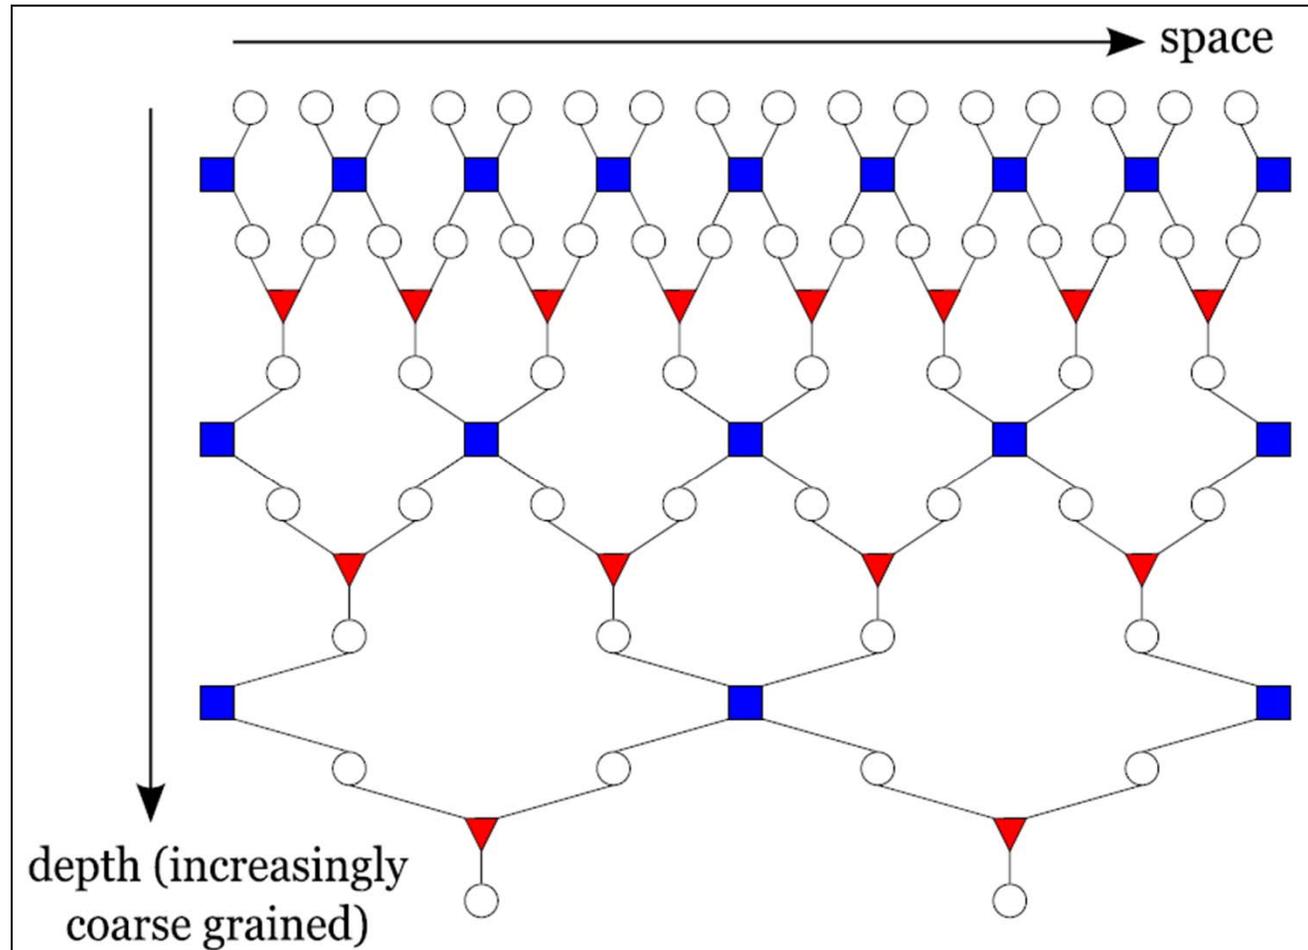


The Entanglement Frontier



 Institute for Quantum Information



John Preskill, Caltech
QuEST, 28 March 2011

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 implications?

Are there deeper principles that illuminate the foundations of quantum physics?

Can we?

Will as yet undiscovered principles of fundamental physics prevent large-scale quantum computers from ever working?

