

BES Roundtable on Quantum Computing Opportunities in Chemical and Materials Sciences

*John Preskill
31 October 2017*




INSTITUTE FOR QUANTUM INFORMATION AND MATTER



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



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*

(2) *Quantum error correction*

This talk emphasizes quantum speedups, including potential near-term applications of quantum computing. I won't say much about quantum error correction and fault-tolerant quantum computing, even though these will be essential for long-term scalability.

Quantum Complexity

n-qubit Hilbert space has dimension 2^n .


- Very complex to describe classically
- Need petabytes to store a 50-qubit state

But measuring n qubits yields no more than n bits in one shot.

And the only states we really care about are those that can be prepared with reasonable (quantum) resources. Only these can arise in Nature, and only these are within reach of the quantum engineers.

A mathematical model of the feasible n-qubit pure states: they can be prepared by a circuit of $\text{poly}(n)$ two-qubit gates, applied to an initial (unentangled) product state. (A fraction $\exp[-\text{poly}(n)]$ of all pure states.) Likewise, feasible measurements are $\text{poly}(n)$ -size circuits followed by single-qubit measurements.

Equivalently, feasible states can be prepared starting from product states through evolution in $\text{poly}(n)$ time governed by a *local* Hamiltonian.



Quantum Complexity

Can we perform super-classical tasks using near-term quantum devices?

Evidence that quantum computing is powerful ...


-- Quantum algorithms that solve problems which are believed to be classically hard --- e.g. [factoring](#).

-- Complexity theory. [Hardness of sampling classically](#) from the probability distribution of measurement outcomes, for a measurement performed on a quantum system which can be prepared by a quantum circuit with relatively low depth.

(1) Hardness of multiplicative approximation follows from very plausible assumptions.

(2) Hardness of additive approximation follows from reasonable assumptions.

-- We have [no good classical algorithm](#) for simulating a quantum circuit classically, or for simulating the time evolution of a quantum state governed by a local Hamiltonian.



Quantum Speedups?

We should compare with **post-exascale hardware**, e.g. 10 years from now (or more).

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

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

- Test e.g. Clifford circuits, which can be simulated.
- Run computation backwards?
- Interactive protocol founded on complexity assumptions (e.g. quantum one-way functions).

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 than PB memory (IBM 2017) --- trading depth and space.

Best approximation ratio for **Max E3LIN2** (with bounded occurrence 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).

Quantum hardware: state of the art

IBM Quantum Experience in the cloud: now 16 qubits (superconducting circuit)

Google 22 qubit device (superconducting circuit), 49 qubits next year.

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.

Google 9-qubit simulation of dynamics in spin system with pulsed random couplings; estimated error rate is 0.4% per qubit and pulse. (Classical machine learning used to optimize Hamiltonian control.) For this error rate, meaningful results should be attainable in a 50-qubit system.

Microsoft: is 2018 the year of the Majorana qubit?

And many other interesting platforms ...

Quantum 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.

Hybrid quantum/classical algorithms. Combine quantum evaluation of an expectation value 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 H . (Much easier than algorithms which require simulation of time evolution governed by H .)

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 (Shor):

Simplex method: experiments showed it’s fast long before theorists could prove that it’s fast.

Turbo and LDPC codes. Experiments showed they come close to achieving Shannon capacity, surprising theorists.

Metropolis algorithm: experiments showed it is useful.

Possible quantum examples:

Quantum annealers, approximate optimizers, variational eigensolvers, ... playing around may give us new ideas.

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

What can we do with, say, 50-100 qubits, depth 25-50? **Talk to users!**

Quantum machine learning?

(Classical) deep learning, e.g. restricted Boltzmann machines with multiple hidden layers between input and output. Millions of couplings, optimized on training set for proper relation between input and output.

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

Quantum deep learning hampered by input/output bottlenecks (need for qRAM).

Potentially faster training, e.g., fewer samples needed, or more efficient thermalization? How sensitive to noise?

Might be achieved by (highly controllable) quantum annealer, or other custom quantum device unsuited for general purpose quantum computing.

Perhaps more natural to consider quantum inputs / outputs; e.g. better ways to characterize or control quantum systems.

Potential quantum applications of *classical* machine learning include e.g. improved control of quantum systems, or decoding quantum error-correcting codes when noise model is unknown.

Power of quantum annealing?

Adiabatic quantum computing (AQC) is powerful (unless $BQP=BPP$). Any problem that can be solved using (noiseless) polynomial-size quantum circuit families can also be solved by adiabatically preparing, and then measuring, the ground state of an (ideal) local Hamiltonian, where the time to prepare the state scales polynomially with input size. AQC fails on instances which are hard for QC (energy gap gets exponentially small).

But in contrast to the circuit model, we don't know whether noisy AQC is scalable. And the circuit-to-Hamiltonian map has high overhead.

The D-Wave machine is a 2000-qubit quantum annealer (QA), attempting to solve optimization problems. We don't have a good theoretical argument for a QA speedup apart from the circuit-to-Hamiltonian map.

Theorists are more hopeful that a QA can achieve speedups if the Hamiltonian has a sign problem (is “non-stoquastic”). Then e.g. Quantum Monte Carlo cannot do a classical simulation.

Assessing the performance of QA may already be beyond the reach of classical simulation. We need more quantum annealers to experiment with. That is the rationale for the IARPA Quantum Enhanced Optimization (QEO) program.

Might also use quantum annealer for quantum simulations rather than classical optimization problems (D-Wave, unpublished).

Noise-resilient quantum circuits

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

Noiseless constant-depth (Clifford) quantum circuits can solve problems beyond reach of constant-depth classical circuits (Bravyi-Gosset-Koenig 2017, no complexity theory assumptions needed). What about quantum circuits with constant error rate per output qubit?

