Simulating quantum field theory with a quantum computer

John Preskill

INSTITUTE FOR QUANTUM INFORMATION AND MATTER

Lattice 2018
28 July 2018

This talk has two parts

(1) *Near-term prospects for quantum computing.*

(2) *Opportunities in quantum simulation of quantum field theory.*

Exascale digital computers will advance our knowledge of QCD, but some challenges will remain, especially concerning real-time evolution and properties of nuclear matter and quark-gluon plasma at nonzero temperature and chemical potential.

Digital computers may never be able to address these (and other) problems; quantum computers will solve them eventually, though I'm not sure when. The physics payoff may still be far away, but today's research can hasten the arrival of a new era in which quantum simulation fuels progress in fundamental physics.

Frontiers of Physics

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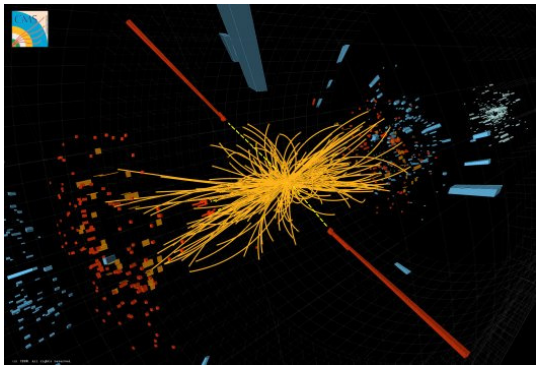		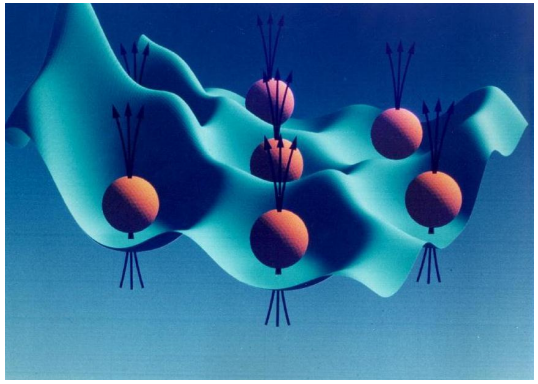
Higgs boson	Large scale structure	“More is different”
Neutrino masses	Cosmic microwave background	Many-body entanglement
Supersymmetry	Dark matter	Phases of quantum matter
Quantum gravity	Dark energy	Quantum computing
String theory	Gravitational waves	Quantum spacetime

