Quantum computing and quantum information

John Preskill

Quantum computing for nuclear physics

Institute for Nuclear Theory
14 November 2017

Frontiers of Physics

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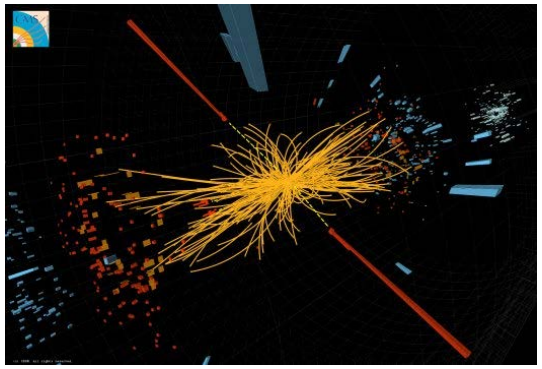		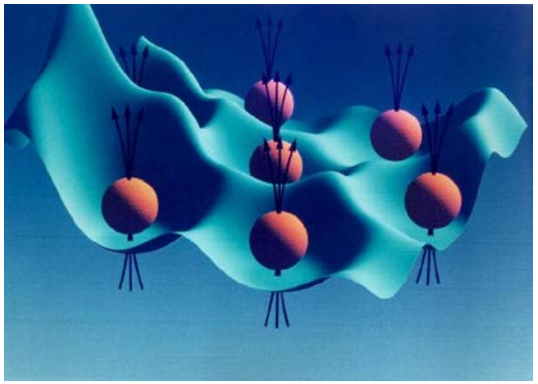
Higgs boson	Large scale structure	“More is different”
Neutrino masses	Cosmic microwave background	Many-body entanglement
Supersymmetry	Dark matter	Phases of quantum matter
Quantum gravity	Dark energy	Quantum computing
String theory	Gravitational waves	Quantum spacetime

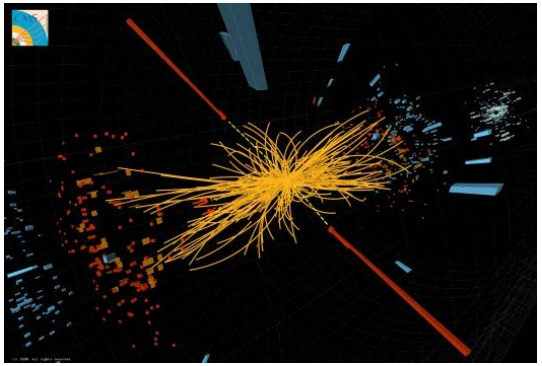

particle collision

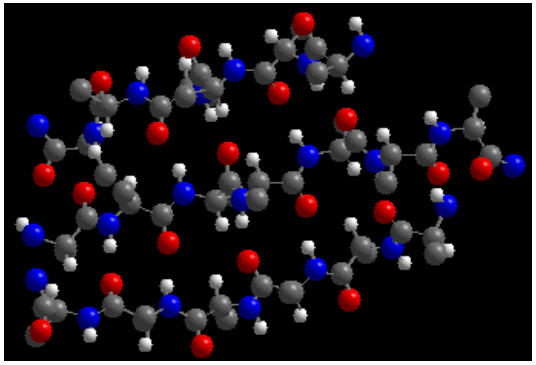

molecular chemistry

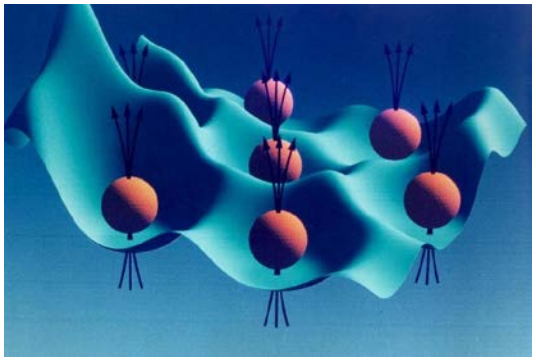

entangled electrons

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

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

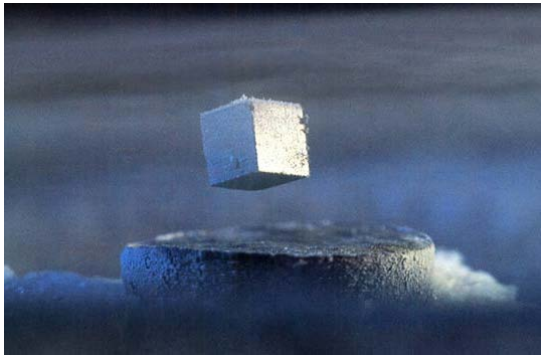

superconductor

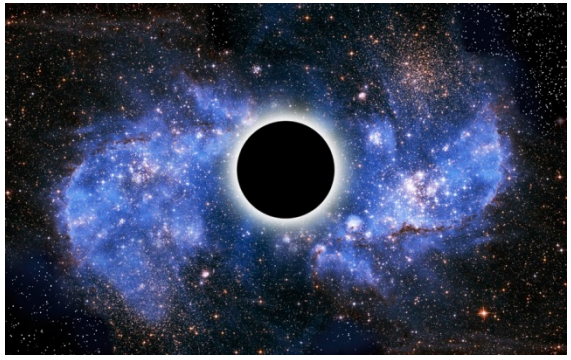

black hole

early universe

Two fundamental ideas

(1) *Quantum complexity*

(2) *Quantum error correction*

I won't say much today about quantum error correction and fault-tolerant quantum computing, even though these will be essential for long-term scalability.