Will fault-tolerant quantum computing work against all noise mechanisms in realizable quantum processors?

Should we?

Can we find new quantum algorithms that achieve exponential speedups relative to classical computers?

What quantum many-body systems of interest in physics and chemistry can be simulated using classical computers? Using quantum simulators? Using large-scale quantum computers?

How?

For example: Can exotic nonabelian topological phases of quantum matter be realized and can the nonabelian statistics of the quasiparticles be confirmed?

And ...?

Can the strength of nonlocal correlations allowed in quantum theory be derived from deeper physical or mathematical principles?

For example: What are the advantages in principle of entangled and other non-classical states for metrology, and can these produce practical gains?

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).

To push against the “entanglement frontier” we must slay the dragon of decoherence ...

Is this merely *really, really hard*?

Or is it *ridiculously hard*?

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.

The “natural” application for quantum computers (Feynman 1982).

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, and what quantum tasks are hard to simulate classically?**

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 resources. Only these can arise in Nature, and only these are within reach of the quantum engineers.

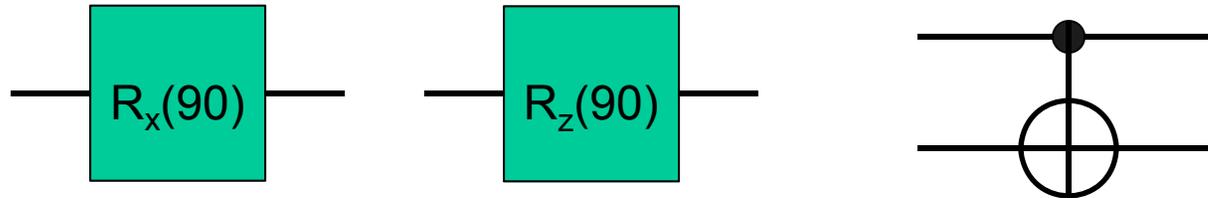
A mathematical model of the feasible n -qubit pure states: they can be prepared by a circuit of $\text{poly}(n)$ two-qubit gates, applied to an initial (unentangled) product state. (A fraction $\exp[-\text{poly}(n)]$ of all pure states.) Likewise, feasible measurements are $\text{poly}(n)$ -size circuits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

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.

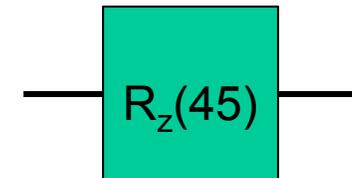
Easy to simulate: Clifford group

We can rotate each qubit on the Bloch sphere by 90° about the x, y, or z axes, and we can perform a CNOT gate on any pair of qubits.



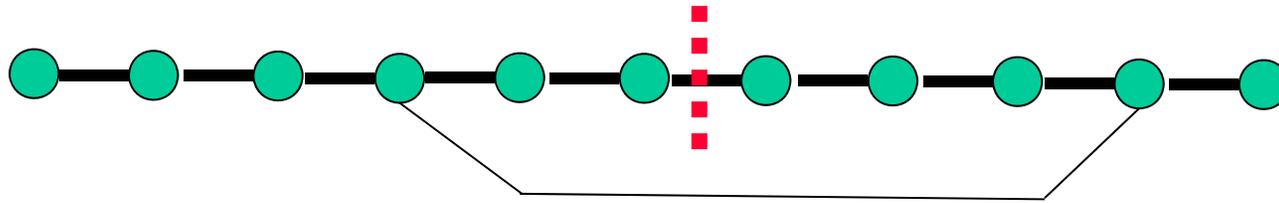
Action of these gates can be described succinctly in the Heisenberg picture: Pauli operators \rightarrow Pauli operators (Gottesman-Knill 1998); therefore circuit is easy to simulate if initial state is a Z eigenstate and we measure in the Z basis.

But ... if we add a rotation about the z axis by 45° ,
Then universal quantum computation is achievable.



