

Ph/CS 219A

Quantum Computation

Instructor: John Preskill

preskill@caltech.edu

theory.caltech.edu/~preskill

TAs: Thom Bohdanowicz and Charles Xu

Topics covered in Ph/CS 219A include density operators, quantum operations, quantum entanglement, quantum circuits, and quantum algorithms. Ph/CS 219B (Winter term, taught by professor Kitaev) will cover quantum error correction and some special topics.

Primary source: Lecture notes


Course materials are on Google Drive: videos, slides, notes, problems, solutions

Course Website: http://theory.caltech.edu/~preskill/ph219/ph219_2020-21

Homework to be submitted and returned via Canvas.

We have a Piazza site for discussions, comments, questions. Connect to Piazza from Canvas


Quantum Information Science



Planck




Turing




Shannon

quantum theory
+ computer science
+ information theory



quantum
information
science



Quantum Information Science

Quantum sensing

Improving sensitivity and spatial resolution.

Quantum cryptography

Privacy founded on fundamental laws of quantum physics.

Quantum networking

Distributing quantumness around the world.

Quantum simulation

Probes of exotic quantum many-body phenomena.

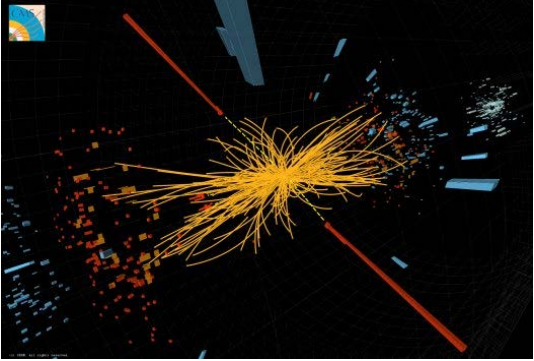

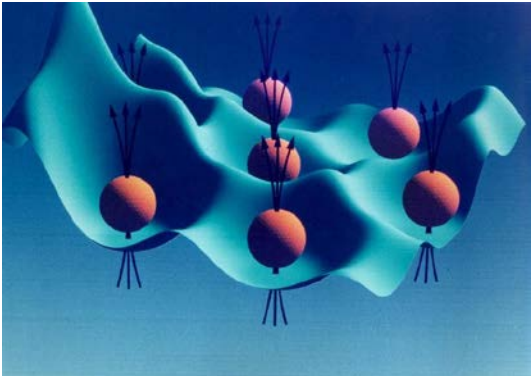
Quantum computing

Speeding up solutions to hard problems.

Quantum information concepts

Entanglement, error correction, complexity, ...

Frontiers of Physics

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

Two fundamental ideas

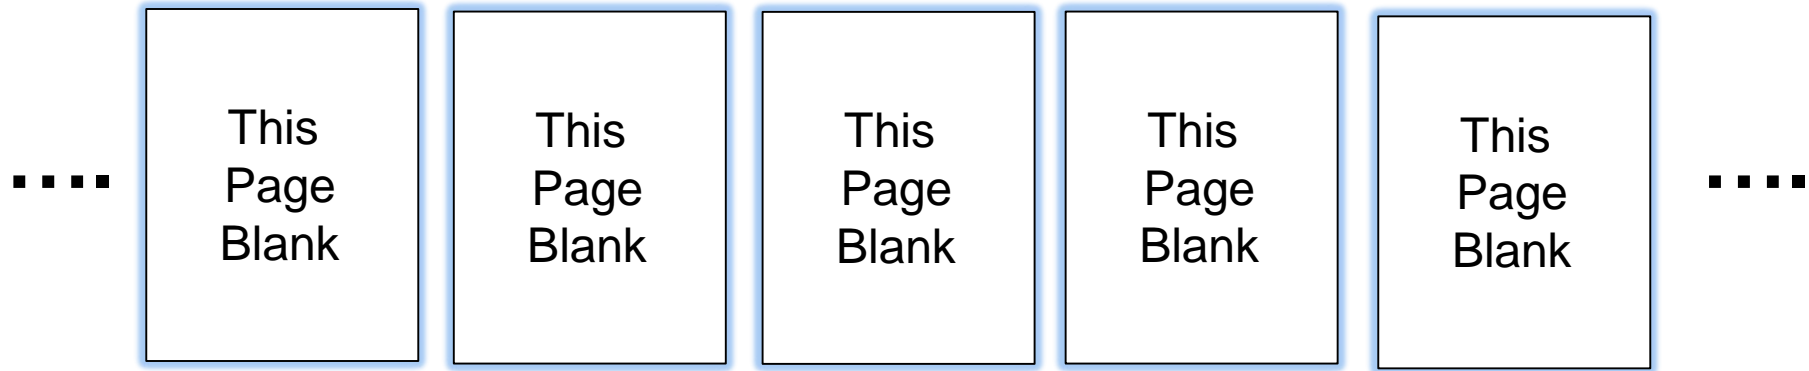
(1) Quantum complexity

Why we think quantum computing is powerful.

(2) Quantum error correction


Why we think quantum computing is scalable.

Quantum entanglement



Nearly all the information in a typical entangled “quantum book” is encoded in the correlations among the “pages”.

You can't access the information if you read the book one page at a time.



