CPSC 447/547 – Introduction to Quantum Computing

Lecture 1 - Overview

Yongshan Ding
Computer Science
Yale Quantum Institute
Fall 2023

Structure of the Course

Four Modules

1. The mathematical formulation of quantum information and computation
2. Building blocks of a universal quantum computer
3. Some quantum algorithms and their applications
4. Practicality of fault-tolerant quantum computation

A Computer Scientist's Guide to Quantum Computing

Theory:
Algorithms and complexity

Systems:
Programming language, compiler, computer architecture

Quantum science meets computer science

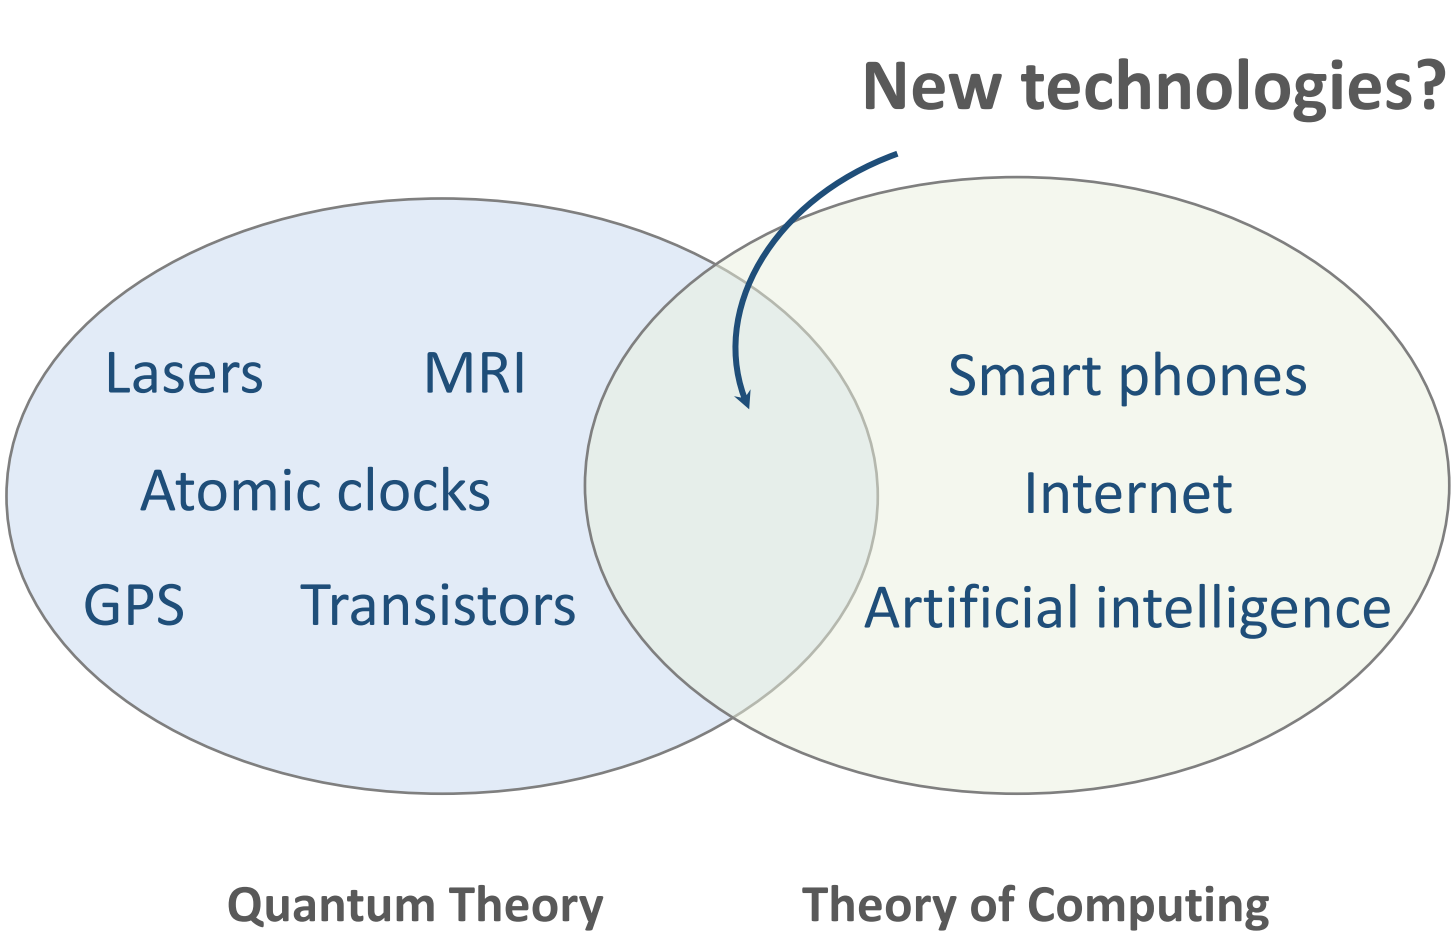

New technologies?

- Quantum Communication
- Quantum Computation
- Quantum Simulation
- Quantum Sensing

Solve Problems Faster with a Quantum Computer

Some problems are hard to compute, in terms of resources in space (memory) or time (steps).
But they might be easier in a quantum world.

Prime Factorization [Shor, 1994]

1. Pick a random number $1 < a < N$.
2. Compute $K = \gcd(a, N)$, the greatest common divisor of a and N.
3. If $K \neq 1$, then K is a nontrivial factor of N, with the other factor being $\frac{N}{K}$ and we are done.
4. Otherwise, use the quantum subroutine to find the order r of a.
5. If r is odd, then go back to step 1.
6. Compute $g = \gcd(N, a^{r/2} + 1)$. If g is nontrivial, the other factor is $\frac{N}{g}$, and we're done. Otherwise, go back to step 1.

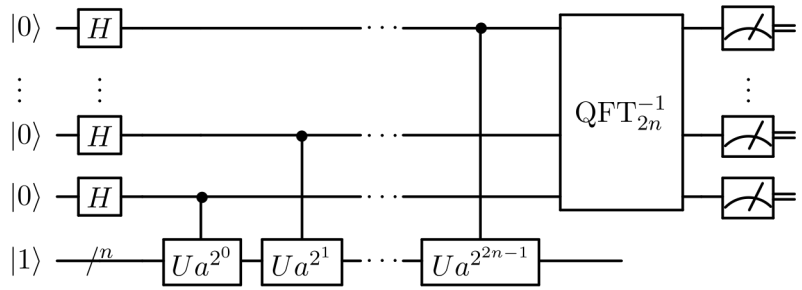

Complexity: $O((\log N)^2(\log\log N))$ quantum gates
Quantum subroutine in an algorithm: encode and process information in quantum systems.

Source: https://en.wikipedia.org/wiki/Shor%27s_algorithm

Computational Hardness

This is similar to using randomness as a computational resource:
e.g., solving problems faster if we allow making random moves in an algorithm.

Minimum Cut [Karger, 1993]

```
begin
    i = 1
    repeat
        repeat
            Take a random edge (u,v) ∈ E in G
            replace u and v with the contraction u'
        until only 2 nodes remain
        obtain the corresponding cut result C_i
        i = i + 1
    until i = m
    output the minimum cut among C_1, C_2, ..., C_m.
end
```

Complexity: $O(mn^2)$

Variant: $\tilde{O}(n^2)$

Source: https://en.wikipedia.org/wiki/Randomized_algorithm#Min_cut

The physics of computation

The laws of physics determines what kinds of computation can be done (efficiently).