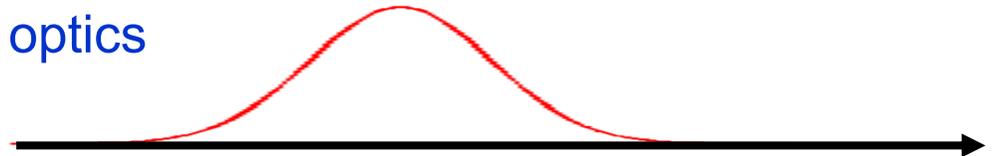
This example is central in the theory of fault tolerance, since encoded Clifford group transformations are relatively simple to realize, but encoded operations beyond the Clifford group are realized by the more complicated gate teleportation trick (Shor 1996, Gottesman and Chuang 1999).

Easy to simulate: “weakly entangled” 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).**

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: “matchgates” = free fermions

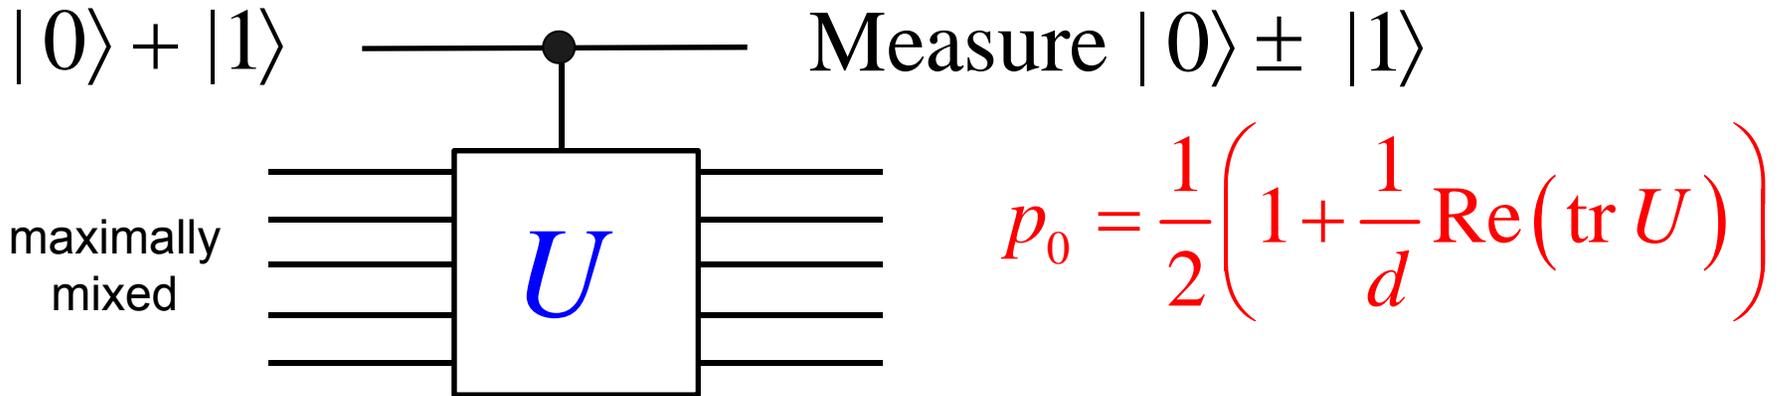
$$V : \{|00\rangle, |11\rangle\}, \quad W : \{|01\rangle, |10\rangle\}, \quad \det V = \det W$$

Arrange n qubits on a line. Nearest neighbor two-qubit gates preserve parity, and transformations on the even and odd parity blocks have the same determinant (Valiant 2001). After a Jordan-Wigner transformation, this is equivalent to a system of free fermions, in which fermionic parity is preserved (Terhal and DiVincenzo 2002).

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

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

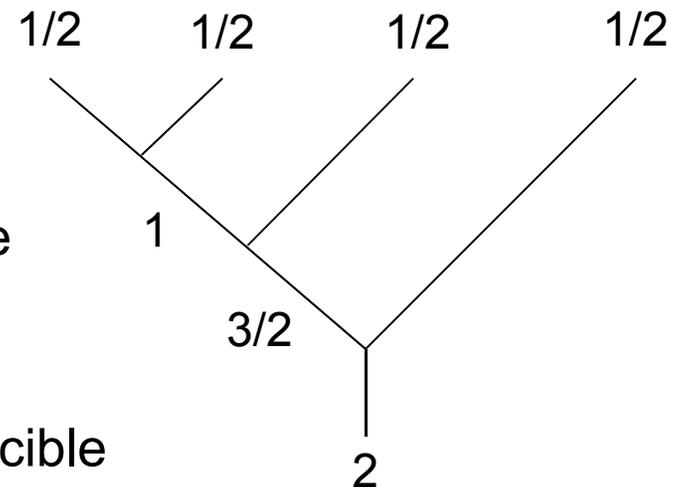


Approximates the Jones polynomial for the trace closure of a braid (Shor and Jordan 2008), which is a complete problem for this complexity class.

