Quantum computing and the entanglement frontier

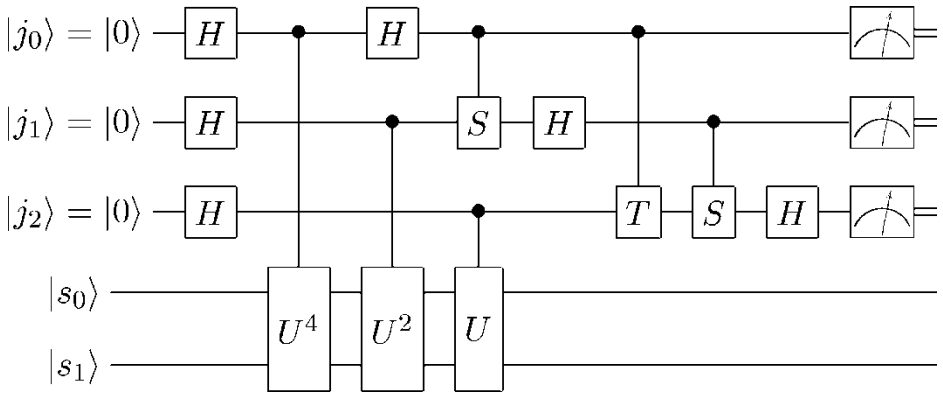

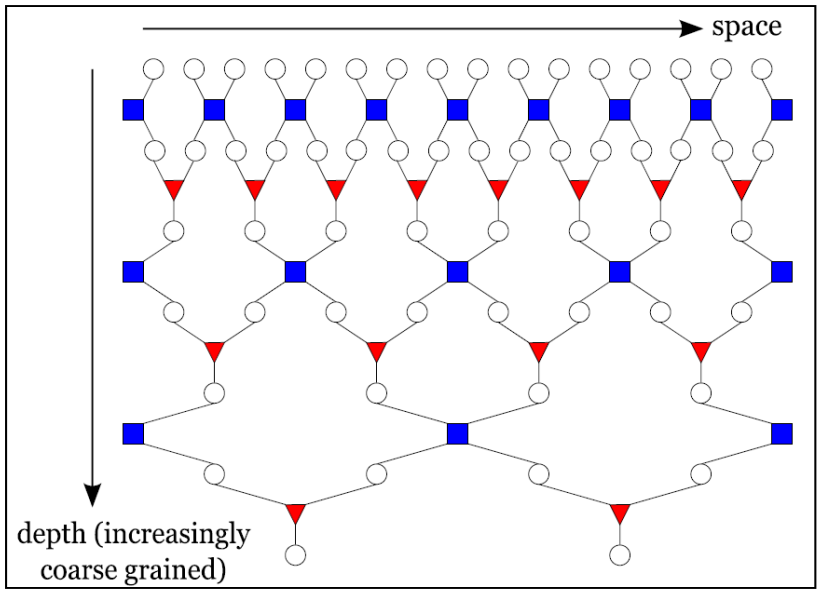

Institute for Quantum Information and Matter

John Preskill, Caltech
Solvay Conference
19 October 2011

Big Questions

HEP:

What underlying theory explains the observed elementary particles and their interactions, including gravity?

QIS:

Can we control complex quantum systems and if so what are the scientific and technological implications?

Not the frontier of short (subnuclear) distances or long (cosmological) distances, but rather the frontier of highly complex quantum states: *The entanglement frontier*

Also: emergence of classicality, security of quantum cryptographic protocols, foundations of statistical mechanics and thermalization, information theoretic principles illuminating the foundations of quantum physics, information processing by e.g. black holes, etc.

Truism:
the macroscopic world is classical.
the microscopic world is quantum.

Goal of QIS:
controllable quantum behavior in scalable systems

Why?

Classical systems cannot simulate quantum systems efficiently (a widely believed but unproven conjecture).

But to control quantum systems we must slay the dragon of decoherence …

Is this merely *really, really hard*?
Or is it *ridiculously* hard?

Quantum entanglement

If you read ten pages of an ordinary hundred-page book, you learn about 10% of the content of the book. But if you read ten pages of a “typical” hundred-page quantum book, you learn almost nothing about the content of the book. That's because nearly all the information in a quantum book is encoded in the correlations among the pages; you can't access it if you read the book one page at a time.

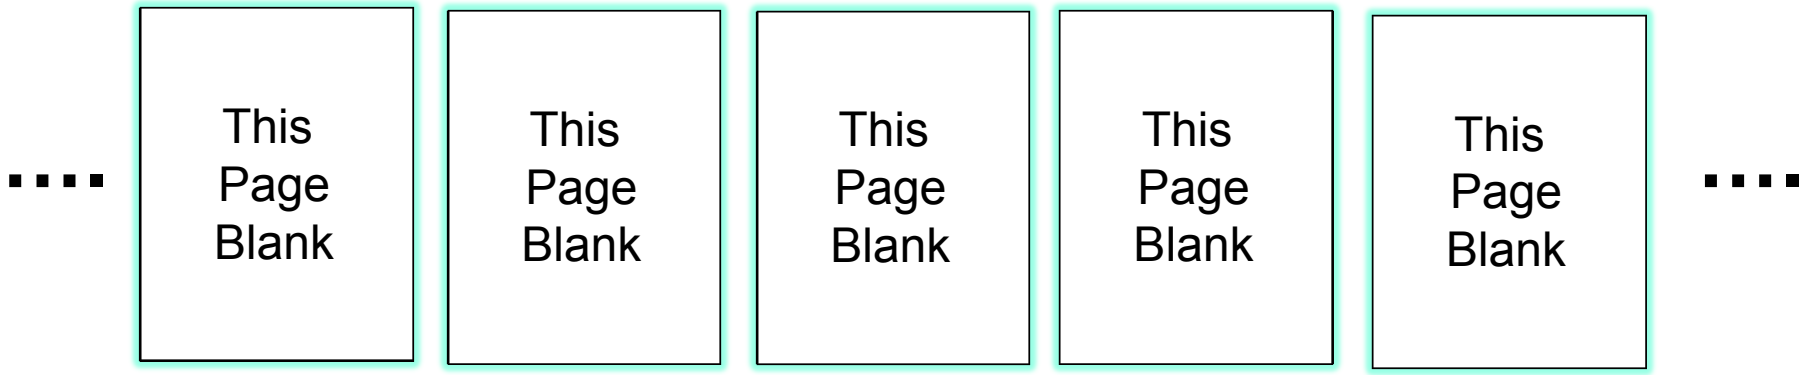

Describing “typical” quantum states with many parts requires extravagant (exponential) classical resources.