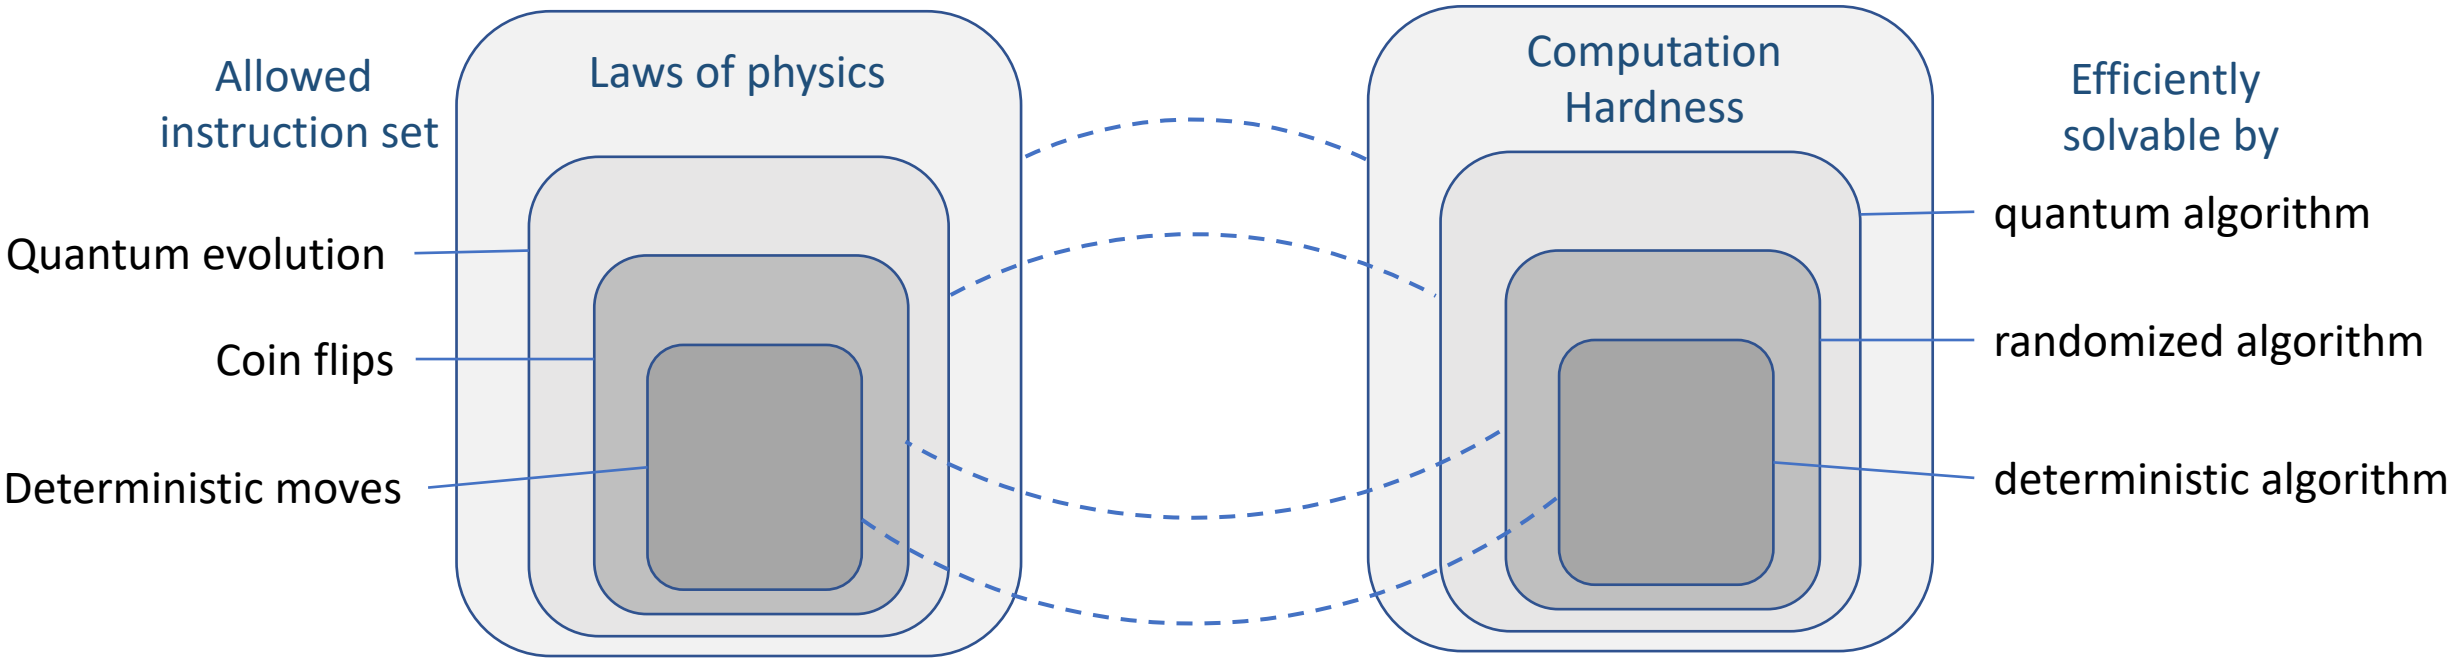

Extended Church-Turing thesis?
The set of problems that can be efficiently computed (in polynomial time) is the same for any realistic (physically realizable in principle) model of computation, such as a probabilistic Turing Machine.

Computational Hardness

Some problems are hard to compute, in terms of resources in space (memory) and time (steps).
But easier in a quantum world.

Quantum Simulation [Manin, Feynman, 1982]

> "The full description of quantum mechanics for a large system with R particles… has too many variables, it cannot be simulated with a normal computer with a number of elements proportional to R…
>
> And therefore, the problem is, how can we simulate the quantum mechanics? … We can give up on our rule about what the computer was, we can say:
>
> **Let the computer itself be built of quantum mechanical elements which obey quantum mechanical laws.** "

Emergence of computational problems that are inherently quantum.

Moore's law hits a wall

In real life, exponential growth cannot be sustained. [~2000]

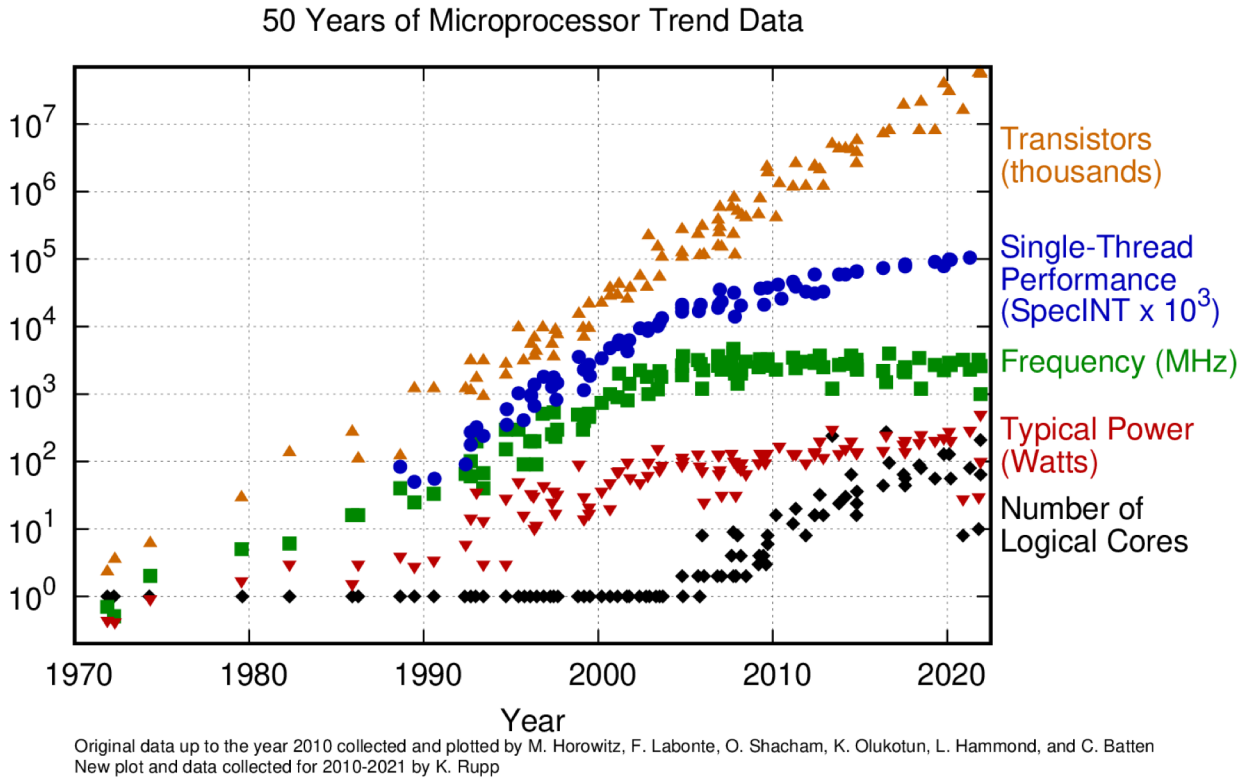

Source: https://github.com/karlrupp/microprocessor-trend-data

CMOS Transistor

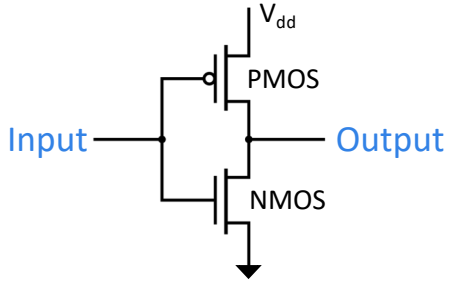

Not gate

Input	Output
0	1
1	0

Dennard's Scaling:

As transistors shrink, power consumption per unit area on chip stays the same.

Transistor dimension↓
energy↓, op frequency↑

This is no longer true, as of early 2000s.

