Why Nuclear Physics Needs Quantum Computers

and What it's doing with them

Daniel Sierra-Sosa

THE CATHOLIC UNIVERSITY OF AMERICA

Outline

- Motivation
- Why quantum computing?
- Hardware landscape: quantum technologies, NISQ devices, and fault tolerance
- Quantum basics: qubits, gates, circuits, measurement, and sampling
- Simulation workflow: encode → prepare → evolve → measure
- Algorithms and noise: QPE, VQE, SQD/QSCI, and error mitigation
- Fragment-based quantum simulation and nuclear-physics applications
- Hybrid QC-ML

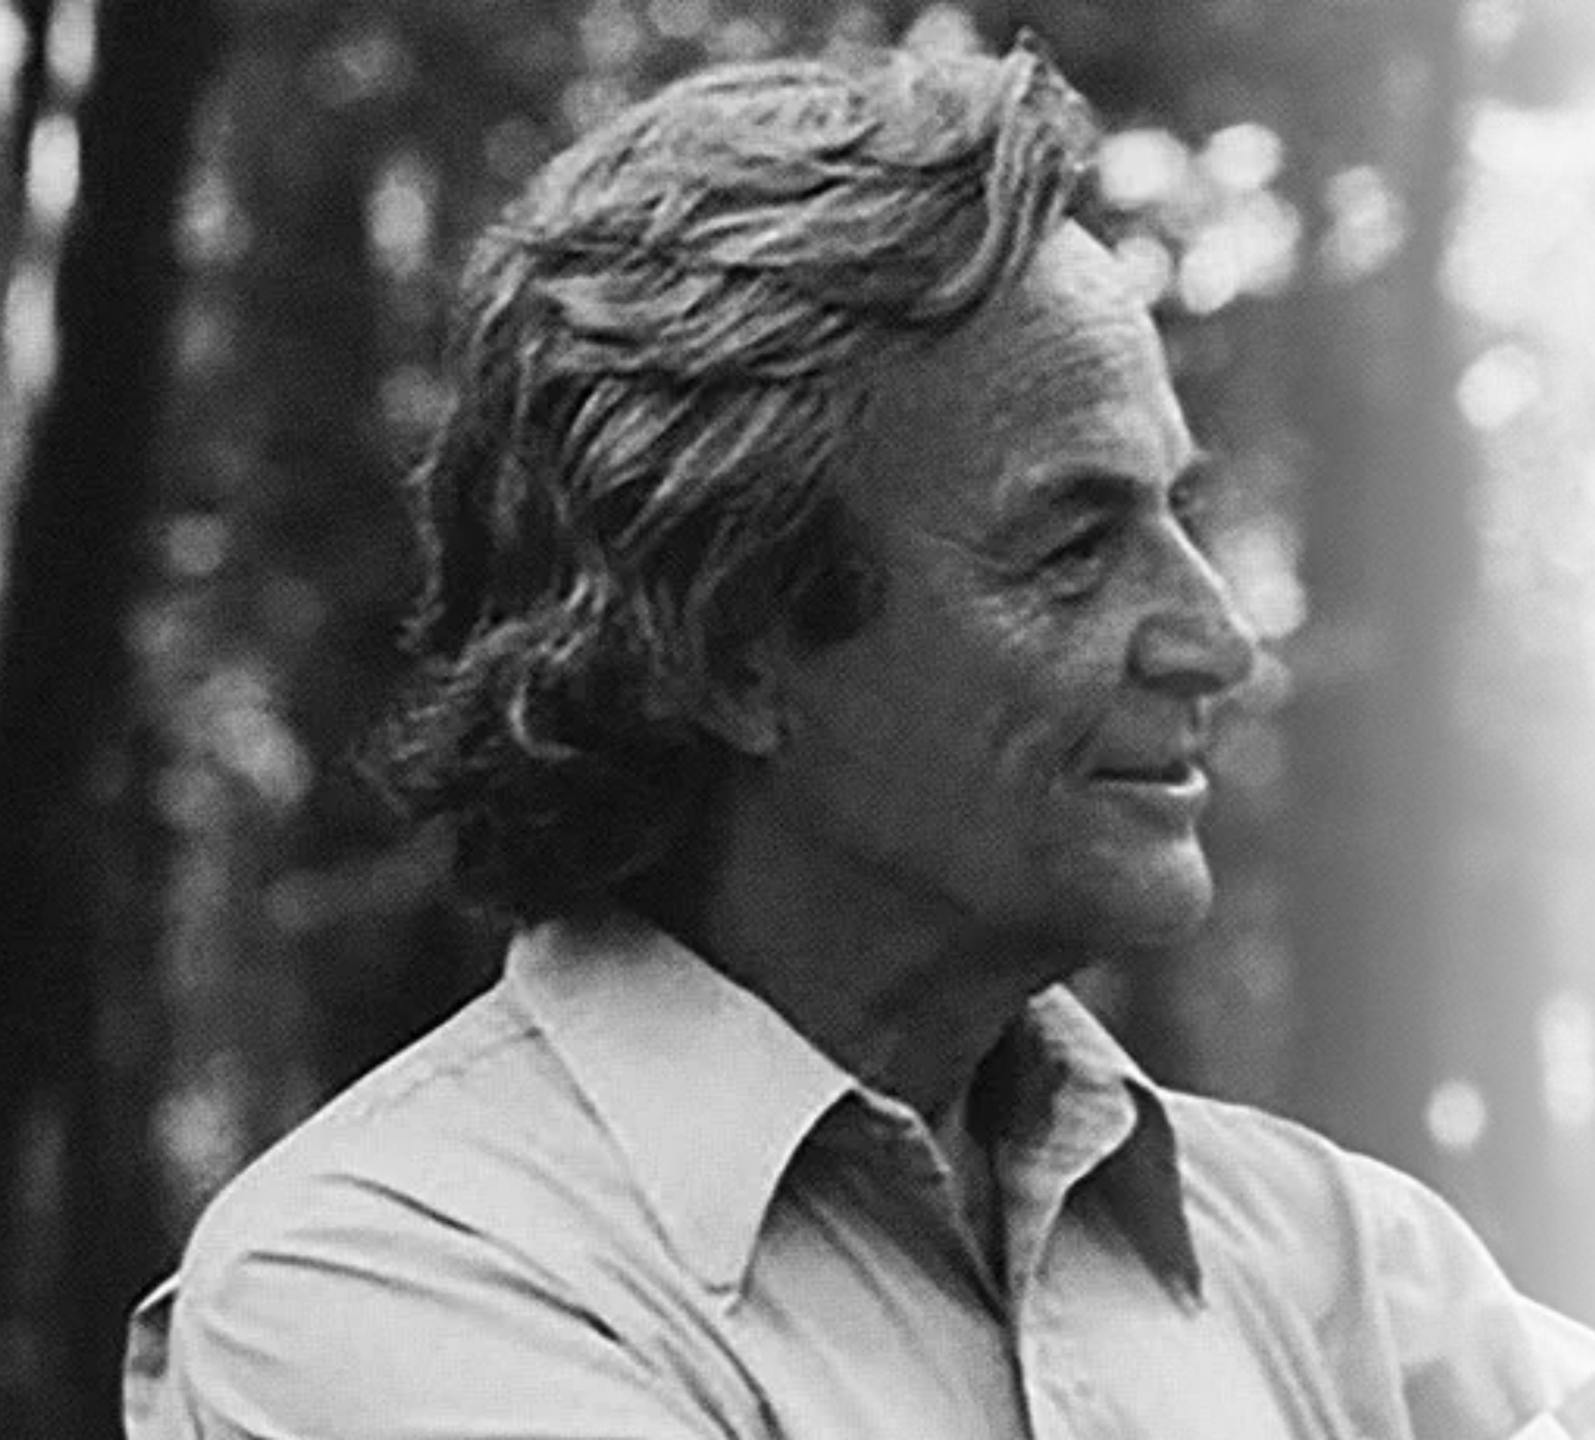

Feynman’s 1982 vision

"… 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, Int. J. Theor. Phys. 21, 467 (1982)

Why Quantum Computers?

- Feynman: a classical computer simulating a quantum system pays an exponential cost in the number of degrees of freedom
- A quantum machine made of N qubits naturally represents 2^N amplitudes. The speedup is not from reading all amplitudes, but from manipulating amplitudes coherently through interference.
- Lloyd (1996): any local Hamiltonian can be simulated efficiently on a quantum computer with polynomial overhead

