Quantum Computing: Transforming the Digital Age

Krysta Svore
Quantum Architectures and Computation (QuArC)
Microsoft Research

Quantum Optimization Workshop 2014

ieee spec

AEROSPACE BIOMEDICAL

COMPUTING // HARDWARE

NEWS

Quantum Chip Help

Experimental chip does

EMAIL PRINT

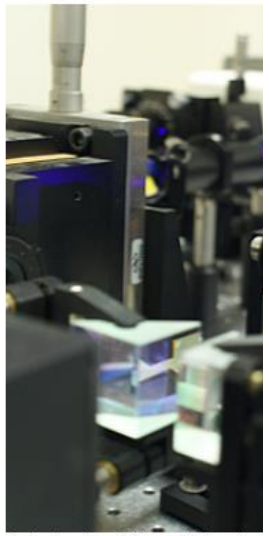

Photo: Jonathan Matthews/University o

BY ANNE-MARIE CORLEY // SEPT

3 September 2009—Modern cryp
have in factoring huge numbers,

The New York Times

Chinese Hackers Pursue Key Data on U.S. Workers

BITS BLOG

Adidas Joins Wearable Stampede With Fitness Tracker

BITS BLOG

IBM Wants to Inven
Future, Not Make Th

TECHNOLOGY

Microsoft Makes Bet Quantum Computing

By JOHN MARKOFF JUNE 23, 2014

EMAIL
FACEBOOK
TWITTER
SAVE
MORE

Top Stories
This article and others like it are part of our new subscription.
Learn More »

SANTA BARBARA, Calif. — Modern computers are not unlike the looms of the industrial revolution: They follow programmed instructions to weave intricate patterns. With a loom, you see the result in a cloth or carpet. With a computer, you see it on an electronic display.

Now a group of physicists and computer scientists funded by Microsoft is trying to take the analogy of interwoven threads to what some believe will be the next great leap in computing, so-called quantum computing.

If the scientists are right, their research could lead to the design of computers that are far more powerful than today's supercomputers and could solve problems in fields as diverse as chemistry, material science, artificial intelligence and code-breaking.

Michael Freedr
2005 that can
Emily Berl for Th

The machine does not
effects

The New York Times Technology | Personal Tech | Business Day

SUBSCRIBE NOW

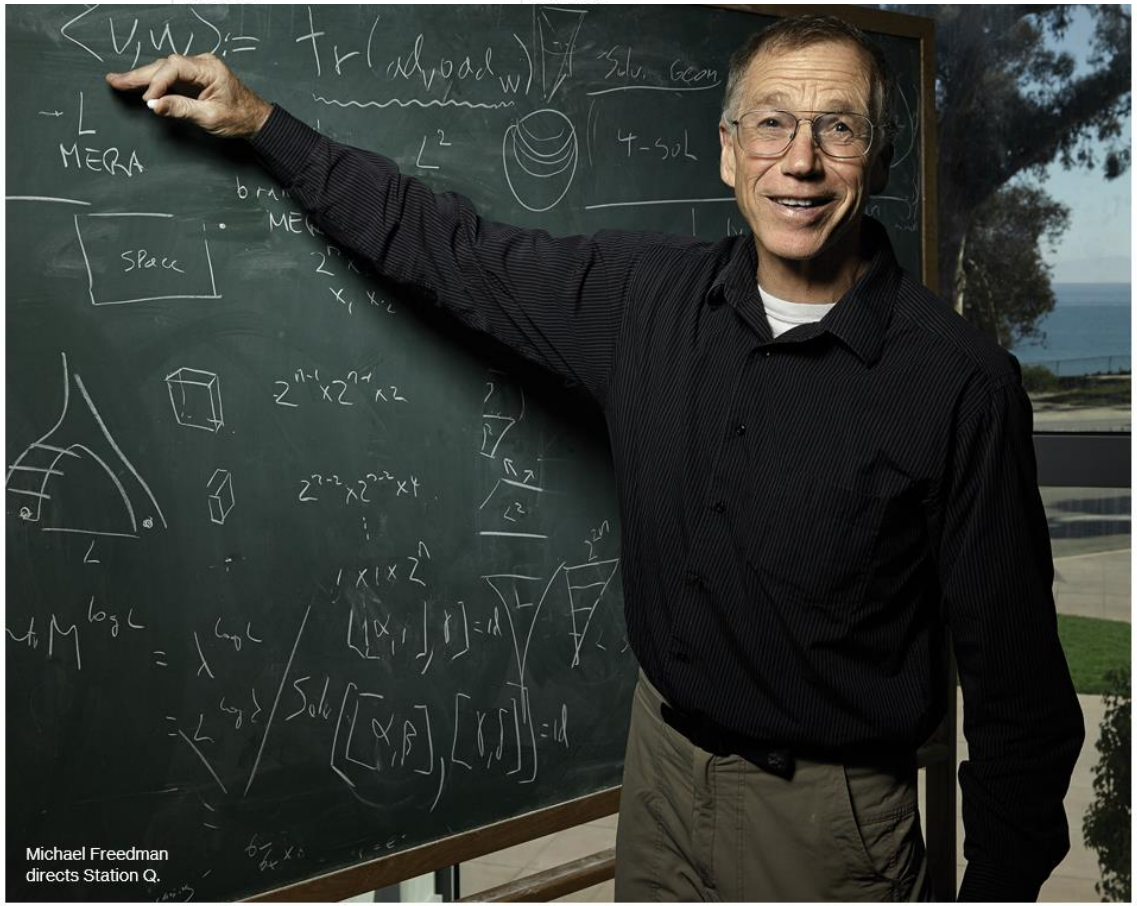

Michael Freedman directs Station Q.

Microsoft's Quantum Mechanics

Can an aging corporation's adventures in fundamental physics research open a new era of unimaginably powerful computers?

By Tom Simonite on October 10, 2014

Computers have come a long way

Antikythera mechanism
(100 BC)

Babbage's Difference Engine
(proposed 1822)

ENIAC
(1946)

Sequoia
(2012)

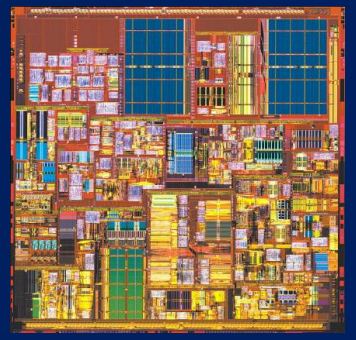

Quantum
(2025?)

Thanks to Matthias Troyer

Success of Digital Computers and Moore's Law

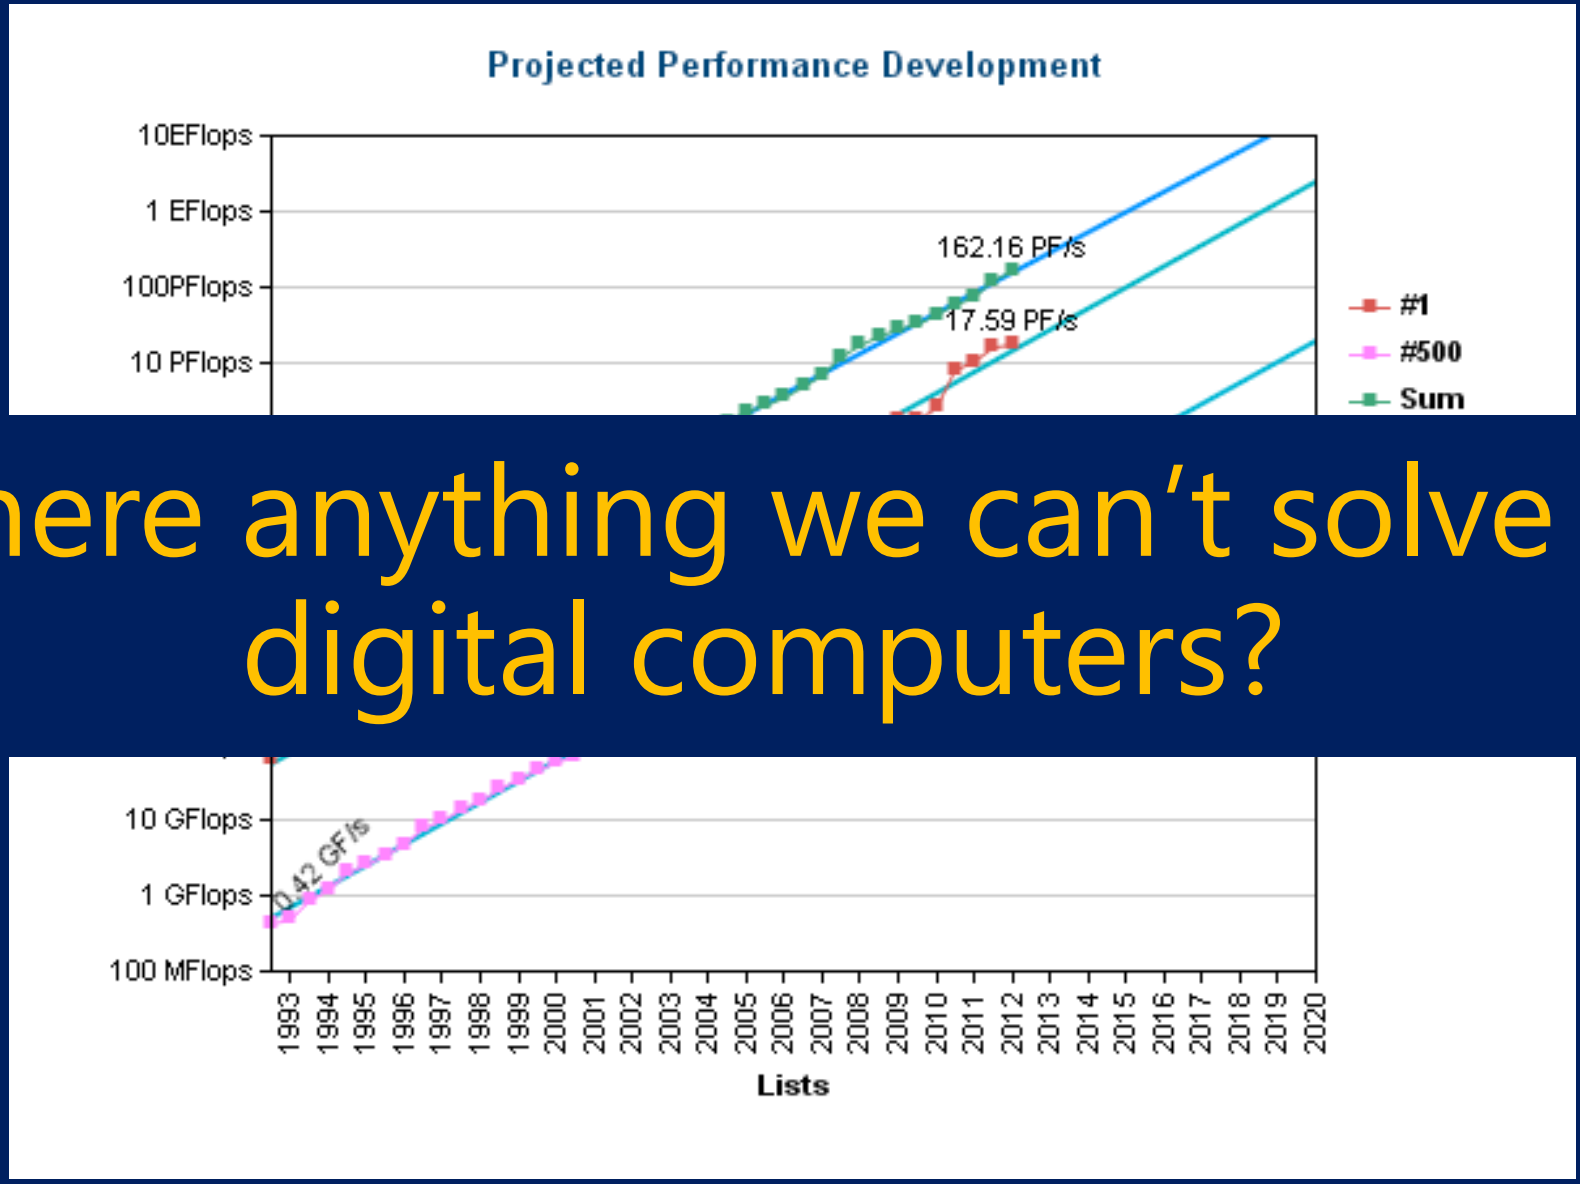

Is there anything we can't solve on digital computers?

Some problems are hard to solve

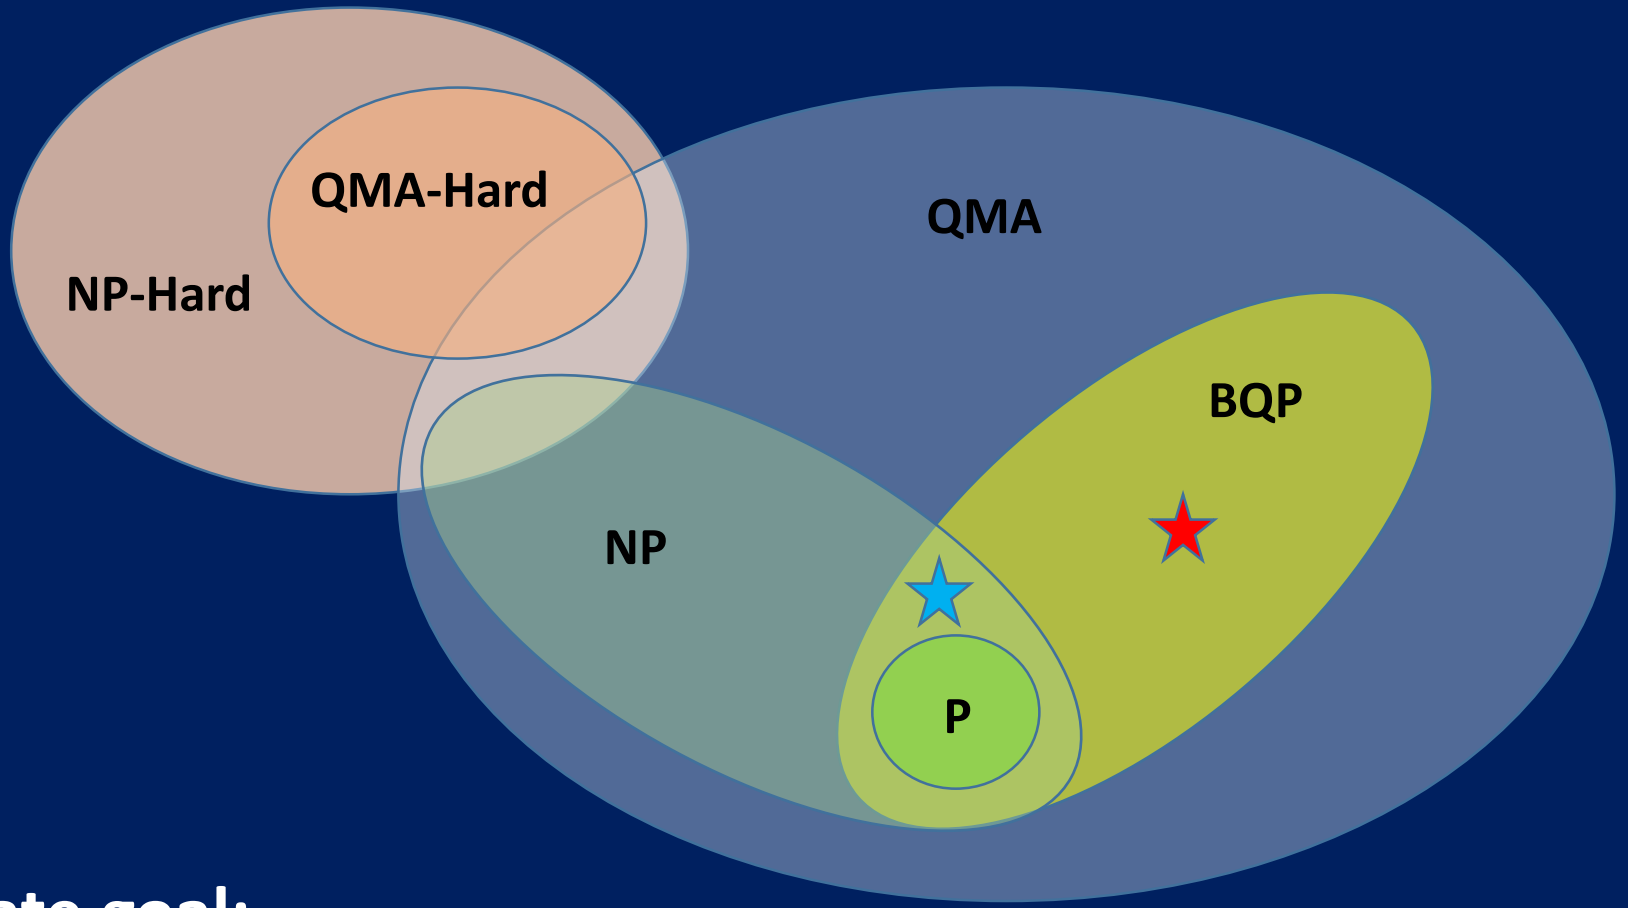

Ultimate goal:

Develop quantum algorithms whose complexity lies in BQP\P

Quantum Magic: Interference

source of particles

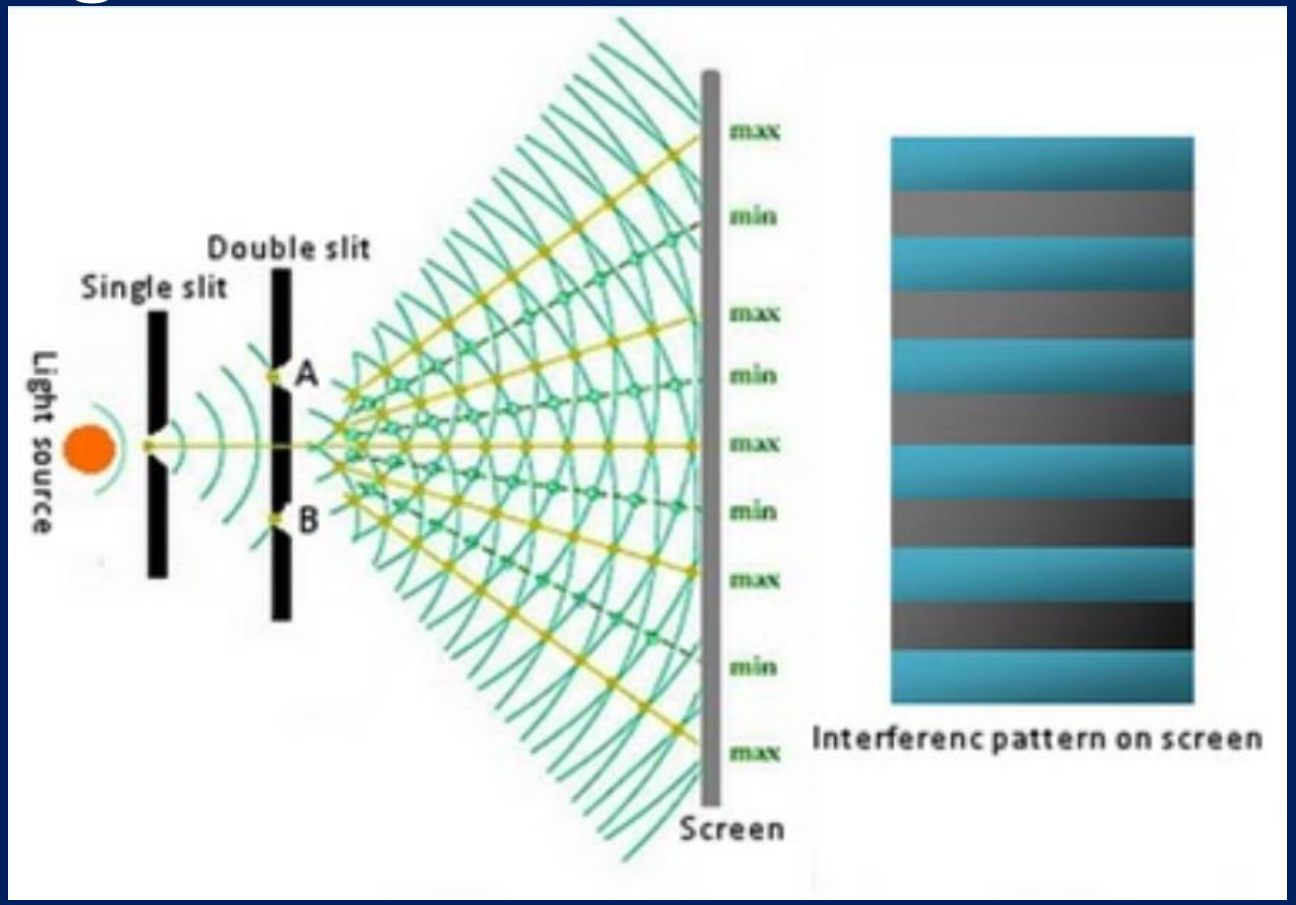

interference pattern = quantum coherence

Classical objects go ***either*** one way or the other.

Quantum objects (electrons, photons) go ***both*** ways.

Gives a quantum computation an inherent type of parallelism!

Quantum Magic: Qubits and Superposition

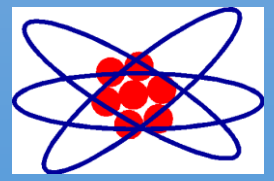

single atom

$|g\rangle = |0\rangle$

$|e\rangle = |1\rangle$

single spin

$|\downarrow\rangle = |0\rangle$

$|\uparrow\rangle = |1\rangle$

$$|\psi\rangle = |0\rangle + |1\rangle \qquad |\psi\rangle = \downarrow + \uparrow$$

Information encoded in the state of a two-level quantum system

$$|\psi\rangle =$$

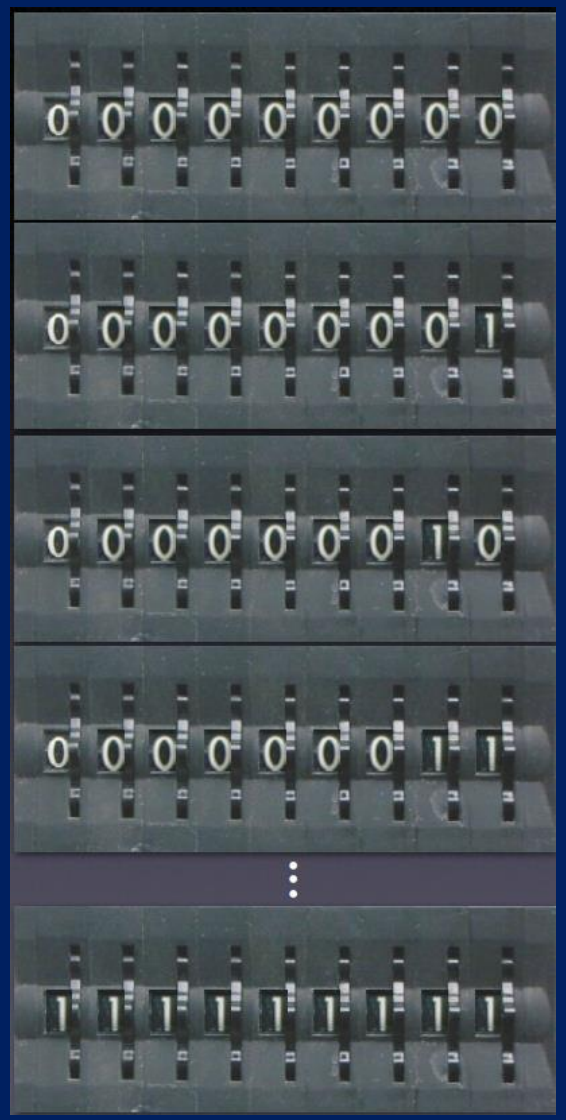

Thanks to Charlie Marcus

Input

Output

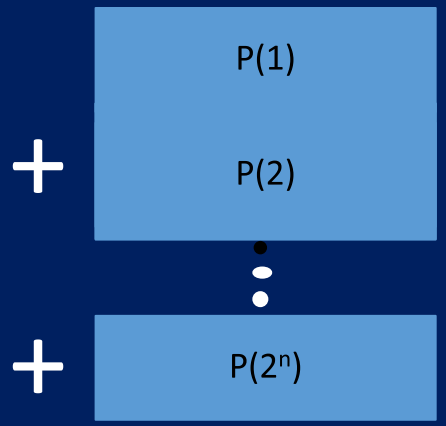

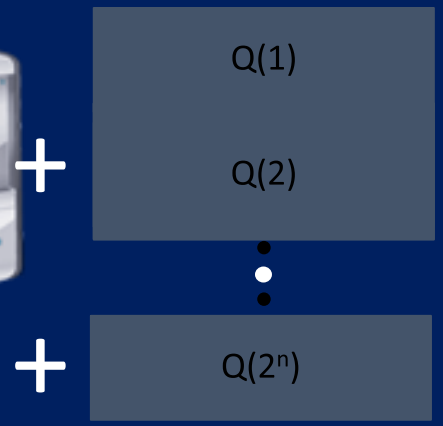

Quantum Magic: Entanglement

Nonlocal Correlations!

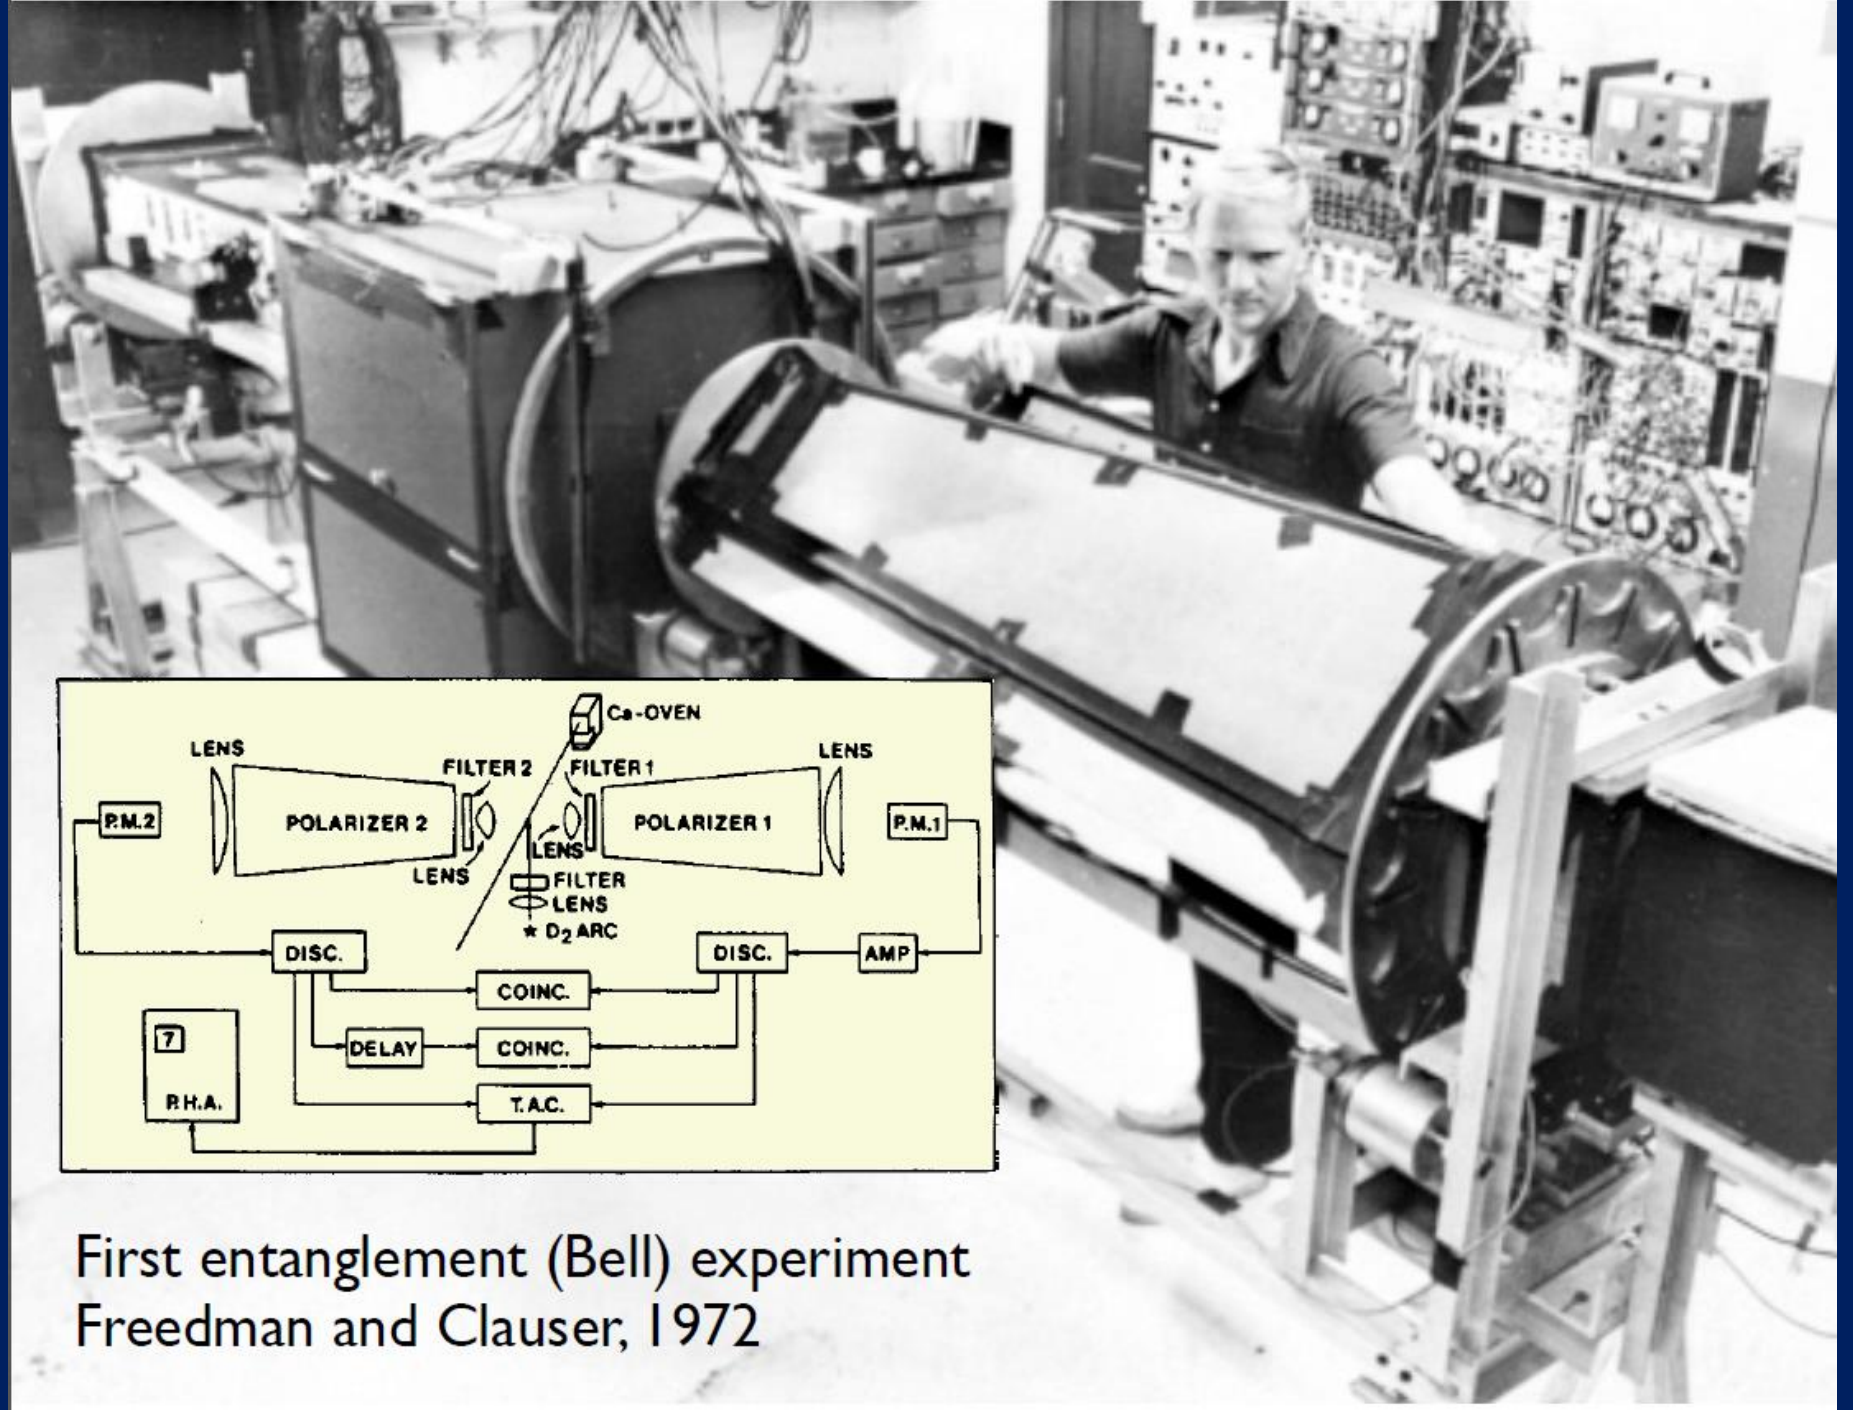

First entanglement (Bell) experiment
Freedman and Clauser, 1972

Quantum Magic: Entanglement

2*5 distinct amplitudes

N non-interacting qubits

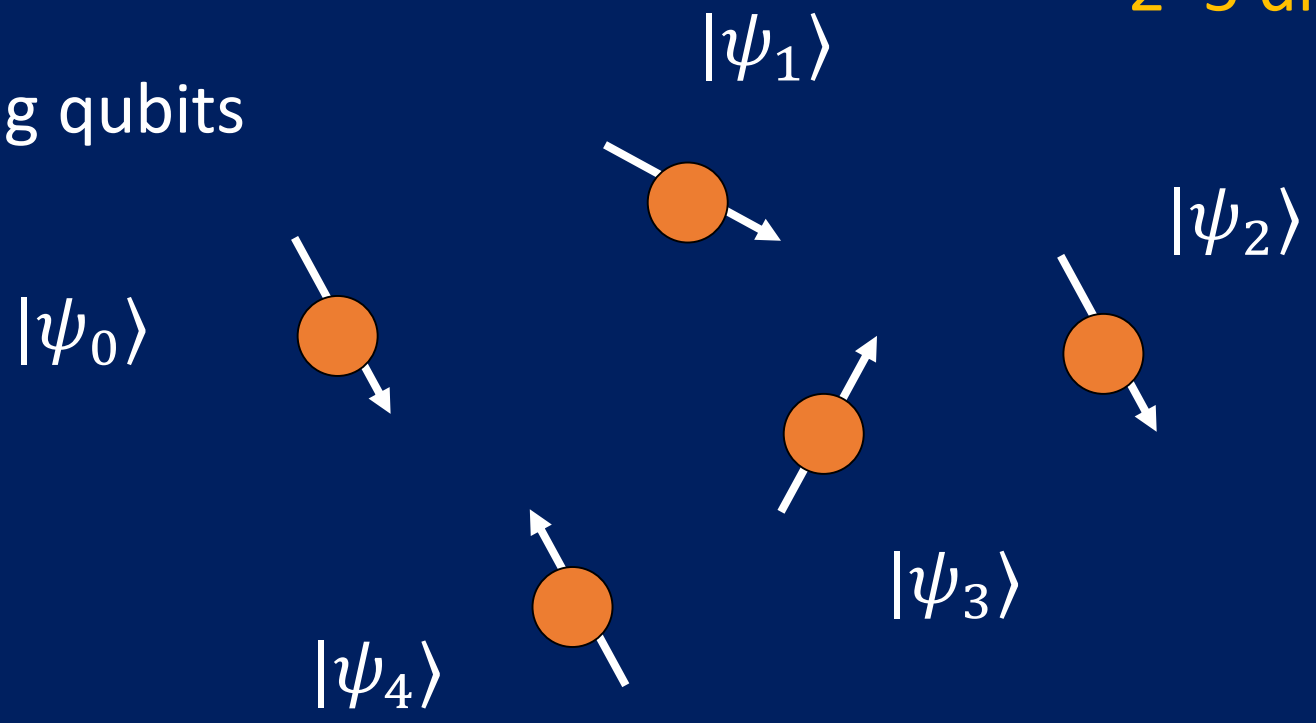

$$|\psi_{total}\rangle = (\alpha_0|0\rangle + \beta_0|1\rangle) \otimes (\alpha_1|0\rangle + \beta_1|1\rangle) \otimes \cdots \otimes (\alpha_{N-1}|0\rangle + \beta_{N-1}|1\rangle)$$

State of N non-interacting qubits: ~ N bits of info

Thanks to Rob Schoelkopf

Quantum Magic: Entanglement

General state of N *interacting* qubits

32 distinct amplitudes!