Implementing a Quantum Bit (Qubit)

Quantum Bit:

Represent 0 and 1 at the same time, as a linear combination: $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle \in \mathbb{C}^2$

Atom/Ion qubits

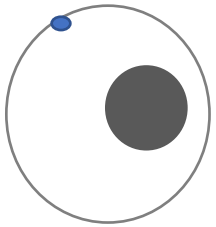

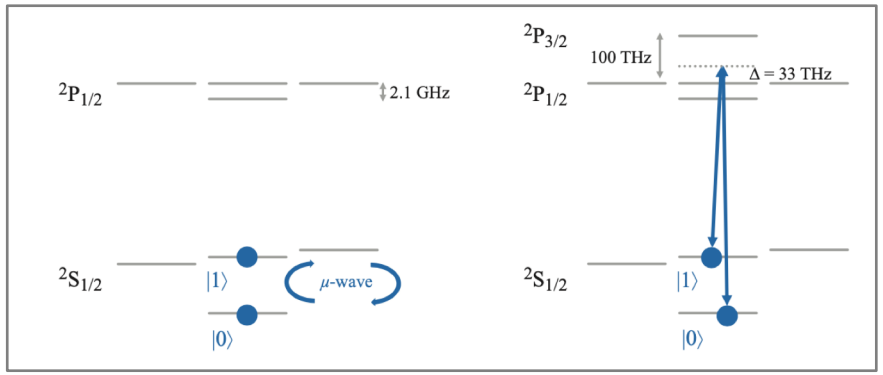

Superconducting Transmon qubit

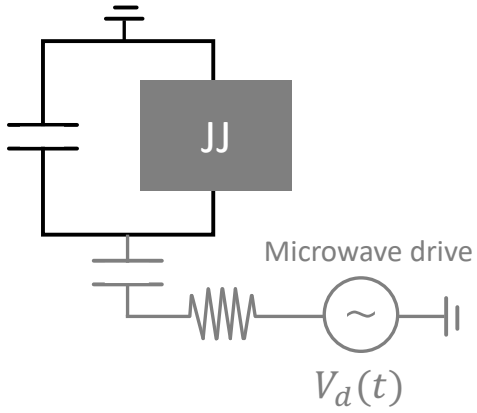

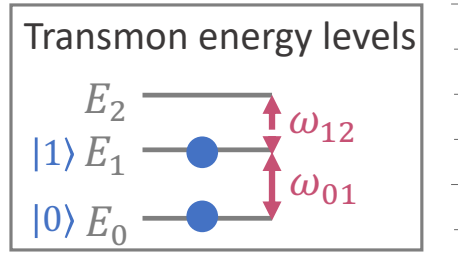

[Schoelkopf, Devoret, Girvin, 2003]

Fighting off decoherence – a natural tendency to reduce to classical behavior

Emerging Quantum Computers

Noisy Intermediate-Scale Quantum (NISQ) Devices: 100-1000 qubits

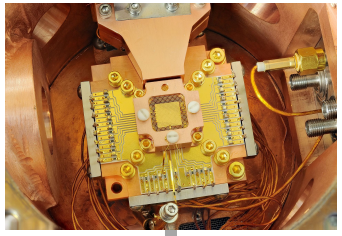

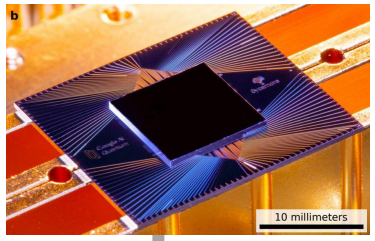

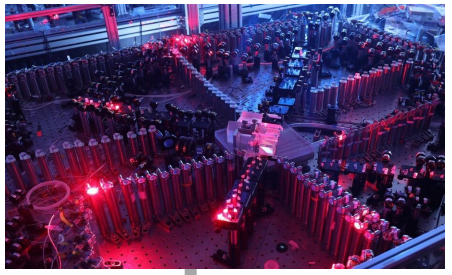

IBM Roadmap

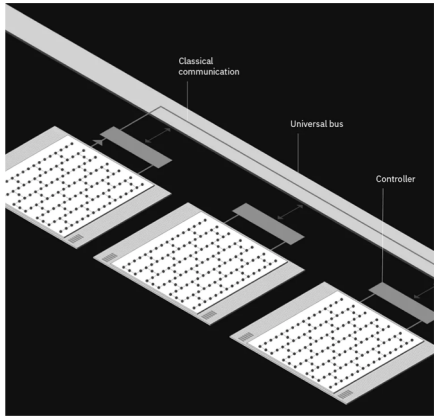

...

2010

2015

2020

2025

...

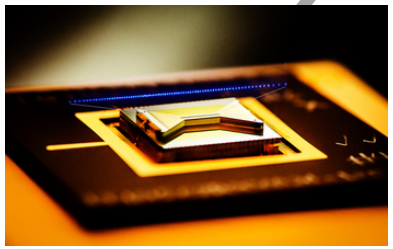

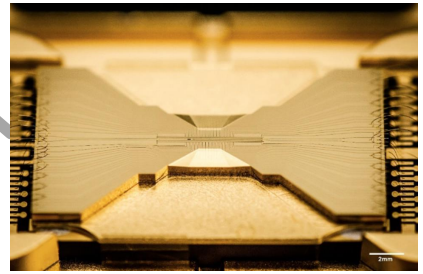

Photos from: NIST, IBM Q Graham Carlow, Google AI Quantum, Kai Hudek IonQ, E. Edwards JQI, Hansen Zhong USTC

Gap between Algorithms and Hardware

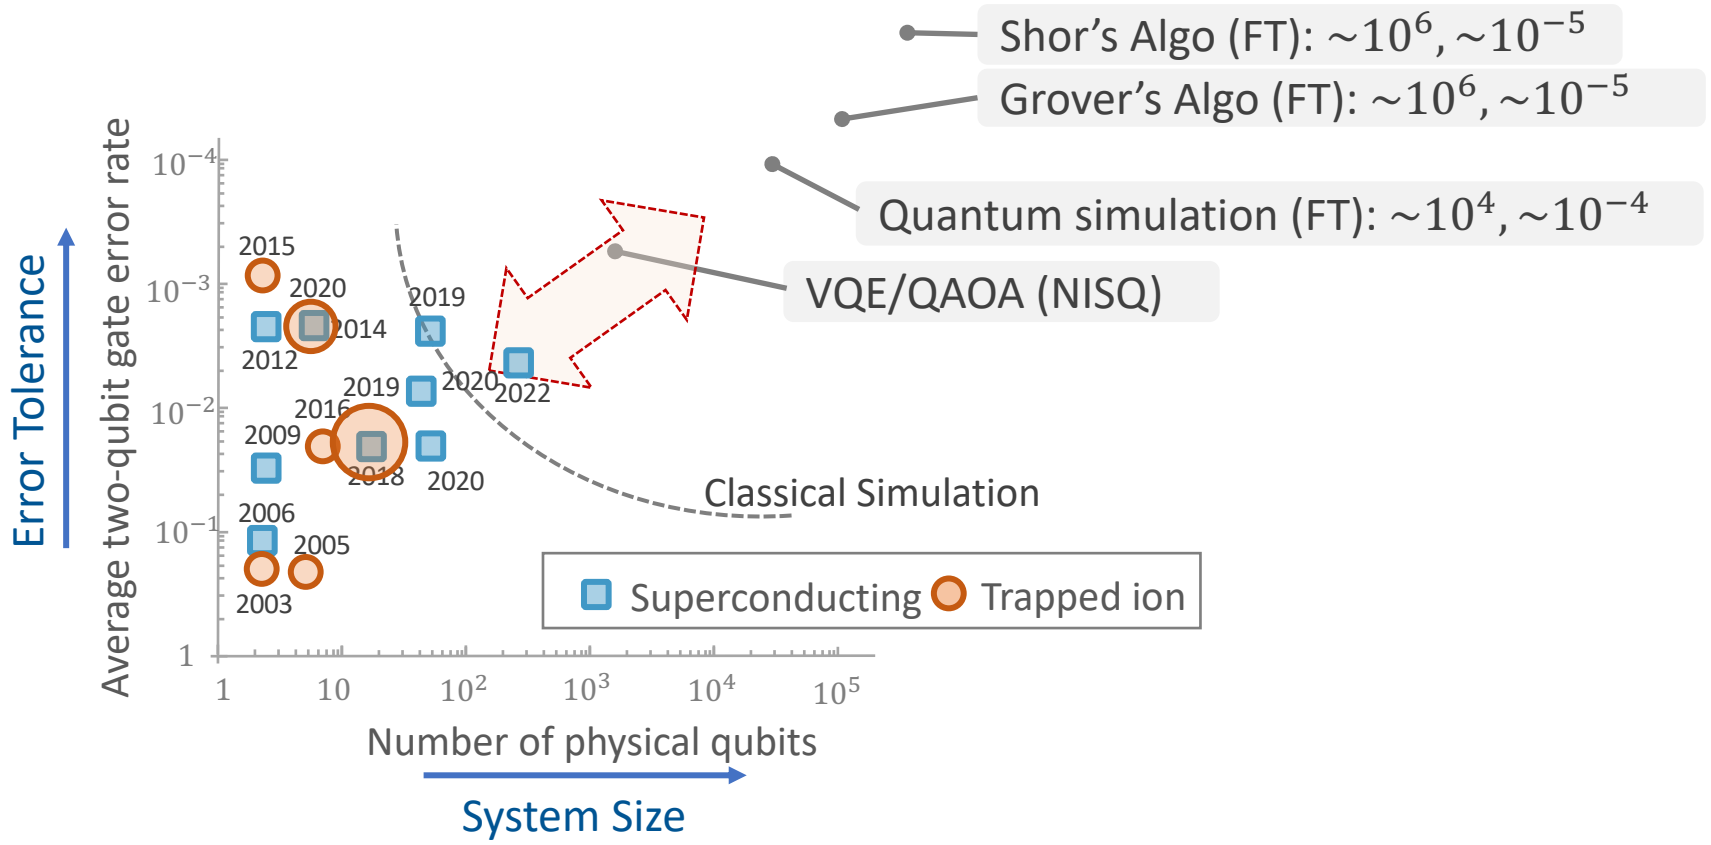

*Size of data point indicates connectivity; larger means denser connectivity.