Hard(?): Permutational model (Jordan 2009)

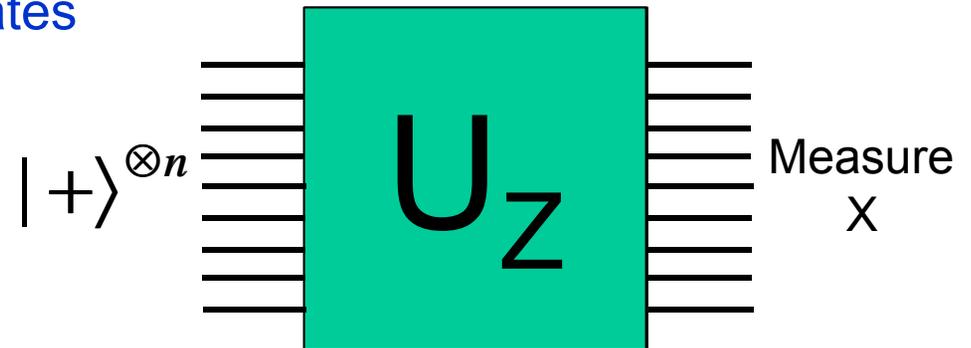
Prepare initial state of spin-(1/2) particles with specified fusing of angular momentum, permute particles, then measure a commuting set of total angular momenta of subsets of particles.

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



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).

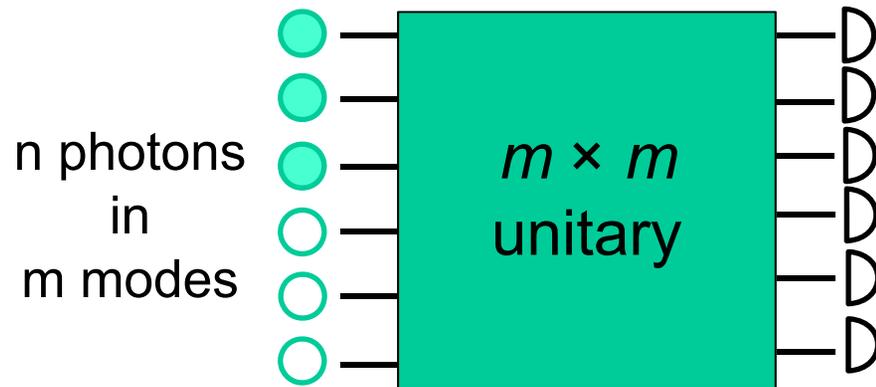


It is not obvious how to simulate this simple quantum circuit classically. If the simulation is possible (in a rather strong sense --- sampling the probability distribution of outcomes in a multiplicative approximation), there would be surprising implications for classical complexity theory: collapse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