Where quantum computing sits

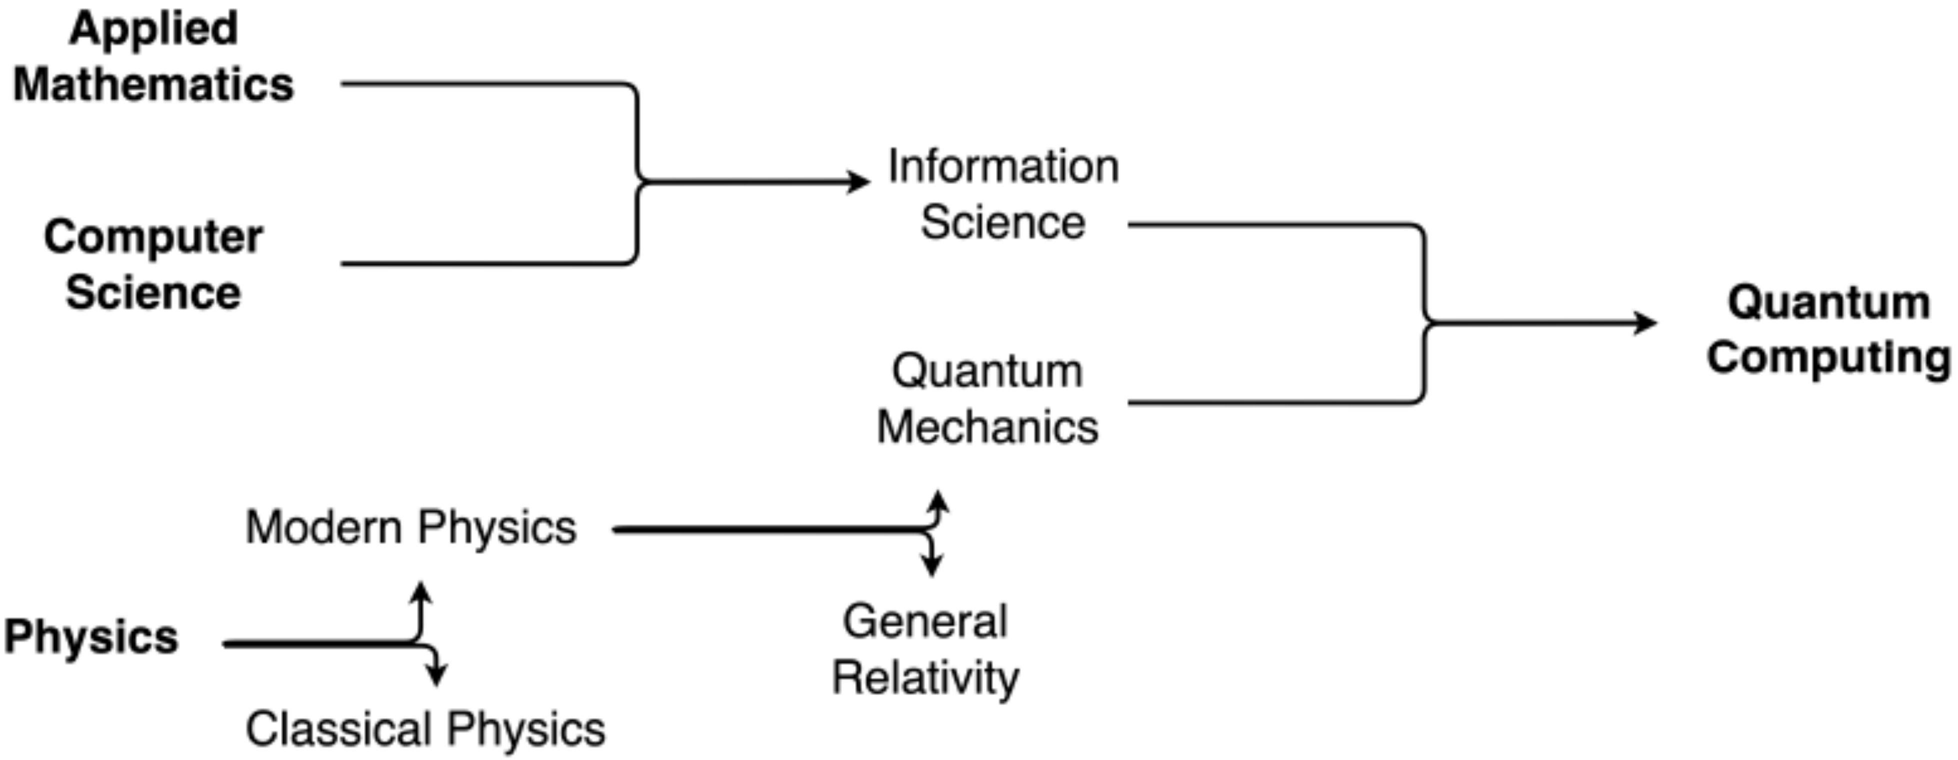

Taken from: Ayoade, O., Rivas, P., & Orduz, J. (2022). Artificial Intelligence Computing at the Quantum Level. Data, 7(3), 28.

“It is no longer a physicist’s dream—it is an engineer’s nightmare.”

Isaac Chuang

Why now?

- Hardware: 100s of physical qubits available, with gate errors ~10^{-3} on best platforms.
- Algorithms: designed for noisy devices (variational methods, Trotter techniques, among others)
- Funding & infrastructure