Data: Ding & Chong (2020)

Lack of Computer Architecture

Hardware Organization

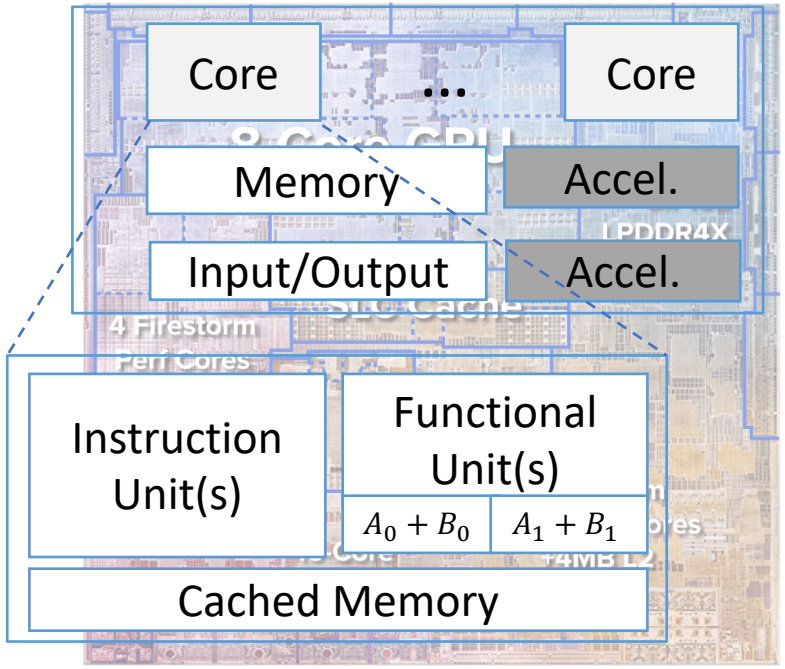

Apple M1 Chip

[Source: Toptal.com]

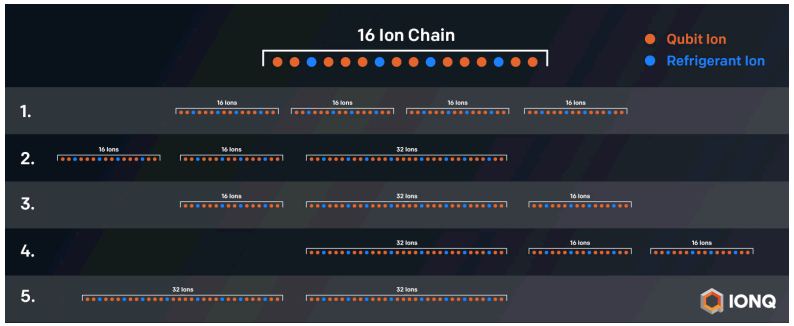

Trapped Ion

[Source: IonQ]

Superconducting Transmon

[Source: IBM Q]

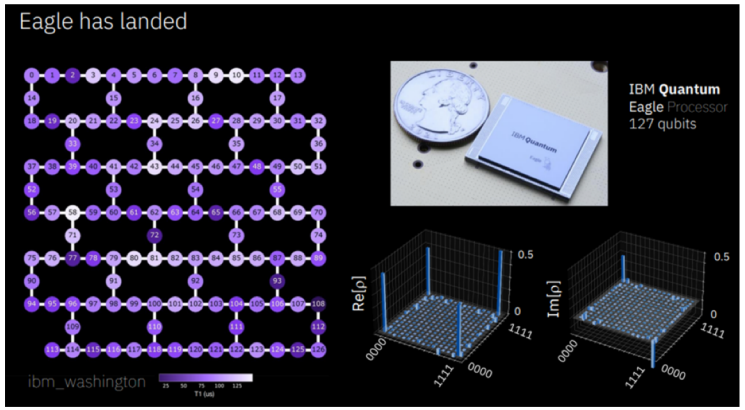

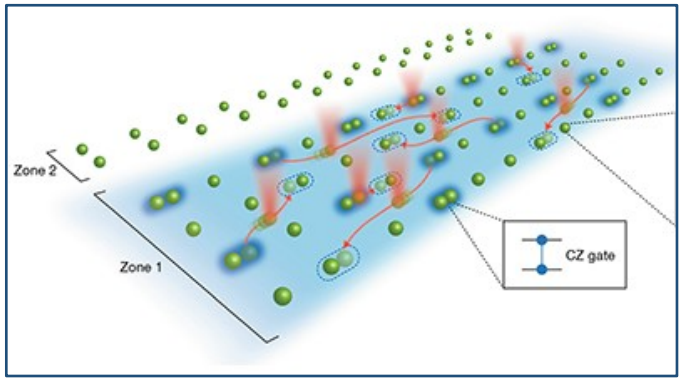

Rydberg Atoms

[Source: Bluvstein et al, Nature, 2022]

CLASSICAL

QUANTUM

Classical v.s. Quantum Information

Classical Bit

Quantum Bit

$|\psi\rangle \propto 0.6|0\rangle + 0.8|1\rangle$

Two Quantum Bit

$|\psi\rangle \propto (0.6|0\rangle + 0.8|1\rangle)(0.6|0\rangle + 0.8|1\rangle)$

$\propto (0.36|0\rangle|0\rangle + 0.48|0\rangle|1\rangle$

$+0.48|1\rangle|0\rangle + 0.64|1\rangle|1\rangle)$

Read & Write

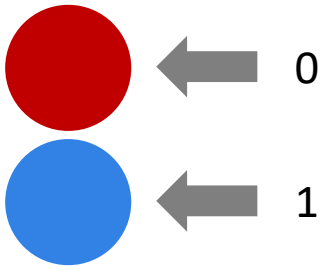

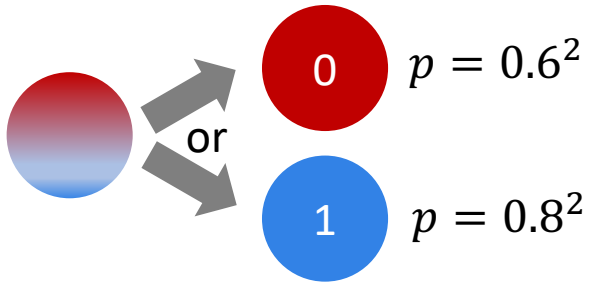

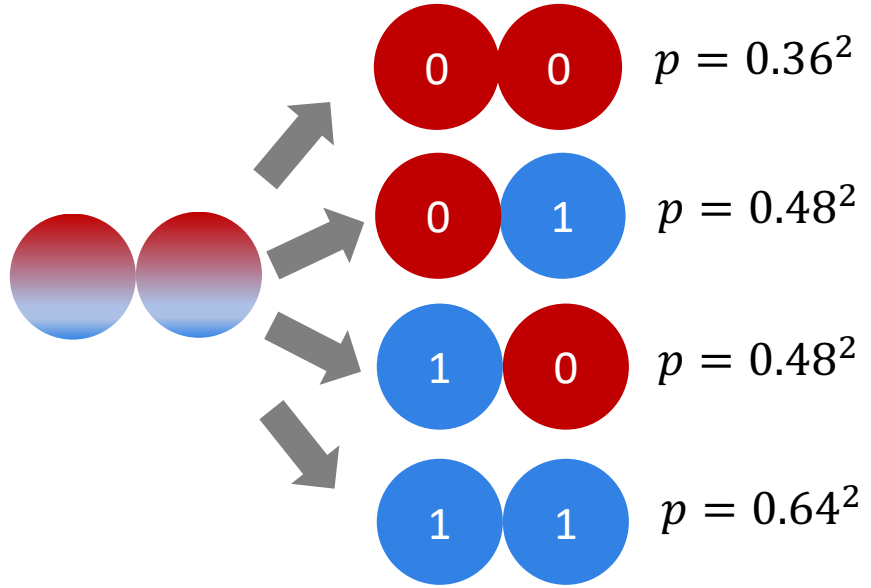

Logic

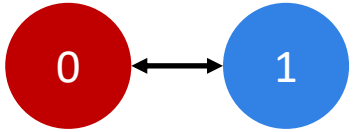

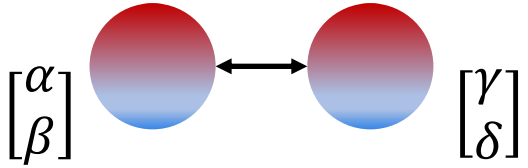

Boolean algebra

Linear algebra

Quantum Computation Model

Goal: compute function f(x) = x + 1

input output

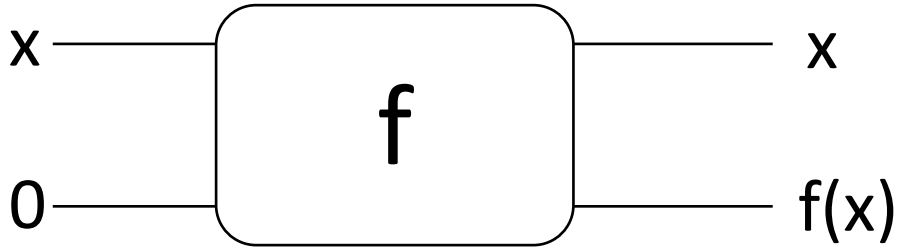

"Writing output out-of-place"
(reversible)