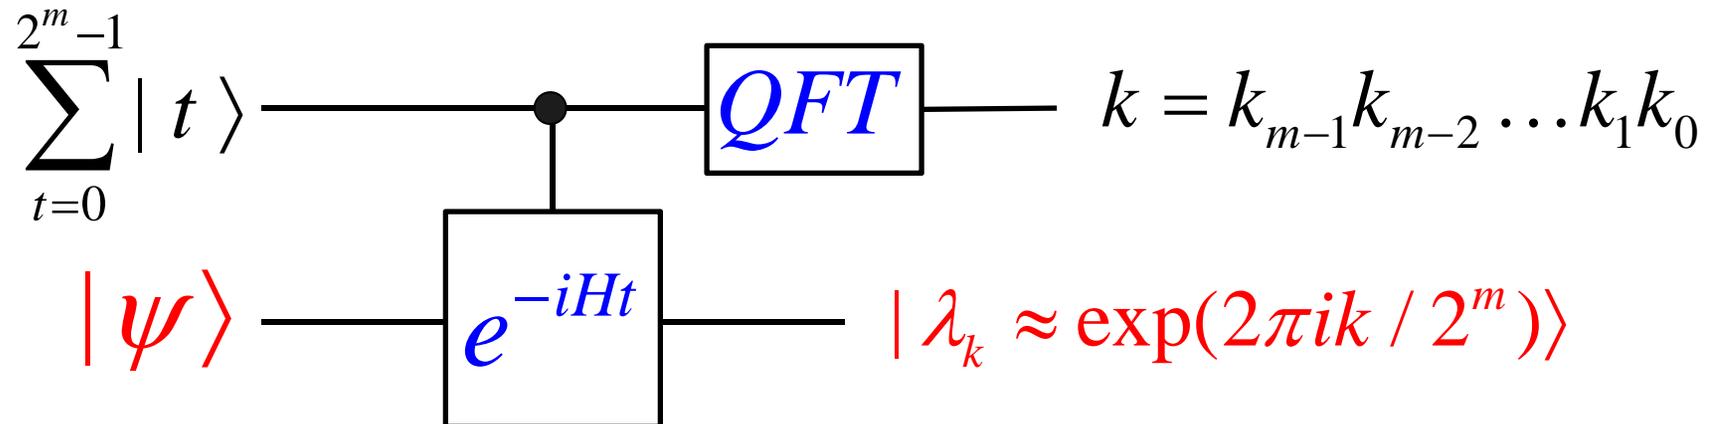
Initially, n photons are prepared in m modes (e.g. the first n modes are occupied by one photon each). A linear optics network executes a unitary on the modes, then the number of photons is counted in each mode (probably 0 or 1 in each mode if $m = \text{constant} \times n^2$).

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?

Quantumly easy: *checking* the ground state energy (polynomial accuracy)



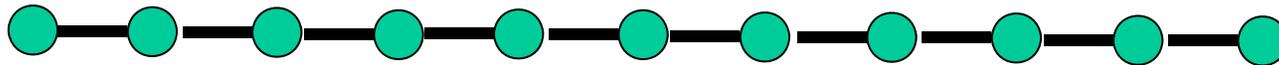
For a local Hamiltonian, measuring the energy to accuracy $1/\text{poly}(n)$ is quantumly easy. We could measure the ground state energy *if we could* prepare a state whose overlap with the ground state is only polynomially small.

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.

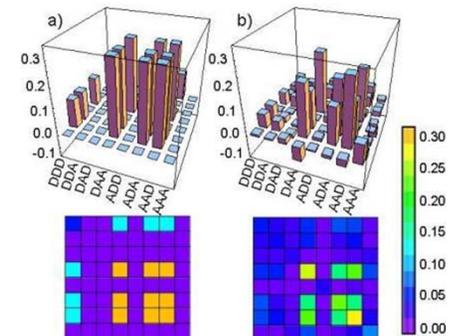
Is Tomography hard?

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 ground state of a gapped local 1D Hamiltonian. 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.

Suppose the state is **close to pure**, i.e. is well approximated by a state with rank $r \ll 2^n$. Then $r n^2 2^n$ measured observables (Pauli operators) suffice, rather than 4^n (Gross et al. 2009). No a priori assumption is needed to certify that the state really is close to pure. The method extends “**compressed sensing**” (reconstructing a sparse vector from a small number of components) to **matrix completion** (reconstructing a sparse matrix from a small number of matrix elements).



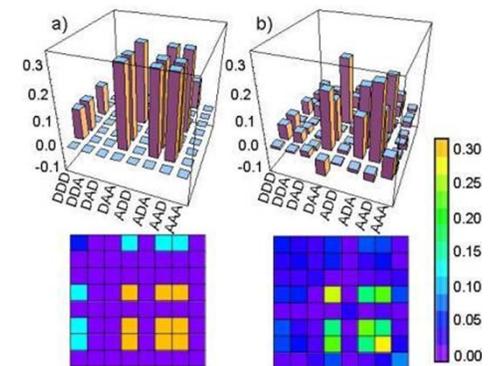
Is Tomography hard?

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.