Can we verify that Nature allows states with no succinct classical description?

Complexity

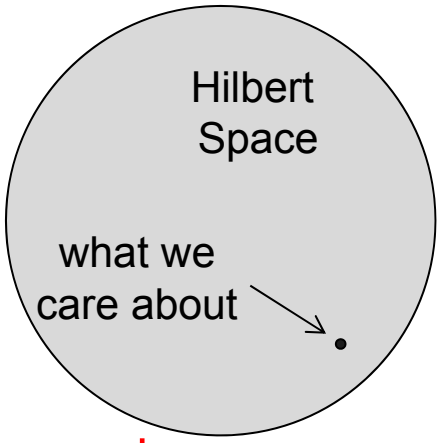

Hilbert space is vast.

But typical quantum states are boring, because
-- they are not useful
-- preparing them is not feasible

The only states we need 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, they can be prepared starting from product states through evolution in poly(n) time governed by a *local* Hamiltonian.

Hubris(?): If Nature can do it, so can we! (Someday…)

But the states and measurements that are quantumly feasible may be hard to simulate classically.

Some quantum algorithms

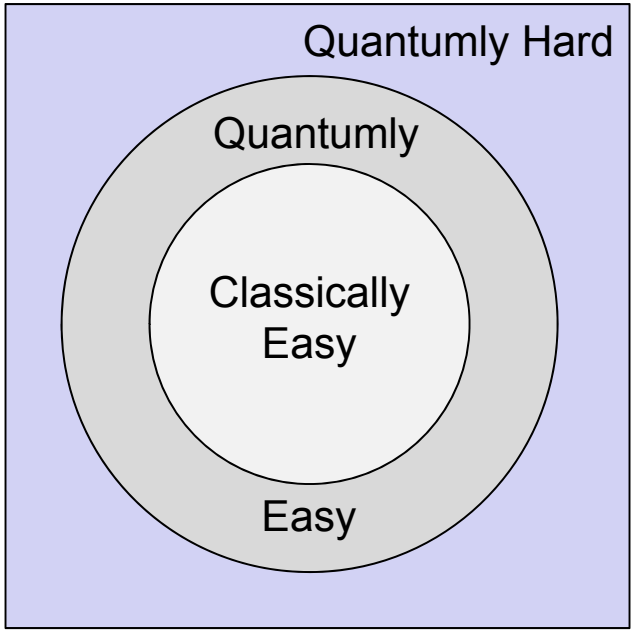

Factoring and finding discrete logarithms (Shor 1994).
Idea: Finding period of a function by Fourier transform.
Application: Breaking classical public key cryptosystems.
Speedup: superpolynomial

Approximating knot invariants (Freedman et al. 2000).
Idea: simulating topological quantum field theory.
Application: Unforgeable quantumly verifiable money.
Speedup: superpolynomial

Exhaustive search in an unstructured space (Grover 1996)
Idea: Increase the angle with marked state by $N^{-1/2}$ in each of many iterations.
Application: finding solutions to NP-hard combinatorial search problems
Speedup: quadratic ($N^{1/2}$ vs. N, where N is the number of possible states)

Evaluation of Boolean formulas (Farhi et al. 2007)
Idea: simulating quantum walk (i.e. scattering) on a tree.
Application: Determining if a two-player game has a winning strategy.
Speedup: polynomial ($N^{.5}$ vs. $N^{.753}$, where N is the number of leaves on the tree)

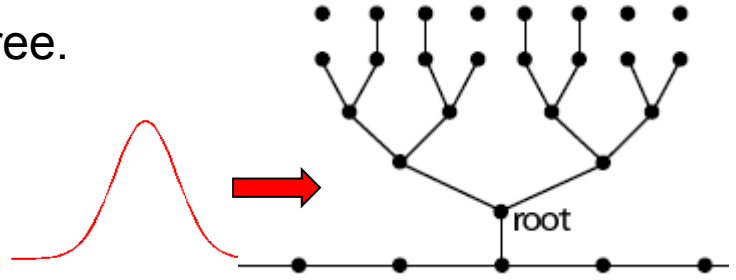

Many more quantum algorithms at math.nist.gov/quantum/zoo/

Quantum algorithms

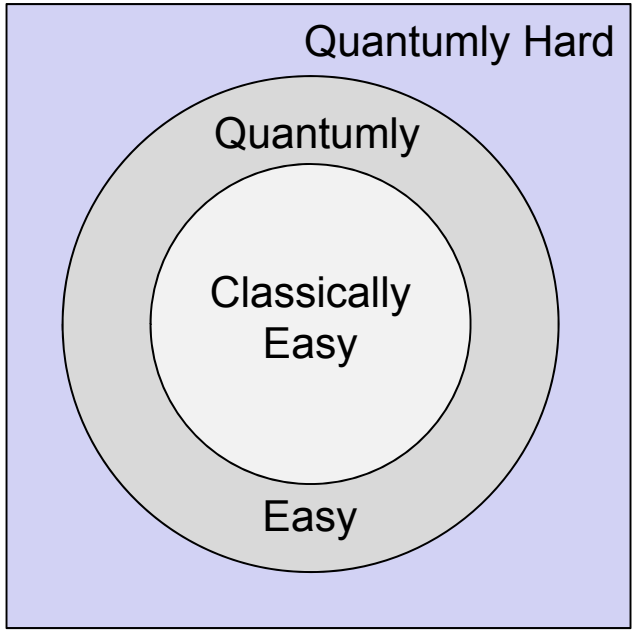

Exploring complexity: We should be able to check (someday) that quantum physics exploits extravagant resources by verifying superpolynomial speedups for (NP) problems where solution can be checked classically, like factoring. (However, there is no *proof* that factoring is hard classically.)