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

Why we think quantum computing is powerful

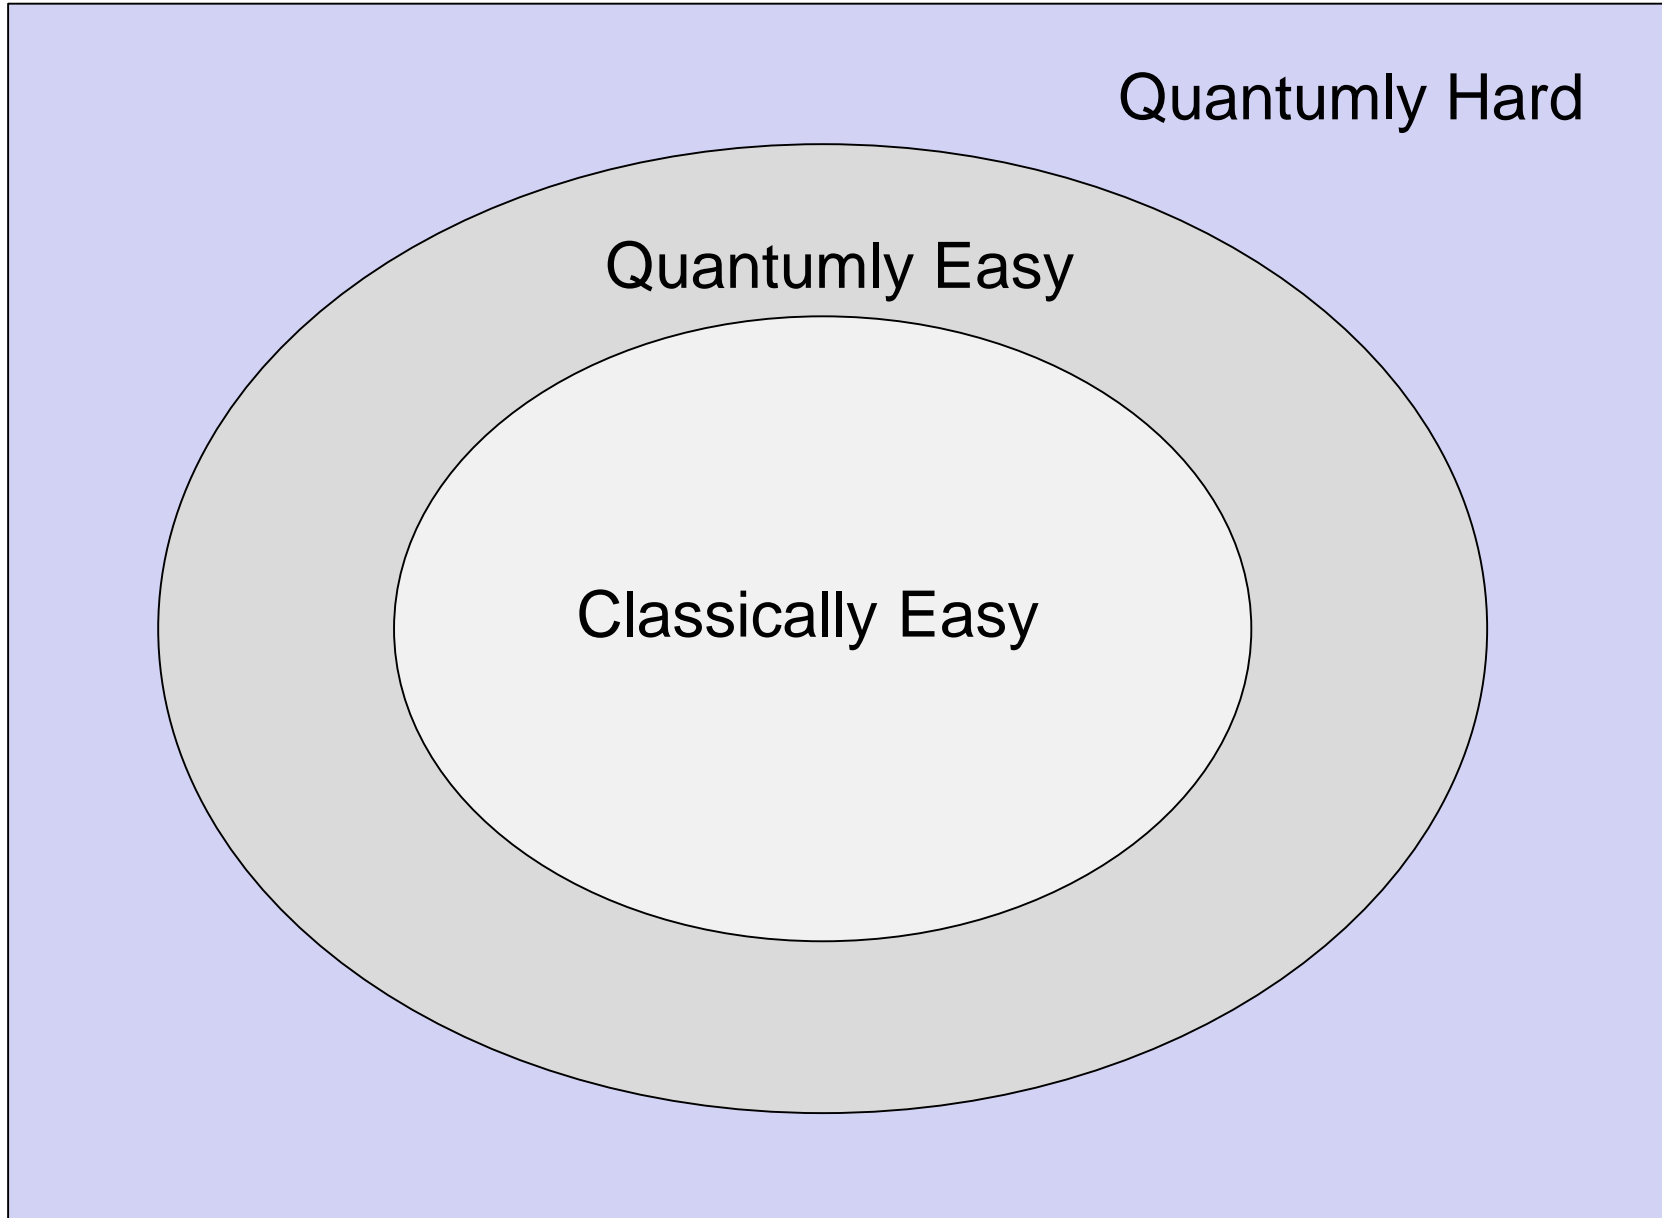
(1) Problems believed to be hard classically, which are easy for quantum computers. **Factoring is the best known example.**

(2) **Complexity theory arguments** indicating that quantum computers are hard to simulate classically.