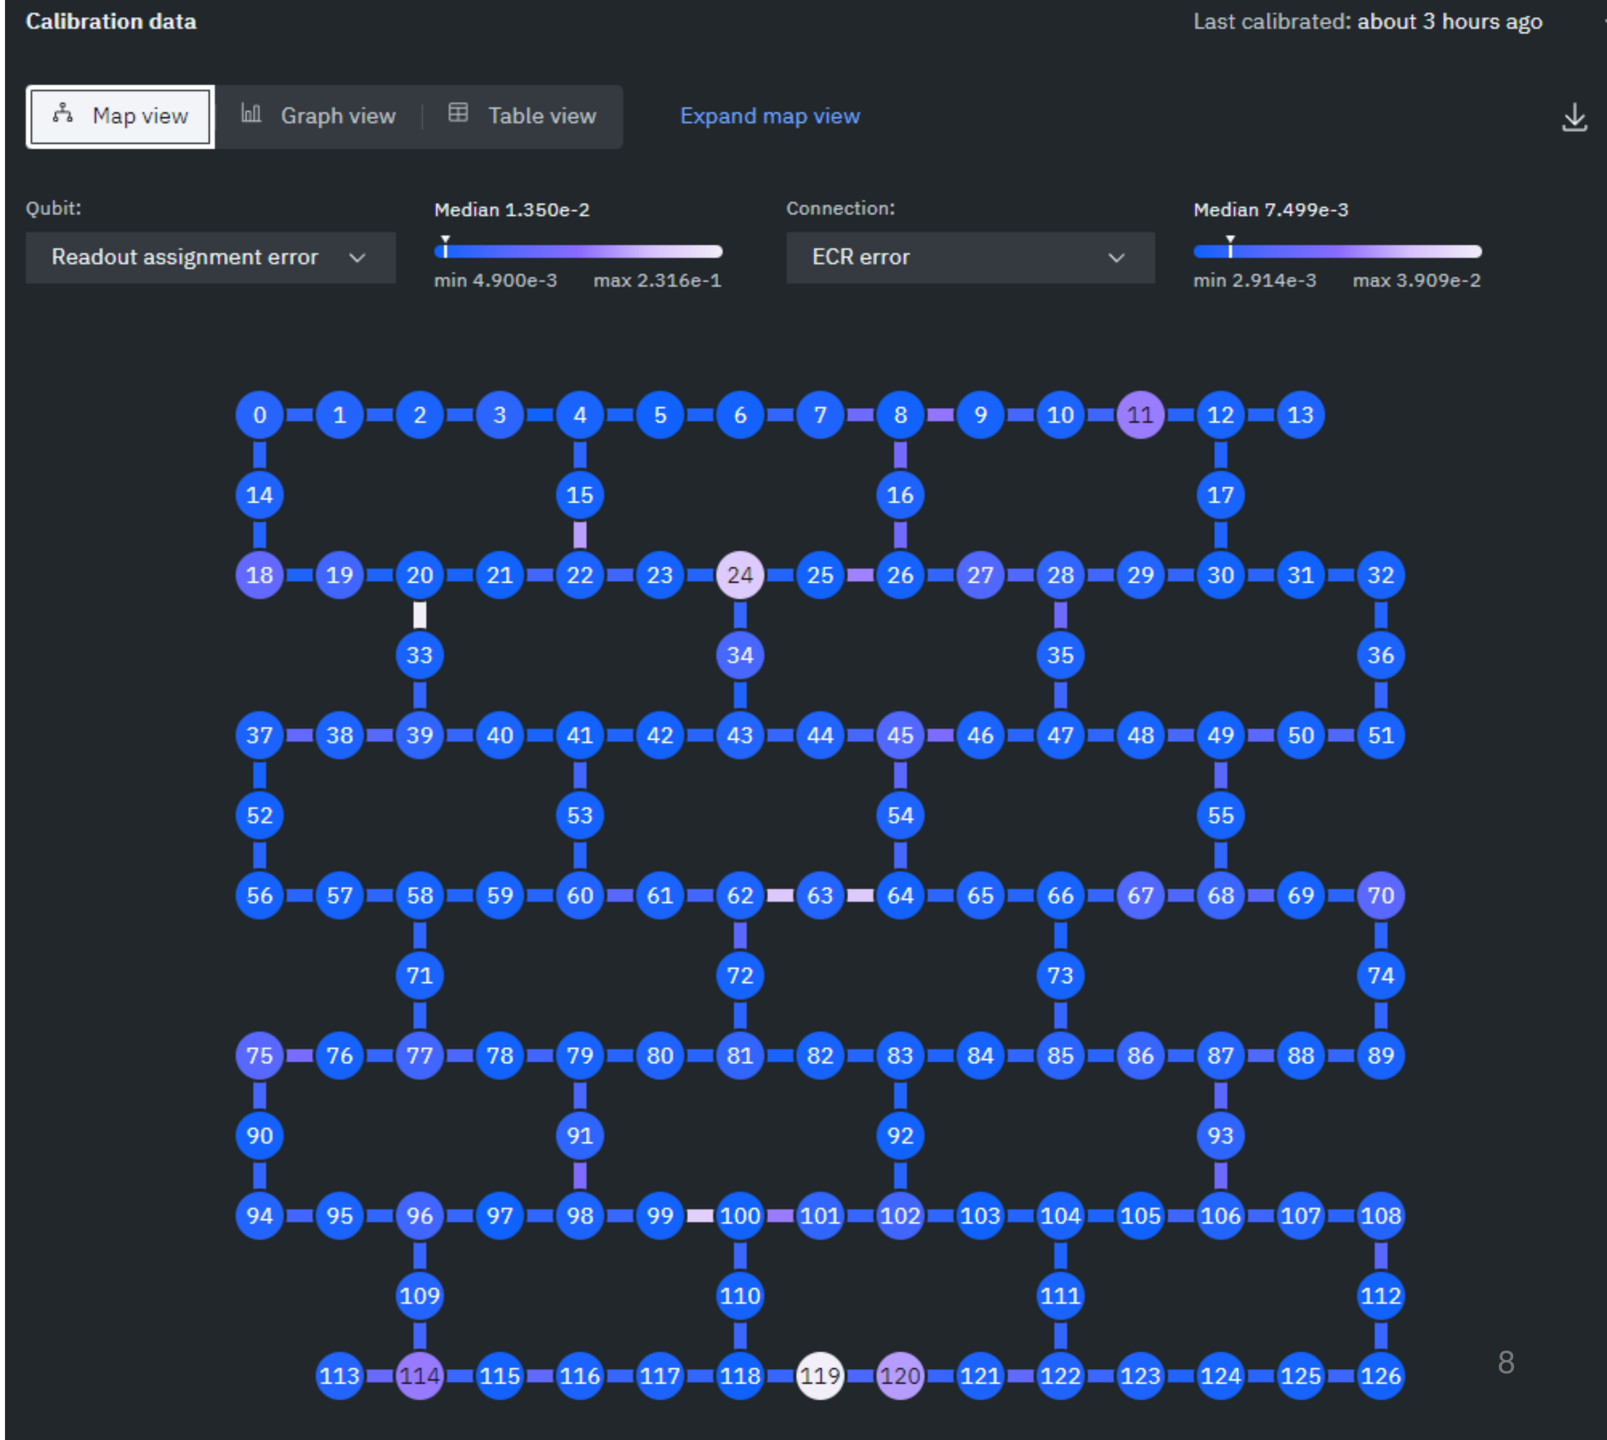

QURECA
WE SPEAK QUANTUM

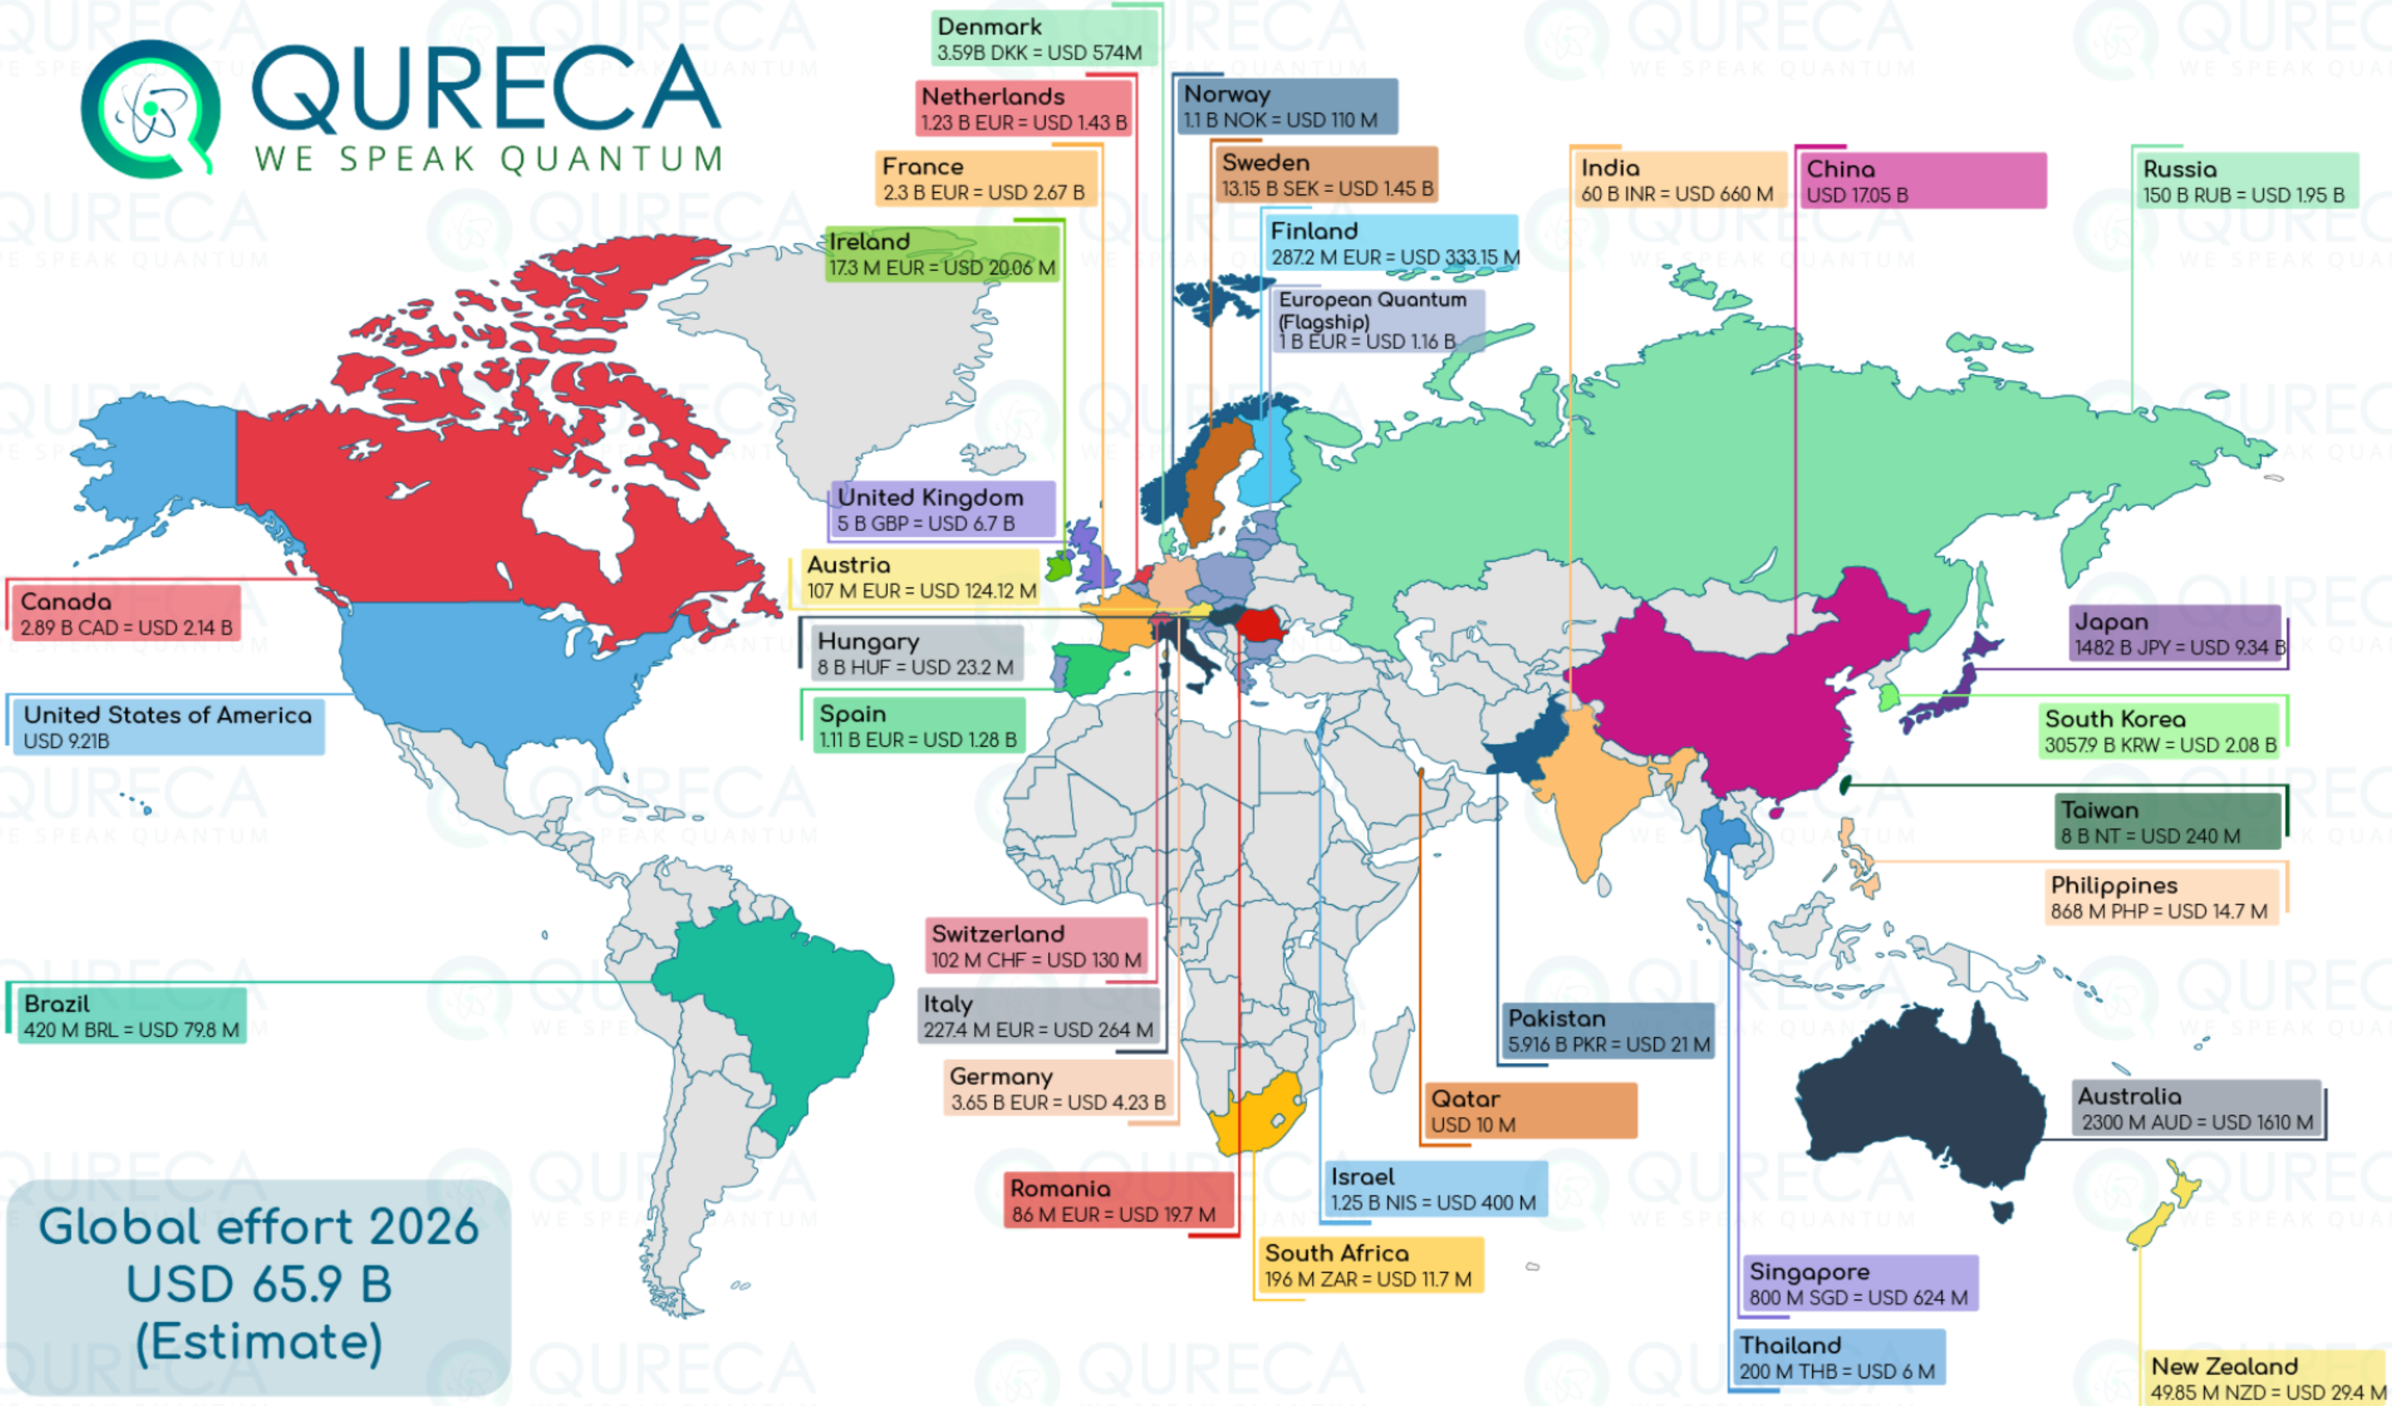

Global effort 2026
USD 65.9 B
(Estimate)