Classical Input

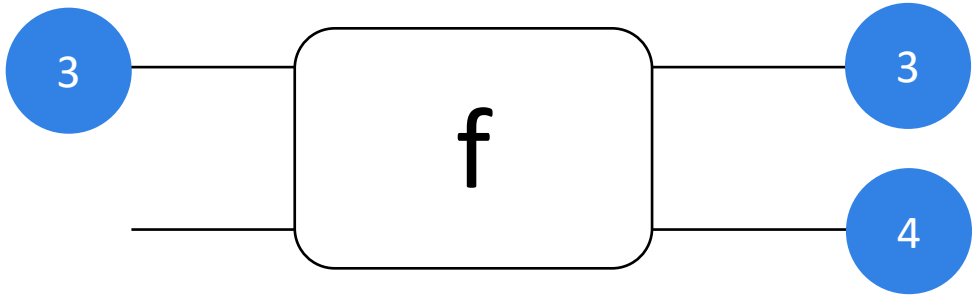

Quantum Input

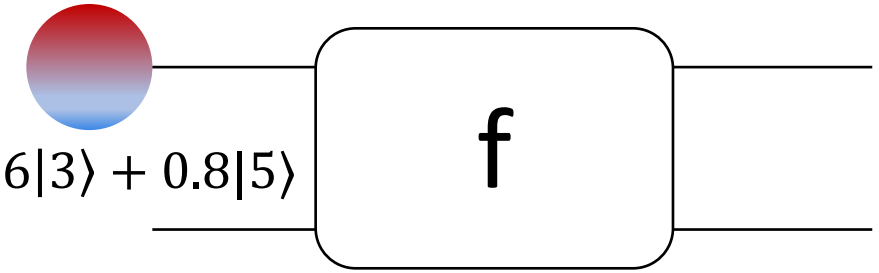

$|\psi\rangle \propto 0.6|3\rangle + 0.8|5\rangle$

What is the output?

$|\psi\rangle \propto (0.6|3\rangle + 0.8|5\rangle)(0.6|4\rangle + 0.8|6\rangle)?$

$|\psi\rangle \propto (0.6|3\rangle|4\rangle + 0.8|5\rangle|6\rangle)?$

Superposition and Interference

The massive quantum parallelism

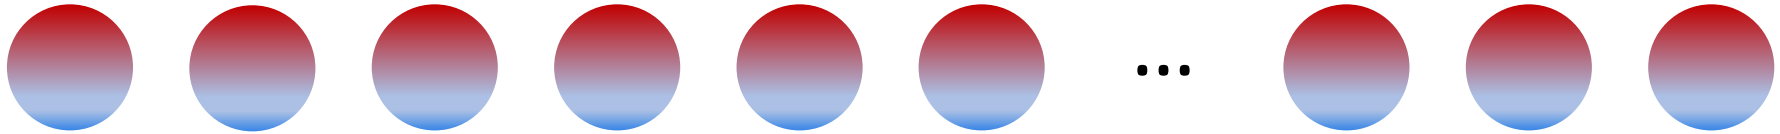

$$\mathbb{C}^{2^n} = \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \cdots \otimes \mathbb{C}^2 \otimes \mathbb{C}^2$$

Feynman paths:
describe the transitions between superposition of states.

