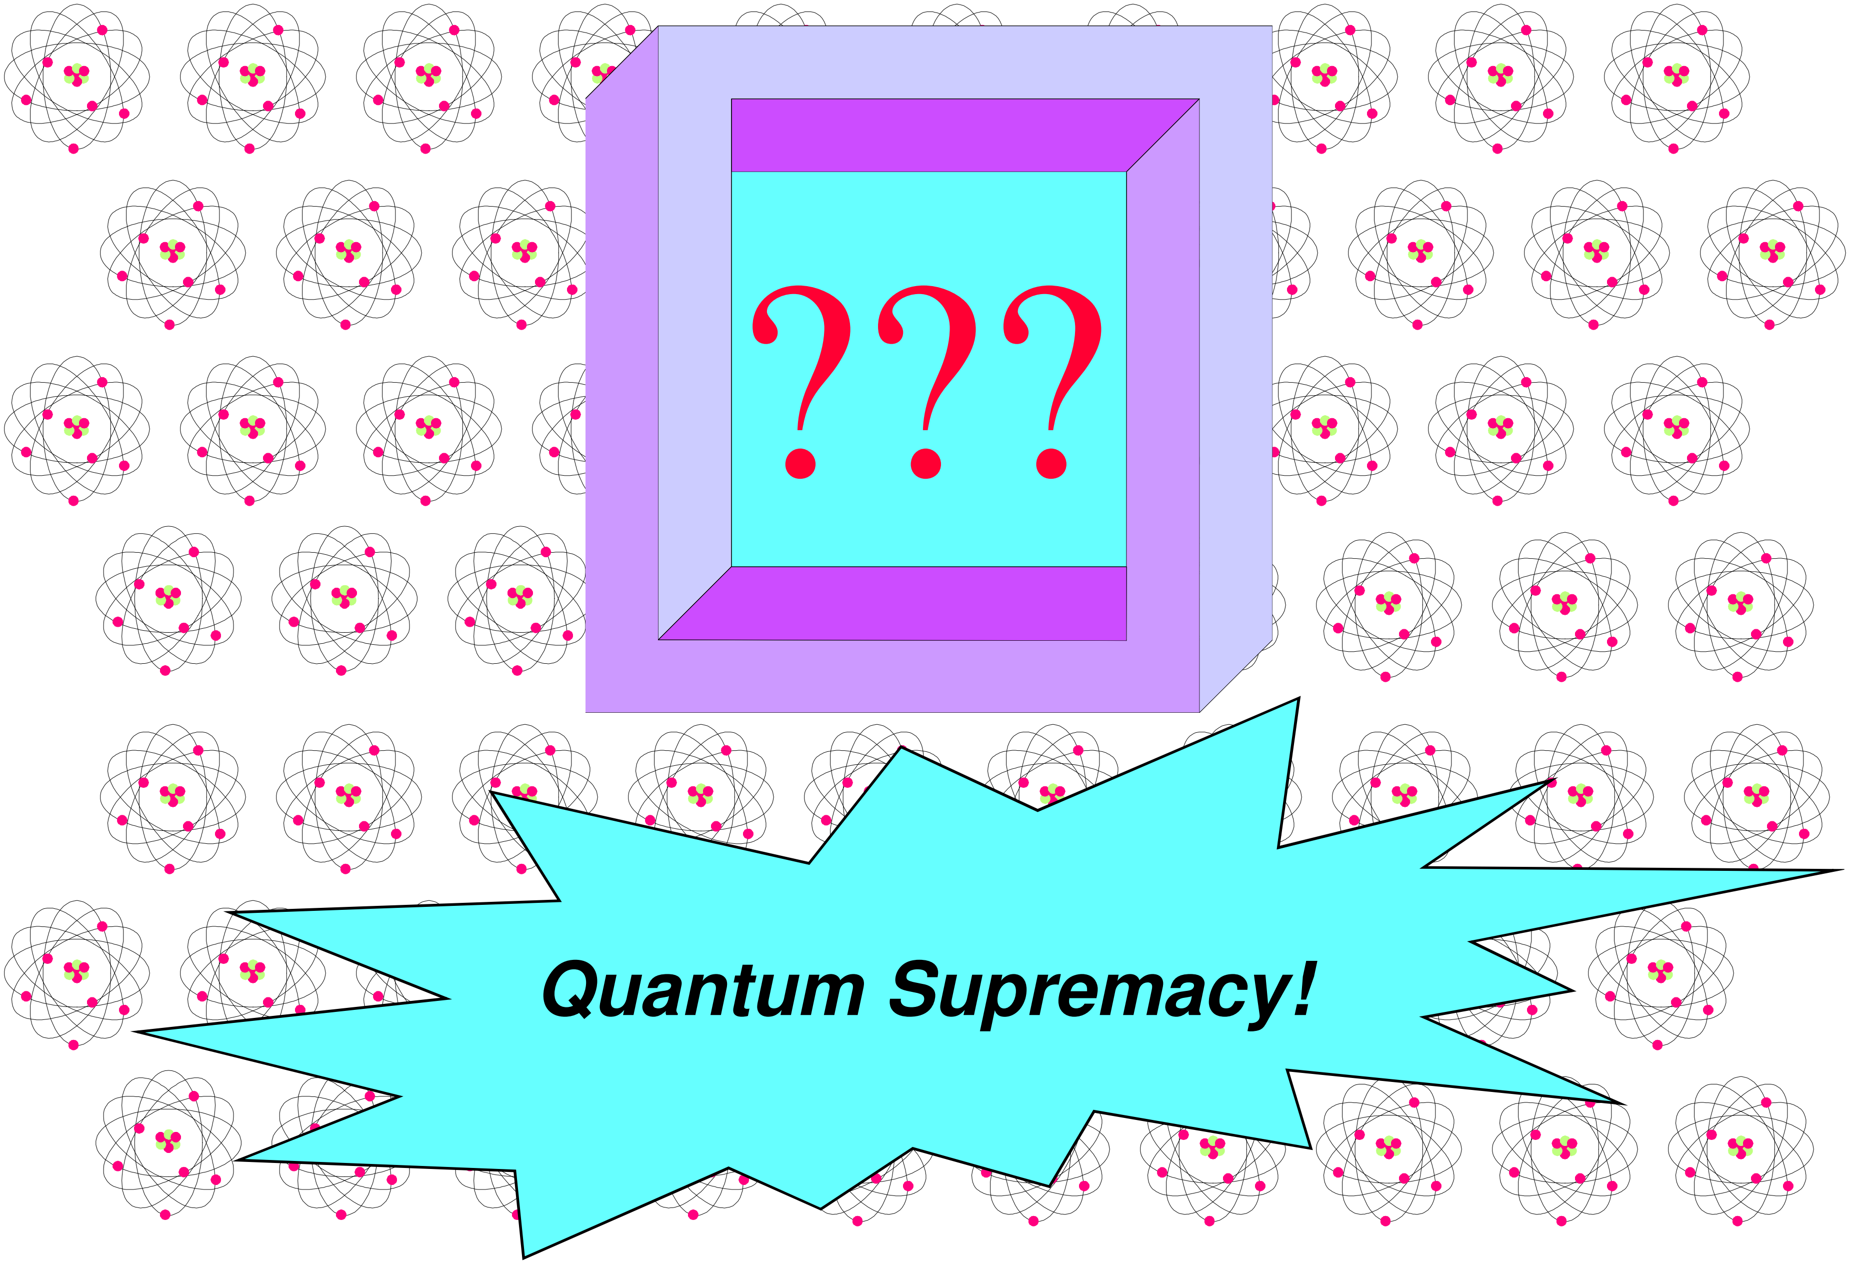
???
Quantum Supremacy!