Quantum Technologies

Superconducting Architecture

IBM rigetti

Google Alibaba Group 阿里巴巴集团

Trapped Ions

IONQ

Topological

Microsoft

Photonic

XANADU

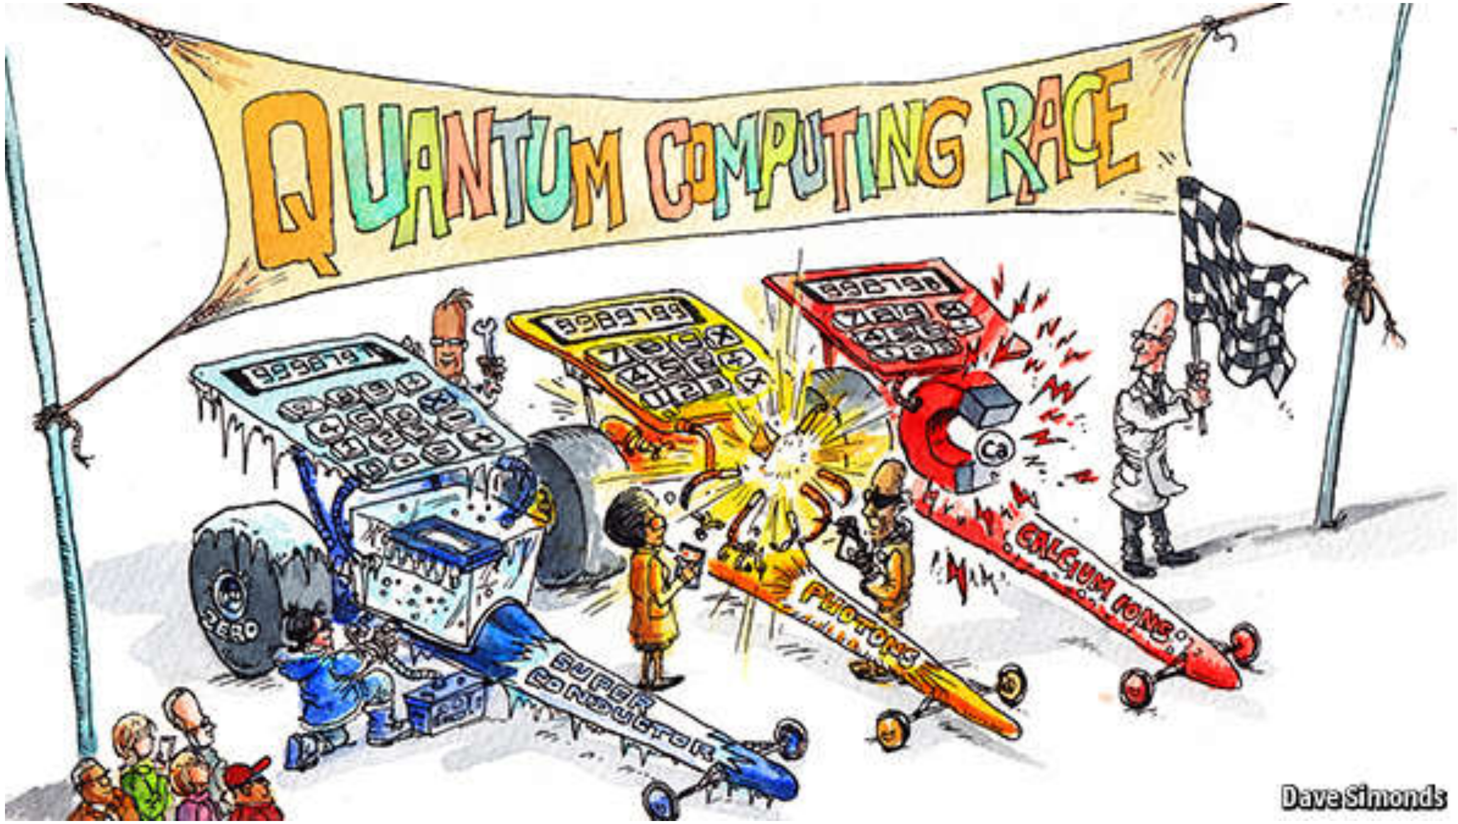

NISQ vs fault-tolerant quantum computing

NISQ (today → ~5 yr)

- $10^2 – 10^3$ physical qubits
- Two-qubit gate error $10^{-3} – 10^{-2}$
- Shallow circuits, hybrid with classical
- Variational, error mitigation
- Demonstrations, benchmarks, small problems

Fault-tolerant (10+ yr horizon)

- 10^6+ physical qubits → 10^3 logical
- Effective error 10^{-10} or better
- Deep circuits with error correction
- QPE, Trotterized dynamics, HHL
- Quantitative ab initio nuclear physics

Roadmaps are targets, not guarantees.

Preskill, Quantum 2, 79 (2018): "Quantum Computing in the NISQ Era and Beyond".