Simulating a 200-qubit interacting system requires ~10^{60} classical bits!

$$|\psi_{total}\rangle = c_0|00\ldots0\rangle + c_1|00\ldots1\rangle + \ldots c_{2^N-1}|11\ldots1\rangle$$

State of N interacting qubits: ~ 2^N bits of info!

Thanks to Rob Schoelkopf

Quantum Magic: What's the catch?

Need coherent control

Need strongly Interacting system

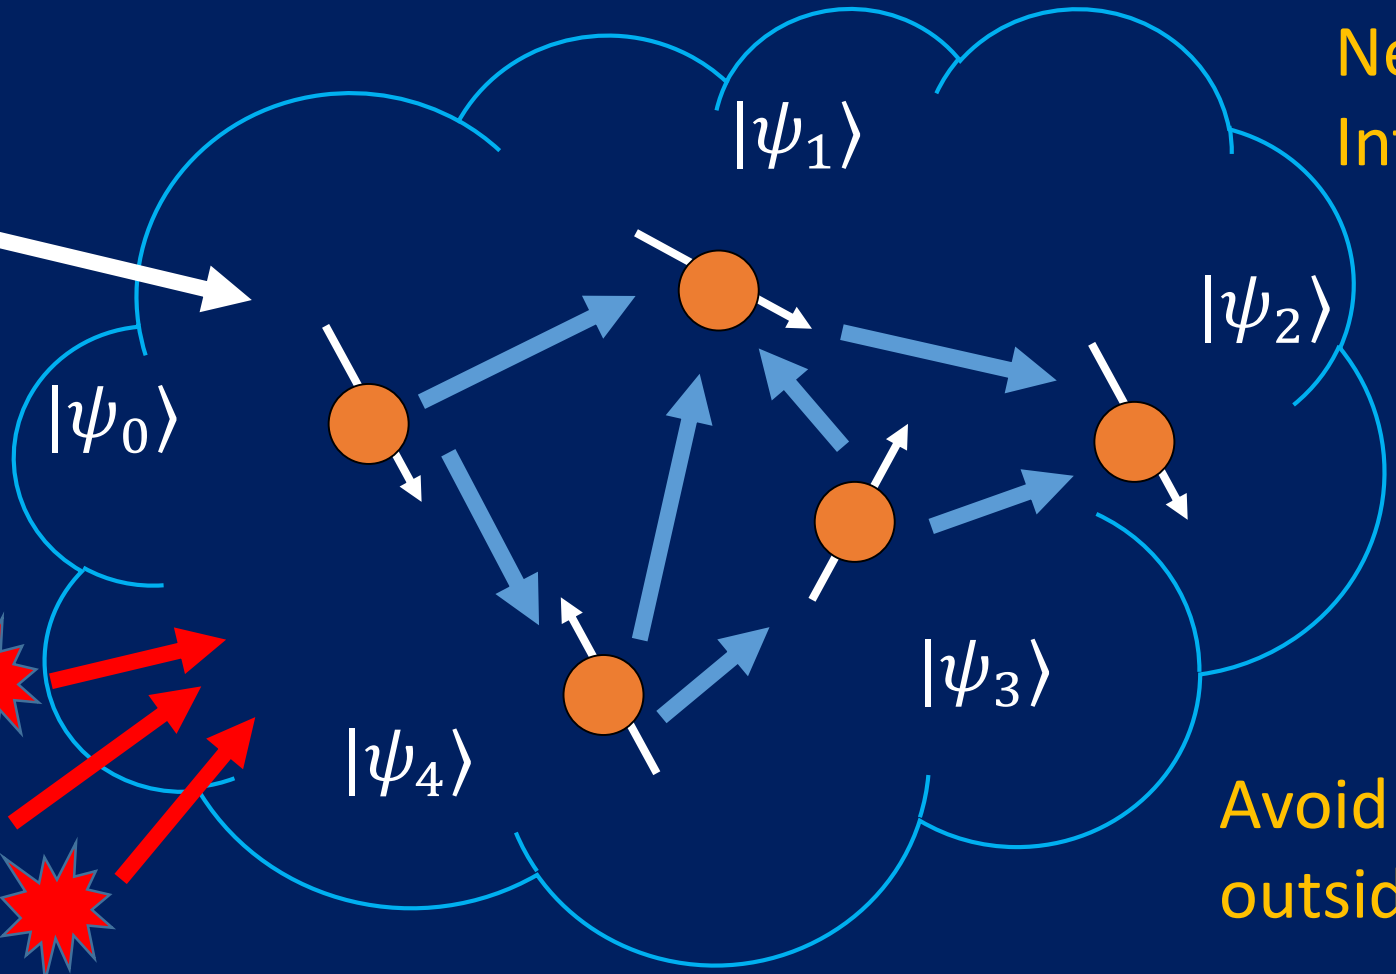

Decoherence and errors!

Avoid interaction with outside environment!

Thanks to Rob Schoelkopf

Quantum Gates: Digital quantum computation

Basic unit: **bit** = 0 or 1

Computing: **logical** operation

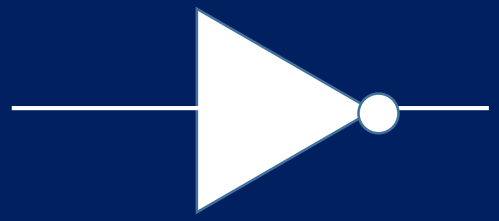

NOT

$0 \rightarrow 1$

$1 \rightarrow 0$

Basic unit: **qubit** = unit vector

$\alpha|0\rangle + \beta|1\rangle$

Computing: **unitary** operation

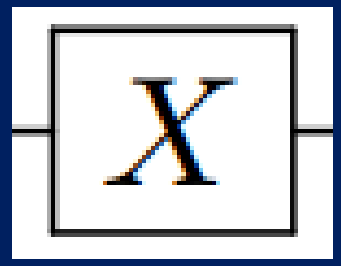

NOT

$$\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} \beta \\ \alpha \end{bmatrix}$$

Quantum Gates: Digital quantum computation

Basic unit: **bit** = 0 or 1

Computing: **logical** operation

Description: **truth table**

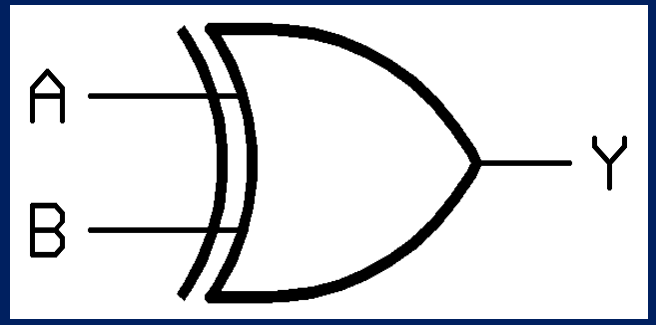

XOR gate

A	B	Y
0	0	0
0	1	1
1	0	1
1	1	0

Basic unit: **qubit** = unit vector

$$\alpha|0\rangle + \beta|1\rangle$$

Computing: **unitary** operation

Description: **unitary matrix**

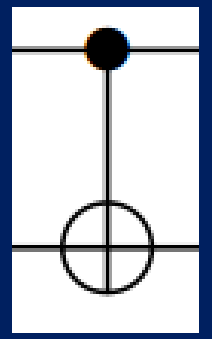

CNOT gate

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

Quantum power unleashed: super-fast FFT

FFT

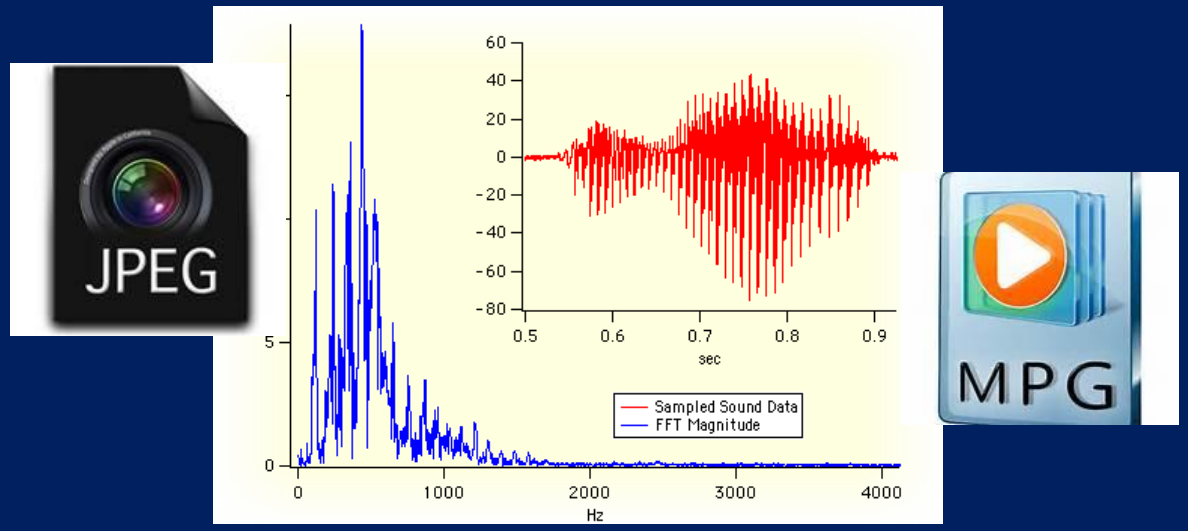

 # ops = N log N

Example:
1GB of data =
10 Billion ops

Quantum FFT

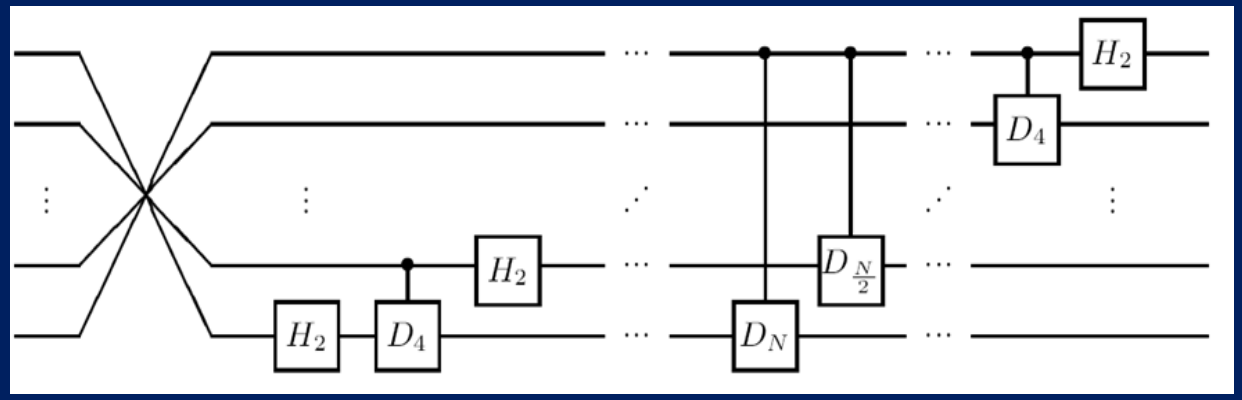

 # ops = log N

Example:
1GB of data =
27 ops (!!!)

Any other catches?

No-cloning principle

Quantum information
cannot be copied

I/O limitations

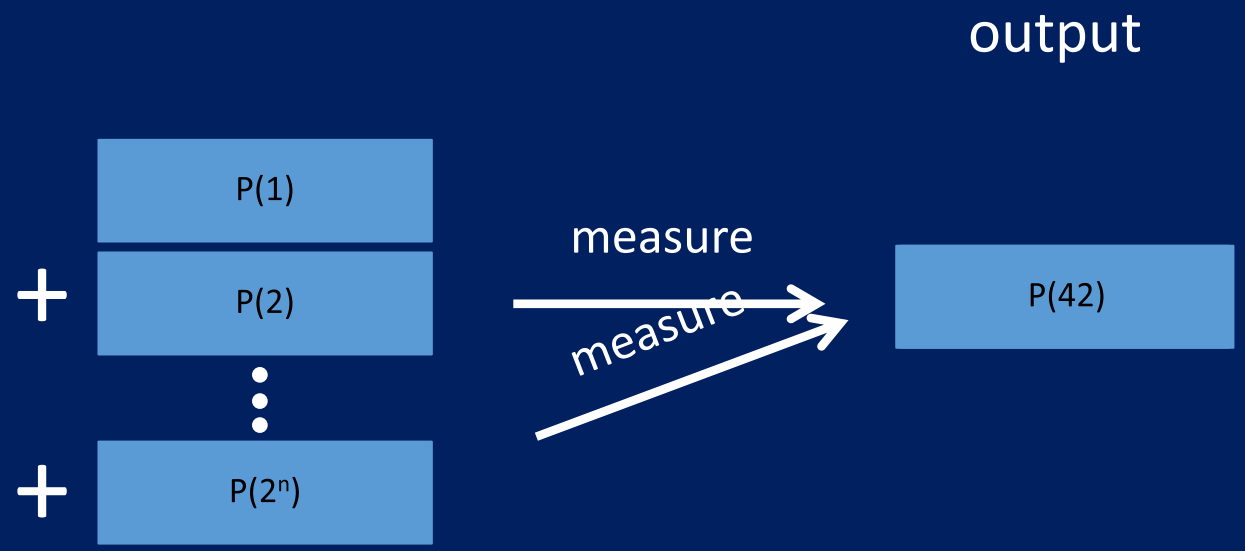

Input: preparing initial state can be costly

Output: reading out a state is probabilistic

Quantum Algorithms Exist!

Shor's Algorithm (1994)	• Breaks RSA, elliptic curve signatures, DSA, El-Gamal • Exponential speedups	
Solving Linear Systems of Equations (2010)	• Applications shown for electromagnetic wave scattering • Exponential speedups	
Quantum simulation (1982)	• Simulate physical systems in a quantum mechanical device • Exponential speedups	

Cryptography

$$15 = \square \times \square$$

$$15 = 5 \times 3$$

$$1387 = \square \times \square$$

$$1387 = 19 \times 73$$

1807082088687
4048059516561
6440590556627
8102516769401
3491701270214
5005666254024 = ■ × ■
4048387341127
5908123033717
8188796656318
2013214880557

1807082088687 3968599 4553449
4048059516561 9459597 8646735

Example: (n=2048 bits)
classically ~7×10^{15} years
quantum ~100 seconds

8188796656318 5551572 9990445
2013214880557 43 99

Breaking RSA and elliptic curve signatures

Classical:

$$O\left(\exp\left(n^{\frac{1}{3}}(\log n)^{\frac{2}{3}}\right)\right)$$

Quantum:

$$O(n^2 \log n)$$

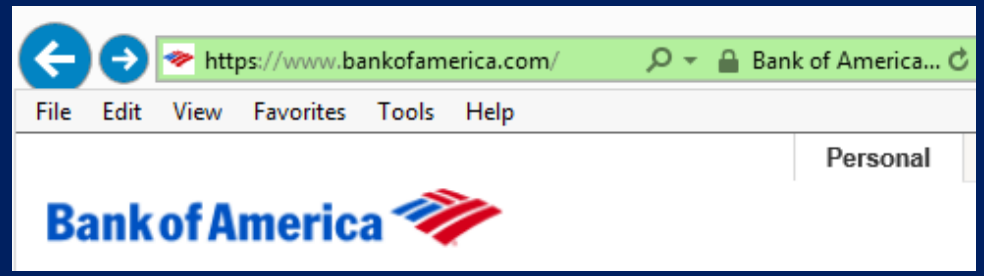

Machine learning

The Problem in Artificial Intelligence

- How do we make computers that *see*, *listen*, and *understand*?