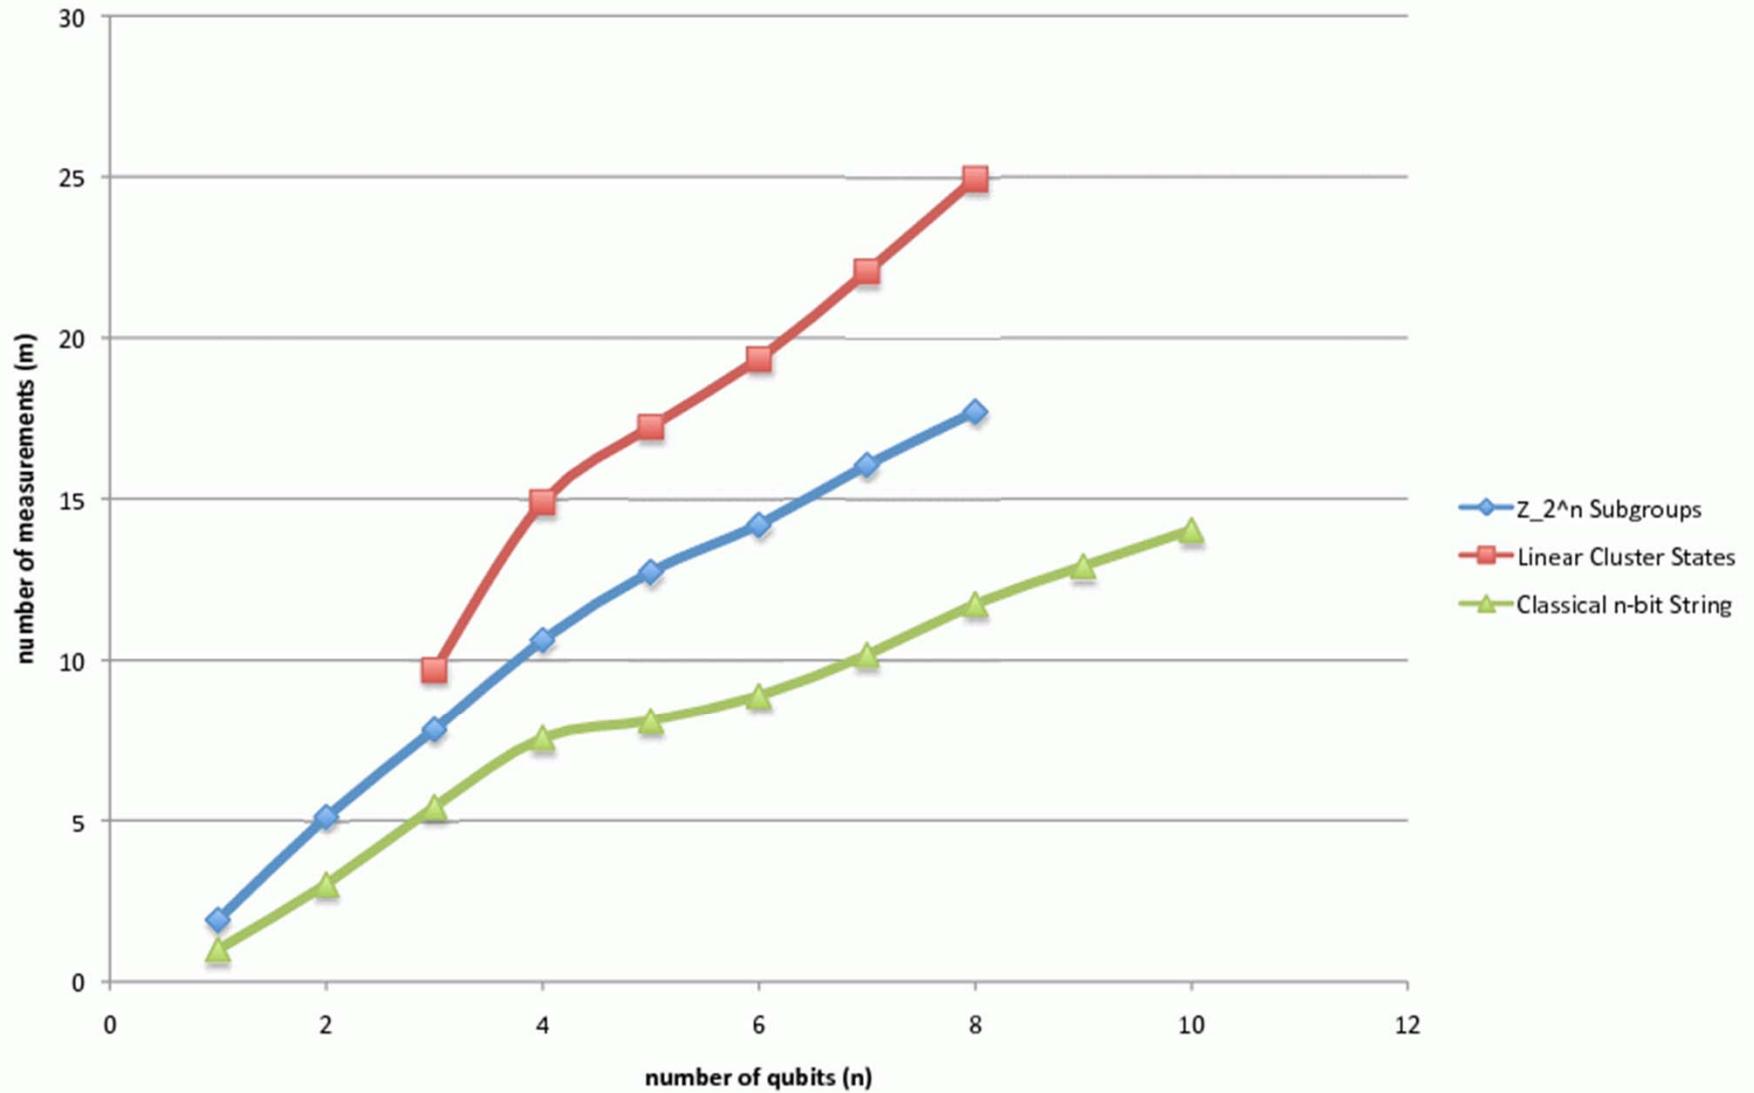
But ... we are not really interested in the expectation values of *arbitrary* observables, only the ones we are likely to want to measure (most measurements are infeasible anyway). And ... Holevo's theorem and its extensions tell us that we can extract only one bit of information per qubit measured --- if quantum states have so limited a capacity to surprise, then *why should it be hard to learn a quantum state?*