Quantum computing in the NISQ Era

The (noisy) 50-100 qubit quantum computer is 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

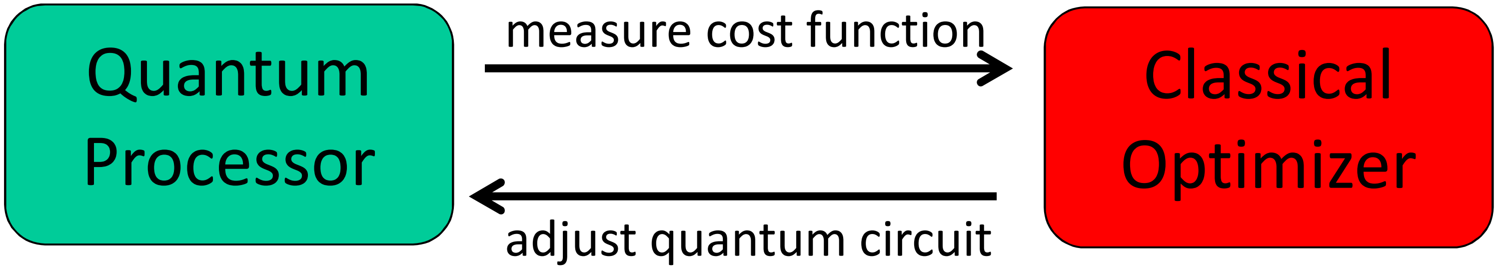

Quantum Processor

measure cost function

adjust quantum circuit

Classical Optimizer

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.

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.

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?

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. The output is a quantum state; 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 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.

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 qubit technology, systems engineering, algorithm design, and theory 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: solid state

NV center = Nitrogen vacancy color center in diamond. *High resolution* scanning probe microscopy. Sensors for nanoscale imaging.

Long coherence time at room temperature. Electric spin for probing, nuclear spin for storage. Optical addressability. Stable in nanostructures.

Noninvasive sensing and imaging of bio-magnetism in living cells. Bacteria, human cells, etc. E.g., action potential of individual neurons.

In vivo mapping of temperature and chemical changes. Monitor (and repair?) damage at cellular or molecular level.

Scanning to guide exploration and development of advanced materials.

Better materials can also enhance sensing platforms. Guiding and amplifying photon signals in multiqubit sensors. Topological materials for robust (against disorder) transport of electrons, spins, photons, phonons.

Quantum sensing: atoms

PNT: position, navigation, and timing. E.g., GPS: 32 satellites in medium-earth orbit. 2 billion users.

Clocks: Current GPS uses microwave frequency standards. Optical atomic clocks can be orders or magnitude more accurate.

General relativity on a tabletop: atomic clock senses gravitational redshift.

What to do when GPS not available? Inertial sensors. Accelerometer (detects linear acceleration), gyrometer (detects angular velocity), …

Gravimeters and gravity gradiometers for geophysical surveying (oil, minerals, water, mapping, ….). Atom interferometers, and also superconducting and optomechanical devices.

Economic impact of atomic clocks and quantum magnetometers already. What's coming? Quantum enhancements from entanglement, squeezing, error correction.

Hybrid quantum technologies for multi-modal function: sensing, storing, processing, communicating quantum information and controlling the environment.

“Next generation” quantum sensing

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

Standard quantum limit (SQL): Sensitivity scales like $N^{-1/2}$ with the number of sensors N, or with the total probing time. In principle, this scaling can be improved to N^{-1} using squeezing/entanglement.

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

Quantum radar (a.k.a. quantum illumination). Create entangled photon pair and bounce one photon off a target. Entanglement enhances signal to noise. Transduction from microwave to visible.

An application for long-baseline quantum networks: Distantly separated optical telescopes which share entanglement can perform interferometry by teleporting photons. Detect a (brightly shining!) elephant on another planet.

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 and robust devices, … and new ideas.

Quantum sensing: LIGO

2 black holes merged 5 billion years ago. Gravitational wave passing by in 2017 stretched the earth by 10^{-15} m.

LIGO (on ground) detects 10^{-18} m change in distance between mirrors 4 km apart. Time scale: LIGO detects 1 ms to 100 ms.

eLISA (in space, proposed) will detect 10 s to 1000 s.

How to cover 100 ms to 10 s? Atom interferometer / optical atomic clocks. (Detect the relative phase shift of two separated ultracold atomic ensembles.)

Meanwhile … At 1 ms, LIGO is limited by *quantum* noise. For neutron star mergers that's the time scale sensitive to nuclear fluid equation of state.

And, by the way … though we know the equations (quantum chromodynamics), we can't compute properties of nuclear matter classically – quantum simulation is needed.