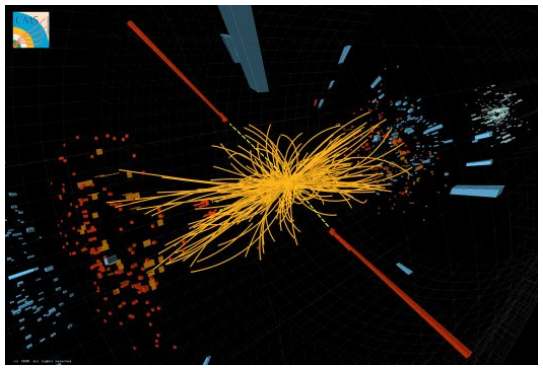

particle collision

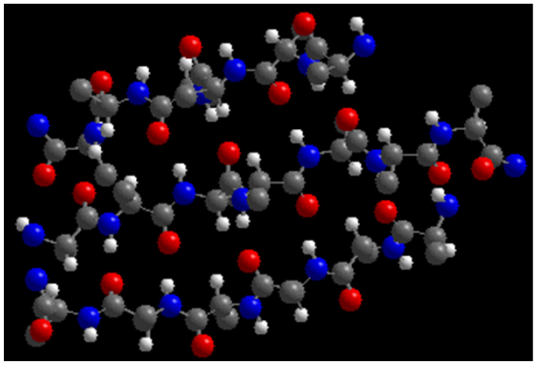

molecular chemistry

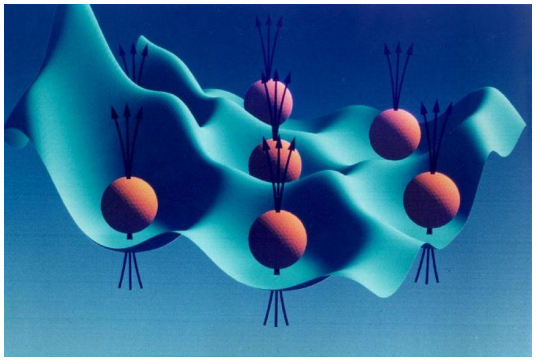

entangled electrons

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

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

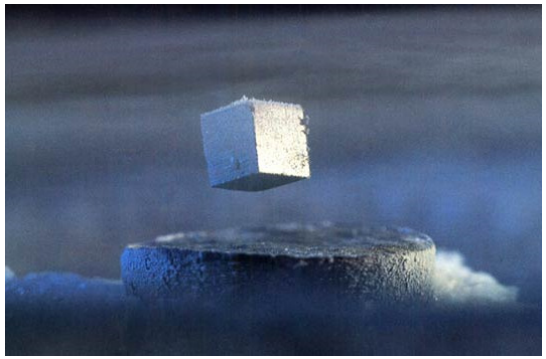

superconductor

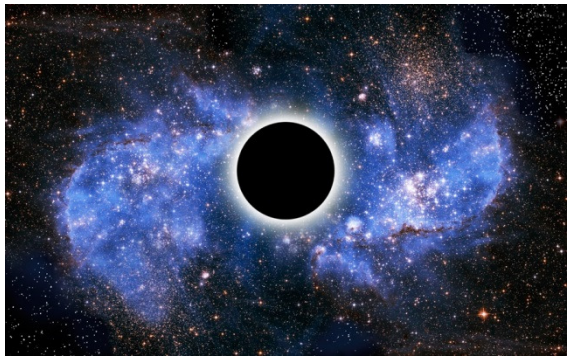

black hole

early universe

Two fundamental ideas

(1) *Quantum complexity*

Why we think quantum computing is powerful.

(2) *Quantum error correction*

Why we think quantum computing is scalable.