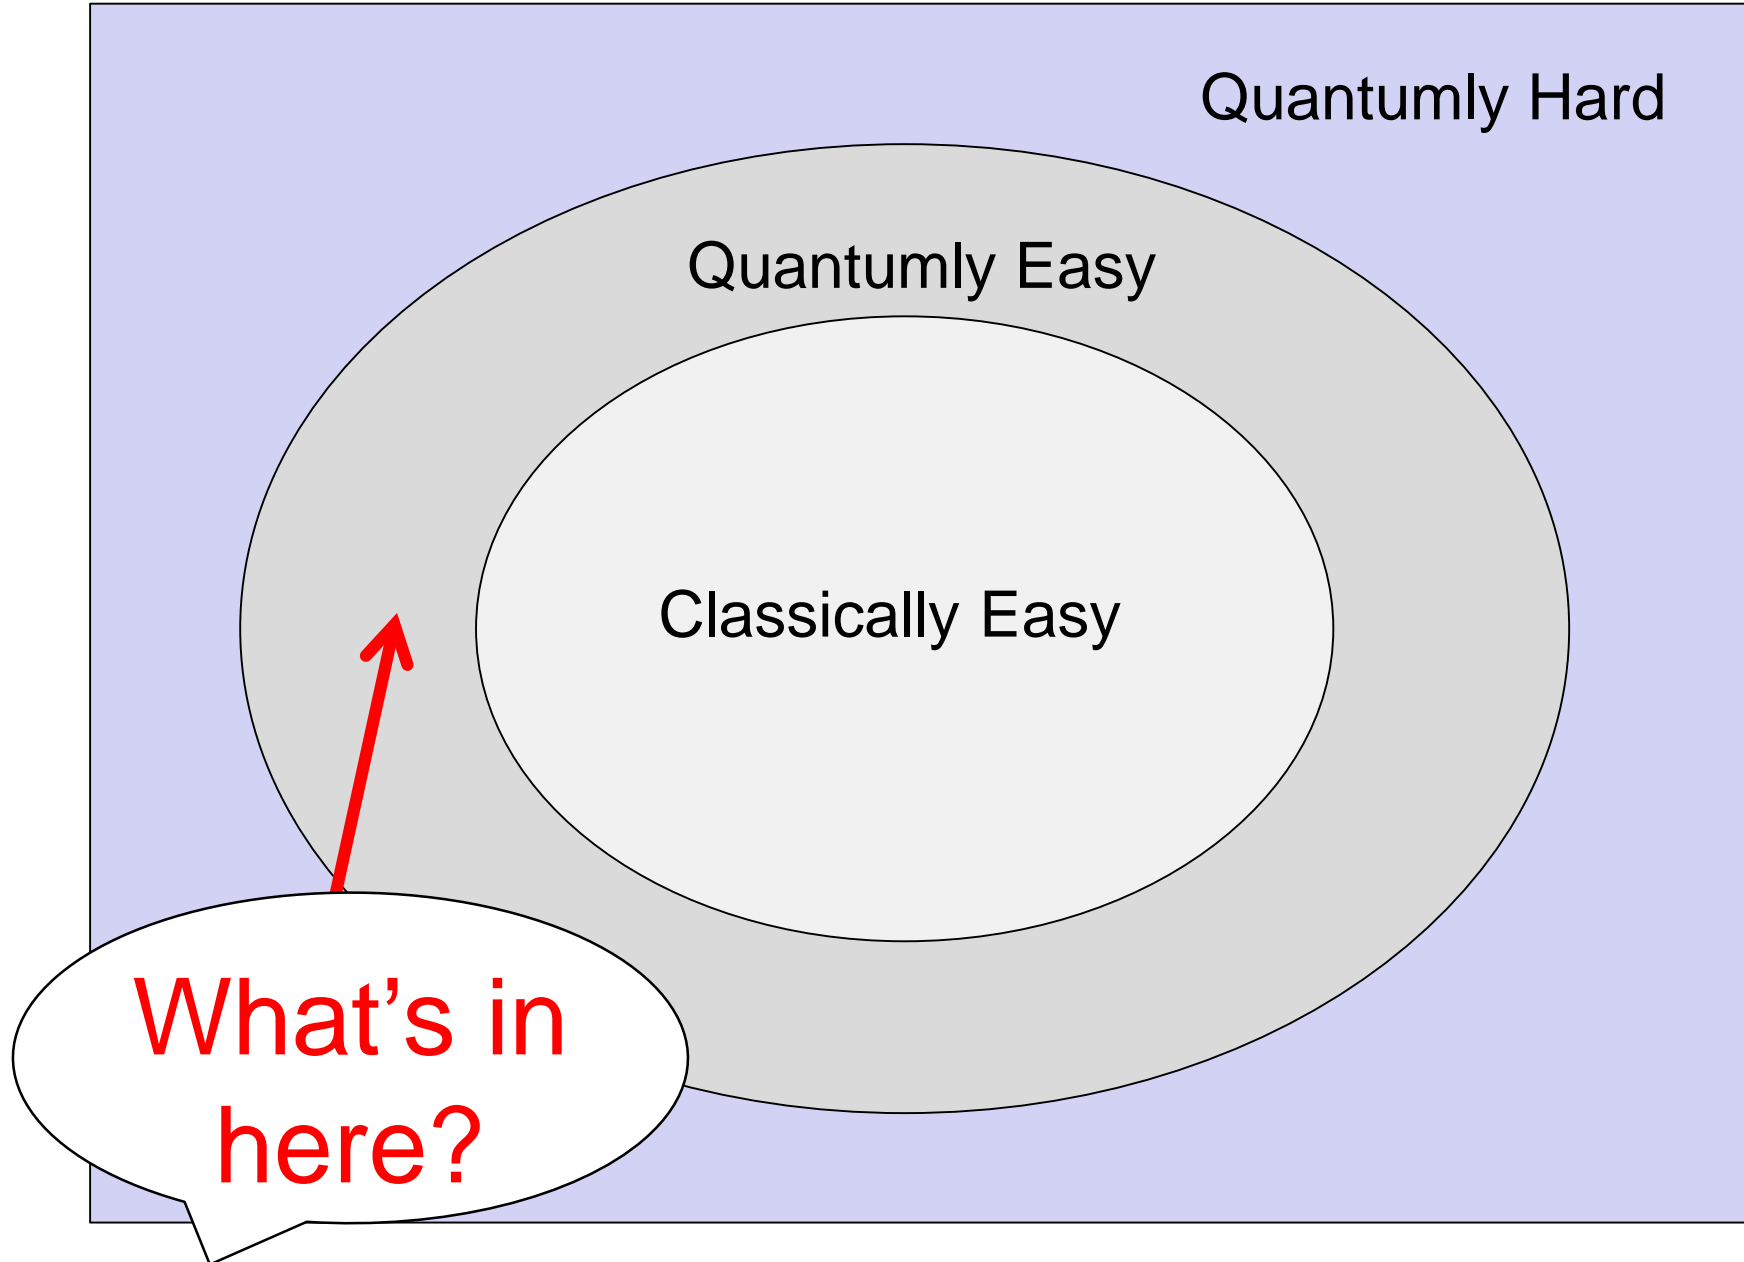
(3) **We don't know how to simulate a quantum computer** efficiently using a digital ("classical") computer. The cost of the best known simulation algorithm rises exponentially with the number of qubits.