Answer: it's not!

Suppose ρ is an n -qubit quantum state, and D is a distribution on two-outcome measurements. We draw m measurements from D , and find any σ that agrees well with ρ regarding the outcome probability for these m measurements. Then *for $m=O(n)$ σ predicts accurately the outcomes of future measurements drawn from D and performed on ρ .* (Aaronson 2006). It works even though σ may be far from ρ in the trace norm. However ... *finding* σ may be a hard classical computation (exponential in n).



Measurement Complexity of n



Aaronson and Dechter 2009

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*).

“Practical” considerations:

Resource requirements, systems engineering issues

Matters of “principle”:

Conditions on the noise model

Hardware

- Robust devices (e.g. “0- Π ” 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 \rightarrow 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?).

Non-Markovian noise

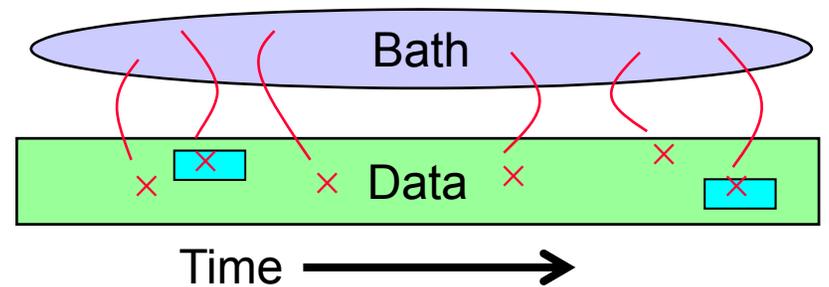
Terhal, Burkard 2005; Aliferis, Gottesman, Preskill 2006; Aharonov, Kitaev, Preskill 2006; Ng, Preskill 2009