Note: 4 binary black hole mergers detected during August 2017, when three interferometers were operating. A factor of 2 increase in strain sensitivity means detection volume 8 times as big. An event every day?

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

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 may be 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).

Self verification: Check the variance of the energy, which should be zero in an eigenstate.

Should remain feasible with ~ 50 ions.

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

Toward applications of quantum technology

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.

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.

The world should respond *now* to the challenge of quantum-safe privacy.

Quantum sensing, networking, and computing will advance together. Next-generation quantum sensors can provide unprecedented capabilities of potential commercial interest.

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.

Additional
Slides

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.

About Hardware and scalability

Ion trap:

Atomic qubits, all-to-all connectivity, entangling gates via Coulomb-coupled vibration.
Optical control (lasers). 2-qubit gates in 10’s of micro-s.
For scalability, shuttle ions between traps or optical interconnects between traps.

Superconducting circuits:

Cryogenic, qubits are anharmonic oscillators with frequency ~ 5 GHz, 2-qubit gates in 10’s of nano-s.
Microwave electronics for control.
For scalability, 3D integration, perhaps eventually cryogenic control electronics.

-- Characterization of devices is challenging and scales badly.
-- Complex noise models: bursts, leakage, non-Markovian noise, crosstalk can cause trouble.
-- Logical failure rate dominated by rare syndromes, and importance sampling is hard.
-- Use small quantum computers to optimize fault tolerance. Machine learning might help.
-- Error detection protocols demonstrated, but not yet a fully fault-tolerant realization of error correction that convincingly achieves “break even”.

Some recent theory developments

An oracle relative to which BQP is not in the polynomial hierarchy (Raz, Tal).

Fault-tolerant quantum computing with constant overhead using quantum expander codes (Fawzi, Grospellier, Leverrier).

Checking a quantum computer using a classical computer (Mahadev).

More efficient algorithms for simulating quantum dynamics (Haah, Hastings, Kothari, Low).

More efficient algorithms for measuring the spectrum of a Hermitian operator (Poulin, Kitaev, Steiger, Hastings, Troyer).

Average case hardness of random quantum circuit sampling (Bouland, Fefferman, Nirkhe, Vazirani).

No (known) exponential quantum speedup for inversion of low-rank matrices, or for recommendation systems, principle component analysis, supervised clustering (Tang; Gilyén, Lloyd, Tang).

Quantum computing: progress and prospects
NAS Report December 2018

1) It is highly unexpected that a quantum computer that can compromise public key cryptosystems will be built within the next decade.
2) If near-term quantum computers are not commercially successful, government funding may be essential to prevent a significant decline in quantum computing research.
3) Research and development into practical commercial applications of noisy intermediate scale quantum (NISQ) computers is an issue of immediate urgency.
4) It is still too early to be able to predict the time horizon for a scalable quantum computer.
5) The research community should adopt clear reporting conventions to enable comparison between devices.
6) Quantum computing is valuable for driving foundational research that will help advance humanity's understanding of the universe.
7) Although the feasibility of a large-scale quantum computer is not yet certain, the benefits of the effort to develop a practical QC are likely to be large.
8) Continued U.S. support is critical if the United States wants to maintain its leadership position.
9) An open ecosystem that enables cross-pollination of ideas and groups will accelerate rapid technology advancement.
10) Prioritization of the development, standardization, and deployment of post-quantum cryptography is critical for minimizing the chance of a security and privacy disaster.

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.

Quantum Speedups?

When will quantum computers solve important problems that are beyond the reach of the most powerful classical supercomputers?

We should compare with post-exascale classical hardware, e.g. 10 years from now, or more (> 10^{18} FLOPS).

We should compare with the best classical algorithms for the same tasks.

Note that, for problems outside NP (e.g typical quantum simulation tasks), validating the performance of the quantum computer may be difficult.

Even if classical supercomputers can compete, the quantum computer might have advantages, e.g. lower cost and/or lower power consumption.

Noise-resilient quantum circuits

For near-term applications, noise-resilience is a key consideration in quantum circuit design (Kim 2017).