But ... **the power of quantum computing is limited.** For example, we don't believe that quantum computers can efficiently solve worst-case instances of NP-hard optimization problems (e.g., the traveling salesman problem).

Problems



Problems



“The theory of everything?”


“The Theory of Everything is not even remotely a theory of every thing ...

We know this equation is correct ... However, it cannot be solved accurately when the number of particles exceeds about 10. No computer existing, or that will ever exist, can break this barrier because it is a catastrophe of dimension

... We have succeeded in reducing all of ordinary physical behavior to a simple, correct Theory of Everything only to discover that it has revealed exactly nothing about many things of great importance.”


R. B. Laughlin and D. Pines, PNAS 2000.






“Nature isn’t classical, dammit, and if you want to make a simulation of Nature, you’d better make it quantum mechanical, and by golly it’s a wonderful problem because it doesn’t look so easy.”


R. P. Feynman, 1981



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

“Understanding strongly interacting quantum systems will be like the transition from alchemy to chemistry.”



Piers Coleman (Rutgers)


Why quantum computing is hard

We want qubits to interact strongly with one another.


We don't want qubits to interact with the environment.

Except when we control or measure them.



Decoherence



Decoherence





Decoherence explains why quantum phenomena, though observable in the microscopic systems studied in the physics lab, are not manifest in the macroscopic physical systems that we encounter in our ordinary experience.



ERROR!

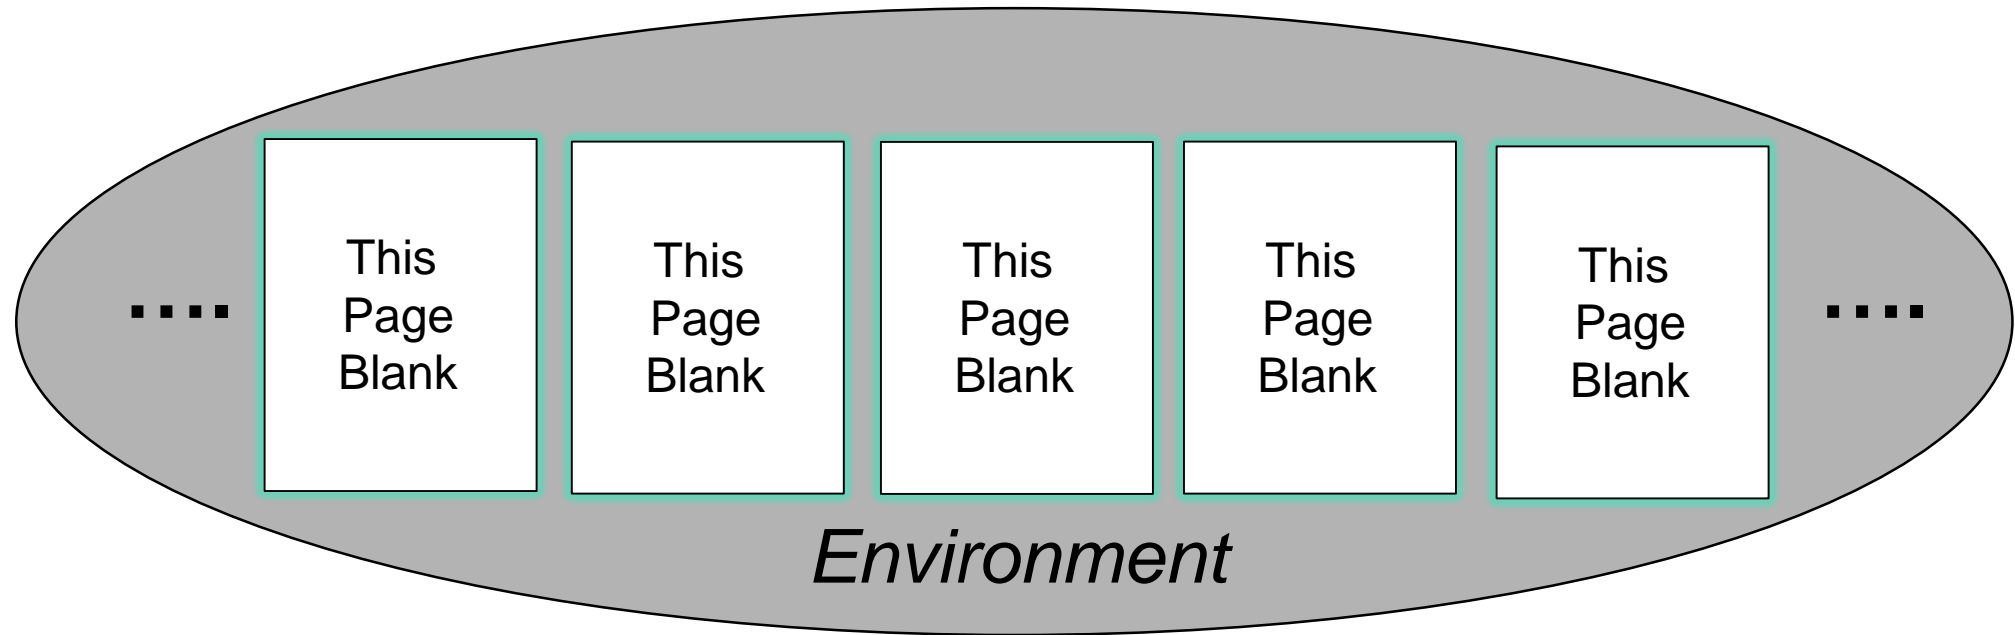
How can we protect a quantum computer from decoherence and other sources of error?