Not NP-hard (in the worst case): Superpolynomial quantum speedups seem to be possible only for problems with special structure, not for NP-complete problems like 3-SAT. Quantum physics speeds up unstructured search quadratically, not exponentially.

Beyond NP: Speedups for problems *outside* NP are also common and important. (Indeed he "natural" application for a quantum computer is simulating evolution governed by a local Hamiltonian, preceded by preparation of a "reasonable" state and followed by measurement of a "reasonable" observable.)

In such cases the findings of a quantum computer might not be easy to check with a classical computer; instead, one quantum computer must be checked by another (or by doing an experiment, which is sort of the same thing).

Toward quantum supremacy

The quantum computing adventure will enter the new, more mature phase of "quantum supremacy" once we can prepare and control complex quantum systems that behave in ways that cannot be predicted using digital computers (systems that "surpass understanding" and surprise us).

To reach that goal, it will be useful to gain a deeper understanding of two questions:

What quantum tasks are feasible?
What quantum tasks are hard to simulate classically?

Might it be that the extravagant "exponential" classical resources required for classical description and simulation of generic quantum states are illusory, because quantum states in Nature have succinct descriptions?

Easy to simulate: “slightly entangling” quantum computation

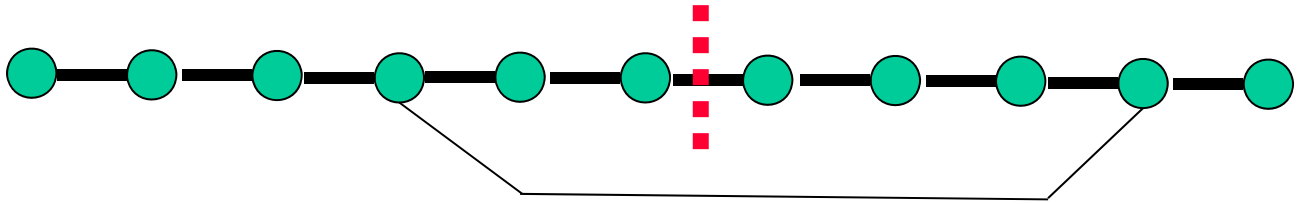

Arrange n qubits on a line. The initial state is a product state, and for any cut into two sets, the number of gates acting across the cut is E. Then if
E = O(log n), the state has a succinct classical description (as a “matrix-product state”) and the quantum computation can be simulated “efficiently” (Vidal 2003, Jozsa 2006).

Is quantum state tomography hard?

The density operator of an n-qubit quantum state has 4^n entries, most of which are exponentially small. Therefore state tomography requires lots of experiments and lots of post-processing. But if the state has a succinct description, how hard is it to find it?

Suppose the state is a matrix product state, e.g. a good approximation to the slightly entangled ground state of a gapped local 1D Hamiltonian (Hastings 2007). Then the quantum complexity of tomography is linear in n. It suffices to do tomography on constant size segments, and then piece the information together using (poly-time) classical post-processing (Cramer et al. 2011). The accuracy can be certified without *a priori* assumptions.

Easy to simulate:
Gaussian linear optics

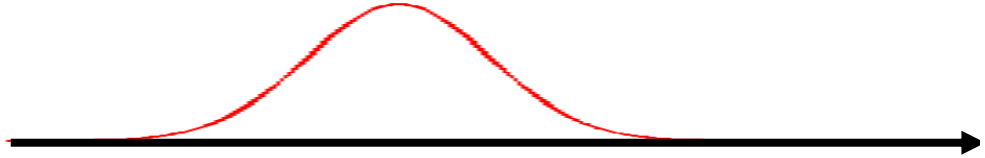

Start with n-mode Gaussian state (e.g. coherent states). Perform linear optics operations (displacement, phase shifting, beam splitting, two-mode squeezing, etc.) and homodyne detection. The state has a succinct (Gaussian) classical description and the quantum computation can be simulated efficiently (Bartlett and Sanders 2003). But add optical nonlinearity, or photon sources and adaptive photon counting measurements, and universal quantum computation is achievable (Knill et al. 2001, Gottesman et al. 2001).

Easy to simulate:
Free fermions

$$H = \frac{i}{4}\sum_{j,k} A_{jk}(t) c_j c_k, \quad \{c_j, c_k\} = 2I\delta_{jk}, \quad c_j = c_j^\dagger$$

Each mode is either empty or occupied in the initial state, and the occupation number is measured in the final readout. In this case, adaptive measurements of the fermion number do not add power, but if we add four-fermion operators to the Hamiltonian, or if we can measure a four-fermion operator nondestructively, universal quantum computation is achievable (Bravyi and Kitaev 2000). (Cf., Majorana fermions as nonabelian anyons in topological insulator / superconductor or superconducting quantum wire.)

Hard to simulate(?): Commuting gates

Qubits are prepared and measured in the X basis. Quantum gates in between are all diagonal in the Z basis. The processing can be done in one time step, by pulsing on and off a diagonal local Hamiltonian ("instantaneous" quantum computing).