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 linear algebra

QRAM: an N-component vector b can be encoded in a quantum state |b ⟩ 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 ⟩ , the HHL (Harrow, Hassidim, Lloyd 2008) algorithm outputs the quantum state |y⟩ = |A^{-1} b⟩, with a small error, in time O(log N). The quantum speedup is exponential in N.

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

If the input b is classical, we need to load |b⟩ 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.

Speeding up semidefinite programs (SDPs)

Given N X N Hermitian matrices C, $\{A_1, \ldots, A_m\}$ and real numbers $\{b_1, \ldots, b_m\}$, maximize $\mathrm{tr}(CX)$ subject to $\mathrm{tr}(A_i X) \leq b_i$, $X \geq 0$.

Many applications, classically solvable in poly(m,N) time.

Suffices to prepare (and sample from) Gibbs state for H = linear comb. of input matrices. Quantum time polylog(N) if Gibbs state can be prepared efficiently (Brandão & Svore 2016). Output is a quantum state $\rho \cong X$.

When can the Gibbs state be prepared efficiently?
-- H thermalizes efficiently.
-- Input matrices are low rank (Brandão et al. 2017).

Can be viewed as a version of quantum annealing (QA) where Hamiltonian is quantum instead of classical, and where the algorithm is potentially robust with respect to small nonzero temperature.

The corresponding Gibbs state can be prepared efficiently only for SDPs with special properties. What are the applications of these SDPs?

Chemistry and materials

Chemical reactions, catalytic pathways. 100-1000 logical qubits, but high circuit depth for full reaction dynamics.

Cuprates and frustrated spin systems.

Topological materials, including driven and far from equilibrium.

Highly entangled materials for higher precision metrology.

Stable coherent behavior in disordered materials (localization).

Can we machine-learn the density functional?

Condensed matter is easier (better locality) than chemistry, but dynamics is still hard!

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.

Compression of Hamiltonian evolution (improved efficiency of quantum simulators).

Supervised quantum machine learning for pharmaceuticals, catalysts, materials.

Quantum machine learning of the density functional?

Applications to quantum biology?

Quantum error correction

-- Superconducting qubits are ready for demonstrations of QECC.

-- Microwave resonators for continuous-variable codes and cat codes.

-- Topological protection is becoming a reality.

-- Other approaches to physical protection (e.g. superinductance).

-- QEC in the AdS/CFT bulk-boundary dictionary.

Benchmarking

-- Describe noise with a small number of parameters; separate gate and SPAM error.

-- Error correction circuits aren’t random.

-- Importance of coherence, non-Markovianity, noise correlations.

-- Gauge invariant measures, e.g. gate set tomography. Certifying small diamond norm.

-- Tailor the noise by twirling / randomized compiling, etc. ?

-- Eventually, *logical* benchmarking, compared with RB expectations. Ancillas for full fault tolerance. Near term: postselected protocols.

-- What will IBM/Google/Yale/LogiQ etc. *not* be doing?