$$\begin{bmatrix} u_1 & u_2 \\ u_3 & u_4 \end{bmatrix}$$

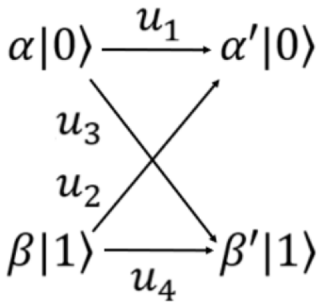

$$\alpha' = u_1\alpha + u_2\beta$$

$$\beta' = u_3\alpha + u_4\beta$$

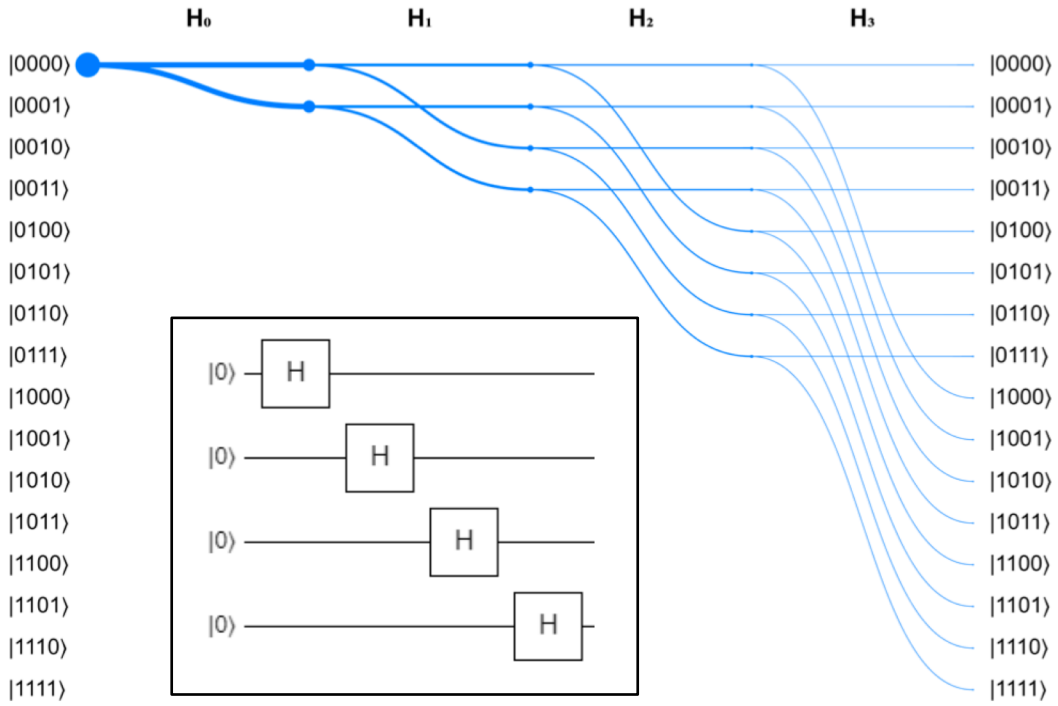

https://github.com/Yale-QCS/feynman-path-visualizer

Superposition and Interference

The massive quantum parallelism

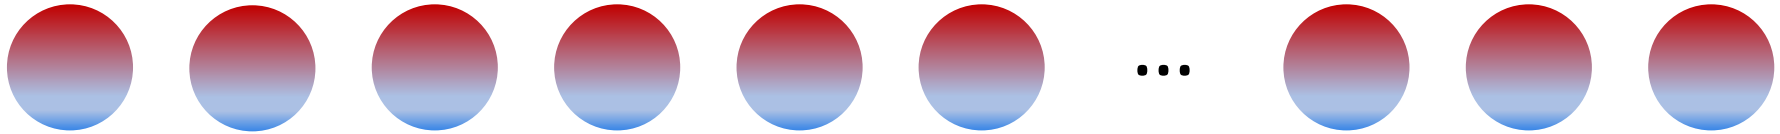

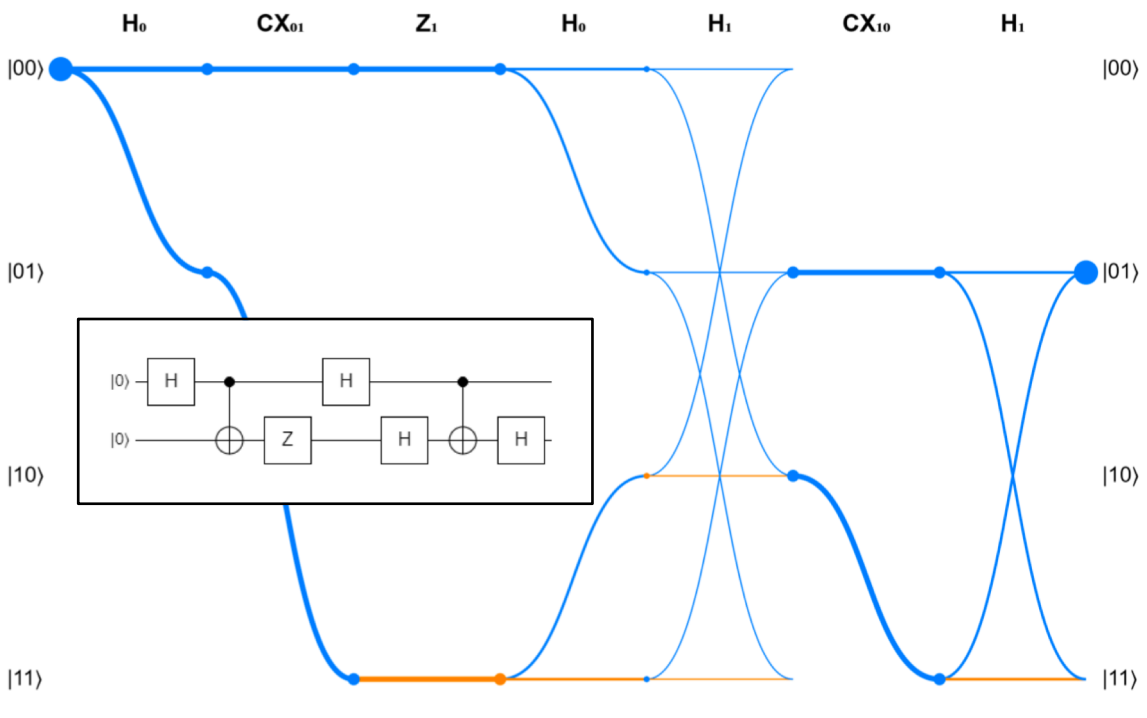

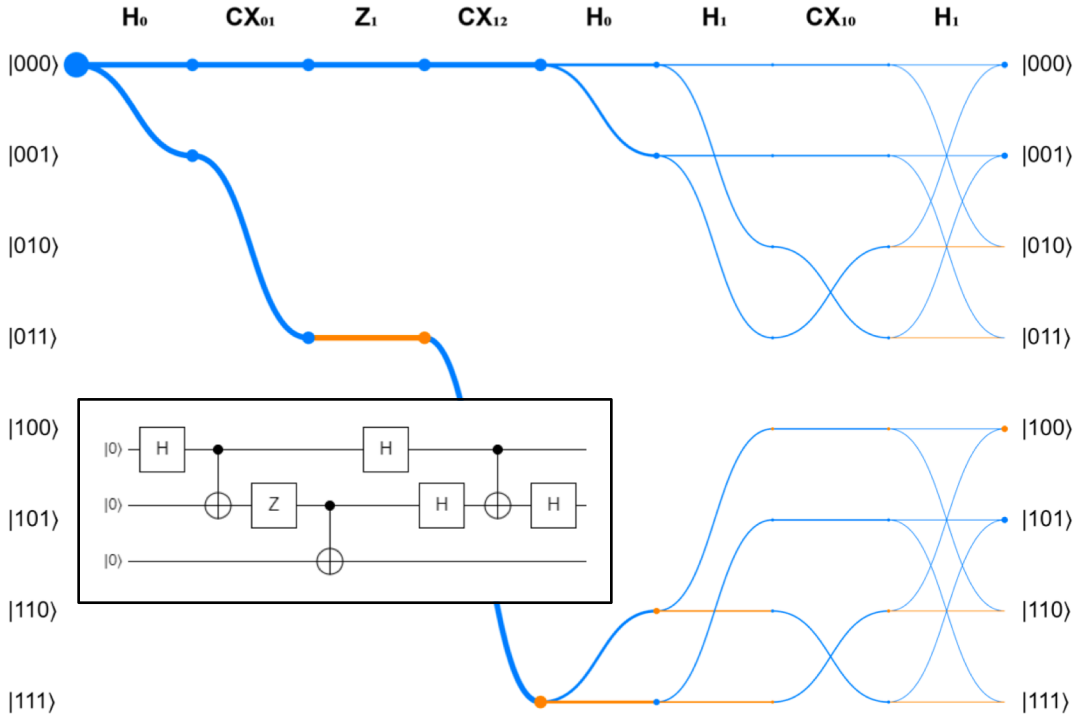

Feynman-path visualization: https://github.com/Yale-QCS/feynman-path-visualizer

From Vacuum Tubes to Modern Computers

IBM System/360 at NASA (1960s)

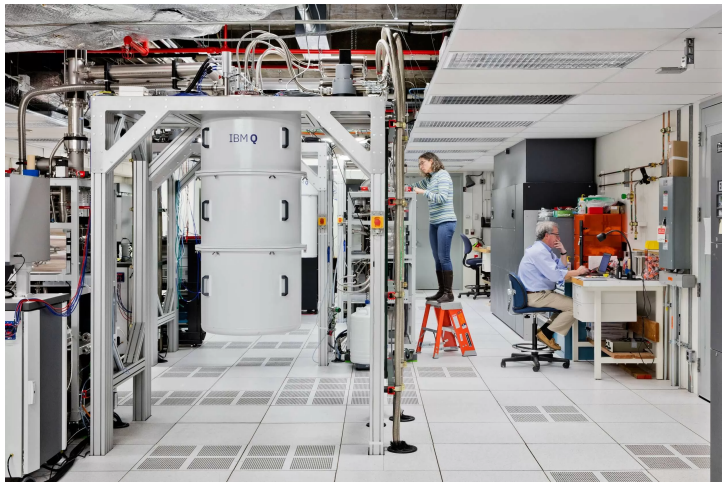

IBM Q (2019)

How do we build practical quantum computers sooner?
The answer has to do with leveraging digital computers and experience of building digital computers.

Photo from: en.wikipedia.org, www.techspot.com

Emerging Applications

Computational tasks that are considered potentially easy on quantum computers:

Comp. Data Science	Numerical Analysis	Simulation	Machine Learning
• Post-Quantum Cryptography • Distributed/blind computation • Secure Communication	• Optimizations • Adiabatic algorithms • Quantum Linear Algebra	• Quantum chemistry • Quantum material science • Quantum gravity	• Quantum Neural Networks • Quantum Learning Theory

...

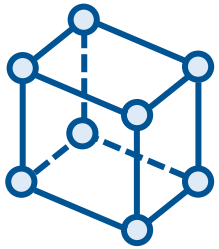

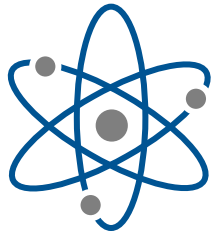

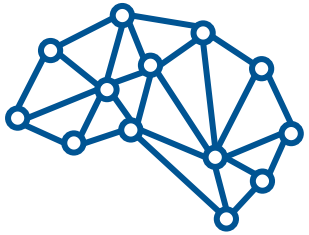

Quantum State

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle \in \mathbb{C}^2$$

basis

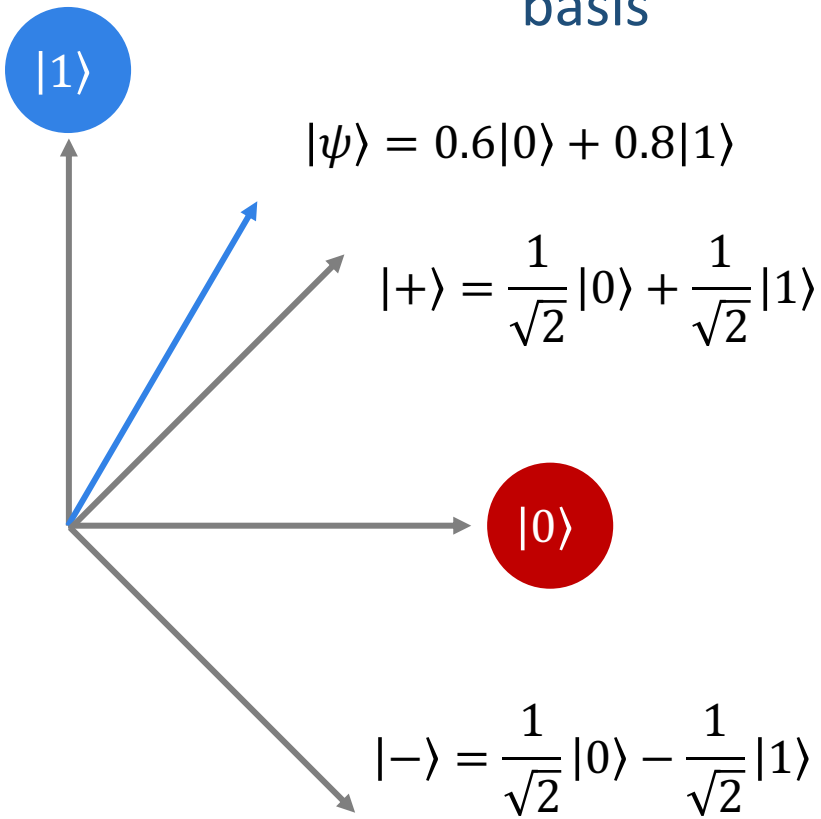

- Normalized: $|\alpha|^2 + |\beta|^2 = 1$
- Global phase does not matter:
  $|\psi\rangle$ and $e^{i\phi}|\psi\rangle$ not distinguishable