The Basics

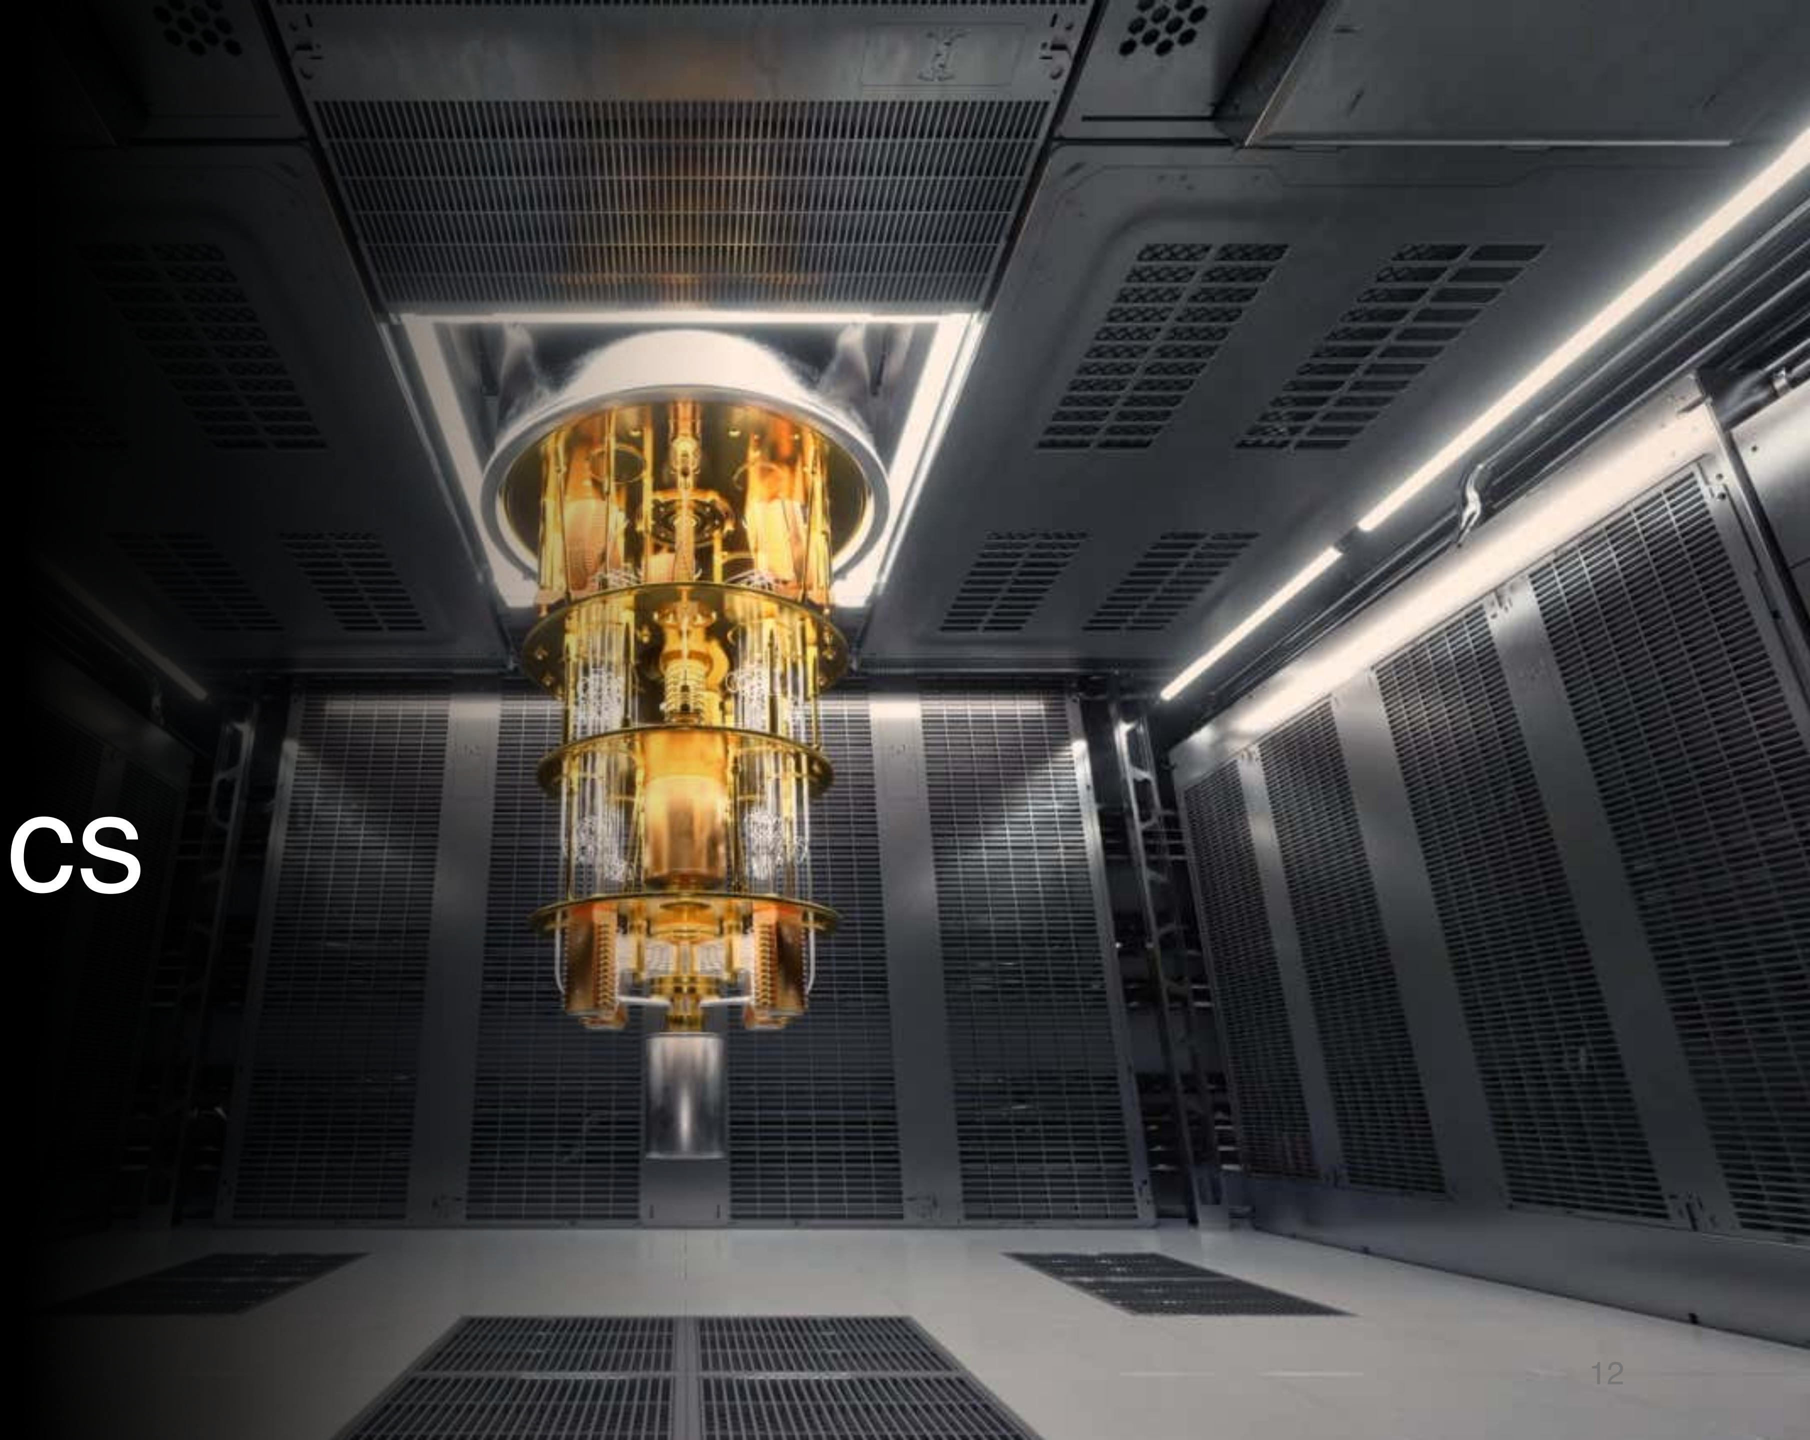

Requirements for a Quantum Computer

1. A scalable physical system with well characterized qubits
2. The ability to initialize the state of the qubits to a simple state
3. Long decoherence times, much longer than the gate operation time
4. A “universal” set of quantum gates
5. A qubit-specific measurement capability

Two criteria requiring the possibility to transmit information:

6. The ability to interconvert stationary and flying qubits.
7. The ability to faithfully transmit flying qubits between specified locations.

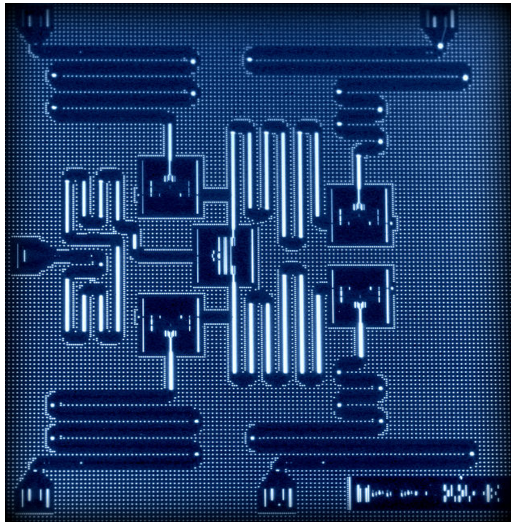

Layout of IBM’s five superconducting quantum bit device. (credit: IBM Research)

DiVincenzo, D. P. (2000). The physical implementation of quantum computation. *Fortschritte der Physik: Progress of Physics*, *48*(9-11), 771-783.

Postulates of Quantum Mechanics

- **State Space:** Describes the state of a closed system.
- **Evolution:** describes the evolution of a closed system.
- **Measurement:** describes how information is extracted from a closed system via interactions with an external system.
- **Composite systems:** describes the state of a composite system in terms of its component parts

The only known photo of Schrodinger's cat.

Quantum algorithms

A quantum algorithm consists of three basic steps:

- Encoding of the data, which could be classical or quantum, into the state of a set of input qubits.
- A sequence of quantum gates applied to this set of input qubits.
- Measurements of one or more of the qubits at the end to obtain a classically interpretable result.