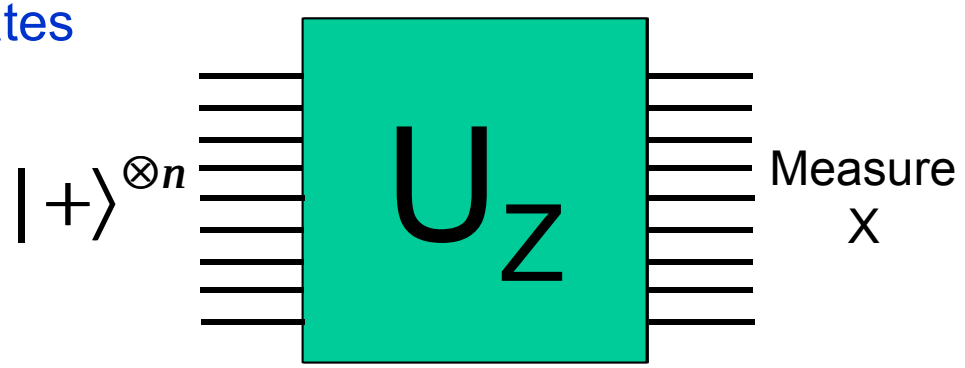

It is not obvious how to simulate this simple quantum circuit classically. If the simulation is possible (in a rather strong sense --- sampling the probability distribution of outcomes in a multiplicative approximation), there would be surprising implications for classical complexity theory: collapse of the polynomial hierarchy to the third level (Bremner, Jozsa, Shepherd 2010).

This model does not seem to have the full power of universal quantum (or even classical) computing, yet may achieve a task beyond the reach of the classical world.

Diagonal gates may be relatively easy to achieve (adiabatically). How robust is the model's "power" against noise?

Hard to simulate(?): Linear optics and (nonadaptive) photon counting

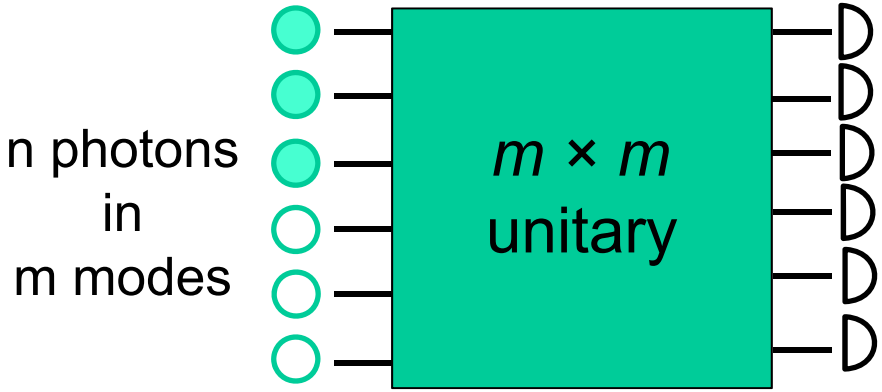

Initially, n photons are prepared in m modes (e.g. the first n modes are occupied by one photon each). A linear optics network executes a unitary on the modes, then the number of photons is counted in each mode (probably 0 or 1 in each mode if $m = \text{constant} \times n^2$.

It is not obvious how to simulate this "simple" quantum experiment classically. If the simulation is possible (with the simulated distribution close to the ideal one) and a plausible conjecture is true, there would be surprising implications for classical complexity theory: collapse of the polynomial hierarchy to the third level (Aaronson and Arkhipov 2010).

This is a theorist's version of the "Hong-Ou-Mandel dip". The classical simulation would be at the limits of current technology for, say, 30 photons. The experiment requires a many-photon coincidence, so photon paths can interfere.

How robust is the model's "power" against noise? (Loss, imperfect photon sources and detectors, …) Fault-tolerant linear optics?

Quantum error correction

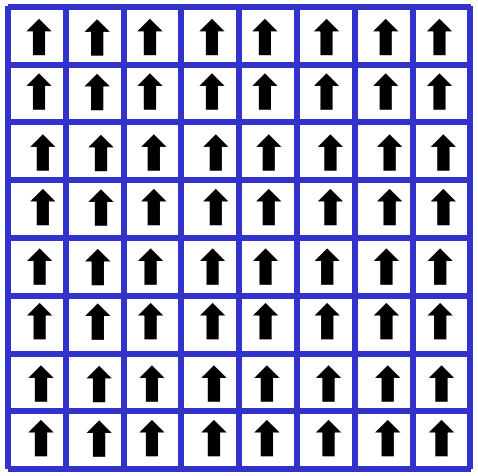

Classical memory ⇔ ferromagnet order

Readout by local measurement (and majority vote).
Errors produce domains walls where spins misalign.
Logical error: domain wall sweeps across sample.
Robust at sufficiently low nonzero temperature; storage time increases exponentially with system size.
Especially simple type of redundant classical storage.

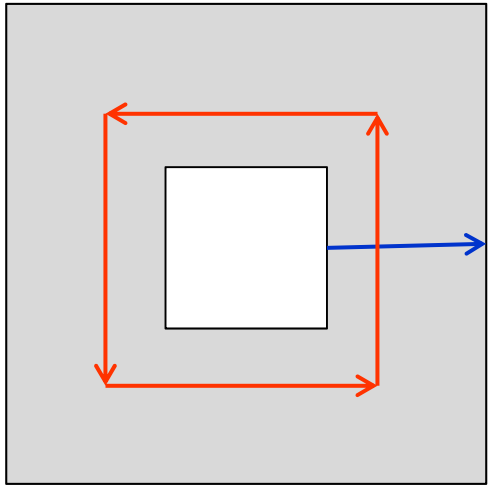

Quantum memory ⇔ topological order

Readout: nonlocal observables X and Z.
Errors produce pointlike excitations (red and blue quasiparticles), with Z_2 relative Aharonov-Bohm phase.
Logical error: X: blue particle escapes from hole. Z: red particle circumnavigates hole.
Protected by energy gap, but storage time does not increase with system size.
Robust if quasiparticles continually monitored.
Especially simple type of redundant quantum storage.