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

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



Threshold condition can be formulated as $\epsilon \leq \epsilon_0 \cong 10^{-4}$, where *noise strength* ϵ can be defined in either of two ways:

$$\epsilon = \max \left\| H_{System-Bath}^{(a)} \right\| t_0$$

over all times and locations

gate execution time

Internal bath dynamics can be strong and nonlocal

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

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 $m = 1, 2, 3, \dots$

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

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

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

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

over all times
and qubits

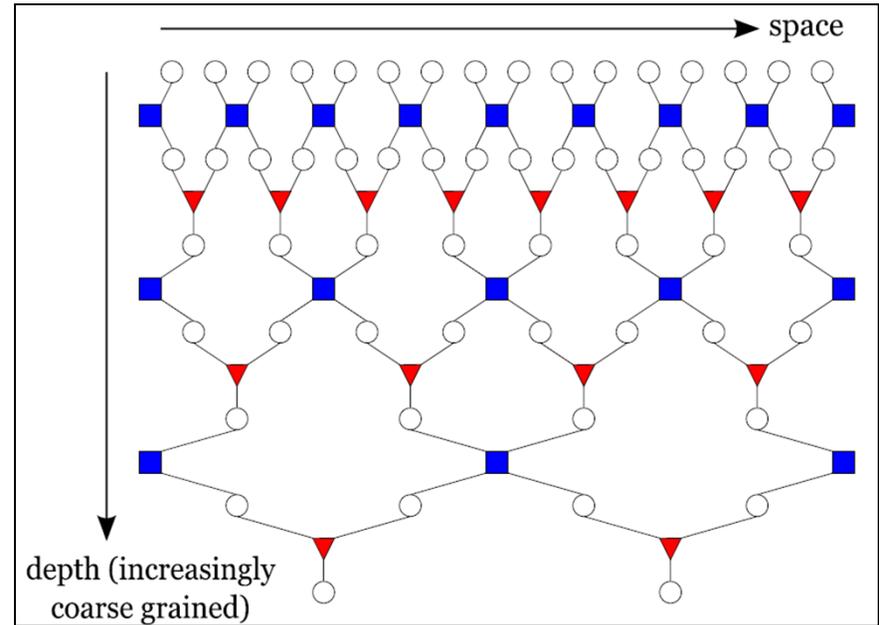
interactions fall
off with distance

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

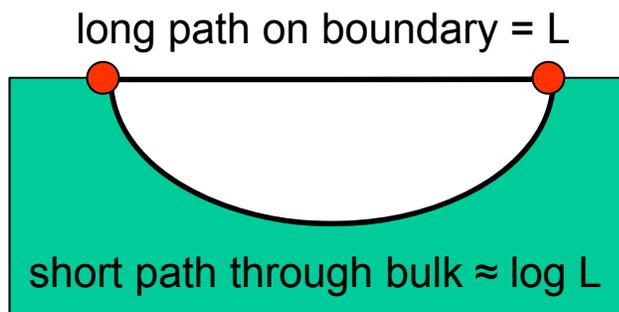
Proofs of the threshold theorem require the noise to be “quasi-local” in the sense that the m -qubit noise term in the Hamiltonian decays exponentially with m . [Can experiments verify this scaling?](#)

Geometry = Entanglement

A succinct description of a quantum state is the description of the quantum circuit that prepares the state. Run backwards, the circuit can be viewed as a coarse graining procedure, or renormalization-group transformation. How entangled a region is with its complement is determined by the structure of this circuit. (Vidal 2007).



We can view the depth of the circuit as an “emergent” spatial dimension. For a scale invariant system, the geometry is hyperbolic, and the entanglement of a region on the “boundary” is the length of a bounding curve that descends into the “bulk” (Ryu and Takayanagi 2006, Swingle 2009).



The AdS/CFT duality (Maldacena 1998) can be viewed as a prescription for building the vacuum of a quantum field theory with a quantum circuit?

What if the circuit (= emergent geometry) *fluctuates*? What about asymptotically flat space? Volume law not area law? (Li and Takayanagi 2010)

Summary: the entanglement frontier

Goal for QIS: perform quantum experiments that cannot be simulated classically. (“Quantum supremacy”)

The vast size of Hilbert space is in some ways misleading, if we confine our attention to **feasible states and measurements**. (Cf., tomography.)

For a quantum computer, simulating the evolution of the excited states of a local quantum system is feasible, but **finding the ground state might not be**.

On the other hand, if a system finds its ground state (or its thermal Gibbs state) in Nature, then **we should be able to prepare the state** in a simulation.

Some **systems that lack the full power of universal quantum computation may still be hard to simulate classically**, and are therefore “interesting”.

Proofs of scalability require **reasonable scaling of noise correlations**, which should be checked experimentally.

Exploring the entanglement frontier should provide valuable insights into highly correlated quantum many body physics, quantum field theory, quantum gravity, and **the difference between the classical and quantum worlds**.