$$\begin{bmatrix}\alpha\\ \beta\end{bmatrix} \sim \begin{bmatrix}a\\ e^{i\theta}b\end{bmatrix} \sim \begin{bmatrix}a\\ e^{i\theta}\sqrt{1-a^2}\end{bmatrix}$$: two real numbers

$$0 \le a \le 1, 0 \le \theta < 2\pi$$

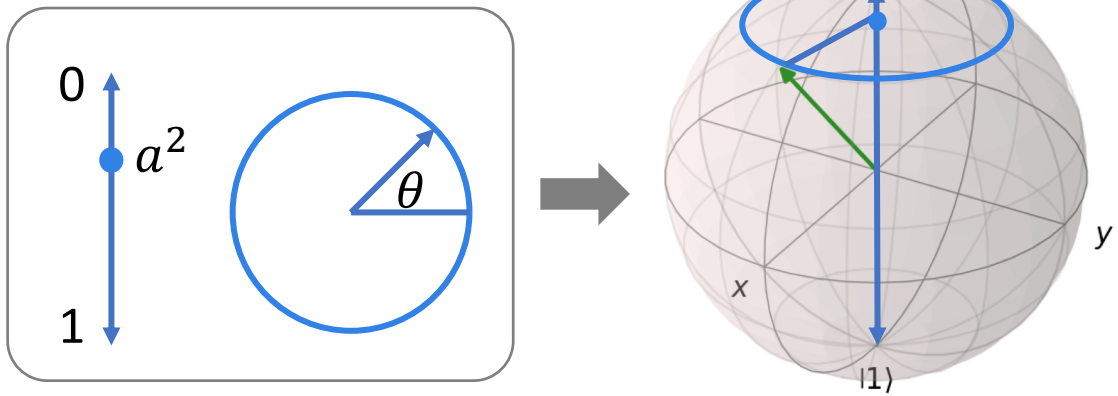

Quantum Measurement

Projection to subspaces H of $\mathbb{C}^{2^n}$.

Orthonormal decomposition $\mathbb{C}^{2^n} = H_1 \oplus H_2 \oplus \cdots \oplus H_m$

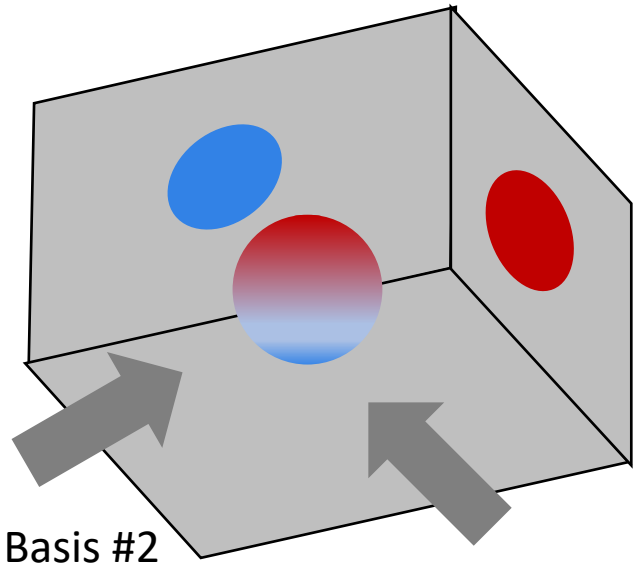

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle = \alpha'|+\rangle + \beta'|-\rangle$$

Randomness:

The measurement event is inherently random, even given full description of the qubits.

Irreversibility:

The measurement operation collapses the quantum state to the associated subspace. This process cannot be reversed.

Non-commutativity:

When A and B do not commute, measuring operator A influences the outcome of the subsequent measurement B.

Entanglement – non-local information

A new notion of shared state.. $\frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle = \frac{1}{\sqrt{2}}|+\,+\rangle + \frac{1}{\sqrt{2}}|-\,-\rangle$

Information is not stored in any subsystems, but as correlations in the entire system.

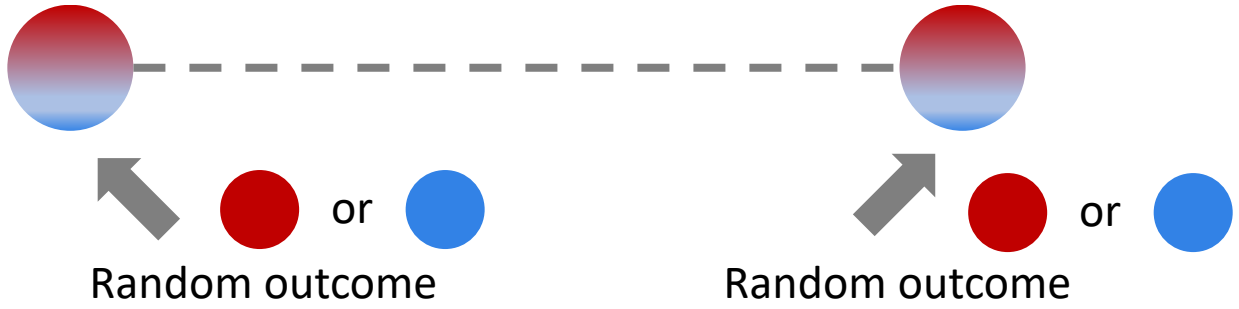

But if measure in an agreed basis, their outcome will always be the same.

Entanglement is stronger than classical correlation

Clauser–Horne–Shimony–Holt (CHSH) Game

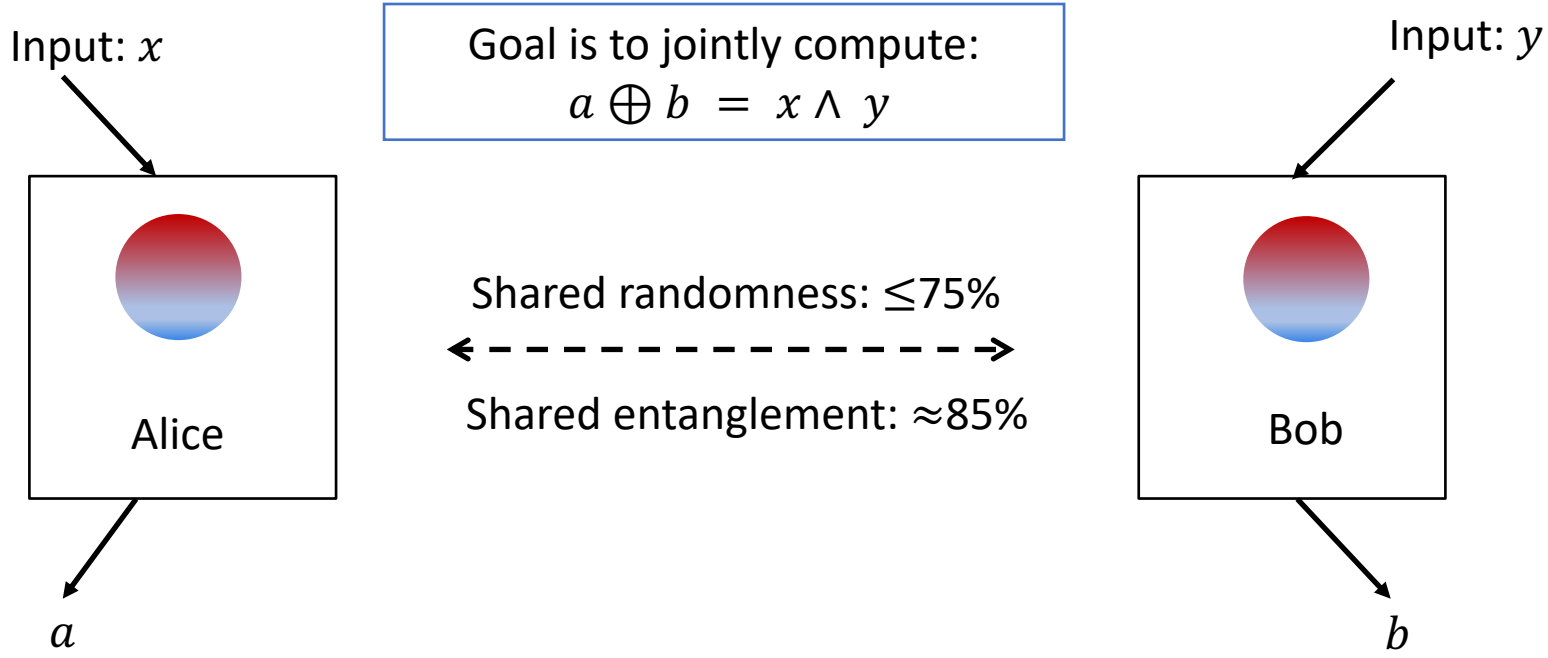

Theory: J. Bell (1964)
Experiment: A. Aspect et al., early ’80s

Entanglement can be used as a resource.
More example: generating certifiable randomness.