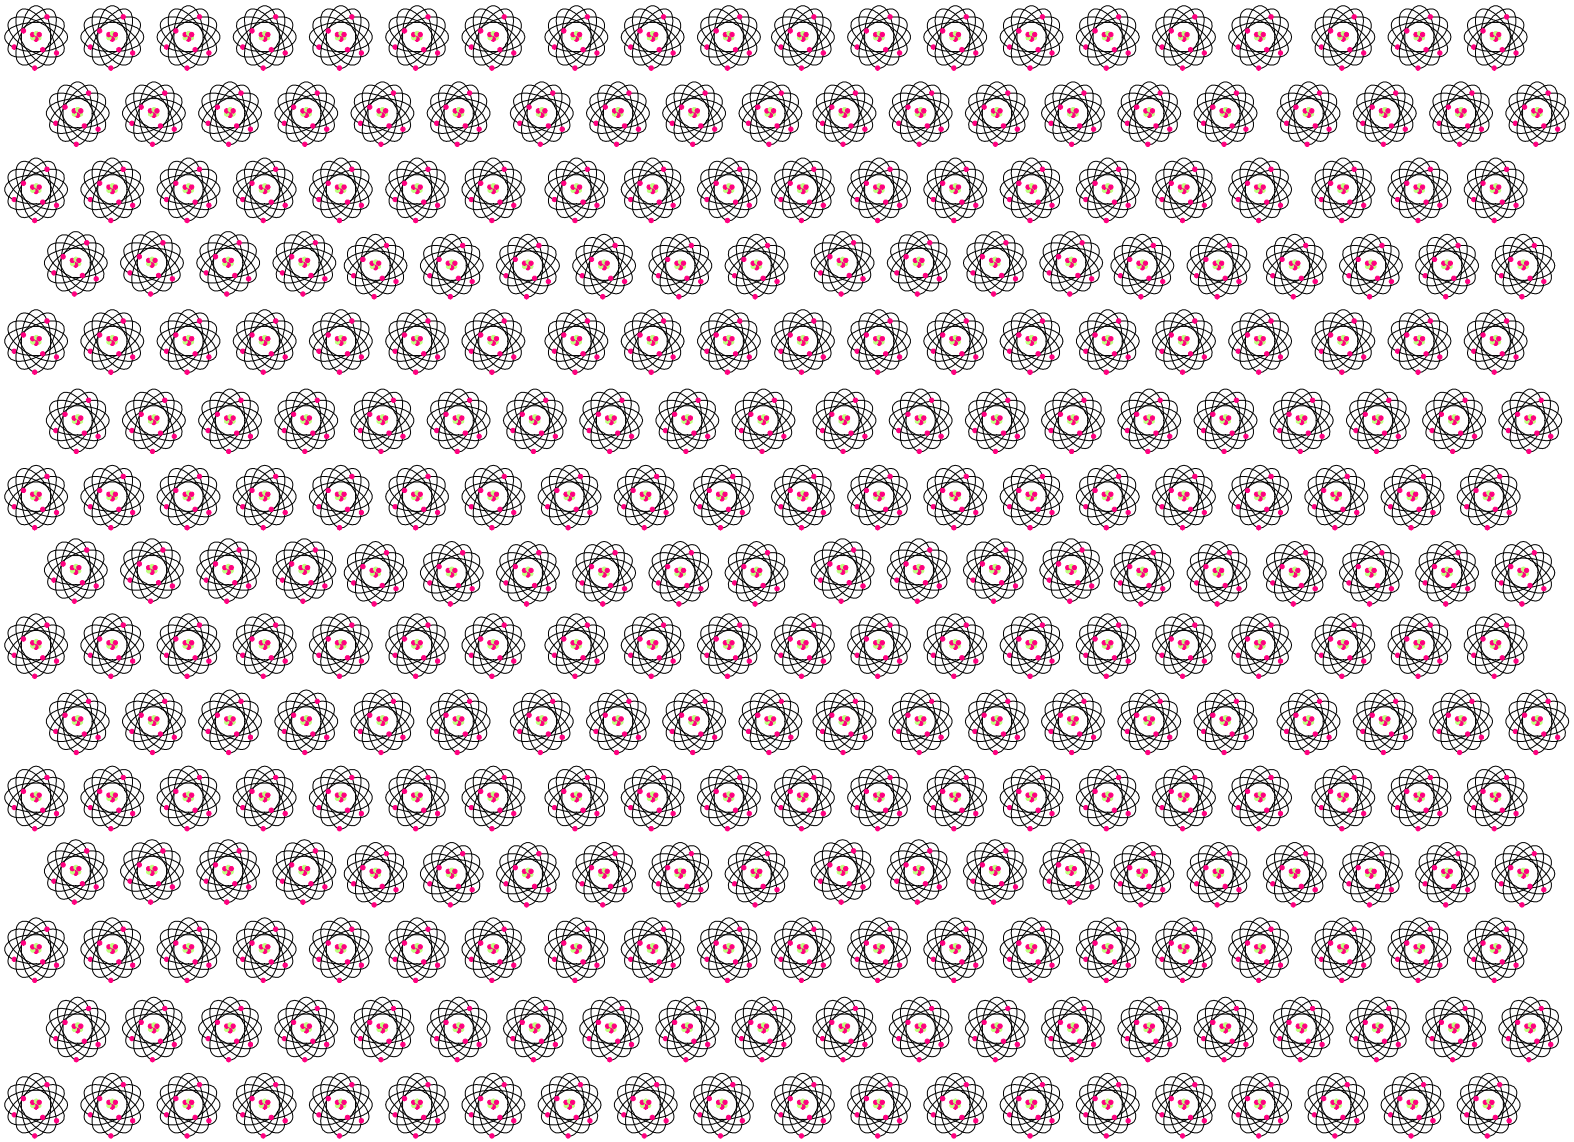

A complete description of a typical quantum state of just 300 qubits requires more bits than the number of atoms in the visible universe.

Why we think quantum computing is powerful

We know examples of problems that can be solved efficiently by a quantum computer, where we believe the problems are hard for classical computers. Factoring is the best known example. No efficient classical algorithm for factoring is known, and not for lack of trying. Factoring numbers which are thousands of bits long is out of reach classically, yet eventually will be feasible quantumly.