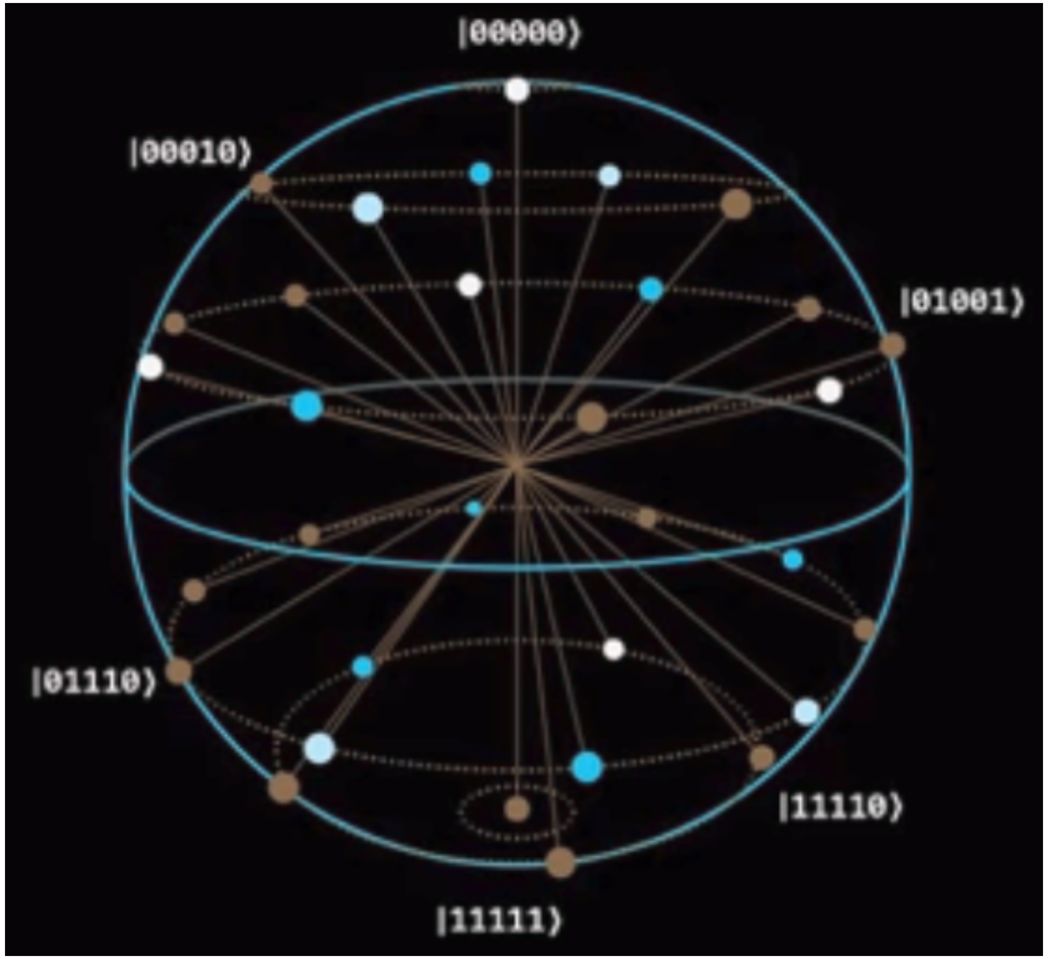

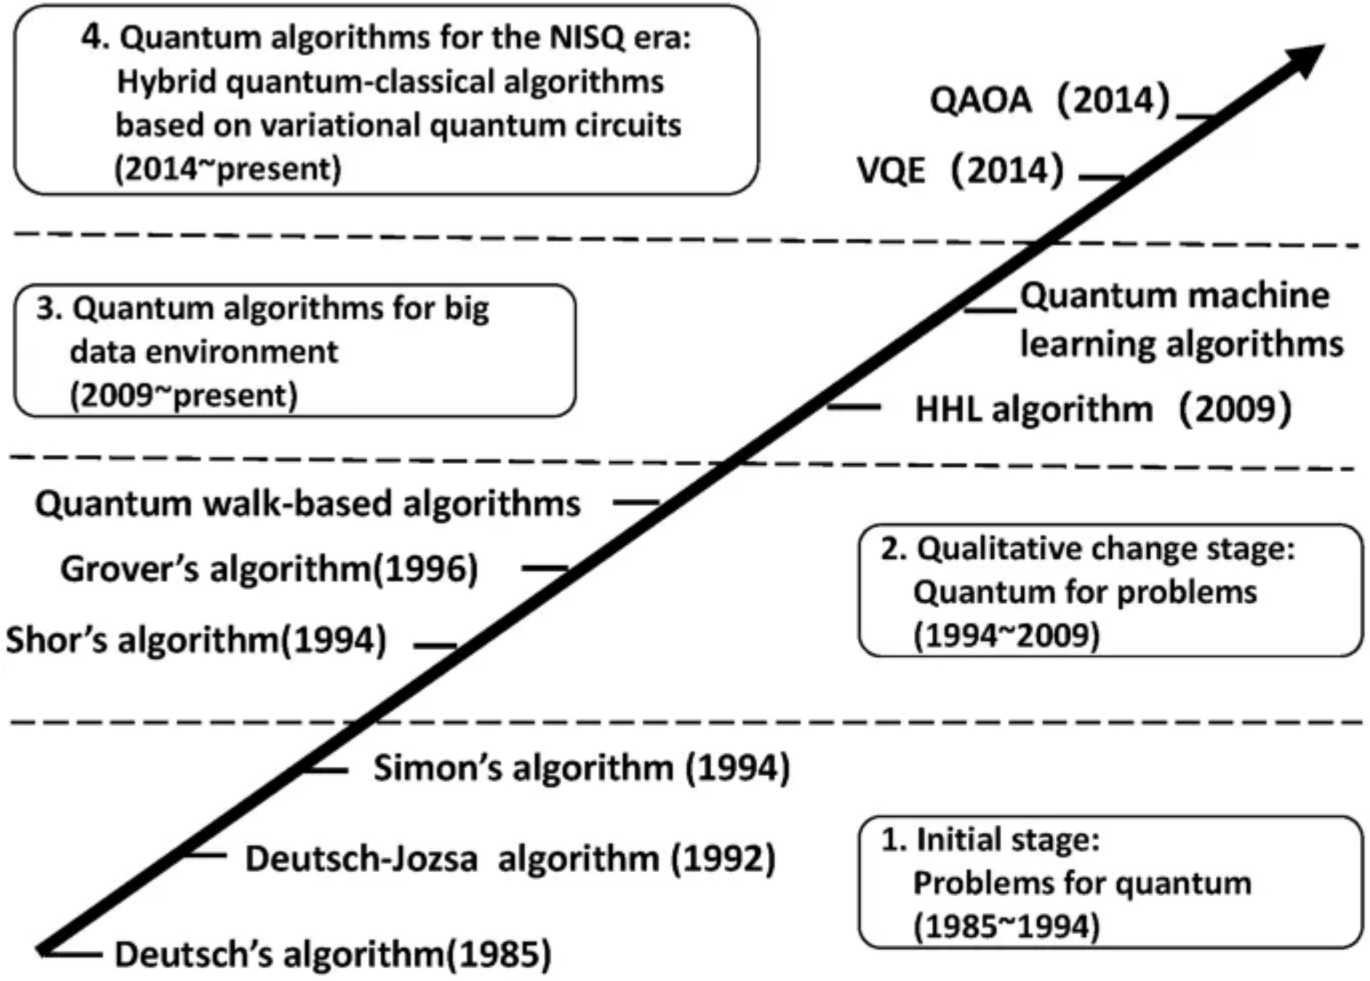

Zhang, S., & Li, L. (2022). A brief introduction to quantum algorithms. CCF Transactions on High Performance Computing, 4(1), 53-62.

Classical vs quantum bits

Classical bit (bit)

Only two possible states

0 and 1

- Like the two faces of a coin
- Heads or tails

Quantum bit (qubit)

$|0\rangle$ or $|1\rangle$

AND any linear combination of the two

e.g. 50% $|0\rangle$ & 50% $|1\rangle$

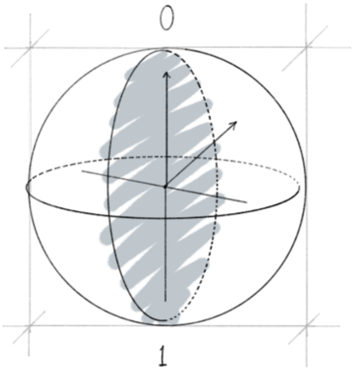

Quantum State (qubit)

Mathematically represented as a vector, or a point on the surface of the Bloch sphere:

$$\left|\psi\right\rangle = \underbrace{cos\left(\frac{\theta}{2}\right)}_{\alpha}\left|0\right\rangle + \underbrace{e^{i\varphi}sin\left(\frac{\theta}{2}\right)}_{\beta}\left|1\right\rangle \quad \left|\alpha\right|^2 + \left|\beta\right|^2 = 1$$

Measurement = projection of state to a basis vector (changes the state – superposition is destroyed)

Quantum gate is a transformation from one qubit state to another. Single-qubit gate = rotation around Bloch sphere. Reversible. Represented by a matrix (unitary, …) acting on the vector.

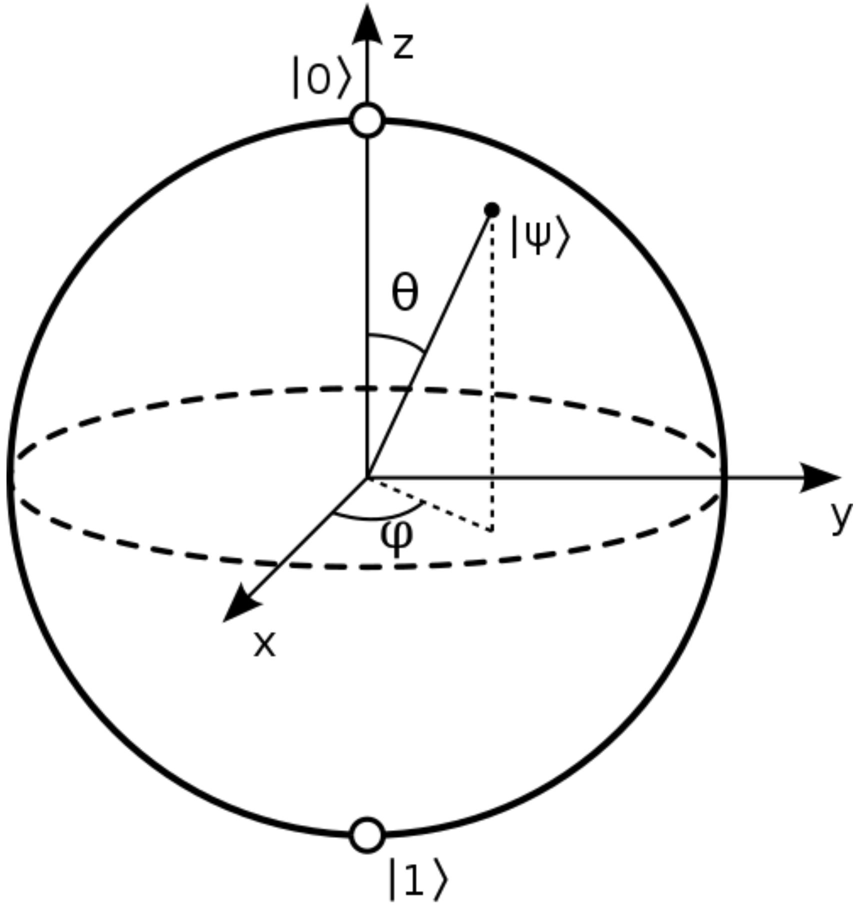

Qubit Gates

- A qubit gate is a black box transforming an input qubit $|\psi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$ into an output qubit $|\varphi\rangle = \alpha'_0|0\rangle + \alpha'_1|1\rangle$
- A gate G is represented by a $2x2$ transfer matrix with complex elements $a_{i,j}$ where $(i,j) \in \{1,2\}$:

$$G = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}$$

Recall $|a_0|^2 + |a_1|^2 = 1$ and $|a'_0|^2 + |a'_1|^2 = 1$

1. Matrix G is a unitary $G^\dagger G = \mathbb{I}$
2. The inverse of a unitary matrix G^{-1} is also unitary

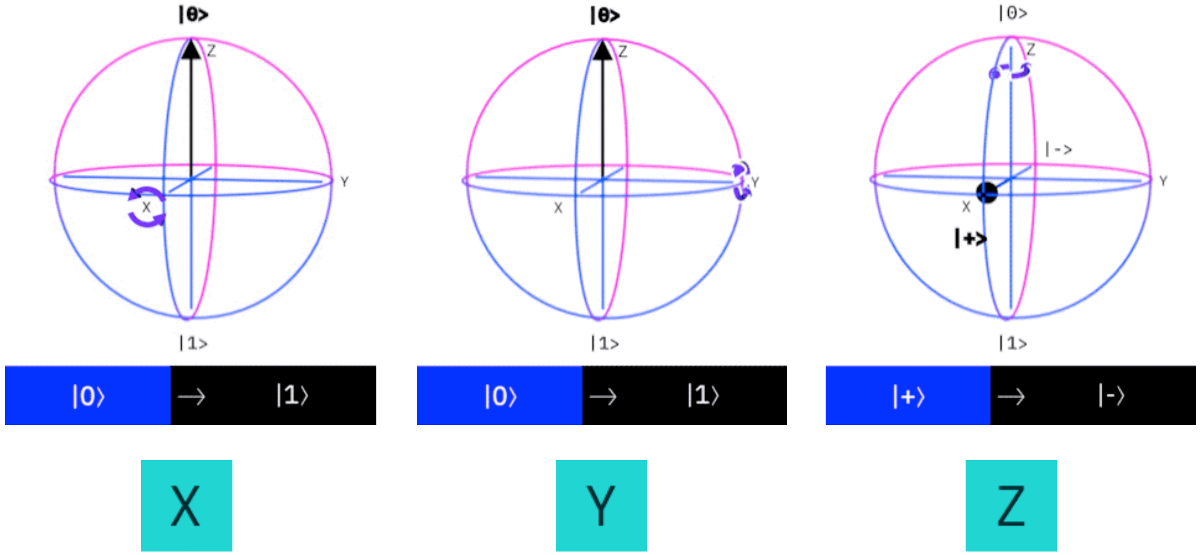

$$\sigma_0 = I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix} = \begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix} \longrightarrow |\phi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle.$$

$$\sigma_1 = X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}\begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix} = \begin{pmatrix} \alpha_1 \\ \alpha_0 \end{pmatrix} \longrightarrow |\phi\rangle = \alpha_1|0\rangle + \alpha_0|1\rangle.$$

$$\sigma_2 = Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}\begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix} = i\begin{pmatrix} -\alpha_1 \\ \alpha_0 \end{pmatrix} \longrightarrow |\phi\rangle = -i\alpha_1|0\rangle + i\alpha_0|1\rangle.$$

$$\sigma_3 = Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}\begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix} = \begin{pmatrix} \alpha_0 \\ -\alpha_1 \end{pmatrix} \longrightarrow |\phi\rangle = \alpha_0|0\rangle - \alpha_1|1\rangle.$$