ERROR!

To resist decoherence, we must prevent the environment from “learning” about the state of the quantum computer during the computation.

Quantum error correction




The protected “logical” quantum information is encoded in a highly entangled state of many physical qubits.

The environment can't access this information if it interacts locally with the protected system.




Persistent current in a superconducting circuit




Electron Magnetic Field

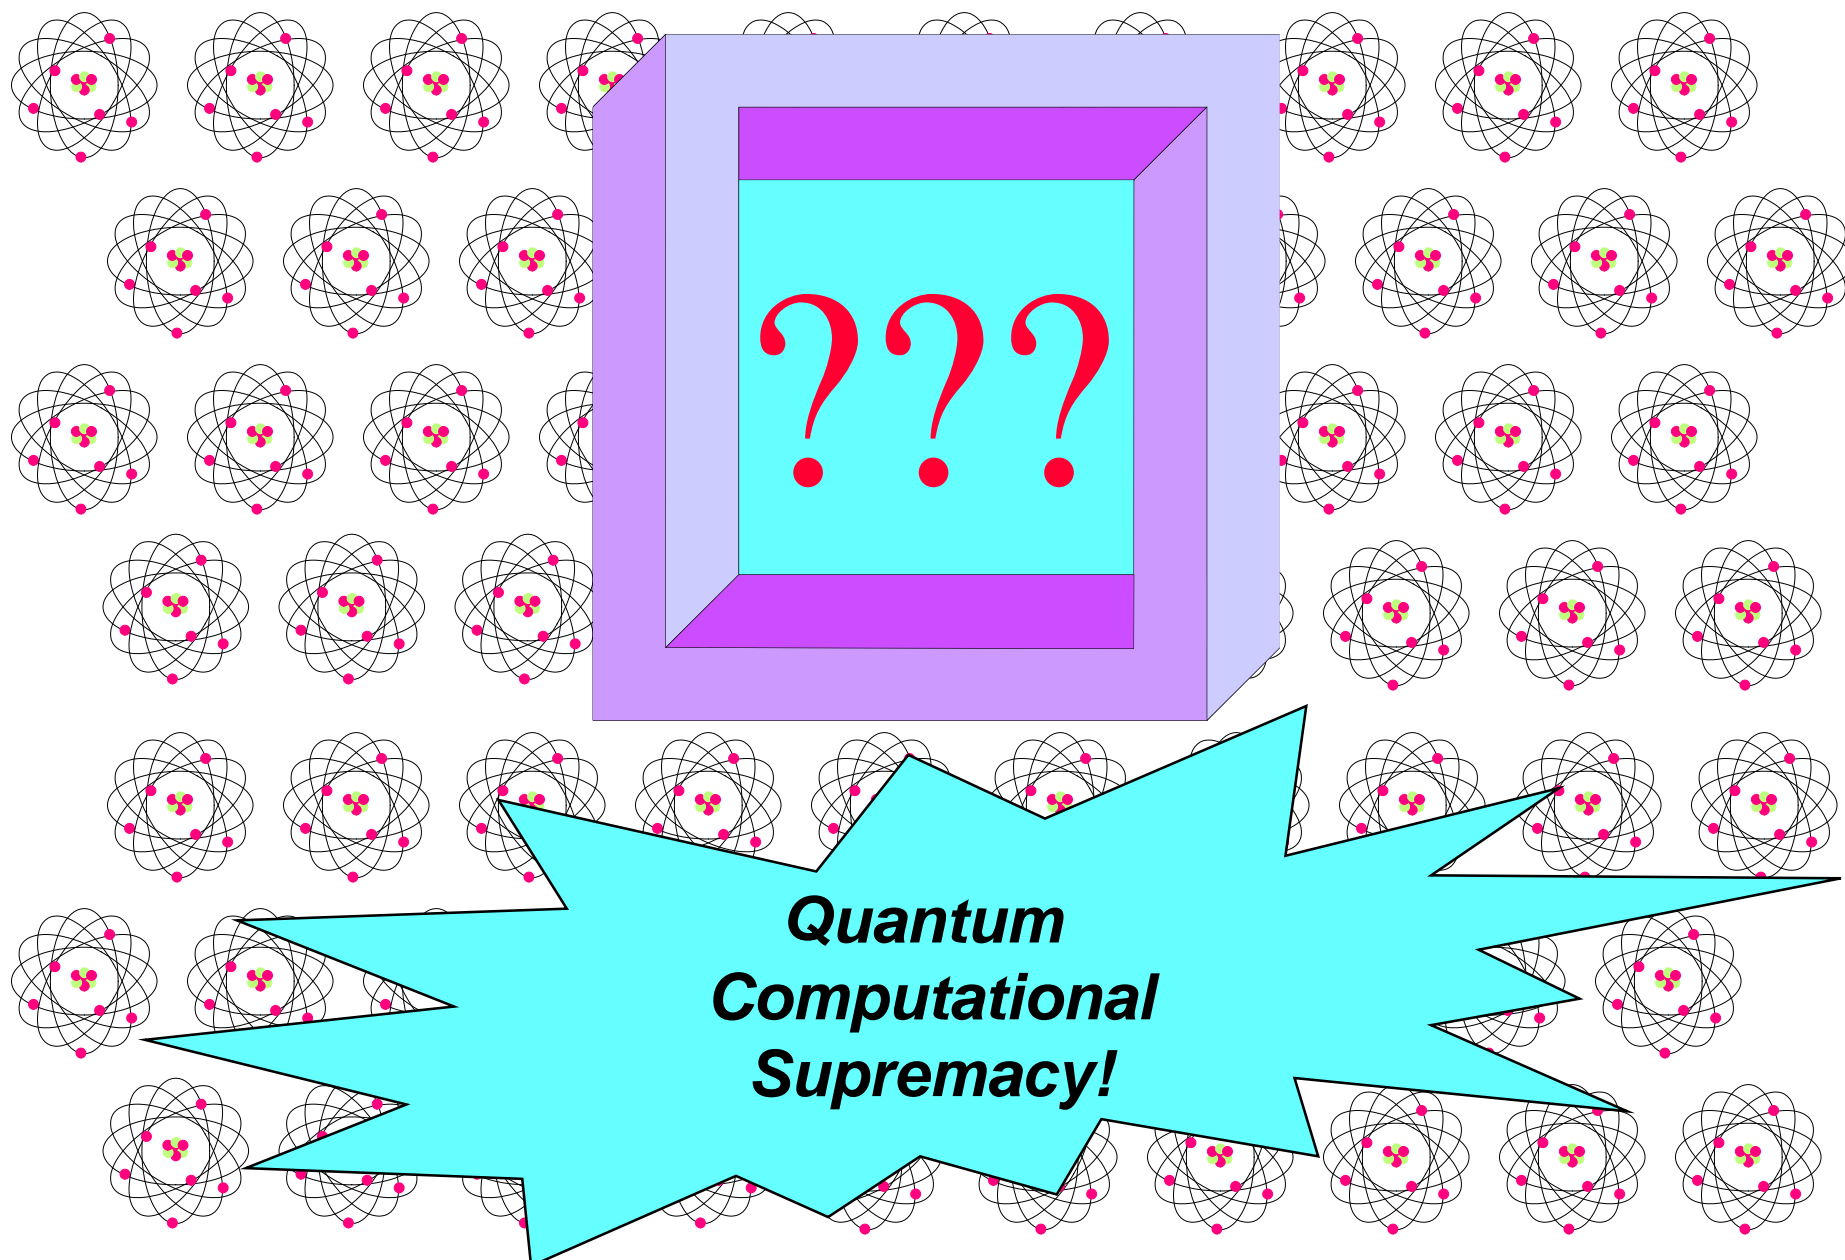
QUBIT



Photon polarization



Atom Internal State



**Quantum
Computational
Supremacy!**

Quantum supremacy using a programmable superconducting processor


<https://doi.org/10.1038/s41586-019-1666-5>

Received: 22 July 2019

Accepted: 20 September 2019

Published online: 23 October 2019

Frank Arute¹, Kunal Arya¹, Ryan Babbush¹, Dave Bacon¹, Joseph C. Bardin^{1,2}, Rami Barends¹, Rupak Biswas³, Sergio Boixo¹, Fernando G. S. L. Brandao^{1,4}, David A. Buell¹, Brian Burkett¹, Yu Chen¹, Zijun Chen¹, Ben Chiaro⁵, Roberto Collins¹, William Courtney¹, Andrew Dunsworth¹, Edward Farhi¹, Brooks Foxen^{1,5}, Austin Fowler¹, Craig Gidney¹, Marissa Giustina¹, Rob Graff¹, Keith Guerin¹, Steve Habegger¹, Matthew P. Harrigan¹, Michael J. Hartmann^{1,6}, Alan Ho¹, Markus Hoffmann¹, Trent Huang¹, Travis S. Humble⁷, Sergei V. Isakov¹, Evan Jeffrey¹,



Credit: Erik Lucero/Google

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


Arguably the most interesting thing we know about the difference between quantum and classical.

About Sycamore

“Quantum David vs. Classical Goliath”