Unitary Transformation

$\begin{bmatrix}\alpha\\ \beta\end{bmatrix}$ 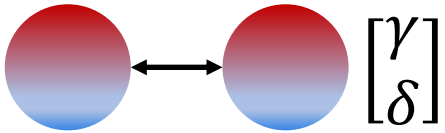 $\begin{bmatrix}\gamma\\ \delta\end{bmatrix}$

Linear algebra: unitary matrices $U^{-1} = U^{\dagger}$ (Prove it!)

Computational universality:
A subset of operations can implement arbitrary transformations.
Claim: Any single-qubit transformation can be implemented by Rx and Rz gates.

$Rz(\pi)$

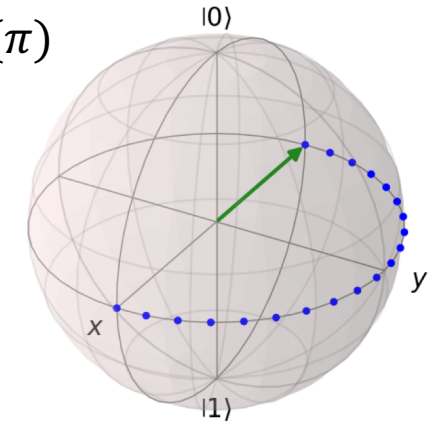

$Rx(\pi)$

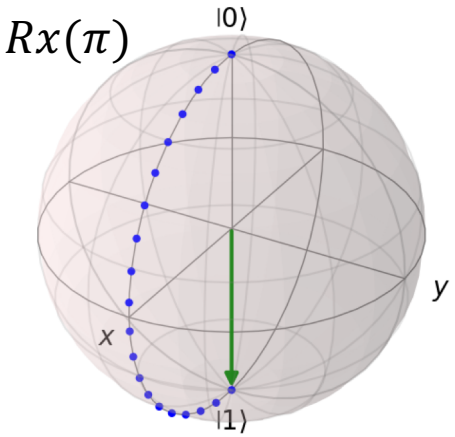

Exercise:

Implement an arbitrary angle, arbitrary axis rotation by three Rx and Rz rotations.

Circuit Synthesis and Quantum Compiling

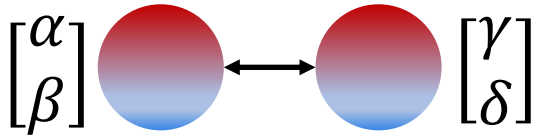

Linear algebra: unitary matrices $U^{-1} = U^{\dagger}$ (Prove it!)

In fact, **Hadamard and $Rz(\pi/4)$ gates** can implement any single-qubit transformations.
Two-qubit gates can implement arbitrary transformation on any number of qubits.

In practice, having a **redundantly universal** instruction set can be helpful: more efficient circuit.

A **quantum circuit/program** specifies a sequence of quantum gates and measurements.

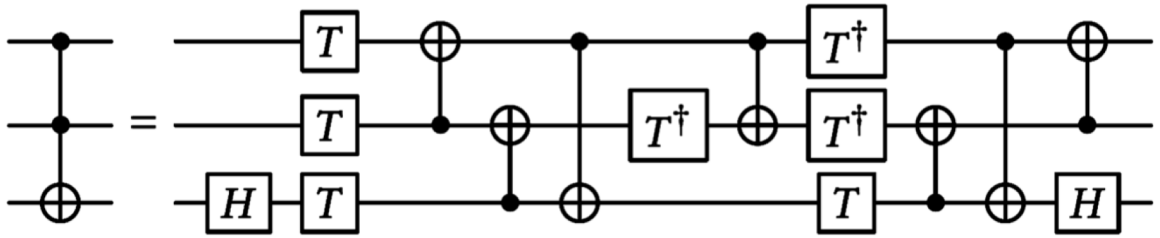

Source: Amy et al, https://arxiv.org/abs/1206.0758

Quantum errors

Decoherence: loss of information to the environment (e.g., bit-type errors, phase-type errors).

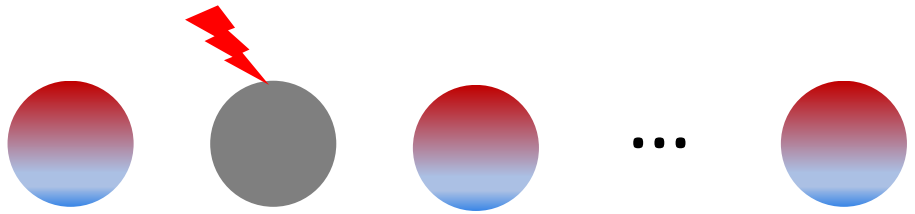

$$\begin{bmatrix}\alpha\\ \beta\end{bmatrix} \sim \begin{bmatrix}a\\ e^{i\theta}b\end{bmatrix} \sim \begin{bmatrix}a\\ e^{i\theta}\sqrt{1-a^2}\end{bmatrix}$$: two real numbers

$$0 \le a \le 1, 0 \le \theta < 2\pi$$

- Bit-type errors: random disturbance to a.
  - Amplitude damping: spontaneous decay from $|0\rangle$ to $|1\rangle$ at a random time.
- Phase-type errors: random disturbance to θ.
  - Dephasing: spontaneous loss of phase information

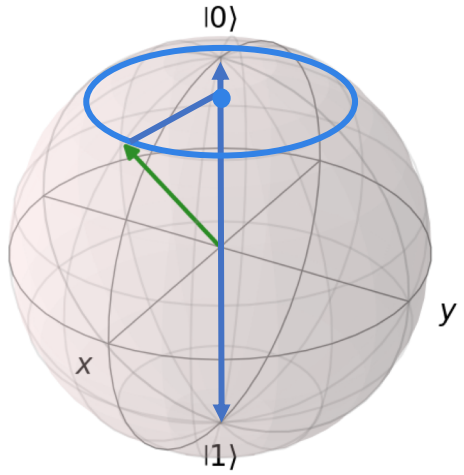

Quantum errors and how to catch them

Protecting information against decoherence:

- Redundancy: encode information non-locally.
- Error detection/correction: frequent checks to restore information

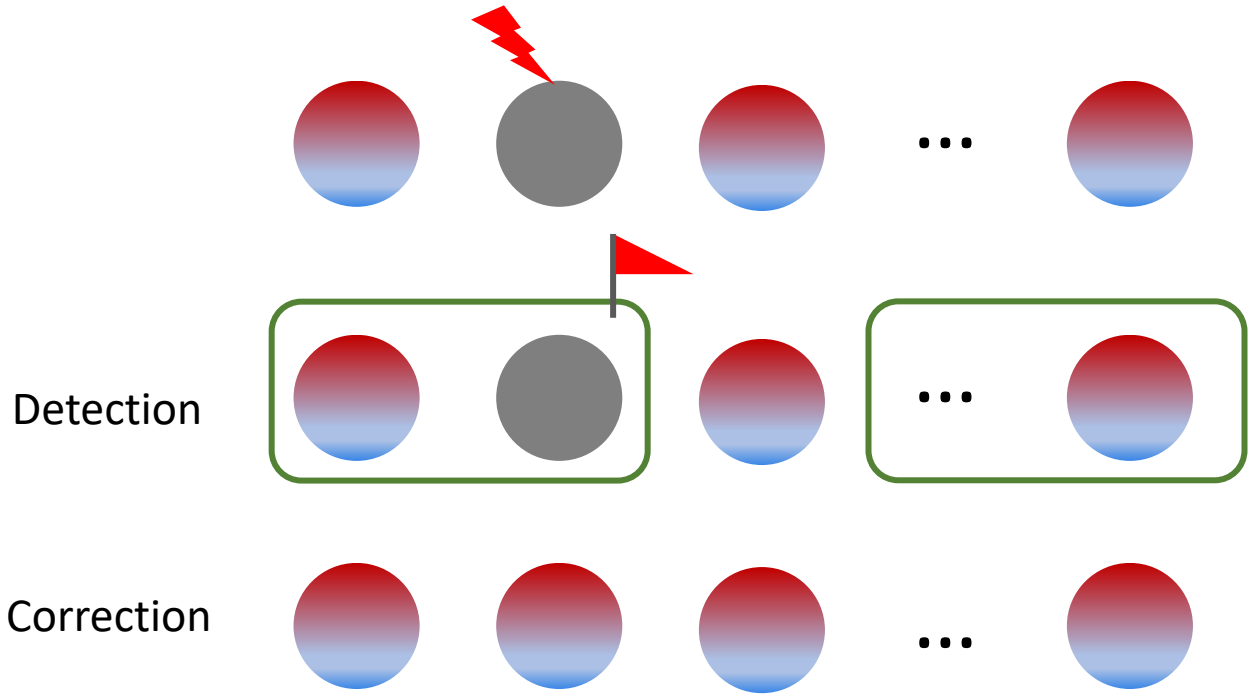

$$\mathbb{C}^{2^n} = \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \cdots \otimes \mathbb{C}^2 \otimes \mathbb{C}^2$$

$$= H_1 \oplus H_2 \oplus H_3 \oplus H_4 \oplus \cdots \oplus H_m$$

Logical subspace H_L | Error subspace H_E