Scalability

Quantum Accuracy Threshold: We can simulate an ideal quantum circuit accurately using a noisy circuit, with a "reasonable" overhead cost [polylog(circuit size)], if the "noise strength" ε is below a critical value ε_0.

Requires: parallel operations, refreshed qubits.

Threshold value ε_0 and overhead cost depends on the scheme and noise model.

In a 2D layout with local gates, it is natural to use topological codes on a punctured plane, with qubits encoded using Z_2 "electric" (or "magnetic") charges placed in the holes.

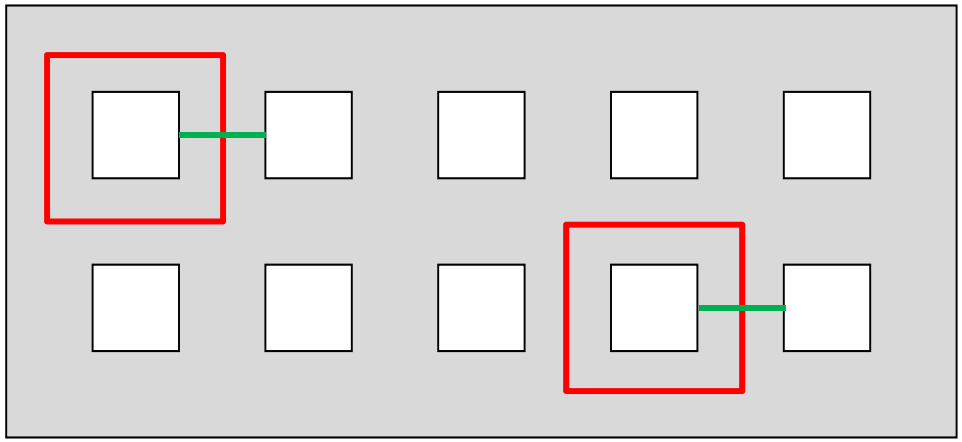

Local gates and independent depolarizing noise: $\varepsilon_0 \sim 7.5 \times 10^{-3}$ Raussendorf et al. Dennis et al.

"Practical" considerations:
Resource requirements, systems engineering issues

Matters of "principle":
Conditions on the noise model, what schemes are scalable, etc.

Non-Markovian noise

Terhal, Burkard 2005; Aliferis, Gottesman, Preskill 2006; Aharonov, Kitaev, Preskill 2006; Ng, Preskill 2009

From a physics perspective, it is natural to formulate the noise model in terms of a Hamiltonian that couples the system to the environment.

$$H = H_{System} + H_{Bath} + H_{System-Bath}$$

where

$$H_{System-Bath} = \sum_{\substack{\text{terms } a \text{ acting} \\ \text{locally on the system}}} H^{(a)}_{System-Bath}$$

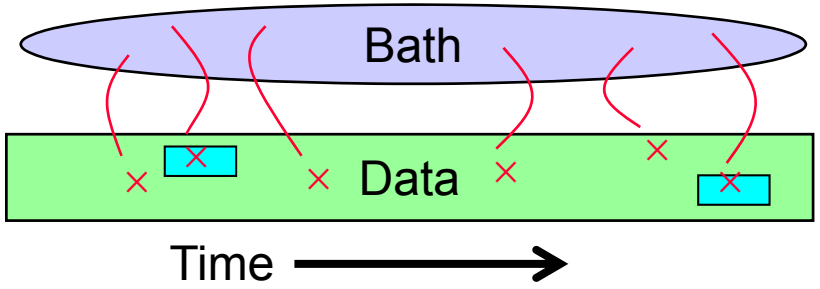

Threshold condition can be formulated as $\varepsilon \le \varepsilon_0 \cong 10^{-4}$, where *noise strength* ε can be defined in either of two ways:

$$\varepsilon = \max \left\| H^{(a)}_{System-Bath} \right\| t_0$$

over all times and locations

gate execution time

Internal bath dynamics can be strong and nonlocal

$$\varepsilon = \max \left(\int_{1,\ circuit\ location} \int_{2,\ all\ spacetime} \left| \Delta_{Bath}(1,2) \right| \right)^{1/2}$$

bath correlation function

applies for a Gaussian (harmonic oscillator) bath

In either scenario, noise Hamiltonian is assumed to act locally on the system

Noise correlations

In general, the noise Hamiltonian may contain terms acting on m system qubits, for m = 1, 2, 3, ….

$$H_{System-Bath} = \sum_i H_i + \sum_{\langle ij \rangle} H_{ij} + \sum_{\langle ijk \rangle} H_{ijk} + \cdots$$

Quantum computing is provably scalable if $\varepsilon \le \varepsilon_0 \cong 10^{-4}$, where

$$\varepsilon = \max_m \eta_m^{1/m}$$ and

$$\eta_m = \max \sum_{j_2, j_3, \ldots, j_m} \| H_{j_1 j_2 j_3 \ldots j_m} \| + \eta_{m+1}$$

over all times and qubits

interactions fall off with distance

term that acts collectively on m system qubits should be exponentially small in m.

Proofs of the threshold theorem require the noise to be “quasi-local” in the sense that the m-qubit noise term in the Hamiltonian decays exponentially with m. Can experiments verify this scaling?

Quantum computing with nonabelian anyons

Kitaev, Freedman et al.