Consider the probability distribution of measurement outcomes for n-qubits in a quantum computer. Complexity theory arguments, based on plausible assumptions, indicate that no efficient classical algorithm can efficiently sample from this distribution.

We don't know how to simulate a quantum computer efficiently using a digital ("classical") computer. It is not for lack of trying. The cost of the best simulation algorithm rises exponentially with the number of qubits.

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).

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.

Microsoft: is 2018 the year of the Majorana qubit?

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 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, …

Quantum Speedups?

When will quantum computers solve important problems that are beyond the reach of the post 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.

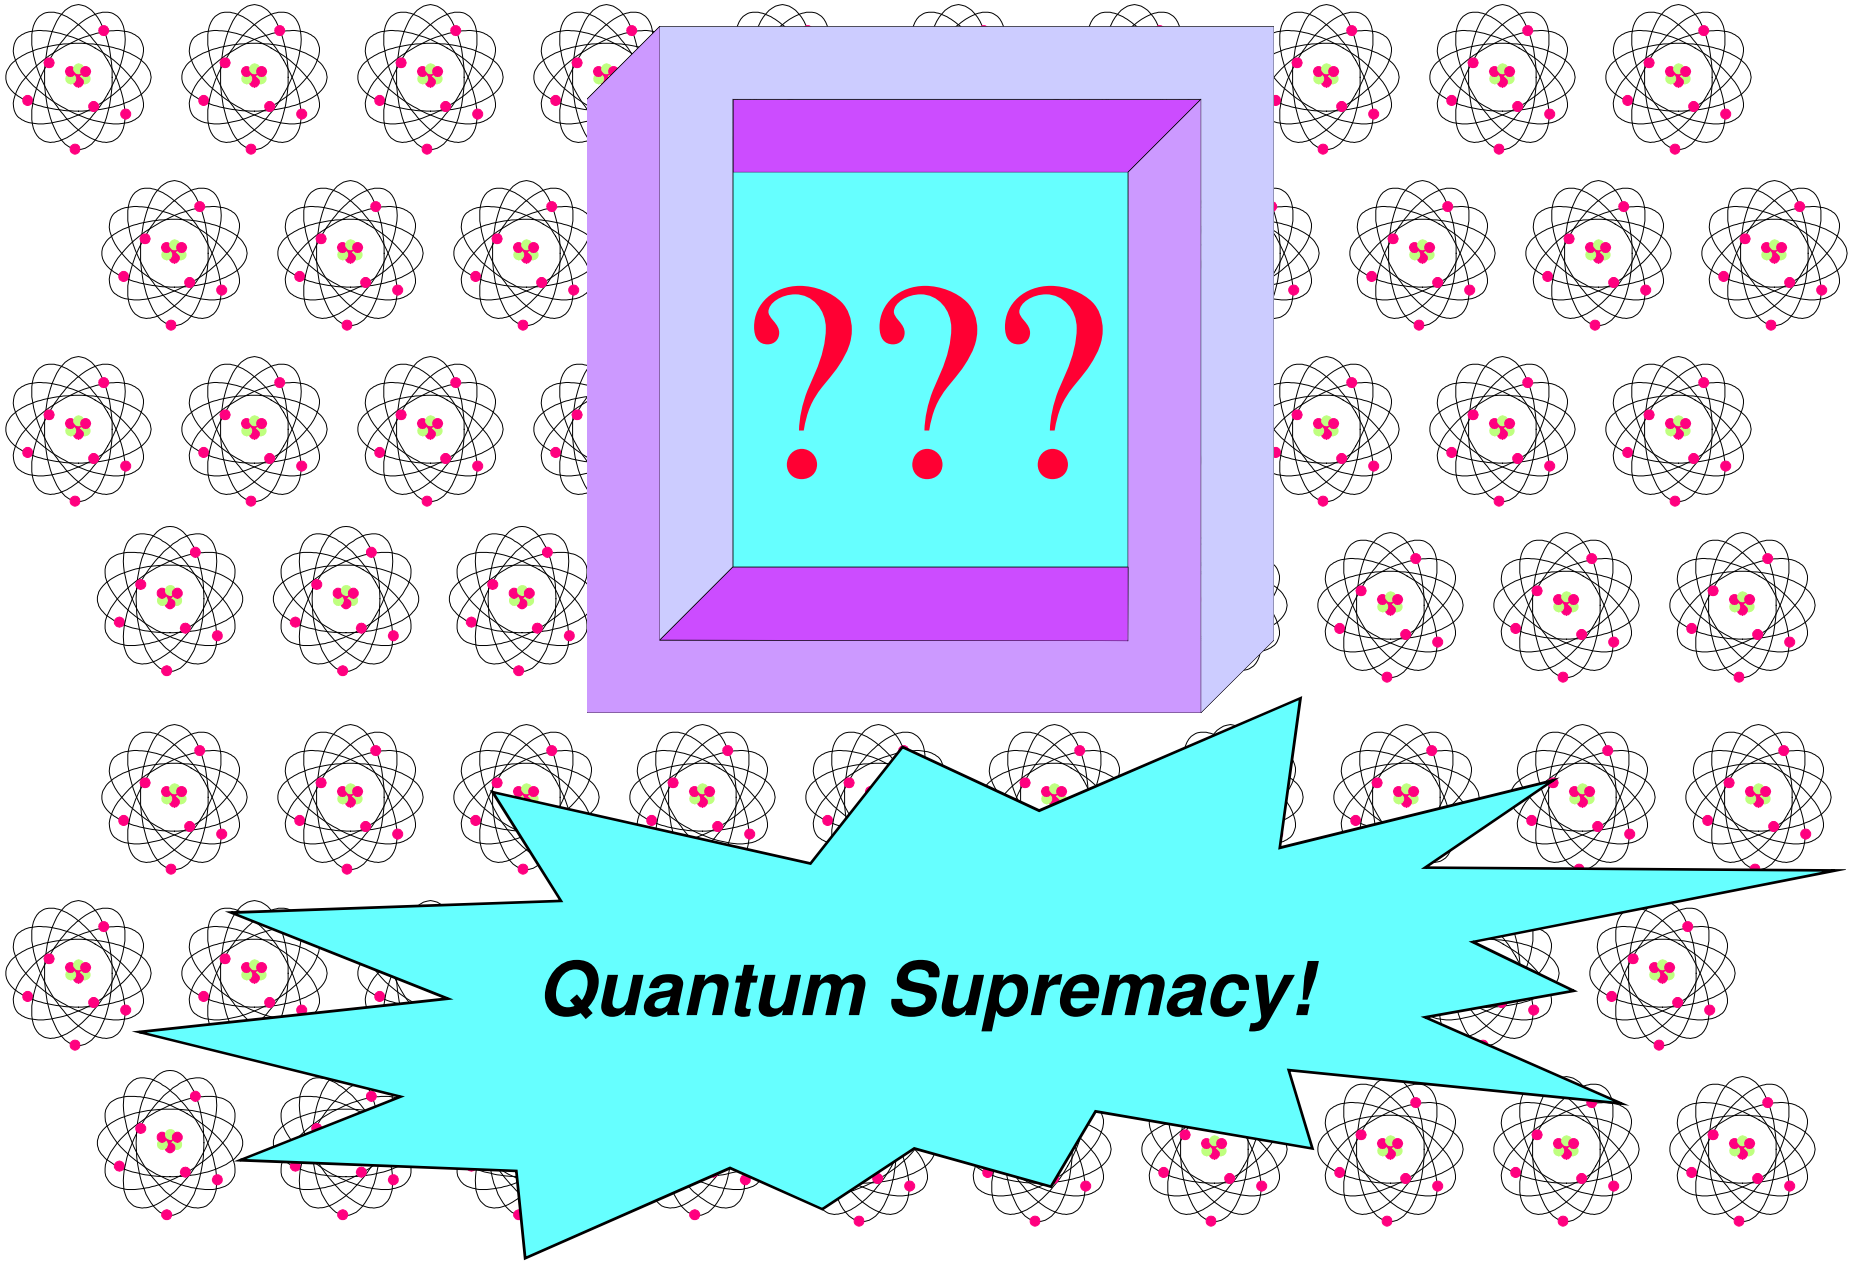
???
Quantum Supremacy!