- **Goal**: Learn <u>complex representations</u> for tough AI problems

- Challenges and
  - Terabytes of s of limited data
  - No “silver bu
  - Good new re 5-10 years

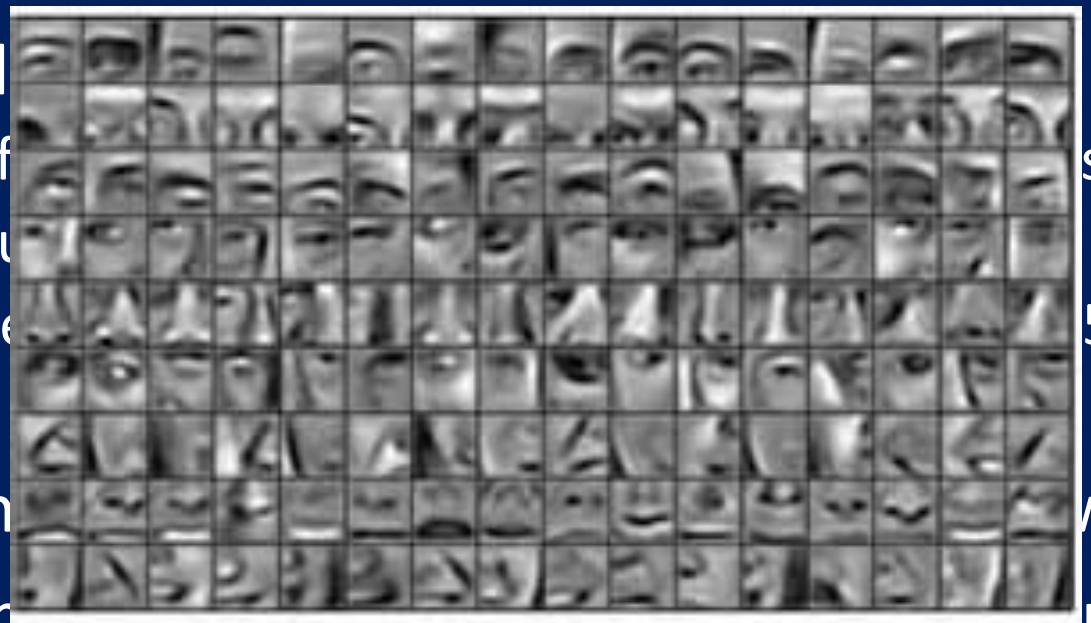

- Can we auton w and high levels?
- Does quantum offer new representations? New training methods?

Deep networks learn complex representations

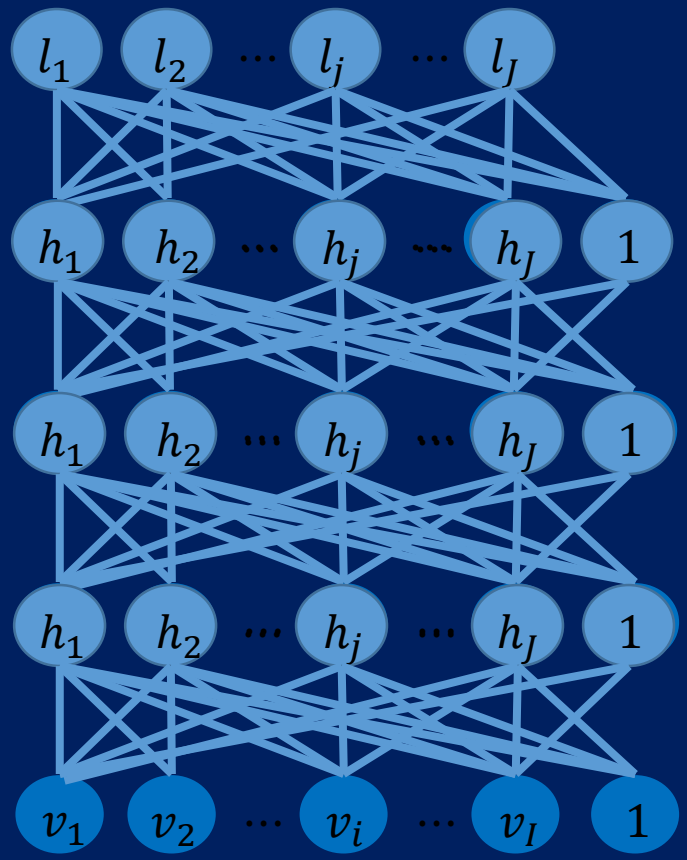

Difficult to specify exactly

Deep networks learn these from data without explicit labels

Analogy: layers of visual processing in the brain

What are the primary challenges in learning?

- Desire: learn a complex representation (e.g., full Boltzmann machine)
- Intractable to learn fully connected graph ⇒
  - Pretrain layers?
  - Learn simpler graph with faster train time?

Can we learn a more complex representation on a quantum computer?

- Desire: efficient computation of true gradient
- Intractable to learn actual objective ⇒
  - Approximate the gradient?

Can we learn the actual objective (true gradient) on a quantum computer?

- Desire: training time close to linear in number of training examples
- Slow training time ⇒
  - Build a big hammer?
  - Look for algorithmic shortcuts?

Can we speedup model training on a quantum computer?

Training RBM - Classical

```
for each epoch                    //until convergence
     for i=1:N                    //each training vector
          CD(V_i, W)              //CD given sample V_i and
                                    parameter vector W
          dLdW += dLdW            //maintain running sum
     end
     W = W + (λ/N) dLdW           //take avg step
end
```

CD Time: # Epochs x # Training vectors x # Parameters

ML Time: # Epochs x # Training vectors x (# Parameters)^2 x $\mathbf{2}^{|v|+|h|}$

Training RBM - Quantum

```
for each epoch                    //until convergence
     for i=1:N                    //each training vector
          qML(V_i, W)             //qML: Use q. computer tp
                                    Approx. to sample P(v,h)
          dLdW += dLdW            //maintain running sum
     end
     W = W + (λ/N)                //take avg step
end
```

qML Time ~ # Epochs x # Training vectors x # Parameters !!!

qML Size (# qubits) for one call ~ |v| + |h| + K, K≤33

Quantum simulation

What does quantum simulation do?

Physical Systems

Quantum Chemistry

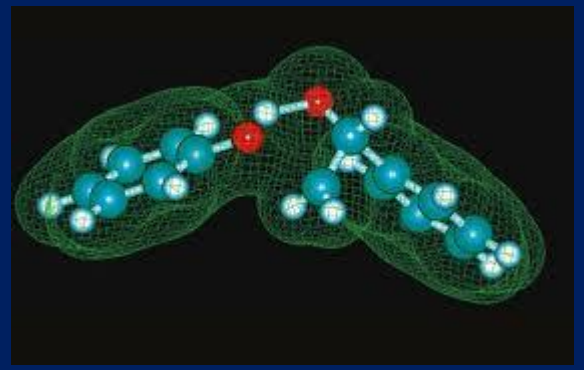

Superconductor Physics

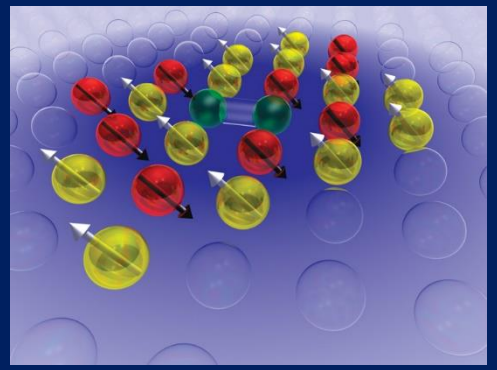

Quantum Field Theory

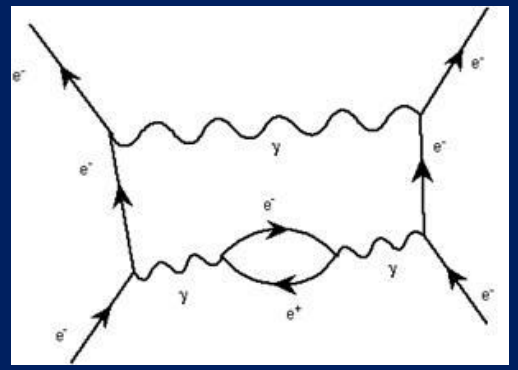

Computational Applications

Emulating Quantum Computers

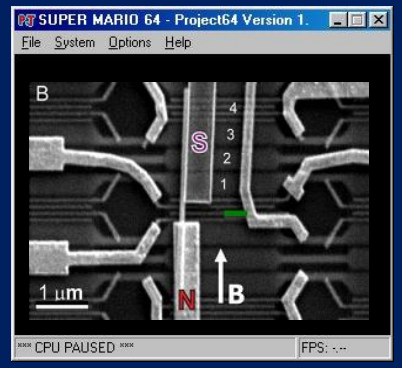

Linear Algebra

Differential Equations

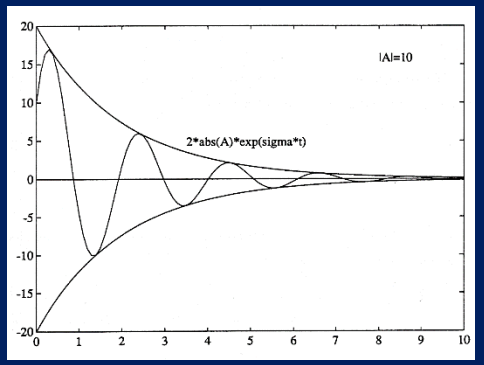

Quantum simulation

Particles can either be spinning clockwise (down) or counterclockwise (up)

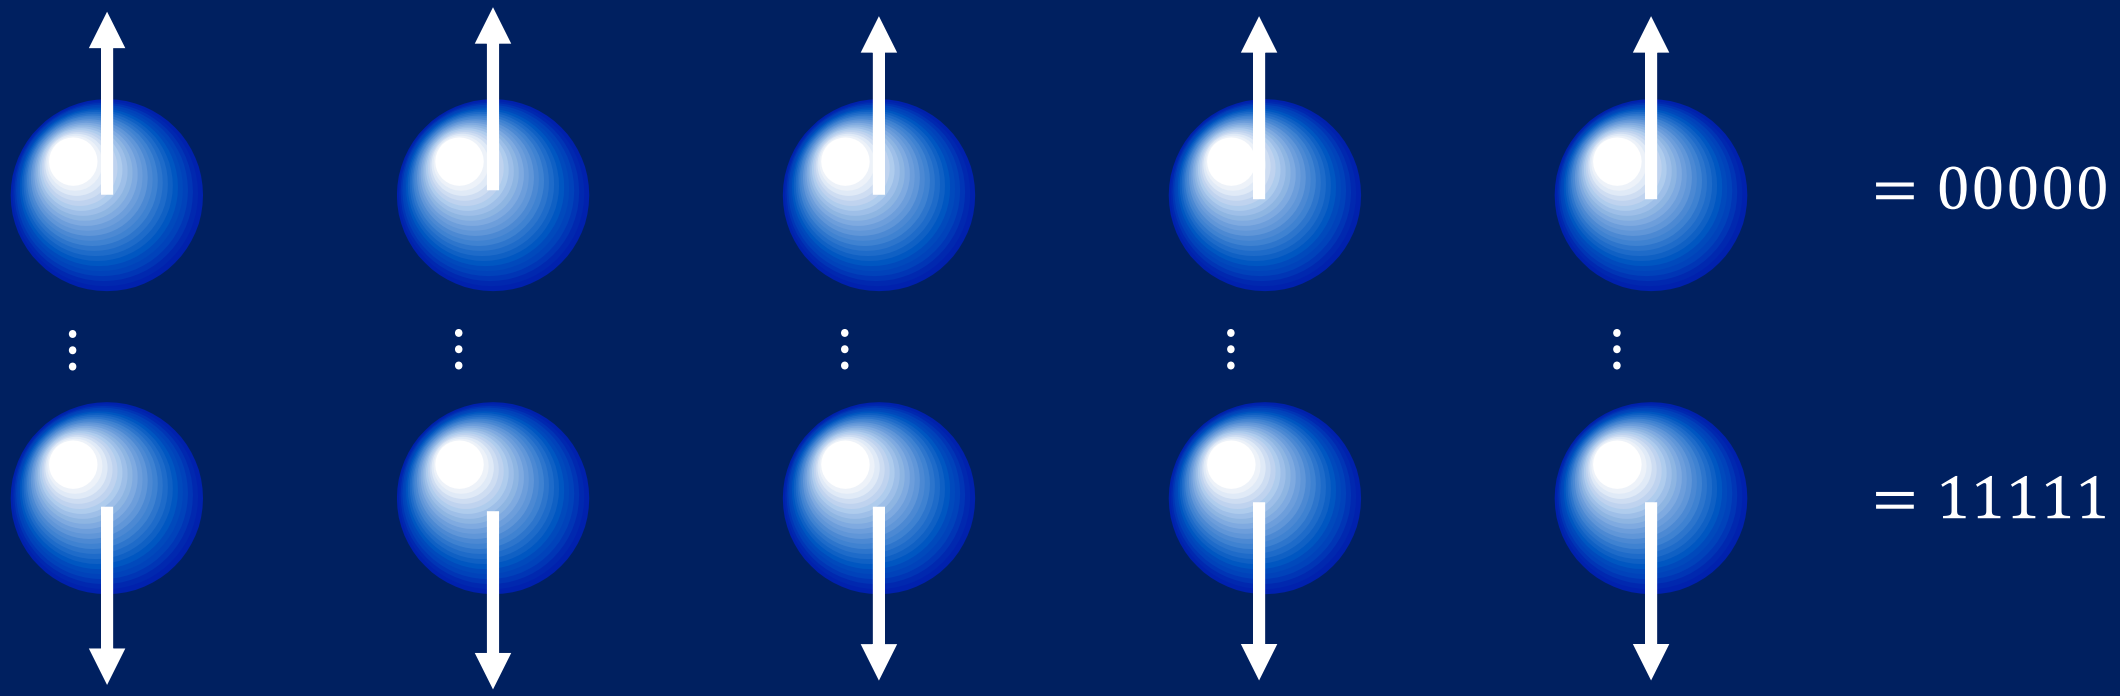

There are 2^5 possible orientations in the quantum distribution.
Cannot store this in memory for 100 particles.

Quantum Simulation for Quantum Chemistry

Ultimate problem:

Simulate molecular dynamics of *larger* systems or to *higher accuracy*

Want to solve system *exactly*

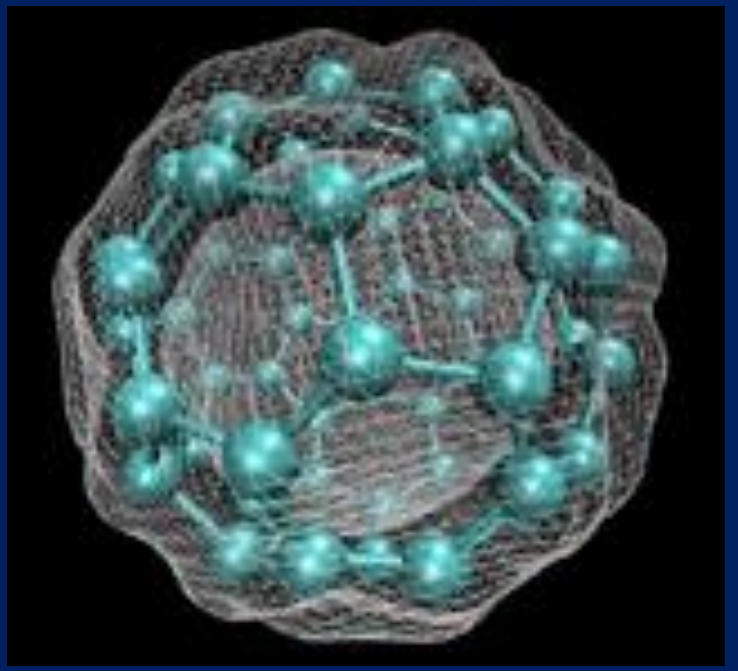

Current solution:

33% supercomputer usage dedicated to chemistry and materials modeling

Requires simulation of exponential-size Hilbert space

Limited to 50-70 spin-orbitals classically

Quantum solution:

Simulate molecular dynamics using *quantum simulation*

Scales to 100s spin-orbitals using only 100s qubits

Runtime recently reduced from $O(N^{11})$ to $O(N^4) - O(N^6)$

Quantum Chemistry

$$H = \sum_{pq} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s$$

Can quantum chem
computer: Dave We
Hastings, Matthias Tr

The Trotter Step Size Required for Accurate Quantum Simulation of Quantum Chemistry
David Poulin, M. B. Hastings, Dave Wecker, Nathan Wiebe, Andrew C. Doherty, Matthias Troyer

m Chemistry: M. B.