A **fully programmable** circuit-based quantum computer. $n=53$ working qubits in a two-dimensional array with coupling of nearest neighbors.

A circuit with 20 layers of 2-qubit gates *can be executed millions of times in a few minutes*, yielding verifiable results.



Simulating this quantum circuit using a classical supercomputer is hard. *It would take at least days*, possibly much longer.

Furthermore, the cost of the classical simulation grows exponentially with the number of qubits.

Conclusion: **the hardware is working well enough** to produce meaningful results in a regime where classical simulation is very difficult.

What quantum computational supremacy means

“Quantum David vs. Classical Goliath”

It’s a programmable **circuit-based** quantum computer.

An impressive achievement in experimental physics and a testament to ongoing **progress** in building quantum computing hardware.

We have arguably entered the regime where the **extravagant exponential resources** of the quantum world can be validated.

This confirmation does not surprise (most) physicists, but it’s a **milestone** for technology on planet earth.

Building a quantum computer is **merely really, really hard, not ridiculously hard**. The hardware is working; we can begin a serious search for useful applications.

But the specific task performed by Sycamore to demonstrate quantum computational supremacy is not particularly useful.

Quantum computing in the NISQ Era

The (noisy) 50-100 qubit quantum computer has arrived.
(*NISQ* = noisy intermediate-scale quantum.)