Qubit Gates Phase: Z, S, '

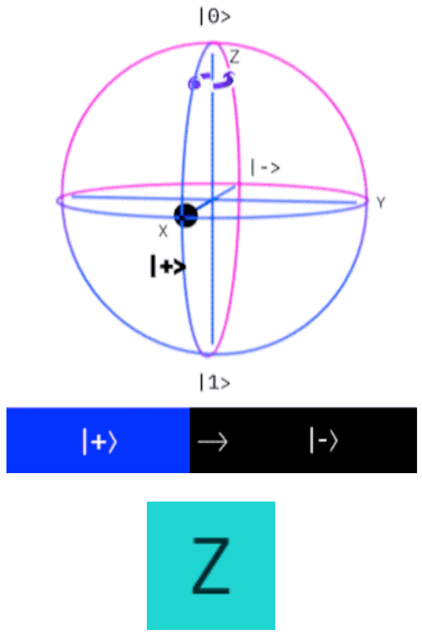

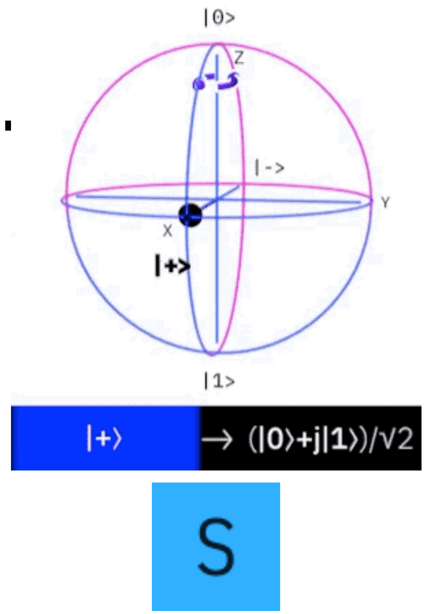

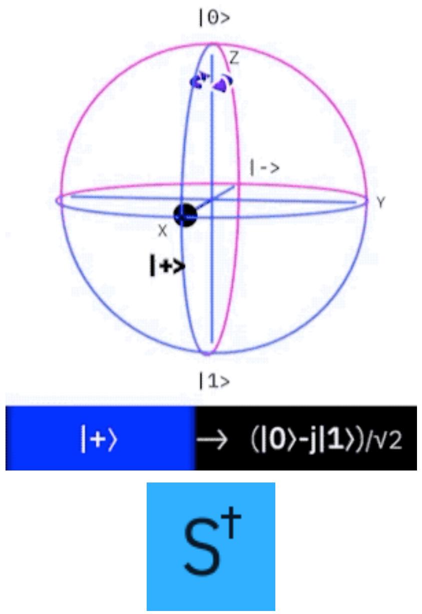

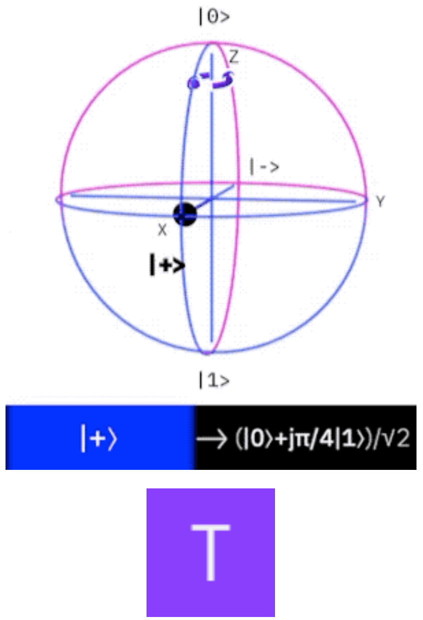

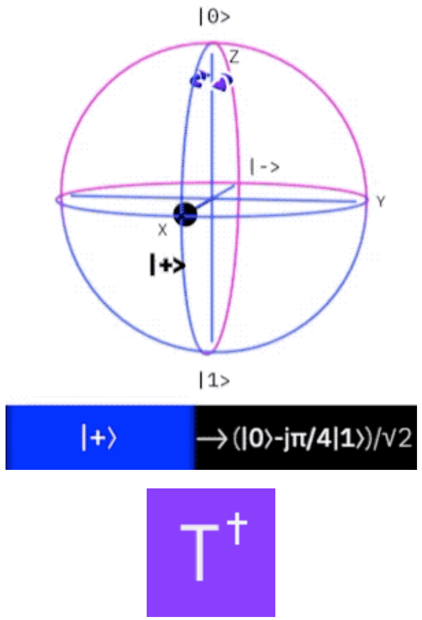

Hadamard (H) Gate: Superposition

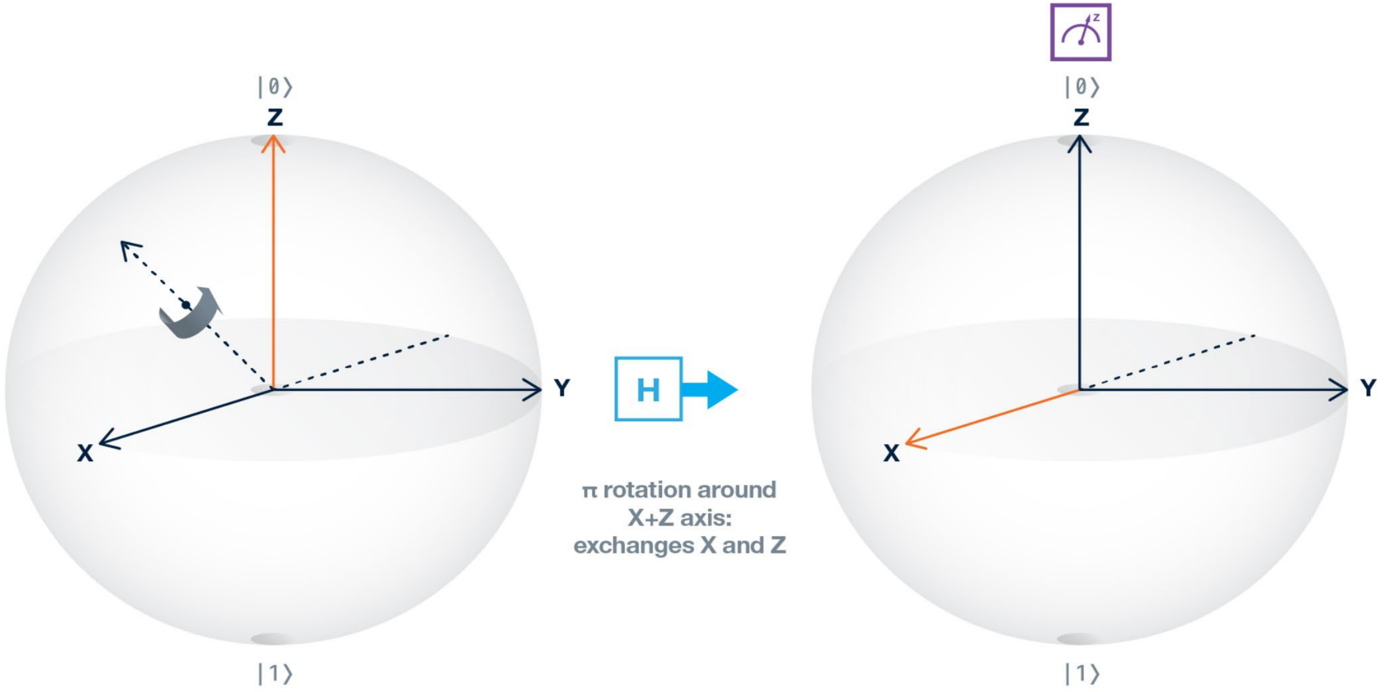

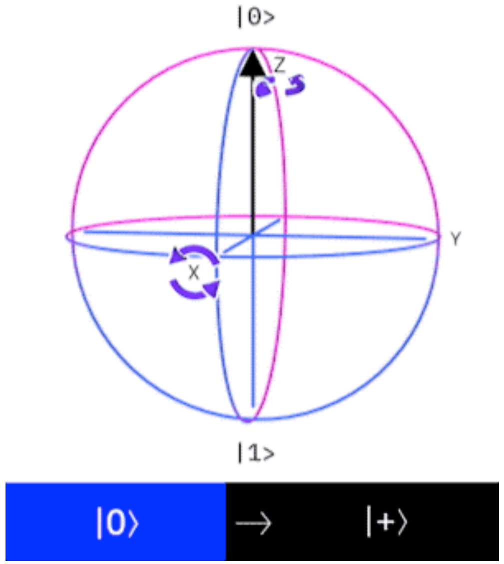

$$H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \qquad H|0\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}\begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}$$