As quantum
computers
appear fea
applicatior
frequently
simulating
of molecul
computatic
perform qu
the quantu
molecule t
exactly. We
increase in
required in
executed is n
quantum computatic
problems, drastic alg

tter-Suzuki based
on a quantum
nations are
mic in the
not require
ny operations
n the parallel
increase in
rder in the
error at given
e Hamiltonian
uki timestep. All
tion and

Ferredoxin (Fe_2S_2) used in many metabolic reactions including energy transport in photosynthesis

- *Intractable on a classical computer*
- *First paper: ~300 million years to solve*
- *Second paper: ~30 years to solve (10^7 reduction)*
- *Third paper: ~300 seconds to solve (another 10^3 reduction)*

http://arxiv.org/abs/1406.4920

Quantum Chemistry

$$H = \sum_{pq} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s$$

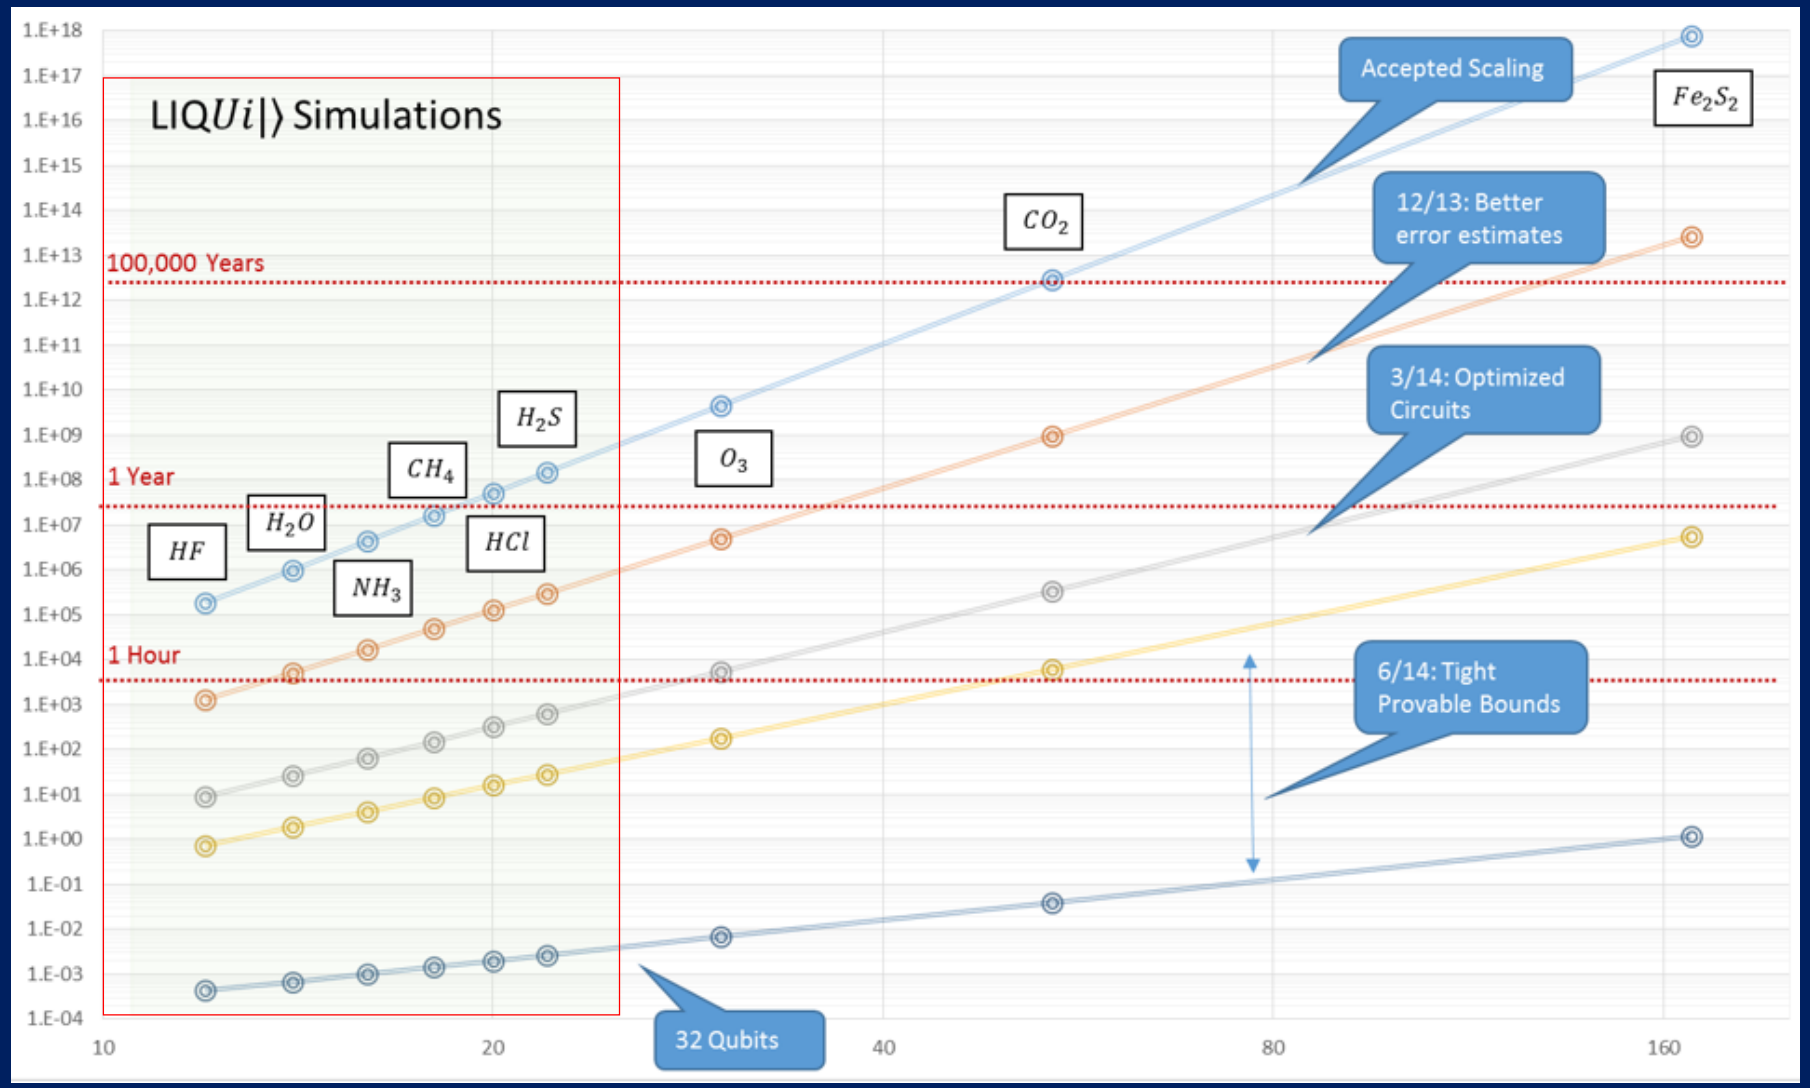

Application: Nitrogen Fixation

Ultimate problem:

Find catalyst to convert nitrogen to ammonia at room temperature

Reduce energy for conversion of air to fertilizer

Current solution:

Uses Haber process developed in 1909

Requires high pressures and temperatures

Cost: 3-5% of the worlds natural gas production (1-2% of the world's annual energy)

Quantum solution:

~ 100-200 qubits: Design the catalyst to enable inexpensive fertilizer production

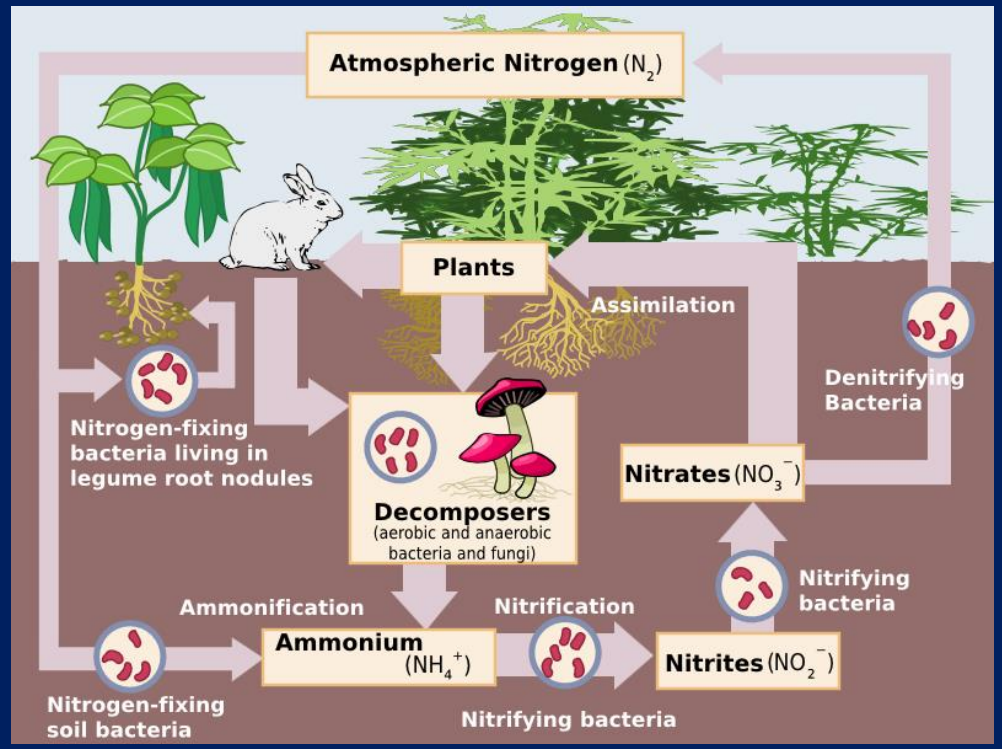

Application: Carbon Capture

Ultimate problem:

Find catalyst to extract carbon dioxide from atmosphere

Reduce 80-90% of emitted carbon dioxide

Current solution:

Capture at point sources

Results in 21-90% increase in energy cost

Quantum solution:

~ 100-200 qubits: Design a catalyst to enable carbon dioxide extraction from air

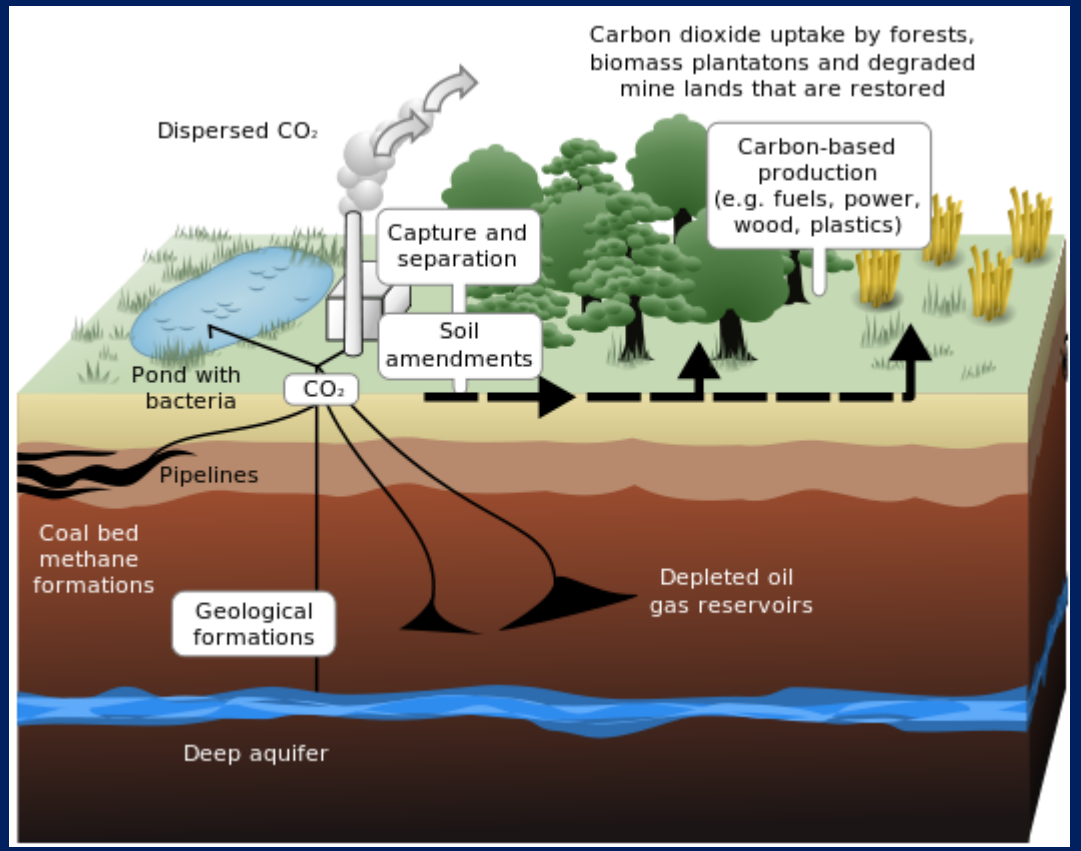

Quantum Algorithm Opportunities

Quantum simulation

- Extend q. chem. method to solid state materials
- E.g., high temp. superconductivity
- ~ 2000 qubits; linear or quad. scaling

Machine learning

- Clustering, regression, classification
- Better model training
- Can we harness interference to produce better inference models?

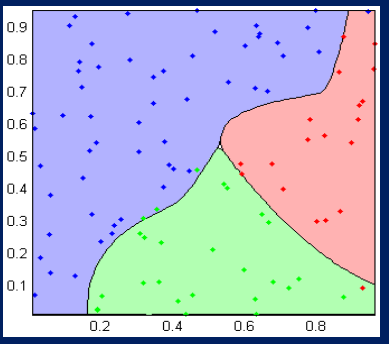

Cryptography

- RSA, DSA, elliptic curve signatures, El-Gamal
- What questions should we pose to a quantum computer?

Requirements for Quantum Computation

Quantum algorithms:

Design real-world quantum algorithms for small-, medium- and large-scale quantum computers

Quantum hardware architecture:

Architect a scalable, fault-tolerant, and fully programmable quantum computer

Quantum software architecture:

Program and compile complex algorithms into optimized, target-dependent (quantum and classical) instructions

Quantum Computing through the Ages

"Age of Scalability"

"Age of Algorithms"

"Age of Quantum Error Correction"

"Age of Quantum Feedback"

"Age of Measurement"

"Age of Entanglement"

"Age of Coherence"

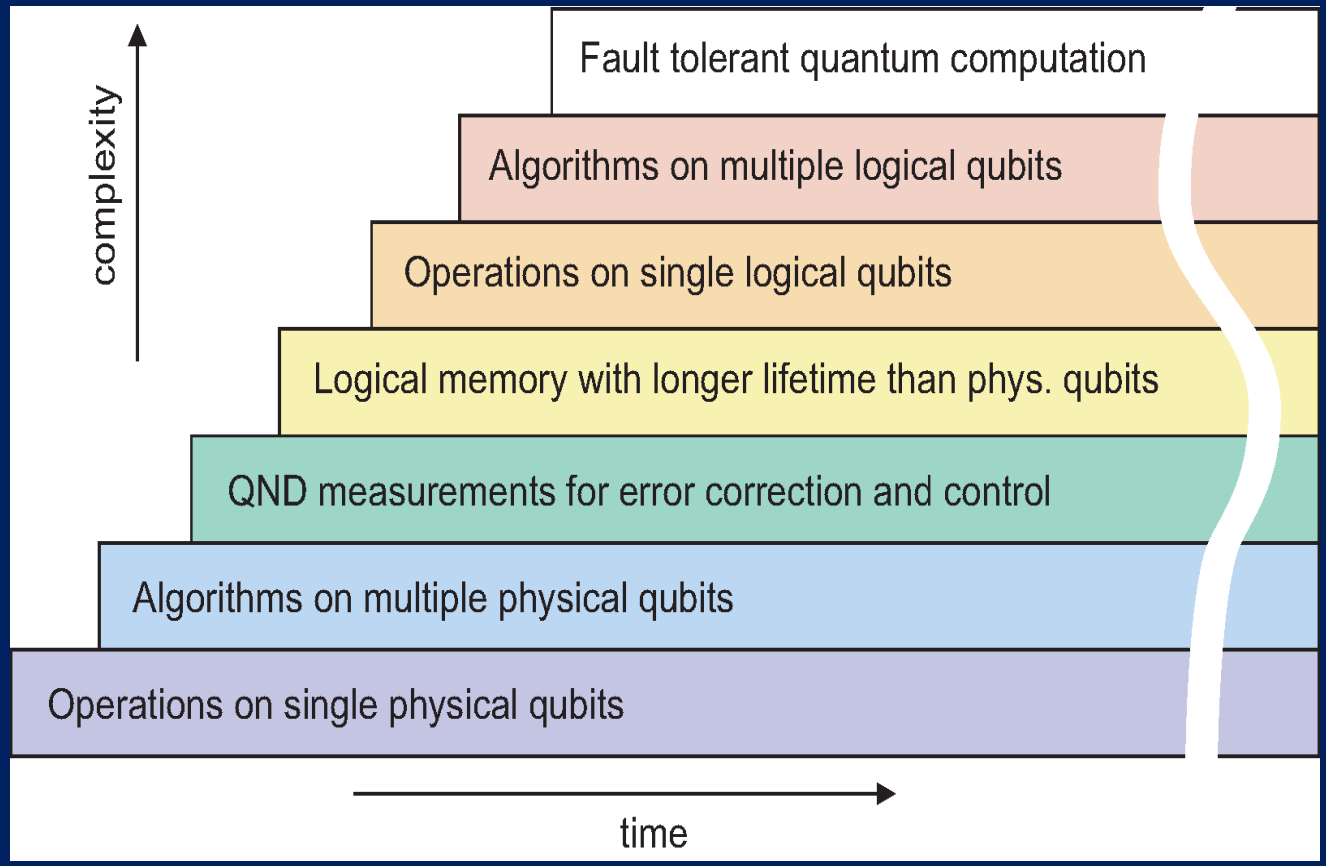

"We" are ~ here

M. Devoret and R. Schoelkopf, Science (2013)

Quantum Hardware Technologies

Ion traps

NV centers

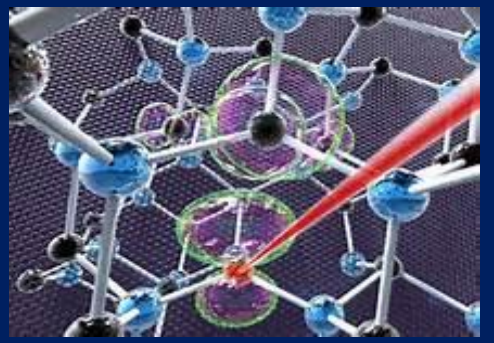

Super-conductors

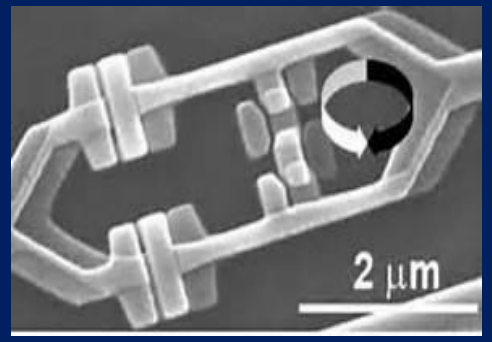

Quantum dots

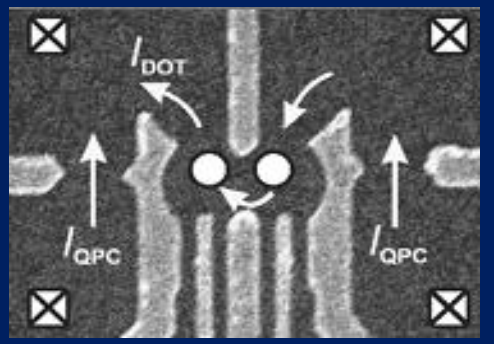

Linear optics

Topological

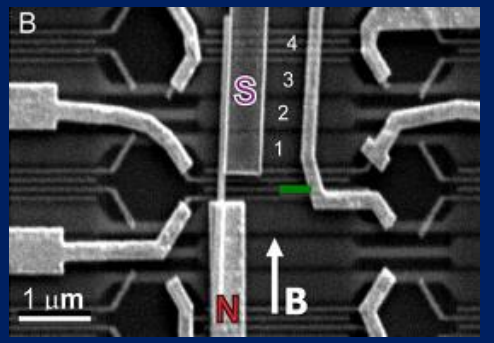

Classical Error Correction

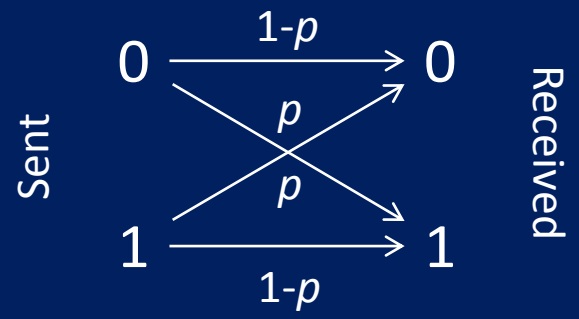

Probability p of having a bit flipped

Repetition code: redundantly encode, majority voting

$0 \rightarrow 000$

$1 \rightarrow 111$

Reduces classical error rate to $3p^2 - 2p^3$

Can we do this for quantum computing? Some reasons to think **no**:

- "No cloning" theorem
- Errors are continuous (or are they?)
- Measurements change the state

Thanks to Rob Schoelkopf

Different Error Correction Architectures

Standard QEC

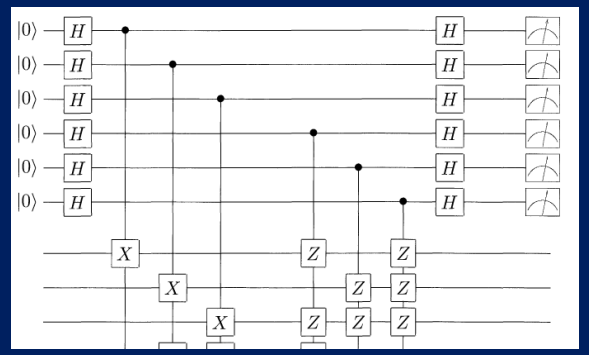

Surface Code

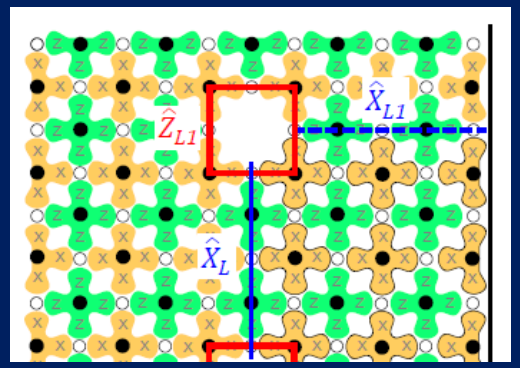

Modular Approach

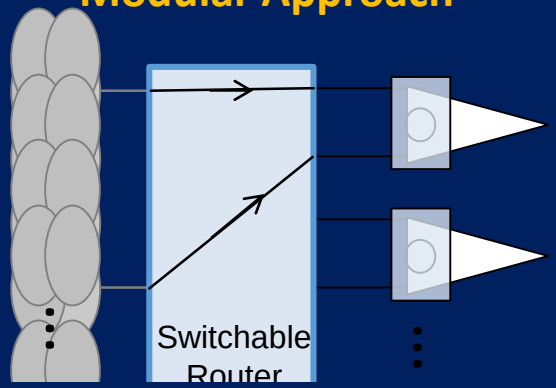

Overhead required in known schemes:

1,000 – 10,000 actual qubits for every logical!!

(+ concatenation!)

- threshold ~ 10^{-4}
- many ops., syndromes per QEC cycle

- threshold ~ 1%
- large system to see effects?

- good local gates (10^{-4}?) remote gates fair (90%?)
- then construct QEC as software layer?

Thanks to Rob Schoelkopf

3 dimensions

$|\psi\rangle \rightarrow |\psi\rangle$ Bosons

$|\psi\rangle \rightarrow -|\psi\rangle$ Fermions

$|\psi\rangle \rightarrow |\psi\rangle$

2 dimensions

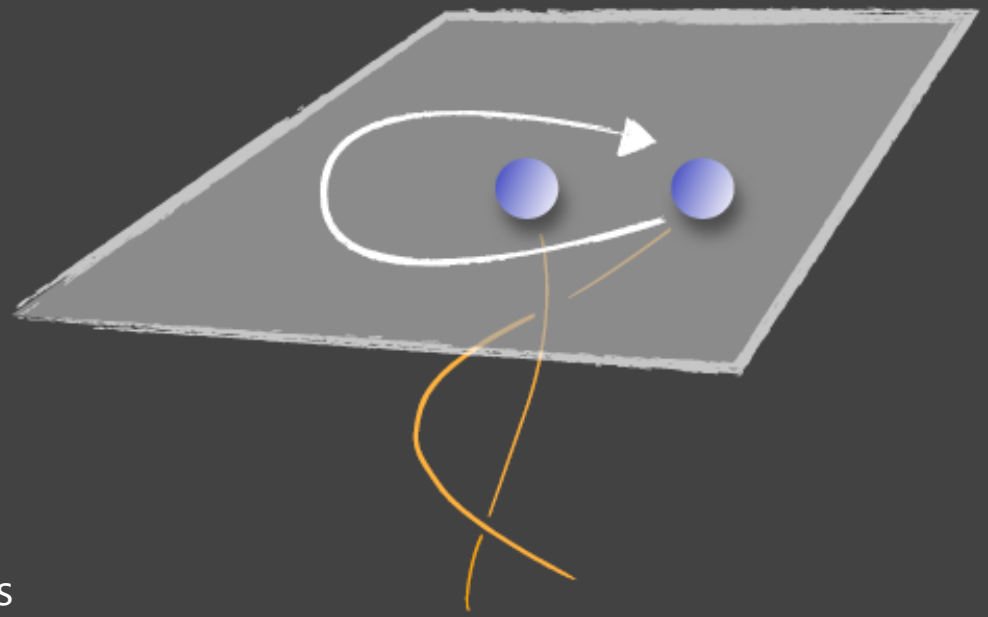

$|\psi_1\rangle \rightarrow |\psi_2\rangle$

Image courtesy of Charlie Marcus

The man and his particle

Chi l'ha visto ?

Ettore Majorana, ordinario di fisica teorica all' Università di Napoli, è misteriosamente scomparso dagli ultimi di marzo. Di anni 31, alto metri 1,70, snello, con capelli neri, occhi scuri, una lunga cicatrice sul dorso di una mano. Chi ne sapesse qualcosa è pregato di scrivere al R. P. E. Majanecci, Viale Regina Margherita 66 - Roma.

1938

Image courtesy of Leo Kouwenhoven

Majorana fermions

(particle) $\gamma^{\dagger} = \gamma$ (anti-particle)

→ Electric charge is zero
→ Energy is zero
→ Everything is zero (except mass)

How to measure the "nikes" particle?

→ That's why it was not yet detected!

Image courtesy of Leo Kouwenhoven

New directions in the pursuit of Majorana fermions in solid state systems

Jason Alicea[1]

[1]*Department of Physics and Astronomy, University of California, Irvine, California 92697*

(Dated: February 8, 2012)

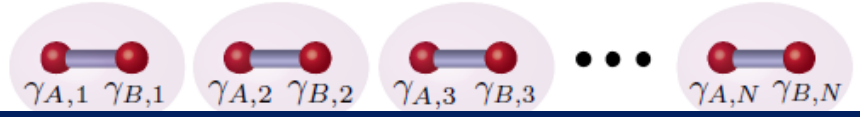

Topology provides natural immunity to noise!

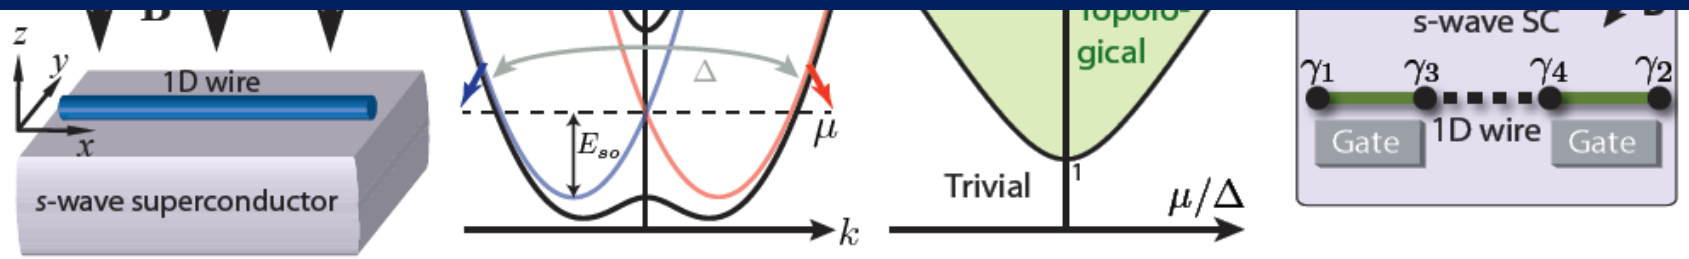

FIG. 6. (a) Basic architecture required to stabilize a topological superconducting state in a 1D spin-orbit-coupled wire. (b) Band structure for the wire when time-reversal symmetry is present (red and blue curves) and broken by a magnetic field (black curves). When the chemical potential lies within the field-induced gap at $k = 0$, the wire appears 'spinless'. Incorporating the pairing induced by the proximate superconductor leads to the phase diagram in (c). The endpoints of topological (green) segments of the wire host localized, zero-energy Majorana modes as shown in (d).

nature physics

ARTICLES

PUBLISHED ONLINE: 13 FEBRUARY 2011 | DOI: 10.1038/NPHYS1915

Non-Abelian statistics and topological quantum information processing in 1D wire networks

Jason Alicea[1]*, Yuval Oreg[2], Gil Refael[3], Felix von Oppen[4] and Matthew P. A. Fisher[3,5]

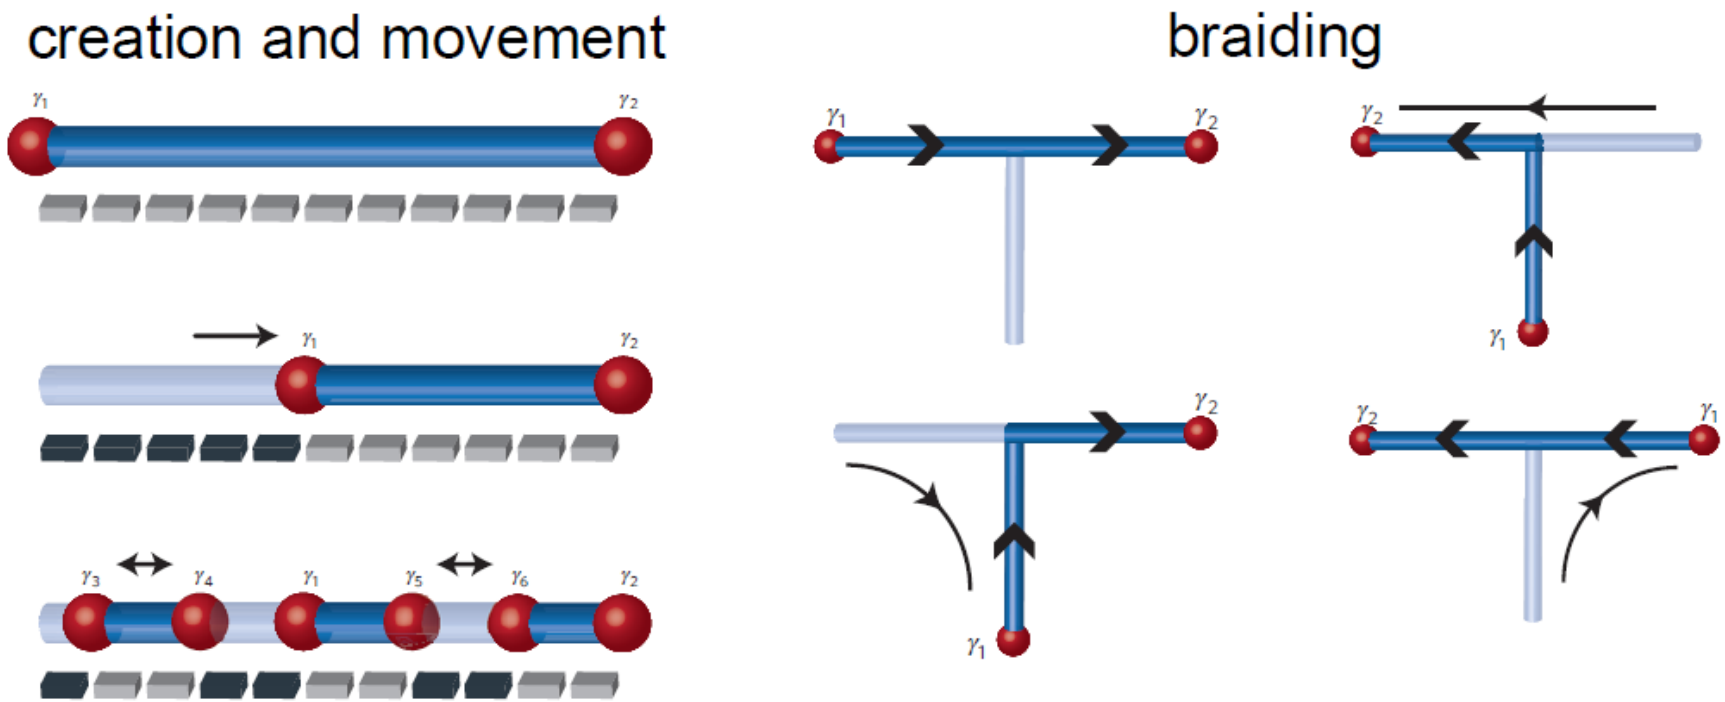

Epitaxial growth of InAs (or InSb) nanowires

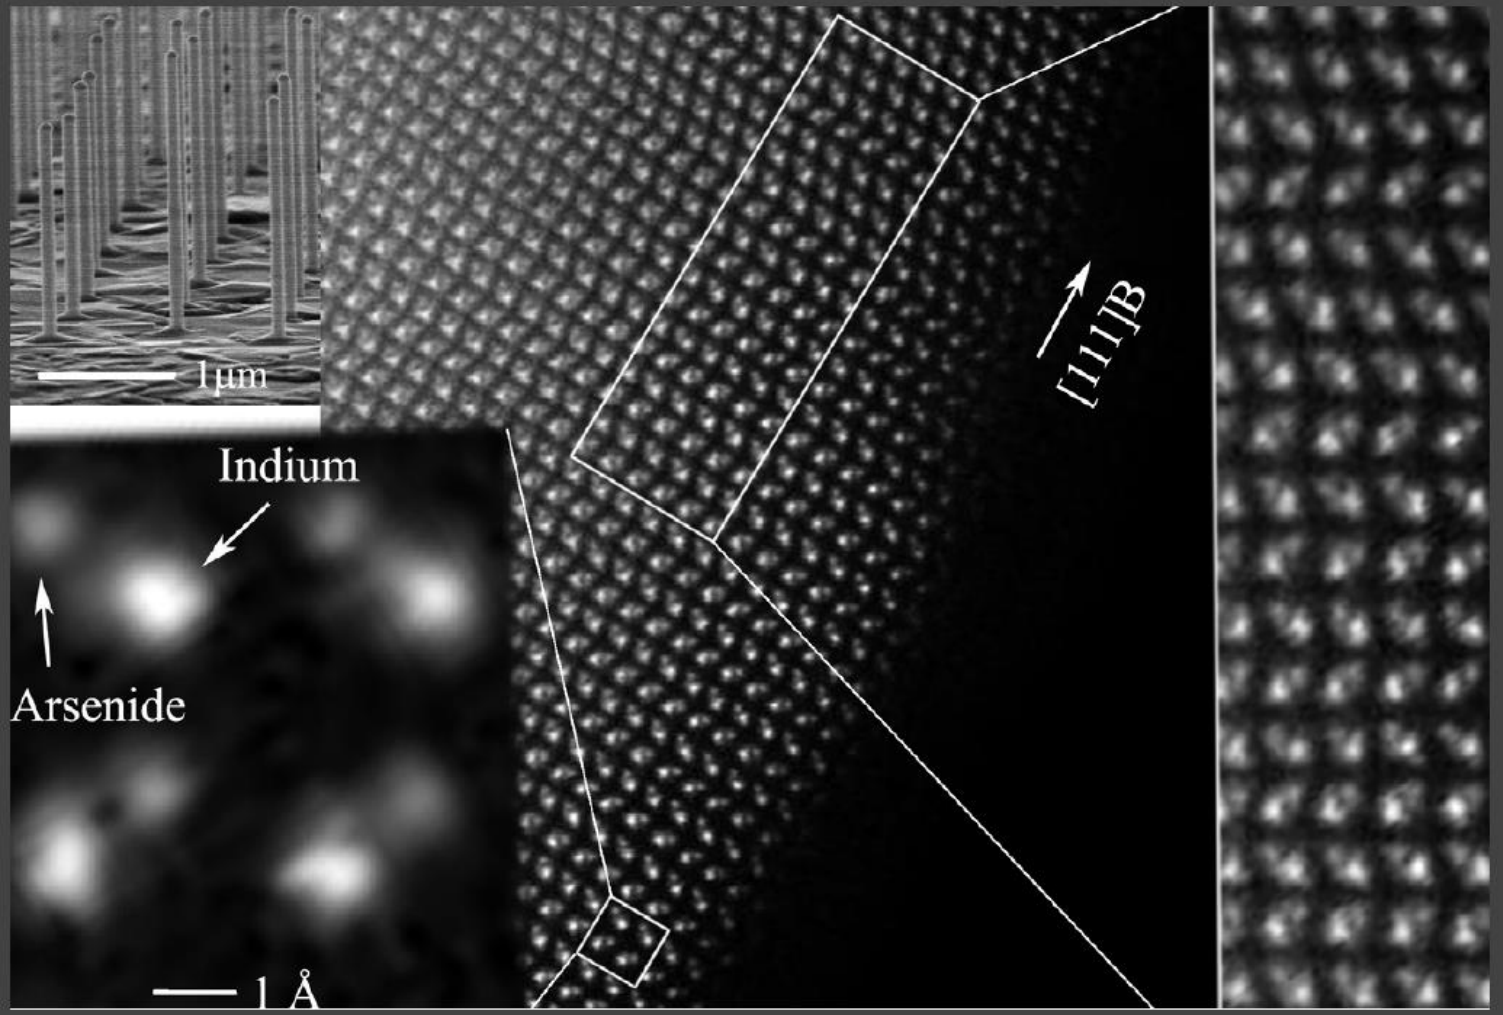

P. Krogstrup, J. Nygård, Univ. of Copenhagen

Signatures of Majorana Fermions in Hybrid Superconductor-Semiconductor Nanowire Devices