NISQ devices *cannot be simulated* by brute force using the most powerful currently existing supercomputers.

Noise limits the computational power of NISQ-era technology.

NISQ will be an interesting tool for exploring physics. It *might* also have other useful applications. But we're not sure about that.


NISQ will not change the world by itself. Rather it is a step toward more powerful quantum technologies of the future.

Potentially transformative scalable quantum computers may still be decades away. *We're not sure how long it will take.*

Quantum 2, 79 (2018), arXiv:1801.00862

Hybrid quantum/classical 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.

Classical optimization algorithms (for both classical and quantum problems) are sophisticated and well-honed after decades of hard work.

We don't know whether NISQ devices can do better, but we can try it and see how well it works.

The era of quantum heuristics

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.”

Sometimes algorithms are effective in practice even though theorists are not able to validate their performance in advance.

Example: Deep learning. Mostly tinkering so far, without much theory input.

Possible quantum examples:

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

What can we do with, say, < 100 qubits, depth < 100 ? **We need a dialog between quantum algorithm experts and application users.**

Maybe we’ll get lucky ...

The steep climb to scalability

NISQ-era quantum devices will not be protected by quantum error correction. Noise will limit the scale of computations that can be executed accurately.

Quantum error correction (QEC) will be essential for solving some hard problems. But QEC carries a high overhead cost in number of qubits & gates.

This cost depends on both the hardware quality and algorithm complexity. With today's hardware, solving (say) useful chemistry problems may require hundreds to thousands of physical qubits for each protected logical qubit.


Recent estimate: 20 million physical qubits to break RSA 2048 (*Gidney, Ekerå 2019*), for gate error rate 10^{-3} .

To reach scalability, we must cross the daunting “quantum chasm” from hundreds to millions of physical qubits. This may take a while.

Advances in quantum gate fidelity, systems engineering, algorithm design, and error correction protocols can hasten the arrival of the fully fault-tolerant quantum computer.




Persistent current in a superconducting circuit




Electron Magnetic Field

QUBIT



Photon polarization



Atom Internal State

*“Quantum computing is
a marathon not a sprint”*



Chris Monroe (UMD/IonQ)

Quantum speedups in the NISQ era and beyond

Quantum computational supremacy demonstrations confirm the extravagant computation resources provided by the quantum world.

In the NISQ era we can explore *heuristic* quantum algorithms. Near-term quantum advantage for useful applications is possible, but not guaranteed.

For truly scalable quantum computing, quantum error correction will be required. But QEC has a dauntingly high overhead cost, and will not be feasible in the near term.

Lower quantum gate error rates will lower the overhead cost of quantum error correction, and also extend the reach of quantum algorithms which do not use error correction.

NISQ will not change the world by itself. Realistically, the goal for near-term quantum platforms should be to *pave the way for bigger payoffs using future devices*. Progress toward fault-tolerant QC must continue to be a high priority for quantum technologists.

Quantum information vs. Classical information

- 1) **Randomness.** Clicks in a Geiger counter are intrinsically random, not pseudorandom. Can't predict outcome even with the most complete possible knowledge of the state.
- 2) **Uncertainty.** Operators A and B do not commute means that measuring A influences the outcome of a subsequent measurement of B .
- 3) **Entanglement.** *The whole is more definite than the parts.* Even if we have the complete possible knowledge of the (pure) state of joint system AB , the (mixed) state of A may be highly uncertain.

Qubit

A vector (actually a “ray” because the normalization is 1 by convention, and the overall phase does not matter) in a two dimensional complex Hilbert space.

$$|\psi\rangle = a|0\rangle + b|1\rangle, \quad a, b \in \mathbb{C}, \quad |a|^2 + |b|^2 = 1, \quad |\psi\rangle \sim e^{i\alpha} |\psi\rangle$$

(two real parameters).

“Classical” in the special case where $|0\rangle$ or $|1\rangle$ is “promised”.

The two orthogonal states $|0\rangle$ and $|1\rangle$ are perfectly distinguishable. If Alice sends one or the other to Bob, he can measure in the $\{|0\rangle, |1\rangle\}$ basis and identify the state.



Qubit

The orthogonal states

$$|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$$

are also perfectly distinguishable.

Suppose Alice sends either $|1\rangle$ or $|+\rangle$ to Bob. Now Bob cannot distinguish the states perfectly. His best measurement, which succeeds with probability $\cos^2(\pi/8) \approx .853$ if the two states are equally likely, projects onto the orthogonal basis shown. (Prove it! Generalize it!) Alternative: a measurement which is sometimes inconclusive, but identifies the state correctly when conclusive.



Information vs. disturbance

Suppose Alice prepares either $|\varphi\rangle$ or $|\psi\rangle$. To distinguish the two possible states, Eve performs a unitary transformation that rotates her probe while leaving Alice's state intact

$$U : \begin{aligned} |\varphi\rangle_A \otimes |0\rangle_E &\rightarrow |\varphi\rangle_A \otimes |e\rangle_E \\ |\psi\rangle_A \otimes |0\rangle_E &\rightarrow |\psi\rangle_A \otimes |f\rangle_E \end{aligned}$$

where $|e\rangle$ and $|f\rangle$ are normalized states. Since U preserves inner products,

$$\langle\psi|\varphi\rangle \cdot \langle f|e\rangle = \langle\psi|\varphi\rangle,$$

and if $|\varphi\rangle$ and $|\psi\rangle$ are nonorthogonal, then $\langle f|e\rangle = 1$;
the states of the probe are the same. Eve's measurement of the probe cannot reveal any information about whether the state is $|\varphi\rangle$ or $|\psi\rangle$. On the other hand if $|\varphi\rangle$ and $|\psi\rangle$ are orthogonal, the probe states can also be orthogonal. Eve can *copy* the info.