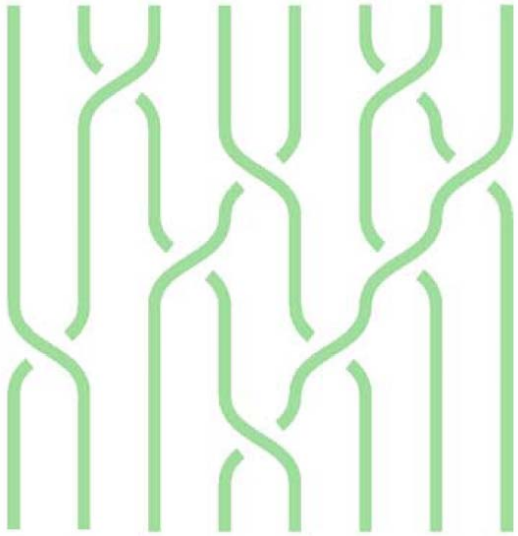

-- Quantum information stored in the exponentially large fusion Hilbert space of n nonabelian anyons.

-- Protected when temperature is small compared to the energy gap (no thermal anyons) and particles are well separated (no tunneling).

-- Read out by measuring charges of anyon pairs.

-- Processing by exchanging (braiding) particles.

Possible realizations:

Fractional quantum Hall states with filling factor = 5/2.

Majorana fermions at ends of quantum wires (topological superconductors), or in topological insulator / superconductor heterostructures.

For the above, braiding is not computationally universal but universality achievable by combining with noisy non-topological operations.

Universal braiding might be realized in other FQH states.

Substantial overhead cost for approximately realizing standard quantum gates as braids when anyons are universal.

Combine with “standard” quantum error correction for scalability (storage time for anyon system does not improve with system size).

Quantum computing vs. quantum simulation

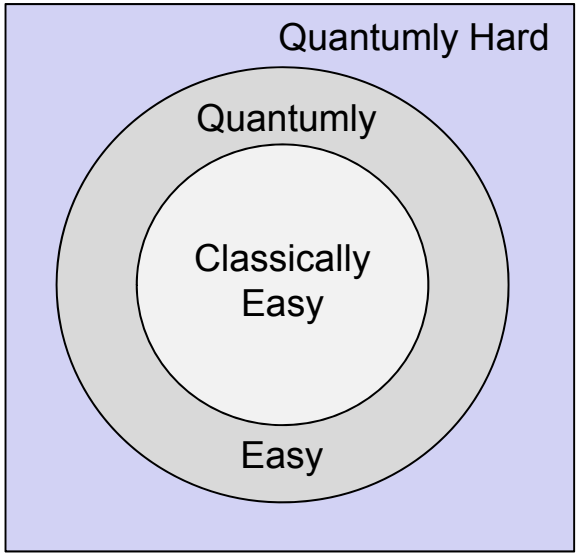

Many of the most challenging problems in physical science concern highly entangled ("strongly correlated") quantum systems: for example, quantum antiferromagnets, exotic superconductors, complex biomolecules, bulk nuclear matter, spacetime near singularities, etc.

A reliable universal quantum computer can simulate efficiently any "reasonable" physical system, while quantum simulators may have intrinsic limitations.

Using either method, the goal should be to learn about quantum phenomena that are hard to simulate classically. We hope to discover previously unsuspected phenomena, not just validate theoretical predictions and models.

Classical hardness may hinge on the *accuracy* of the simulation. Universal quantum computers can be made fault-tolerant, though with a daunting resource blowup.

Novel properties of interest may be robust and universal, hence accessible through crude simulations.

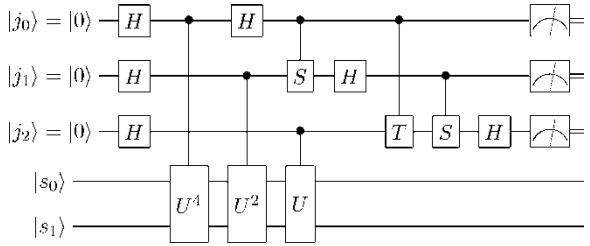

VS.

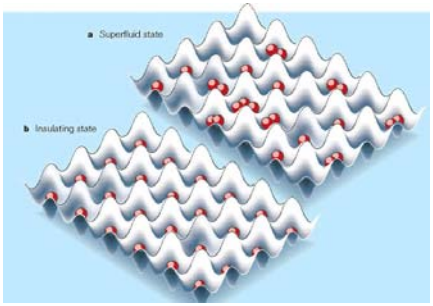

Questions for discussion:

Quantum supremacy: super-classical behavior of controllable quantum systems

Do we *already* know that Nature performs tasks that cannot be simulated efficiently by classical computers (strongly correlated materials, complex molecules, etc.)?

Is quantum simulation (with atoms, molecules, etc.) a feasible path to quantum supremacy, despite difficulties in controlling such systems precisely?

How best to achieve quantum supremacy with relatively small and accessible systems (~100 physical qubits)? Simulation tasks? Short of universal quantum computing?

Can quantum computers simulate quantum gravity? (Yes/No are both interesting.)

Quantum fault tolerance: scalable protection against decoherence

Do e.g. FQH systems already demonstrate that large-scale quantum error-correcting codes are realizable. Will “anyon interferometry” make this case persuasive?

Scalability: anyons vs “standard” qubits protected by quantum codes.

What experiments studying noise in quantum systems will strengthen the case that scalable fault-tolerant quantum computing is feasible?

Is an intrinsically stable quantum memory possible, e.g. in a three dimensional system?

Is nature fault tolerant (does unitary quantum mechanics emerge at long distances)?