Hybrid quantum/classical optimizers

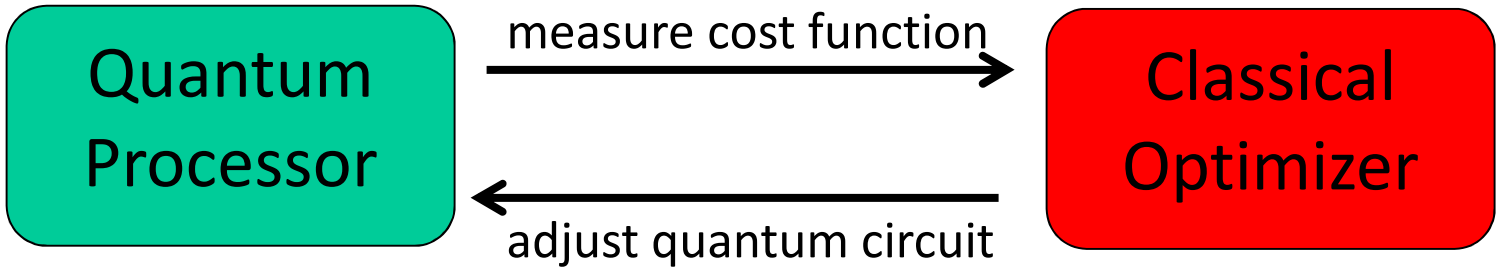

We don't expect a quantum computer to solve worst case instances of NP-hard problems, but it might find better approximate solutions, or find them faster.

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.

Quantum machine learning?

(Classical) deep learning, e.g. restricted Boltzmann machines with multiple hidden layers between input and output. Millions of coupling parameters, optimized on a training set to achieve the proper relation between input and output.

Deep learning may be either unsupervised (unlabeled training set), or supervised (e.g. learning to identify photos).

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 deep learning applications.

However, quantum deep learning is hampered by input/output bottlenecks.

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.

Might be achieved by a (highly controllable) quantum annealer, or other custom quantum device unsuited for general purpose quantum computing. How robust to noise?

Perhaps 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.

Classical deep learning has many applications to quantum science and technology.

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.

Example: Solving (monochromatic) Maxwell's equations in a complex 3D geometry; e.g., for antenna design (Clader et al. 2013). Discretization and preconditioner needed. How else can HHL be applied?

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

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.

QA is a noisy version of adiabatic quantum computing (AQC), and we believe AQC is powerful. Any problem that can be solved efficiently by noiseless quantum computers can also be solved efficiently by noiseless AQC, using a "circuit-to-Hamiltonian map."

But in contrast to the quantum circuit model, we don't know whether noisy AQC is scalable. Furthermore, the circuit-to-Hamiltonian map has high overhead: Many more qubits are needed by the (noiseless) AQC algorithm than by the corresponding quantum circuit algorithm which solves the same problem.

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 (D-Wave, unpublished).

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 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.

Quantum simulation of quantum field theories. Why?

QFT encompasses all fundamental interactions, possibly excluding gravity.

Can quantum computers efficiently simulate any process that occurs in Nature? (Quantum Church-Turing thesis.)

YES and NO are both exciting answers!

Event generators for QCD, etc.

Simulations of nuclear matter, etc.

Exploration of other strongly coupled theories.

Stepping stone to quantum gravity.

Characterizing computational complexity of quantum states.

New insights!

Quantum computing “solves the sign problem”!

Good News / Bad News

Quantum computers simulate quantum systems in real time, not imaginary time.

That's a shame, because imaginary time evolution (in some cases) is an efficient way to prepare ground states and thermal states.

But it's okay, because Nature evolves in real time, too.

And simulation of real time evolution for highly entangled quantum many-body systems (including quantum field theories) is presumed to be hard classically.