Quantum Complexity

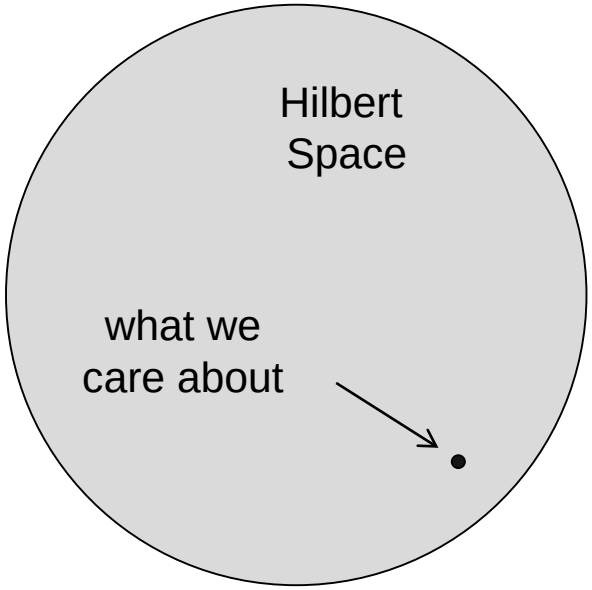

n-qubit Hilbert space has dimension 2^n.

-- 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 poly(n) two-qubit gates, applied to an initial (unentangled) product state. (A fraction exp[-poly(n)] of all pure states.) Likewise, feasible measurements are poly(n)-size circuits followed by single-qubit measurements.

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

Quantum Complexity

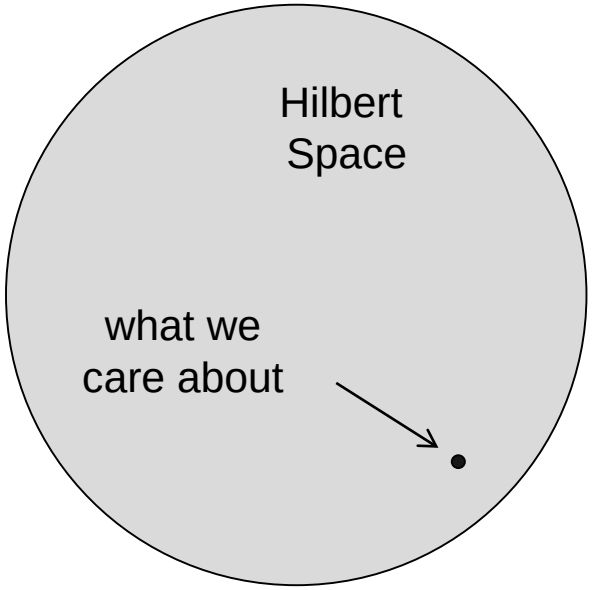

Can we perform super-classical tasks using near-term quantum devices?

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).

Quantum speedups for physics and chemistry problems?

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). 20 qubits by end of 2017, 50-qubit device "built and measured."

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.

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

Microsoft: is 2018 the year of the Majorana qubit?

And many other interesting platforms …

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 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!

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 χ 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.

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 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!

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 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.

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.

Digital vs. Analog quantum simulation

Analog (not based on a compiled quantum circuit) 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.

Our work (with Stephen Jordan and Keith Lee) concerns (error corrected) digital quantum simulation.

Challenges and Opportunities in quantum simulation

Improving resource costs, greater rigor.

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

Simulations with near-term quantum devices?

Scattering of topological defects as probe of nonperturbative physics.

Gauge fields, QCD, standard model, nuclear matter.

Massless particles, chiral fermions, SUSY.

Conformal field theory, holography, chaos.

Alternative paradigms, e.g. conformal bootstrap.

Fresh approaches to noise suppression in quantum simulation.

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.

Additional Slides

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).

Quantum machine learning?

(Classical) deep learning, e.g. restricted Boltzmann machines with multiple hidden layers beween 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.

Matrix inversion on a quantum computer

Given classical input A (N X 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 ~ $s \kappa^2 [\log(\kappa/\varepsilon)]^{2.5} \log N$ [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 X 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 X 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 poly(k) polylog(mn).

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?

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).

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).

$$H = \Sigma\, J_{abcd}\, \psi_a \psi_b \psi_c \psi_d\,, \qquad N >> 1, \qquad \beta J >> 1,$$

-- 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?

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 (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.

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.

How resources scale with precision

Consider, e.g., a spin system on a lattice. Hamiltonian is a sum of k-local terms, each with O(1) operator norm.

Simulation cost for evolution operator U(t) is optimal in error (~ log(1/error), nearly linear in evolution time t.

But QFT has (formally) unbounded number of sites and unbounded operators at each site.

Our cost scales like poly(1/error). The obstacle is the lattice spacing, not the Trotter error.

Can we do better?

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 $ma << 1$.

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!