Additional Slides

Hard to simulate(?): One clean qubit model (Knill and Laflamme, 1998)

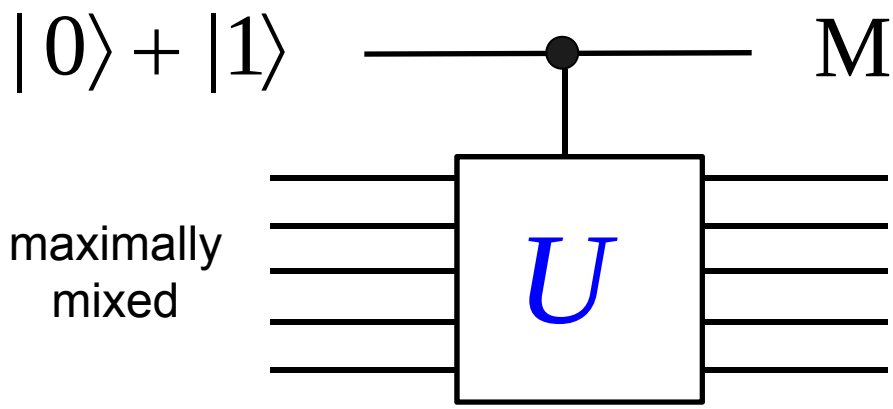

$$p_0 = \frac{1}{2}\left(1 + \frac{1}{d}\,\mathrm{Re}\left(\mathrm{tr}\,U\right)\right)$$

Approximates the Jones polynomial for the trace closure of a braid (Shor and Jordan 2008), which is a complete problem for this complexity class.

Hard(?): Permutational model (Jordan 2009)

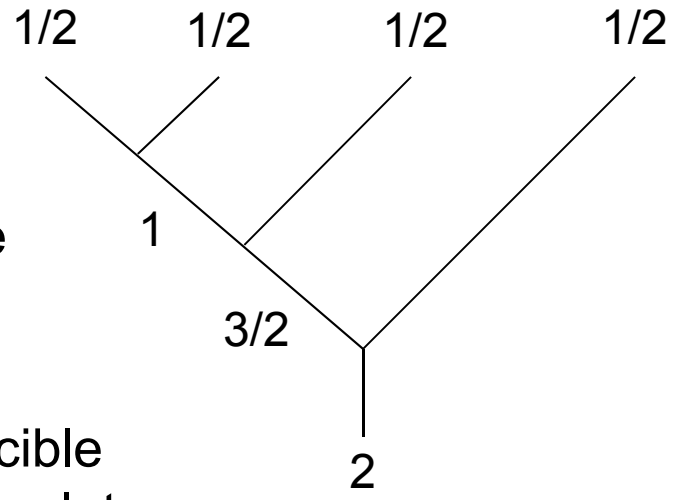

Prepare initial state of spin-(1/2) particles with specified fusing of angular momentum, permute particles, then measure a commuting set of total angular momenta of subsets of particles.

Approximates matrix elements of certain irreducible representations of the symmetric group, and simulates the Ponzano-Regge spin foam toy model of quantum gravity.

Quantumly easy: *measuring* the energy for a local Hamiltonian

(polynomial accuracy)

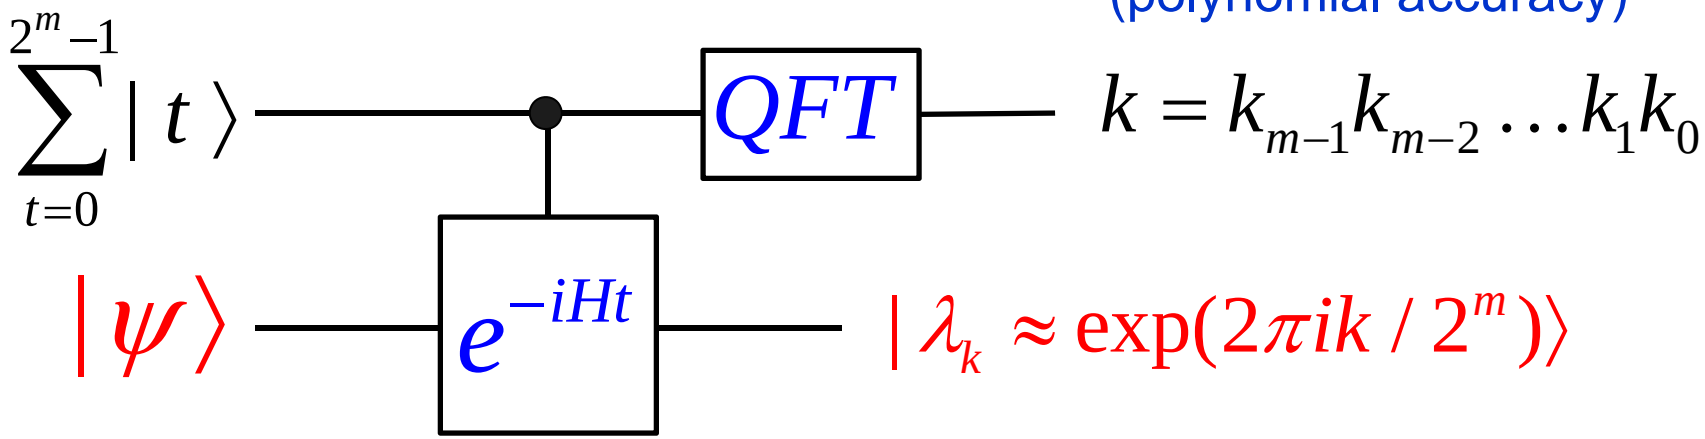

For a local Hamiltonian, measuring the energy to accuracy 1/poly(n) is quantumly easy. We could measure the ground state energy *if* we could prepare a state whose overlap with the ground state is only polynomially small.

But … the state preparation problem seems to be hard in general (Kitaev 2002). This is the quantum version of the P/NP conjecture --- it is easy to *verify* a solution to a constraint satisfaction problem, but hard to find the solution.

A good trick for preparing ground states is adiabatic evolution. But adiabatic state preparation fails in hard cases because the at some point during the evolution the energy gap becomes superpolynomially small.

Quantum error correction and topological order

A “logical qubit” is encoded using many “physical qubits.” We want to protect the logical qubit, with orthonomal basis states $|0\rangle$ and $|1\rangle$, from a set of possible error operators $\{ E_a \}$.

For protection against bit flips:
$E_a |0\rangle \perp E_b |1\rangle$.