Applications include real-time dynamics in strongly correlated quantum many-body systems, quantum chemistry, strongly-coupled relativistic quantum field theory, QCD, nuclear physics, …

We work with the Hamiltonian (not the action), so Lorentz covariance is not manifest. We have to pick an inertial frame, but can obtain frame-independent results (if we're careful).

What problem does the algorithm solve?

Scattering problem: given initial (incoming) state, sample accurately from the distribution of final (outgoing) states.

Vacuum-to-vacuum probability in the presence of spacetime-dependent sources coupled to local observables.

Other S-matrix elements, in cases where particles can be “dressed” adiabatically.

Real-time correlation functions, e.g., for insertions of unitary operators.

Correlation functions and bulk observables at nonzero temperature and chemical potential.

To probe, e.g., transport properties, formulate a simulation that models an actual experiment.

For quantum simulation, no “sign problem” prevents us from performing these tasks efficiently.

Prototypical quantum simulation task

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

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

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

Goal: sample accurately from probability distribution of outcomes.

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

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

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

Need an efficient preparation of initial state.

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

What should we simulate, and what do we stand to learn?

Preparing the ground (or thermal) state of a local Hamiltonian

Can be NP-hard even for a classical spin glass.

And even harder (QMA-hard) for some quantum systems, even in 1D.

But if the state/process exists in Nature, we can hope to simulate it (Quantum Church-Turing thesis).

Same goes for Gibbs states (finite temperature and chemical potential) and states far from equilibrium.

Where did the observed state of our universe come from? That's a question about cosmology …

Prototypical ground-state preparation: prepare ground state for an easy case (e.g., free theory or strong-coupling limit), then adiabatically change the Hamiltonian. Alternatively, we might find a tensor-network approximation via a classical variational algorithm, which can then be compiled as a quantum circuit.

For thermal states, there are *quantum Gibbs sampling* algorithms. Or simulate a thermal bath. Or follow time evolution until equilibration.

Sources of error?

Nonzero lattice spacing a.

Finite spatial volume V.

Discretized fields and conjugate momenta.

Nonzero Trotter step size for simulation of time evolution.

(Diabatic) errors during (adiabatic) state preparation.

These sources of error determine how resources scale with accuracy for the case of an ideal (noiseless) quantum circuit. We also need to worry about noise in (logical) gates.

Where are we now?

Resource scaling (number of qubits and gates) for scattering simulations in scalar and Yukawa theories.

BQP hardness of weakly-coupled 1D scalar theory (with multiple scattering).

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.

In progress: Classical and quantum simulations of nonabelian gauge symmetry, higher dimensions.

Quantum field theory in one spatial dimension: When can it be simulated classically?

For a theory with a mass gap, low energy states are not highly entangled. These states admit a succinct classical description as matrix-product states (MPS) with a relatively low bond dimension.

After a quench, a slightly entangled state becomes more entangled as it evolves. The bond dimension needed in an MPS description grows, and classical simulation may become infeasible if the energy density is nonzero.

What about scattering states? At relatively low energy, only a few particles can be created in a collision; the entanglement remains small, and a time-dependent MPS (TEBD) classical computation may provide an accurate description, even at strong coupling (where Feynman diagram perturbation theory fails).

At high energy, many particles, and much entanglement, can be created; classical simulation may become intractable, but quantum simulation is still efficient.

Even for non-relativistic particles and at weak coupling, multiple scattering events can build up enough entanglement for classical simulation to be hard (the problem becomes “BQP-complete”).

Challenges and Opportunities in quantum simulation

Improving resource costs, greater rigor.

Better regulators: e.g., smearing, improved lattice Hamiltonian, tensor network methods, …

Gauge fields, QCD, standard model, nuclear matter.

Topological defects, massless particles, chiral fermions, SUSY.

Conformal field theory, holography, chaos.

Alternative paradigms, e.g. conformal bootstrap.

Simulations with near-term quantum devices? Hybrid quantum-classical methods. Defining reachable physics goals.

Fresh approaches to noise resilience in quantum simulation.

Collaboration of quantumists and field theorists will be needed to achieve progress!

Additional Slides

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.

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.