We say a noisy circuit is “k-noise-resilient” if any k-local operator has an expectation value close to its ideal value. That’s good enough to access physically useful information (like energy per site, magnetization, response functions) and may apply for non-constant circuit depth (Kim 2017).

In some circuits, an error occurring at an earlier times may decay away by a later time. This can occur for locality-preserving MERA-like circuits. And for MERA prepared by a quantum circuit, $\log \chi$ qubits suffice to realize a bond dimension of χ , a big quantum advantage. (Kim and Swingle, unpublished).

A variational quantum eigensolver can use these quantum tensor-network states, potentially achieving more accurate results than classical methods.

Matrix inversion on a quantum computer

Given classical input A ($N \times N$ matrix, sparsity s and condition number κ) and N -qubit quantum input $|b\rangle$, algorithm **outputs**

$|y\rangle = |A^{-1} b\rangle$ with error ε . Quantum run time is fast! (HHL 2008)

Quantum time $\sim s \kappa^2 [\log(\kappa/\varepsilon)]^{2.5} \log N$ [Speedup exponential in N .]

--- No exp. speedup if A not sparse (more precisely, we need to be able to simulate $\exp[-i A t]$) or not well-conditioned.

--- Input $|b\rangle$ and output $|y\rangle = |A^{-1} b\rangle$ are quantum!

--- We can sample from measurements of $|y\rangle$.

--- If input b is classical need to load $|b\rangle$ into “qRAM” in polylog time.

If matrix A is not well-conditioned, we can invert a well-conditioned part of A (at a modest cost), or apply a “preconditioner” B – e.g. invert AB instead of A if AB has a lower condition number.

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

Applications of quantum linear algebra

Given classical input A ($N \times N$ matrix, sparsity s and condition number κ) and N -qubit quantum input $|b\rangle$, algorithm outputs $|y\rangle = |A^{-1} b\rangle$ with error ε .

It is more promising if the **input b is computed** rather than entered from a database.

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

Alternative method for solving classical scattering problems: **quantum simulation of $N \times N$ Laplacian** using $O(\log N)$ qubits (Jordan et al. 2017). Need efficient preparation of initial state (e.g. input Gaussian wavepacket).

Recommendation systems (e.g. Netflix/Amazon with $m=10^8$ users and $n=10^6$ products). Sample rapidly from preference matrix with *low-rank* $k \cong 100$ (Kerenidis & Prakash 2016). Quantum queries to a classical data structure: Linear-time offline preprocessing, online processing of quantum queries in time $\text{poly}(k)$ $\text{polylog}(mn)$.

Speeding up semidefinite programs (SDPs)

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

Many applications, classically solvable in $\text{poly}(m, N)$ time.

Suffices to prepare (and sample from) Gibbs state for $H = \text{linear comb. of input matrices}$. Quantum time $\text{polylog}(N)$ 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?

Quantum simulation of dynamics: 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.

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

And simulation of real-time evolution for highly entangled quantum many-body systems 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, ...

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?

Quantum chaos

Classical chaos advanced rapidly with onset of numerical simulation of classical dynamical systems.

Goal: advancing theory of quantum chaos via real-time quantum simulation experiments.

Simulating thermalization to elucidate foundations of quantum statistical mechanics.

How rapidly mixing are quantum circuits with randomness?
Hamiltonian systems with high connectivity?

What time-independent Hamiltonians are “fast scramblers?”

Simulations with ~ 100 qubits could be revealing, if not too noisy.

Simulation: fast scrambling

-- Sachdev-Ye-Kitaev model (N Majorana modes, randomized four-Fermi coupling J).

$$H = \sum J_{abcd} \Psi_a \Psi_b \Psi_c \Psi_d, \quad N \gg 1, \quad \beta J \gg 1,$$

-- Out-of-time-order correlators indicate scrambling time $t \sim (\beta/2\pi) \log N$, matching a 1+1 dimensional “AdS black hole”.

-- Classically simulated for up to N=32 (16 qubits), by exact diagonalization (Shenker et al. 2016)

-- *Highlights the importance of nonlocal couplings.*

-- What other simple models have holographic duals?

-- *Robustness to noise? Might be very sensitive.*

-- Dynamics of small black holes. Seeing the black hole interior?

Digital vs. Analog quantum simulation

Analog is very active now, in both experiment and theory. Digital is more aspirational.

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

High connectivity among qubits highly desirable (e.g., for probing scrambling of quantum information).

There are ambitious proposals for simulating gauge field theories with existing experimental tools, e.g., using ultracold atoms.

Analog simulations are limited by imperfect control. What robust universal properties are hard to access in classical simulations, yet accessible in noisy quantum simulations?

Eventually, digital quantum simulation will surpass analog, but perhaps not until fault tolerance is feasible.

The steep climb to scalability

We tend to be too optimistic about the short run, too pessimistic about the long run.

Long-lived logical qubits, protected by quantum error correction, are likely to be realized in the next few years.

But near-term quantum algorithms will need to tolerate noise. **Fully fault-tolerant quantum computing may still be decades away.** We don't really know ...

Lower gate error rates will substantially reduce the overhead cost of fault tolerance, and also extend the reach of quantum algorithms which do not use error correction. (Topological quantum computing is one aspirational approach to achieving much lower error rates.)

Platforms with faster gates have shorter time to solution, all else being equal. This may become more important in the longer term.

Very daunting systems engineering issues!

A long way to go! New insights and developments could substantially alter the outlook for scalability.

Quantum speedups

Can near-term quantum computers compete with 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.

Realistically, the goal for near-term quantum platforms should be to pave the way for bigger payoffs using future devices.

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.