For protection against phase errors:
$E_a (|0\rangle + |1\rangle) \perp E_b (|0\rangle - |1\rangle)$.

In fact, these conditions suffice to ensure the existence of a recovery map.

It follows that
$\langle 0| E_b^\dagger E_a |0\rangle = \langle 1| E_b^\dagger E_a |1\rangle$.

Compare the definition of topological order: if V is a (quasi-)local operator and $|0\rangle$, $|1\rangle$ are ground states of a local Hamiltonian, then $\langle 1| V |0\rangle = 0$, and $\langle 0| V |0\rangle = \langle 1| V |1\rangle$.
up to corrections exponentially small in the system size. (Ground states are locally indistinguishable.)

Scalability

Quantum Accuracy Threshold Theorem: Consider a quantum computer subject to quasi-independent noise with strength ε. There exists a constant $\varepsilon_0 > 0$ such that for a fixed $\varepsilon < \varepsilon_0$ and fixed $\delta > 0$, any circuit of size L can be simulated by a circuit of size L^* with accuracy greater than $1-\delta$, where, for some constant c,

$$L^* = O\left[L (\log L)^c \right]$$

Aharonov, Ben-Or
Kitaev
Laflamme, Knill, Zurek
Aliferis, Gottesman, Preskill
Reichardt

assuming:

parallelism, fresh qubits (*necessary* assumptions)

nonlocal gates, fast measurements, fast and accurate classical processing, no leakage (*convenient* assumptions).

Local gates and independent depolarizing noise: $\varepsilon_0 \sim 7.5 \times 10^{-3}$

"Practical" considerations:
Resource requirements, systems engineering issues

Matters of "principle":
Conditions on the noise model, what schemes are scalable, etc.

Local fault tolerance with 2D topological codes

Qubits are arranged on a two-dimensional lattice with holes in it. Protected qubits are encoded (in either of two complentary bases) by placing “electric” charges inside “primal” holes or “magnetic” charges inside “dual” holes. The quantum information is well protected if the holes are large and far apart.

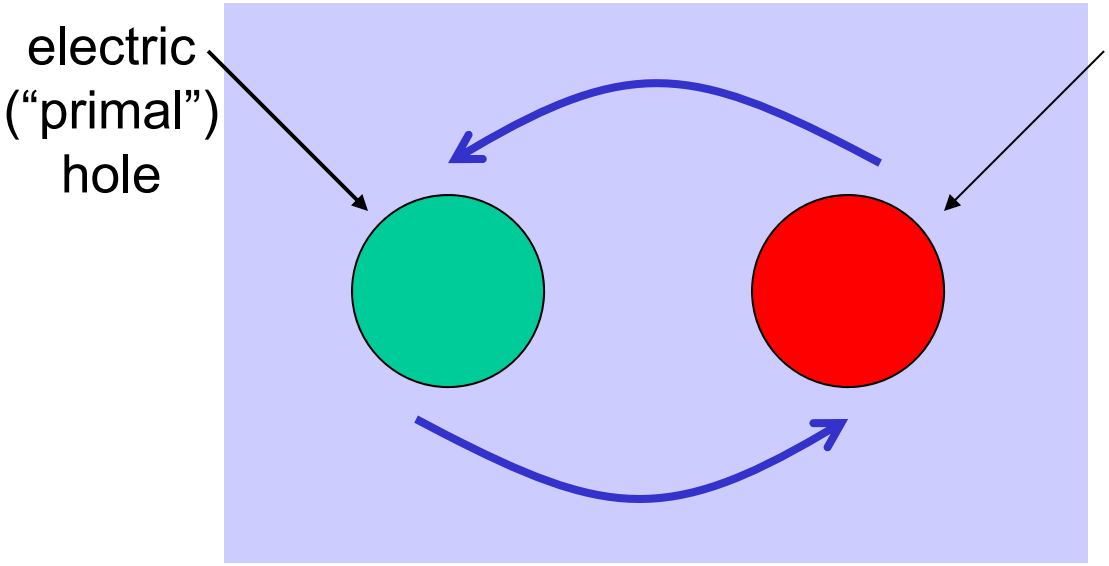

A controlled-NOT gate can be executed by “braiding the holes” which is achieved by a sequence of local gates or measurements.

The protected gates and error syndrome extraction can be done with *local* two-qubit gates or measurements. Numerical studies indicate an upper bound on the threshold for independent depolarizing noise:

$\varepsilon_0 \sim 7.5 \times 10^{-3}$

Raussendorf, Harrington, Goyal
Dennis, Kitaev, Landahl, Preskill

Hardware

-- Robust devices (e.g. “0-Pi” superconducting qubit).

-- Topological protection and processing (e.g. Majorana fermions in quantum wires).

Software

-- Optimized threshold and overhead.

-- Adapting fault tolerance to noise.

-- Dynamical decoupling.

Systems engineering (wires, power, cooling, etc.)

Matters of principle

-- Justifying error phase randomization → error probabilities (e.g., relating randomized error benchmarking to fault tolerance requirements).

-- Limitations on noise correlations

-- Self-correcting hardware (favorable scaling of storage time with system size, in fewer than four dimensions?).

-- Other scalable schemes besides concatenated codes and topological codes (perhaps fault-tolerant adiabatic quantum computing?).

Some themes of quantum information science

- Quantum entanglement. Correlations among parts of a quantum system are different than classical correlations. A quantum system with two parts is entangled when its joint state is more definite and less random than the state of each part by itself.
- Quantum cryptography. Unknown quantum states cannot be copied perfectly. Eavesdropping on quantum communication can be detected.
- Quantum computing. Quantum computers can solve problems (like factoring) that are believed to be hard for classical computers.
- Quantum error correction. Quantum systems are highly vulnerable to noise, but we know in principle how to stabilize them.
- Quantum hardware. We can really do this stuff in the physics lab (though so far only on a modest scale).