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 tr(CX) subject to 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?

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).

Best classical algorithms: cautionary tales

Boson sampling: From 30 photons and 500 modes to 50 photons and 2500 modes (Neville et al. 2017).

Random circuits: Simulating 49 qubits with TB rather then PB memory (IBM 2017) --- trading depth and space.

Best approximation ratio for Max E3LIN2 (with bounded occurence D) achieved by QAOA at level p=1 (Farhi et al. 2014), later surpassed by classical all-star team.

D-Wave evidence for constant factor speedup weakens when quantum annealer is compared with better classical algorithms.

Randomized classical matrix inversion can compete with quantum in *some* parameter regimes (Le Gall).

Tensor network methods for quantum many-body physics and chemistry keep improving (MPS, PEPS, MERA, tensor RG).

Are physically relevant quantum problems really classically hard, even if BQP $\neq$ BPP?

Dynamics seems promising, but MBL (many-body localization) may be classically easy, and ETH (eigenstate thermalization hypothesis = strong quantum chaos) may be physically boring (wisecrack by Frank Verstraete).

How to regulate?

Momentum space is natural for diagonalizing free field theory Hamiltonian, and formulating perturbation theory.

Renormalization group can also be formulated in momentum space.

But real space is better suited for simulation, because the Hamiltonian is spatially local.

We define fields on lattice sites, with lattice spacing a, a source of error. “Bare” parameters. Smaller lattice spacing means better accuracy, but more qubits to simulate a specified spatial volume.

Fields and their conjugate momenta are unbounded operators. We express them in terms of a bounded number of qubits, determined by energy of the simulated process.

Doing better: RG-improved lattice Hamiltonians? Tensor network constructions, e.g., c-MPS, c-MERA, wavelets?

What to simulate?

For example, a self-coupled scalar field in D=2, 3, 4.

$$H = \int d^{D-1}x \left(\frac{1}{2}\Pi^2 + \frac{1}{2}(\nabla\phi)^2 + \frac{1}{2}m_0^2\phi^2 + \frac{1}{4!}\lambda_0\phi^4 \right)$$

Without the ϕ^4 term, a Gaussian theory which is easy to simulate classically (noninteracting particles).

With this interaction term, particles can scatter. The dimensionless coupling parameter is λ/m^{4-D}. Classical simulations are hard when the coupling is strong.

Hardness persists even at weak coupling, if we want high precision, or if particles interact long enough to become highly entangled.

Summing perturbation theory is infeasible, and misses nonperturbative effects.

We assume theory has a mass gap.

A simulation protocol

Input: a list of incoming particle momenta.
Output: a list of (perhaps many) outgoing particle momenta.

Procedure:
(1) Prepare free field vacuum ($\lambda_0 = 0$).
(2) Prepare free field wavepackets (narrow in momentum).
(3) Adiabatically turn on the (bare) coupling.
(4) Evolve for time t with Hamiltonian H.
(5) Adiabatically turn off interaction.
(6) Measure field modes.

Assume no phase transition blocks adiabatic state preparation.

Alternative: create particles with spacetime dependent classical sources (better if there are bound states). Simulate detector POVM.

Lorentz invariance brutally broken in lattice theory, recovered by tuning bare H. (Ugh.) Also tune for achieve $ma \ll 1$.

Example: ϕ^4 theory in D=3 spacetime dimensions

Error ε scales with lattice spacing a as $\varepsilon = O(a^2)$.

Number of qubits Ω needed to simulate physical volume V *is* $\Omega = V/a^2 = O(1/\varepsilon)$.

Gaussian state preparation (matrix arithmetic) uses $\Omega^{2.273}$ gates. (though a customized algorithm exploiting translation invariance does better).

Scaling with energy E: number of gates = $O(E^6)$.

Factor E from Trotter error, E^2 from lattice spacing $a \sim 1/E$, E^3 from diabatic error.

Dominant diabatic error comes from splitting of $1 \rightarrow 3$ particles, for which energy gap $\sim m^2/E$.

Thousands of logical qubits for $2 \rightarrow 4$ scattering with 1% error at $E/m = O(1)$. Yikes!