V. Mourik,[1*] K. Zuo,[1*] S. M. Frolov,[1] S. R. Plissard,[2] E. P. A. M. Bakkers,[1,2] L. P. Kouwenhoven[1†]

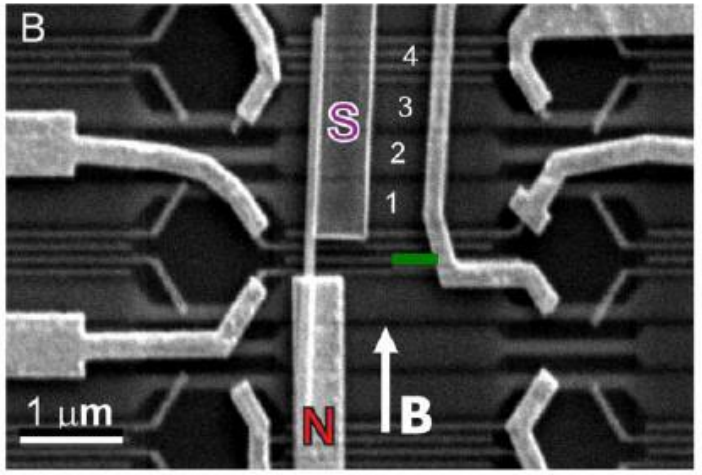

Hardware provides error correction:
Only ~10-100s for every logical???

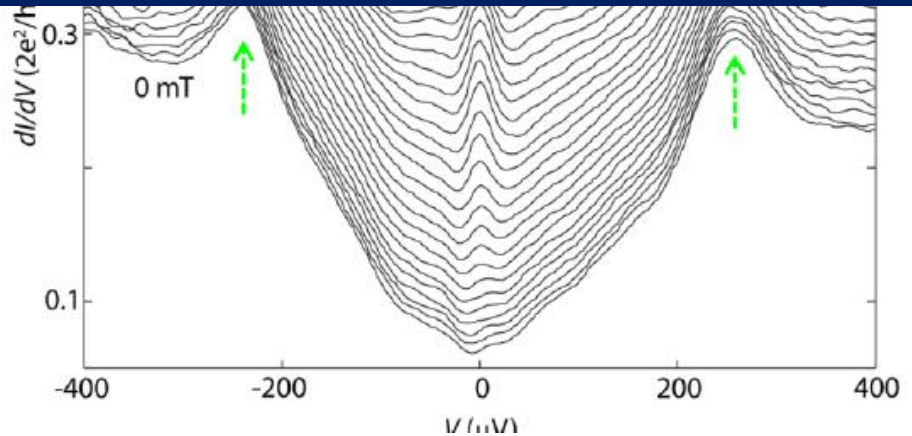

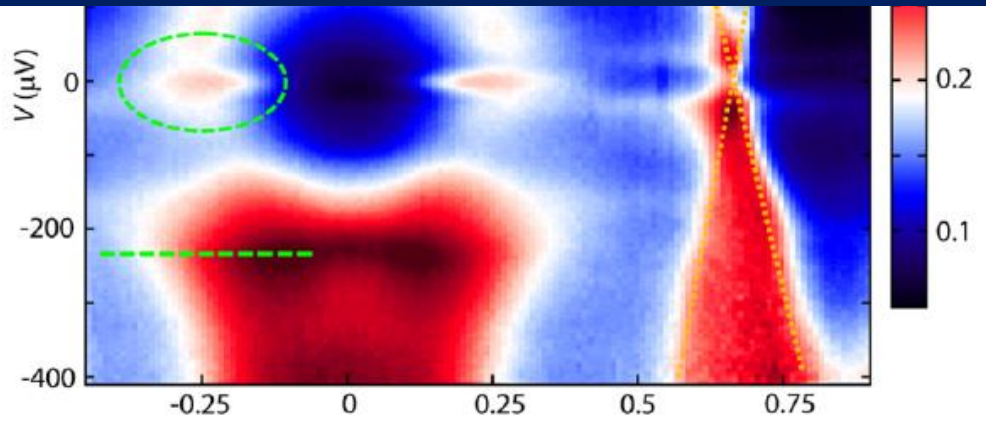

A Software Architecture for Quantum Computing

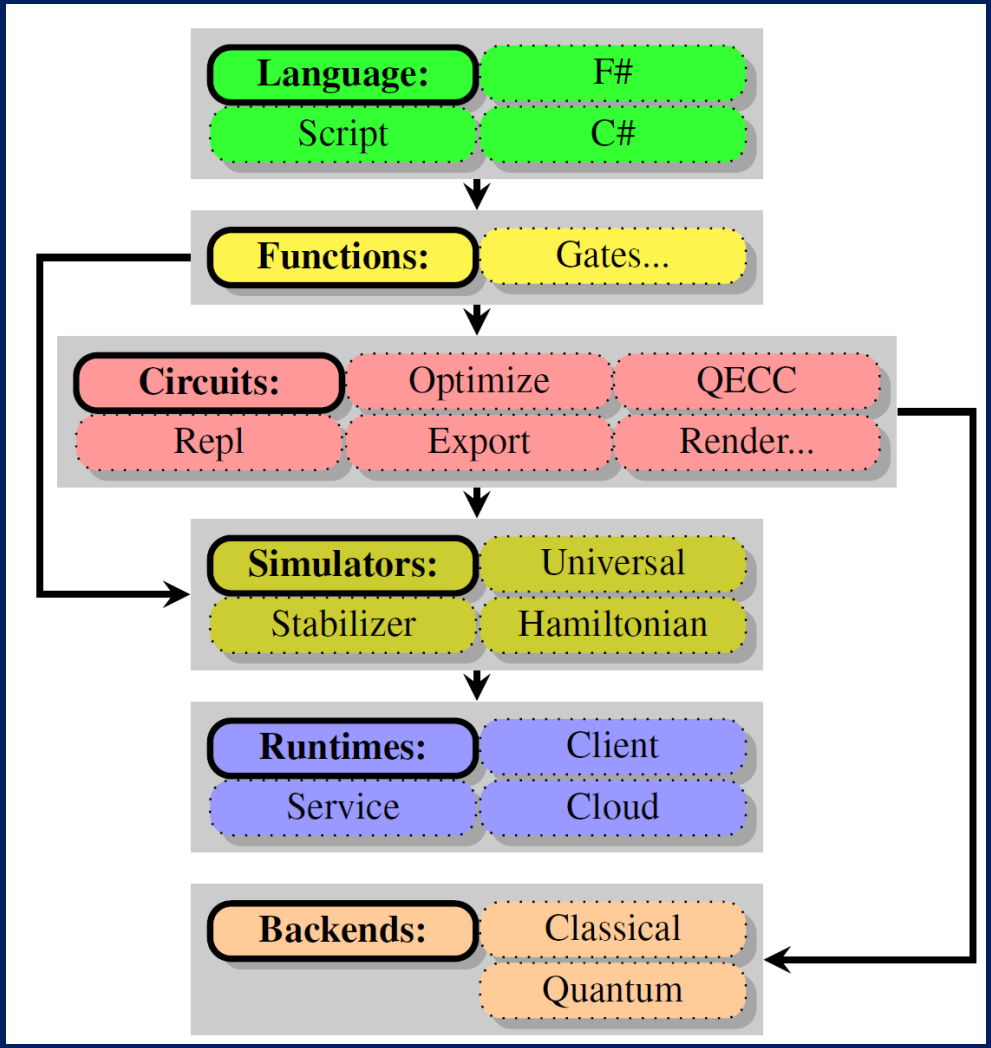

The LIQ$Ui|\rangle$ platform
[Wecker, Svore, 2014]

High-level Quantum Algorithm

Quantum Gates
Quantum Function Implementation

Quantum Circuit Decomposition
Quantum Circuit Optimization
Quantum Error Correction