Quantum Internet

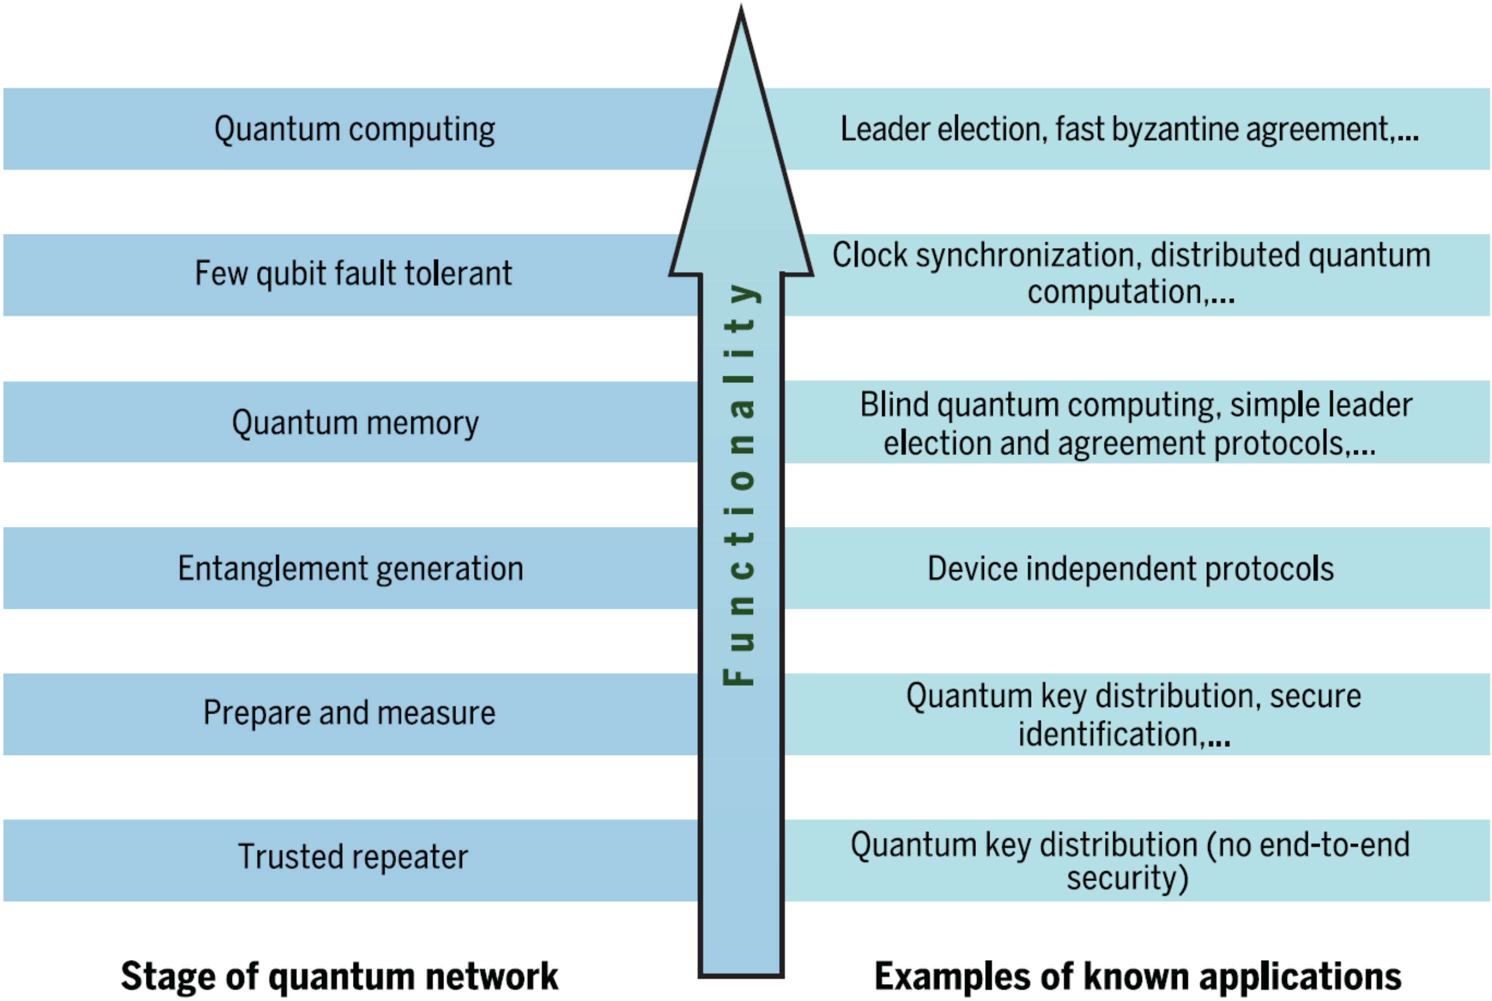

Stage of quantum network	Examples of known applications
Quantum computing	Leader election, fast byzantine agreement,...
Few qubit fault tolerant	Clock synchronization, distributed quantum computation,...
Quantum memory	Blind quantum computing, simple leader election and agreement protocols,...
Entanglement generation	Device independent protocols
Prepare and measure	Quantum key distribution, secure identification,...
Trusted repeater	Quantum key distribution (no end-to-end security)

Wehner, Elkouss, Hanson 2018