Tensor Product



$$\{|i\rangle_A, \quad i = 1, 2, \dots, d_A\} \qquad \{|a\rangle_B, \quad a = 1, 2, \dots, d_B\}$$

Basis states of the composite system are distinguishable if they can be distinguished on either Alice's or Bob's side:

$$(\langle j | \otimes \langle b |)(|i\rangle \otimes |a\rangle) = \delta_{ij} \delta_{ab}$$

If Alice and Bob both have qubits, the basis states

$$\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$$

are all distinguishable.

Many qubits



$$\mathbb{C}^{2^n} = \underbrace{\mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \dots \otimes \mathbb{C}^2}_{n \text{ times}}$$

spanned by ...

$$|x\rangle = |x_{n-1}\rangle \otimes |x_{n-2}\rangle \otimes \dots \otimes |x_1\rangle \otimes |x_0\rangle, \quad x \in \{0, 1\}^n$$

where $\langle x | y \rangle = \delta_{x,y}$

-- complex dimension $d = 2^n$, $2^{n+1} - 2$ real parameters.

$$|\psi\rangle = \sum_{x=0}^{2^n-1} a_x |x\rangle$$

For 300 qubits, vector in a space with dimension $2^{300} \sim 10^{90}$, more than the

the number of atoms in the visible universe. No succinct classical description of the quantum state, in general.

Which decomposition into subsystems?



$$\mathbb{C}^{2^n} = \underbrace{\mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \dots \otimes \mathbb{C}^2}_{n \text{ times}}$$

Typically dictated by spatial locality. The qubits may be in different cities, or encoded in distinct atoms, ...

A *product state* $|\psi\rangle = |\psi_1\rangle_{A_1} \otimes |\psi_2\rangle_{A_2} \otimes \dots \otimes |\psi_n\rangle_{A_n}$

has a succinct description; only $2n$ real parameters. It can be created by n parties each acting locally in his/her city.

If a (pure) state is not a product state, it is *entangled*.


Entanglement cannot be created by remote parties acting locally, even if they communicate classically.

Entanglement



If a (pure) state is not a product state, it is *entangled*. Entanglement cannot be created by remote parties acting locally, even if they communicate classically.

However, by bringing the qubits together pairwise, or by sending qubit messages between parties, arbitrary entangled states can be constructed.



Two qubit gates are universal. In general, though, this construction is *inefficient*. For most entangled states, an exponentially large number of two-qubit gates are needed to create the state, starting with a product state.

Quantum computer: the circuit model

(1) Hilbert space of n qubits: $\mathfrak{H} = \left(\mathbb{C}^{2^n} \right)$
spanned by

$$|x\rangle = |x_{n-1}\rangle \otimes |x_{n-2}\rangle \otimes \cdots \otimes |x_1\rangle \otimes |x_0\rangle, \quad x \in \{0,1\}^n$$

Important: the Hilbert space is equipped with a natural tensor-product decomposition into subsystems.

$$\mathbb{C}^{2^n} = \underbrace{\mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \cdots \otimes \mathbb{C}^2}_{n \text{ times}}$$

Physically, this decomposition arises from spatial locality. Elementary operations (“quantum gates”) that act on a small number of qubits (independent of n) are “easy;” operations that act on many qubits (increasing with n) are “hard.”


(2) Initial state: $|000\dots 0\rangle = |0\rangle^{\otimes n}$

Quantum computer: the circuit model

(3) A finite set of fundamental quantum gates:

$$\{U_1, U_2, U_3, \dots, U_{n_G}\}$$

Each gate is a unitary transformation acting on a bounded number of qubits. The gates form a universal set: arbitrary unitary transformations can be constructed, to any specified accuracy, as a quantum circuit constructed from the gates:



(Universal gates are generic.)

Important: One universal set of gates can simulate another efficiently, so there is a notion of complexity that is independent of the details of the quantum hardware.

Quantum computer: the circuit model

(4) Classical control:

The construction of a quantum circuit is directed by a classical computer, *i.e.*, a Turing machine. (We're not interested in what a quantum circuit can do unless the circuit can be designed efficiently by a classical machine.)

(5) Readout:

At the end of the quantum computation, we read out the result by measuring σ_z , *i.e.*, projecting onto the basis $\{|0\rangle, |1\rangle\}$

(We don't want to hide computational power in the ability to perform difficult measurements.)

Quantum computer: the model

- (1) n qubits
- (2) initial state
- (3) quantum gates
- (4) classical control
- (5) readout

Clearly, the model can be simulated by a classical computer with access to a random number generator. But there is an exponential slowdown, since the simulation involves matrices of exponential size ($2^n \times 2^n$).

The quantum computer might solve efficiently some problems that can't be solved efficiently by a classical computer. ("Efficiently" means that the number of quantum gates = polynomial of the number of bits of input to the problem.)

Quantum computer: the model

- (1) n qubits
- (2) initial state
- (3) quantum gates
- (4) classical control
- (5) readout

Is this “circuit model” sufficiently powerful to *simulate efficiently* any physical process that occurs in Nature (the “strong quantum Church-Turing thesis”)?

Some challenges are:

-- Local quantum field theory (standard model of particle physics). Field modes of arbitrarily small wavelength, hence an infinite number of degrees of freedom per unit volume. Need to include energy as well as time and space in resource accounting.

-- Quantum gravity. May be described by local quantum field theory or (large) matrix quantum mechanics. But subtle resource accounting because short distance in the bulk spacetime corresponds to long distance in the dual boundary field theory.

Three crucial lessons

- (1) Quantum computers can solve hard problems.
- (2) Quantum errors can be controlled.
- (3) Quantum hardware can be constructed.

Lots of questions

- (1) What problems can be solved efficiently with quantum computers?
What problems cannot be?
- (2) Can we build a “quantum Internet” and how would we use it?
- (3) Can we construct much more reliable quantum hardware than we have now?
- (4) How can quantum information concepts illuminate issues of fundamental interest in the physical science?