Simulation

Runtime Environments

Target-dependent Optimization
Layout, Scheduling, Control

Target-dependent Representation
Classical Control/Quantum Machine Instructions

Shor's algorithm: 4 bits ≅ 8200 gates

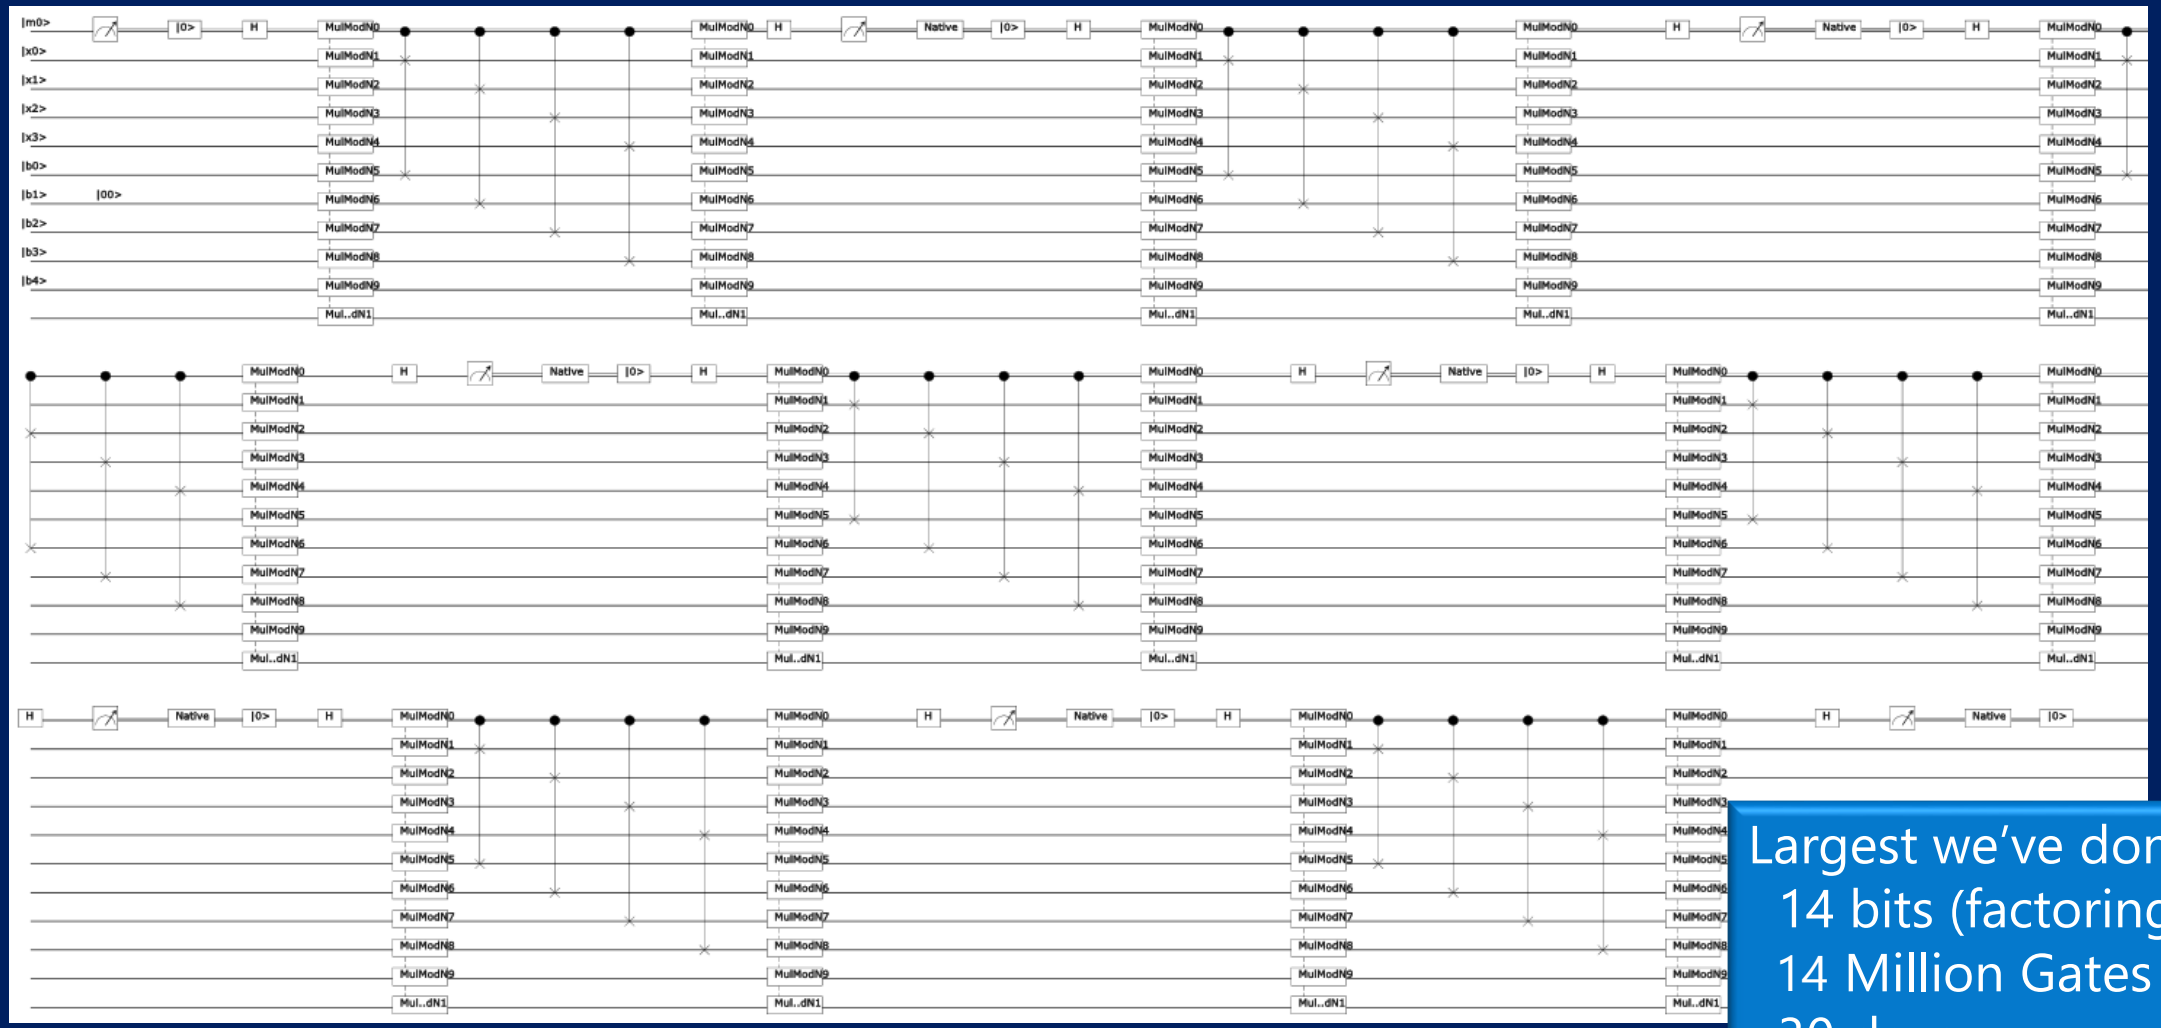

Circuit for Shor's algorithm using 2n+3 qubits – Stéphane Beauregard

Largest we've done:
- 14 bits (factoring 8193)
- 14 Million Gates
- 30 days

Shor's algorithm: Modular Adder

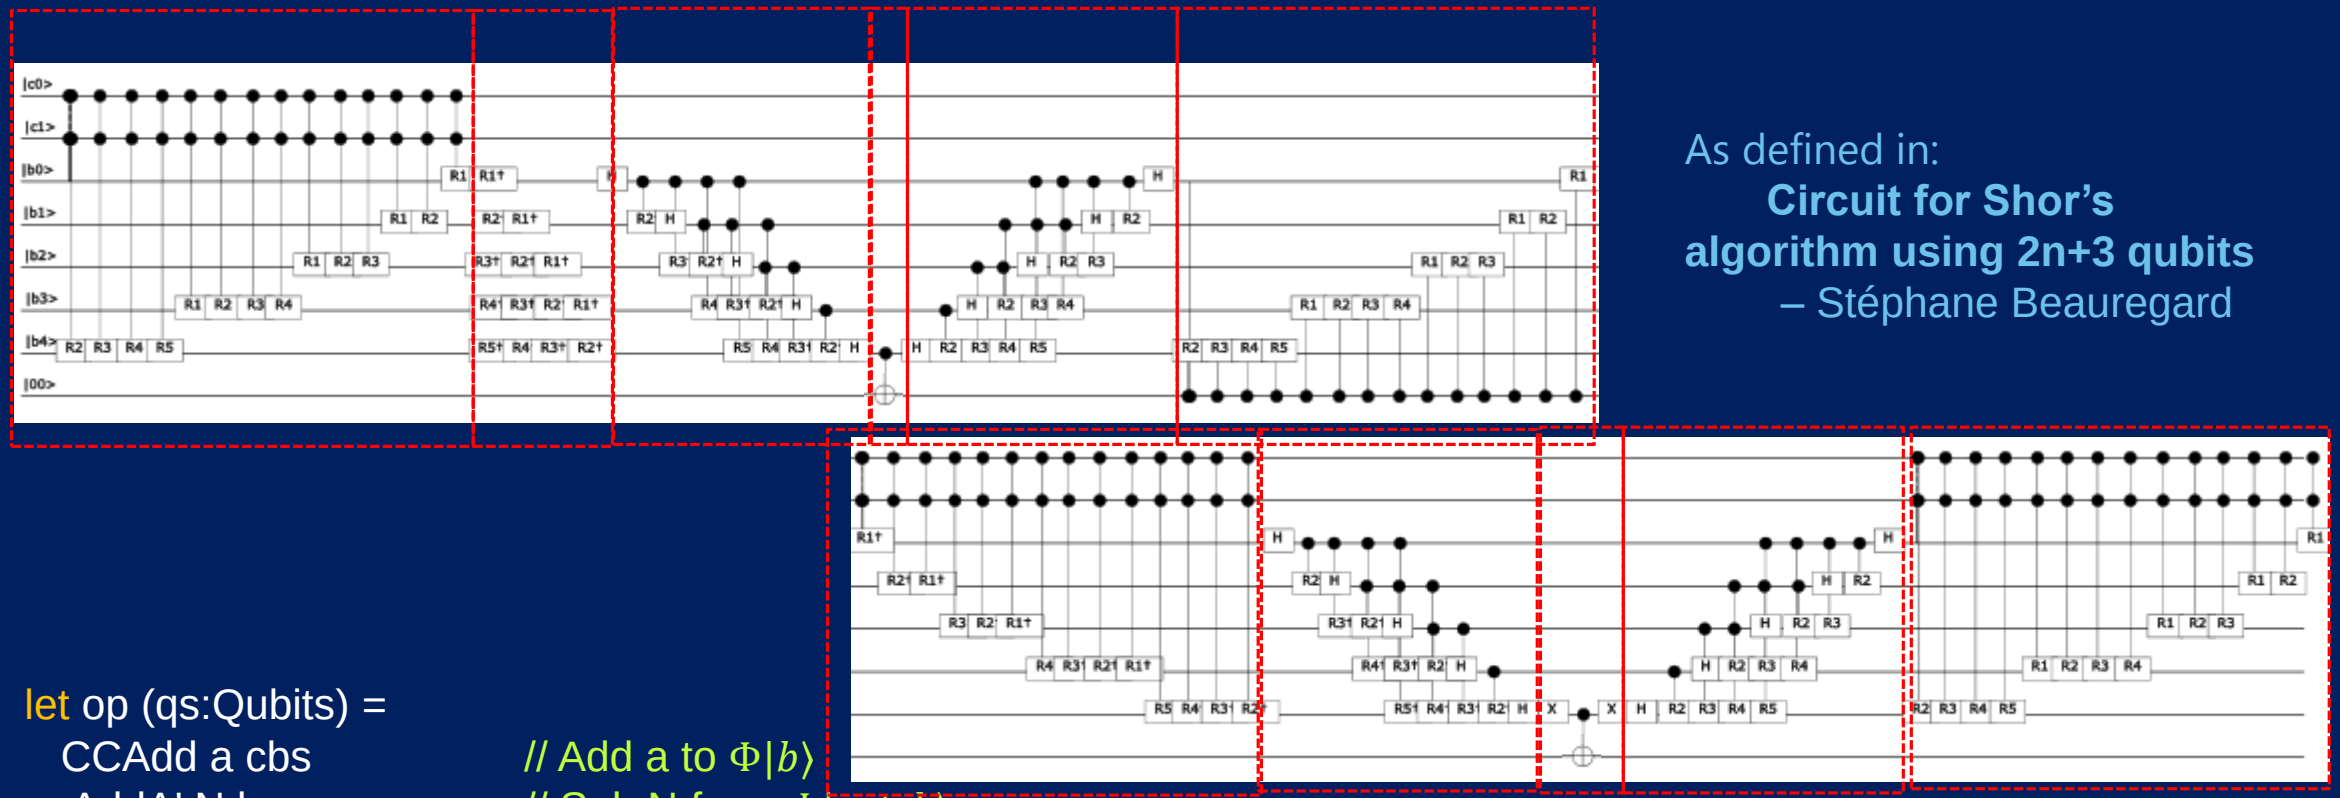

As defined in:
Circuit for Shor's algorithm using 2n+3 qubits
– Stéphane Beauregard

```
let op (qs:Qubits) =
    CCAdd a cbs
     AddA' N bs
     QFT' bs
     CNOT [bMx;anc]
     QFT bs
     CAddA N (anc :: bs)
     CCAdd' a cbs
```

// Add a to $\Phi|b\rangle$
// Sub N from $\Phi|a+b\rangle$
// Inverse QFT of $\Phi|a+b-N\rangle$
// Save top bit in Ancilla
// QFT of a+b-N
// Add back N if negative
// Subtract a from $\Phi|a+b\ mod\ N\rangle$

QFT' bs
X [bMx]
CNOT [bMx;anc]
X [bMx]
QFT bs
CCAdd a cbs

// Inverse QFT
// Flip top bit
// Reset Ancilla to $|0\rangle$
// Flip top bit back
// QFT back
// Finally get $\Phi|a+b\ mod\ N\rangle$

Simulating Shor's Algorithm

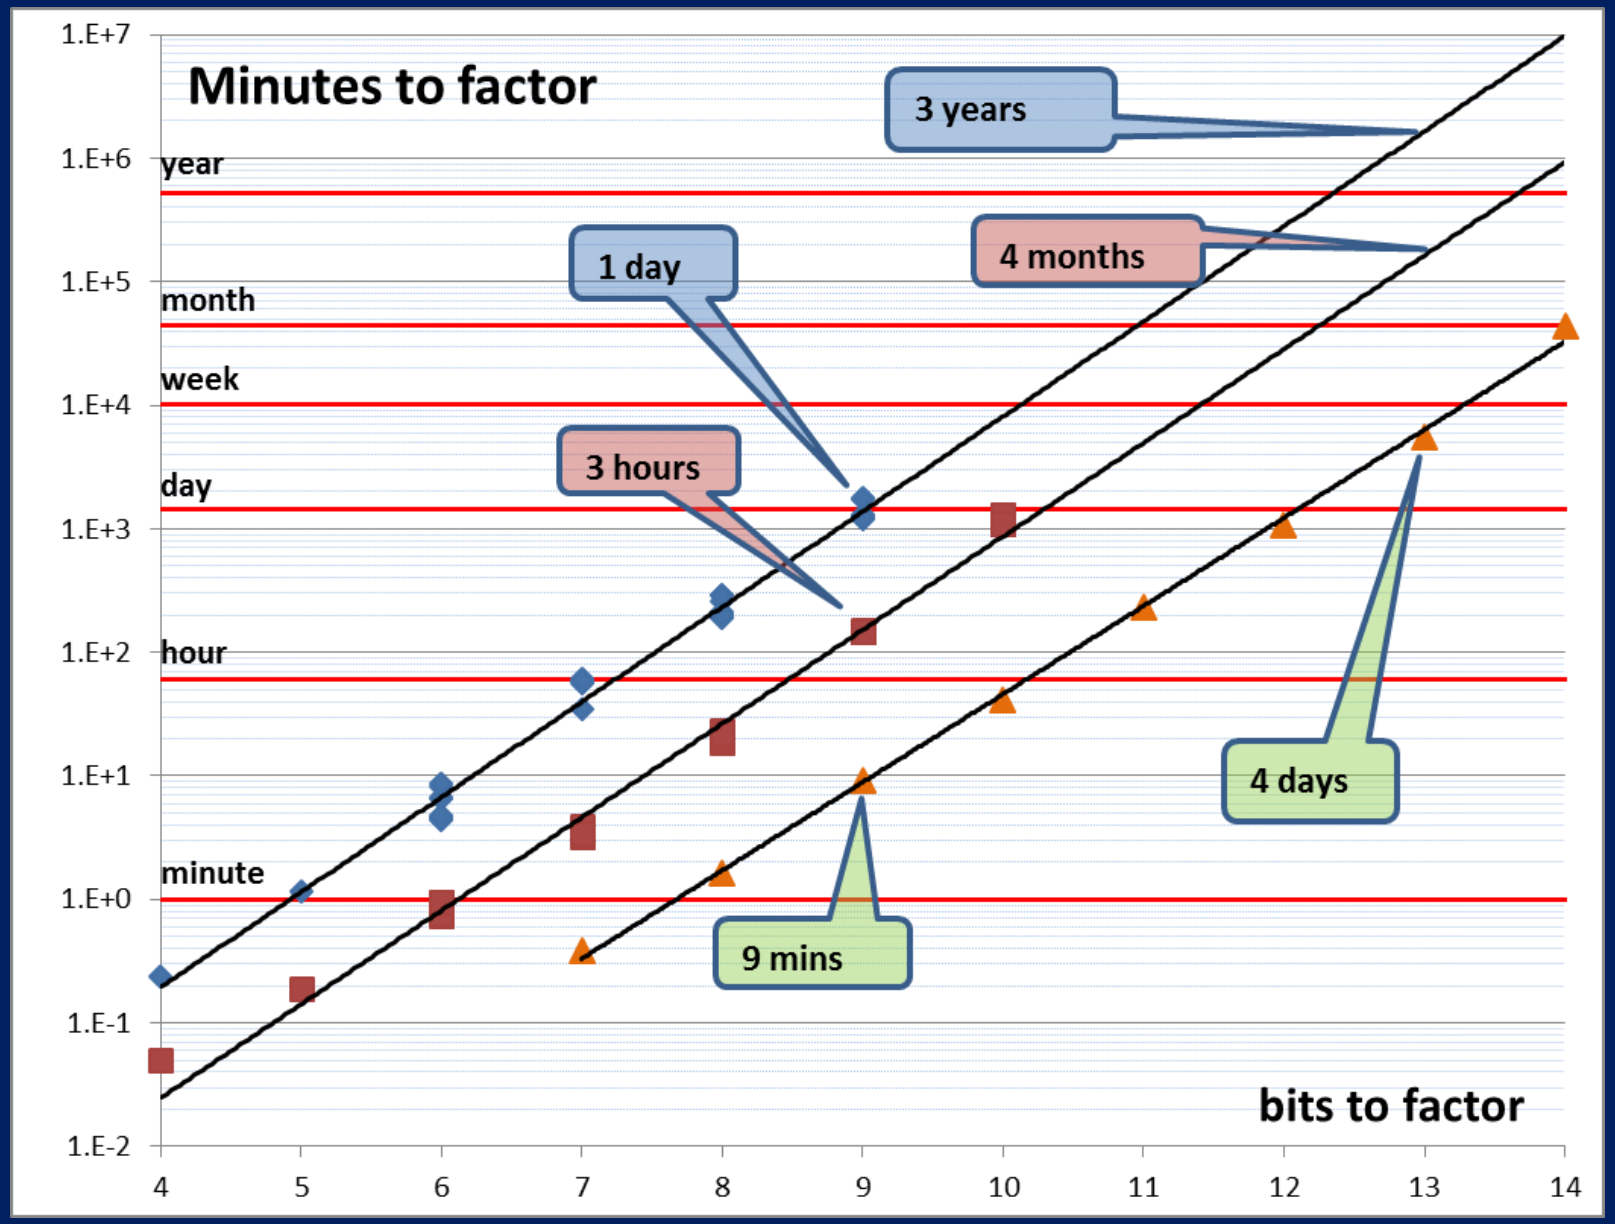

LIQ$Ui|\rangle$ for Compilation onto Hardware

```
let QFT (qs : Qs) =
        let n = qs.Length - 1
        for i = 0 to n do
                let q = qs.[i]
                H q
                for j = (i + 1) to n do
                        let theta = 2.0 * Math.PI /
                                float(1 <<< (j - i + 1))
                        CRz theta qs.[j] q
        for i = 0 to ((n - 1) / 2) do
                SWAP qs.[i] qs.[n - i]

let QftOp  = compile QFT
let QftOp' = adjoint QftOp
```

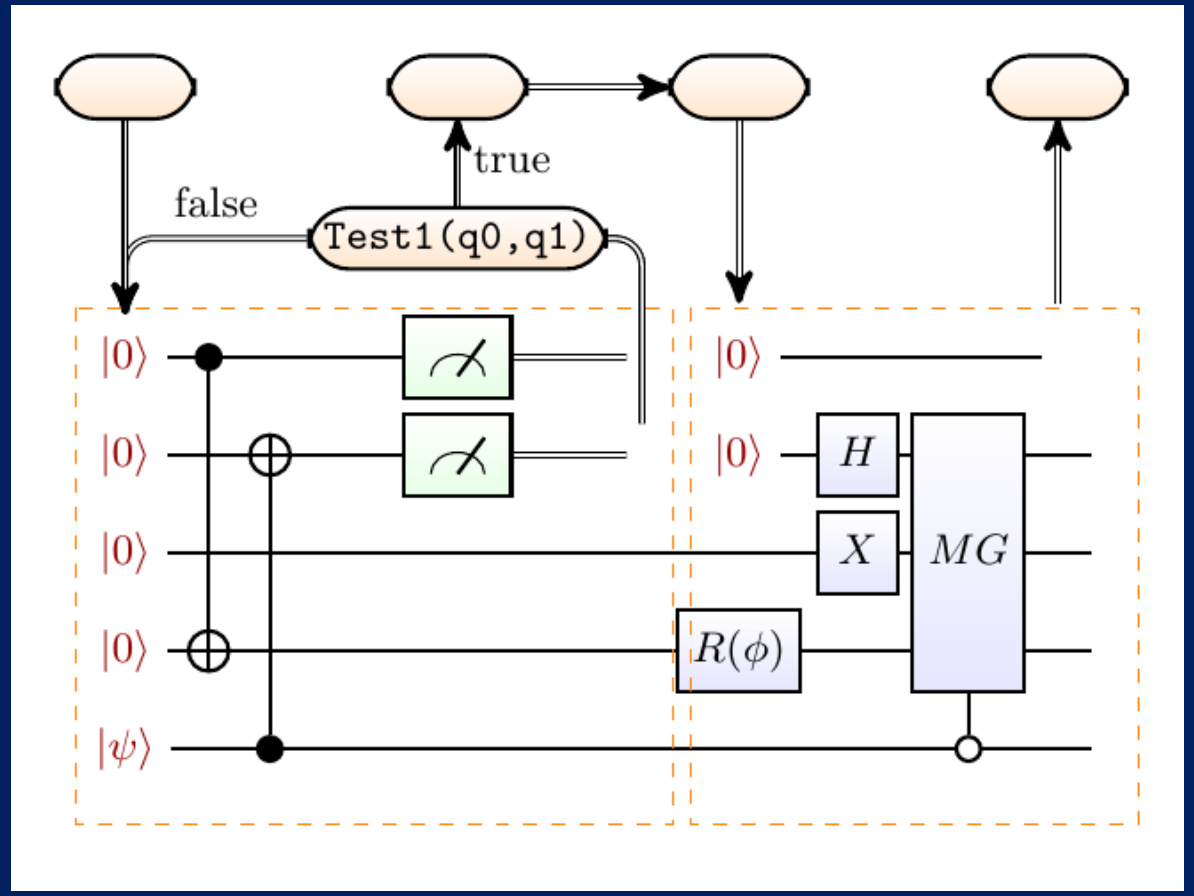

Conclusions

Quantum computers exploit interference and superposition to solve problems.

How big/fast does a quantum computer have to be to have an advantage?

[Boixo, Ronnow et al '13]
[Wecker, Bauer et al '14]

What are the right questions to ask a quantum computer?

[Wiebe, Braun, Lloyd '12]
[Wiebe, Grenade et al '13]

Exponential speedups for *certain* simulation, cryptography, linear algebra problems.

How do you compile, test, and debug quantum algorithms?

[Wiebe, Kliuchnikov'13]
[Bocharov, Gurevich, Svore'13]
[Wecker, Svore Geller' 14]

What other problems does a quantum computer solve better or faster?

QuArC

Doug Carmean

Vadym Klichnikov

Alex Bocharov

Yuri Gurevich

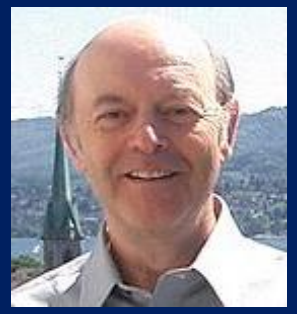

Alan Geller

Dave Wecker

Krysta Svore

Ken Reneris

Martin Roetteler

Burton Smith

Nathan Wiebe

Station Q

Maissam Barkeshli

Michael Freedman

Chetan Nayak

Bela Bauer

Matthew Hastings

Kevin Walker

Parsa Bonderson

Roman Lutchyn

Zhenghan Wang

Meng Cheng

Mike Mulligan

Jon Yard

University Partners

David Reilly
U. Sydney

Charlie Marcus
NBI

Leo Kouwenhoven
Delft

Amir Yacoby
Harvard

Dale Van Harlingen
UIUC

Mike Manfra
Purdue

Chris Palmstrom
UCSB

Bert Halperin
Harvard

Sankar Das Sarma
U Maryland

Matthias Troyer
ETH Zurich

http://research.microsoft.com/groups/quarc/

http://research.microsoft.com/en-us/labs/stationq/

ksvore@microsoft.com