Quantum Circuits

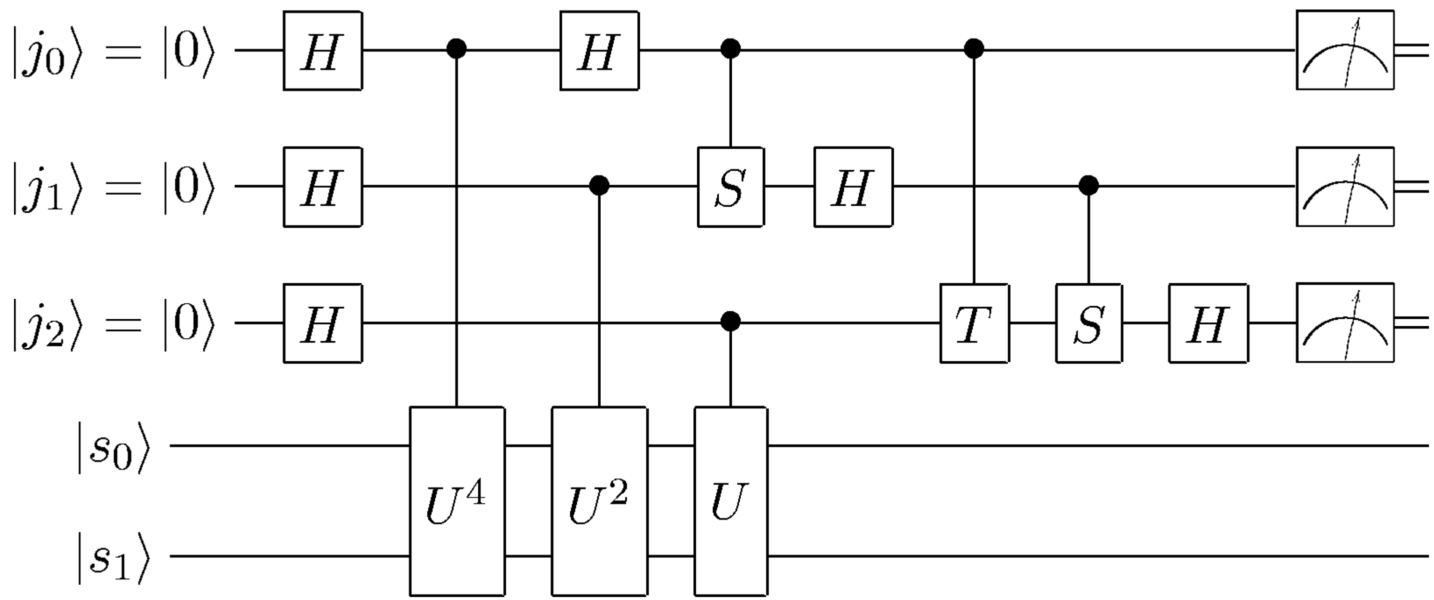

- Time flows left to right.
- Quantum gates (operators) are applied sequentially to qubit states, with result shown on the right.
- Measurement:
  - Double line represents classical bit.
  - Measurement is in the standard basis

Any quantum transformation can be realized in terms of the basic gates of the standard circuit model.

Notation for quantum circuits

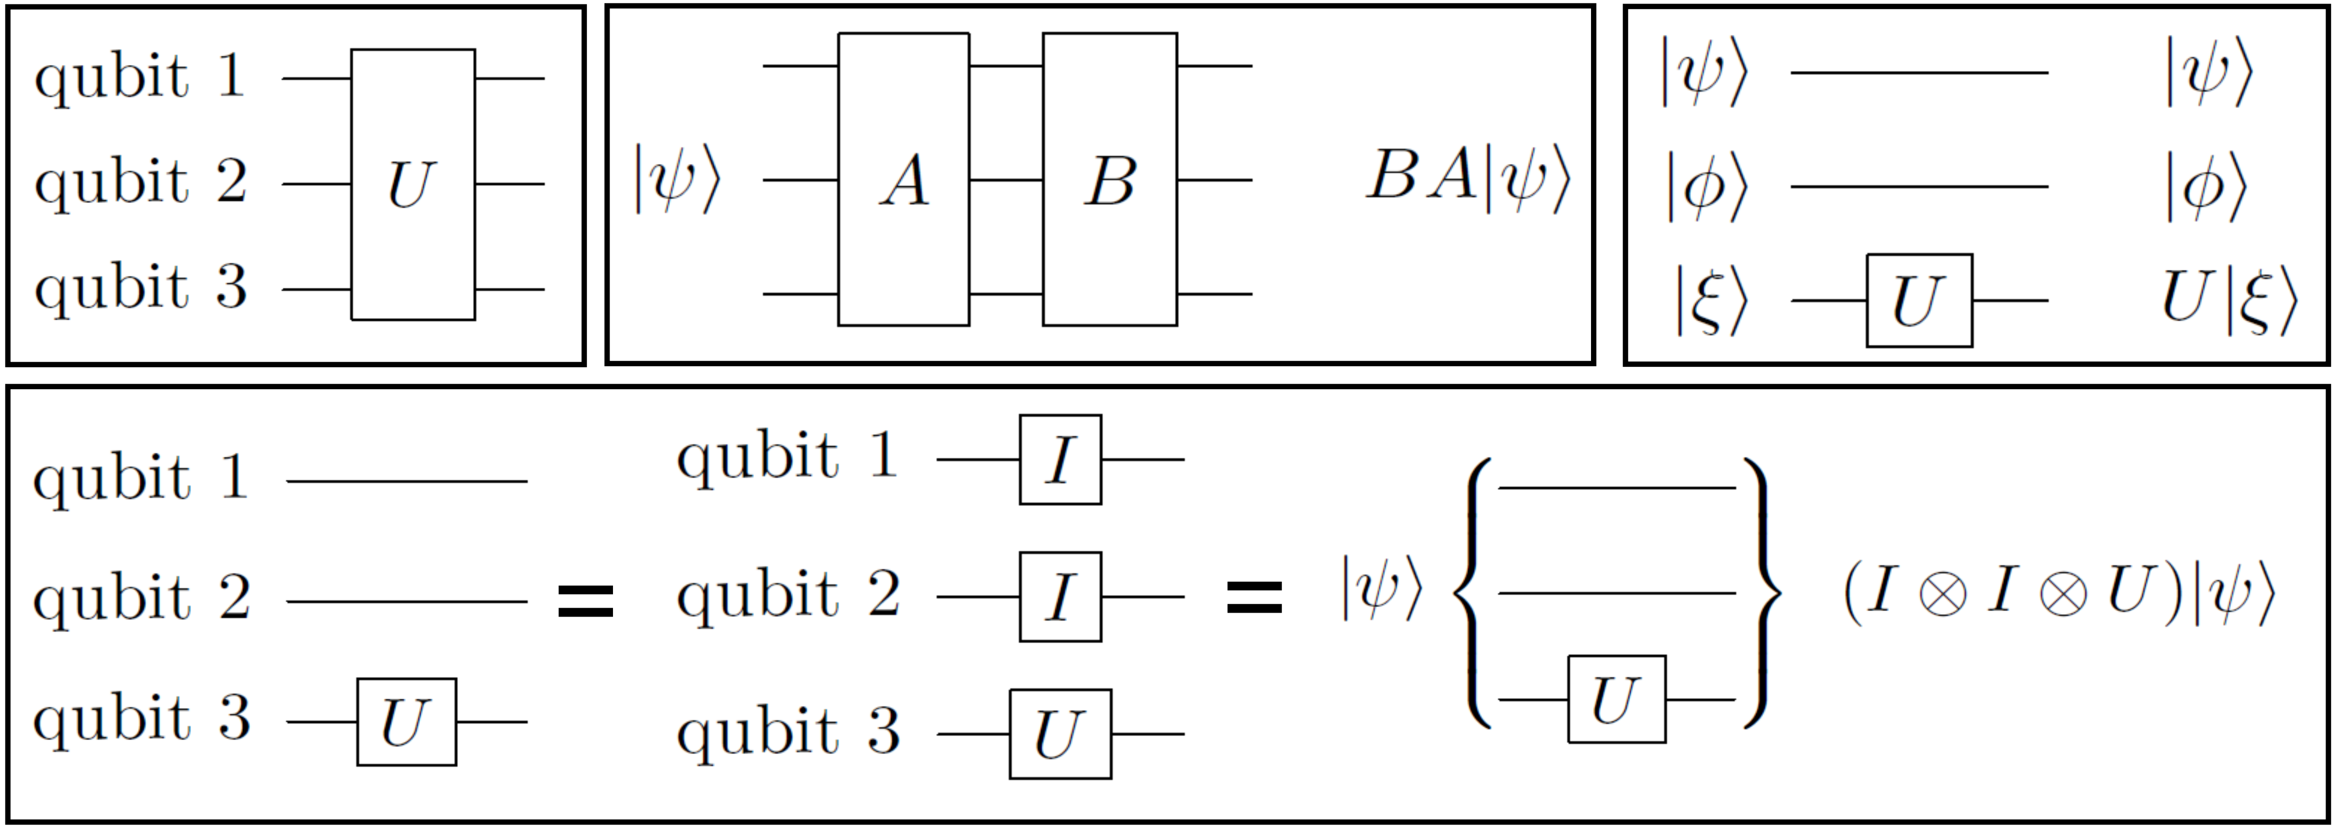

Walsh-Hadamard Transform

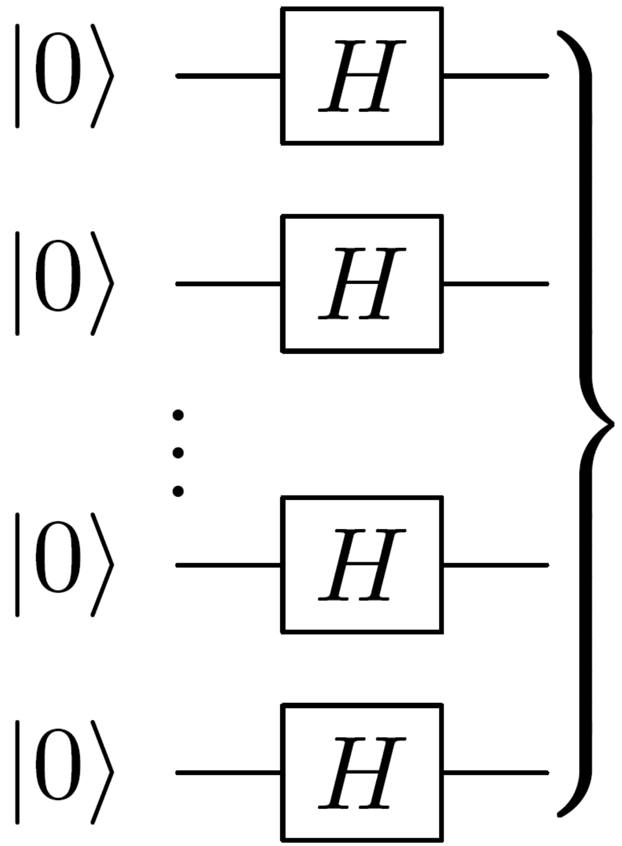

$$\frac{1}{\sqrt{2^n}} \sum_{x=0}^{2^n-1} |x\rangle$$

Used in the setup phase of algorithms, to create a superposition of all inputs.

Transformations occur on all components of the superposition. This is the source of quantum parallelism.

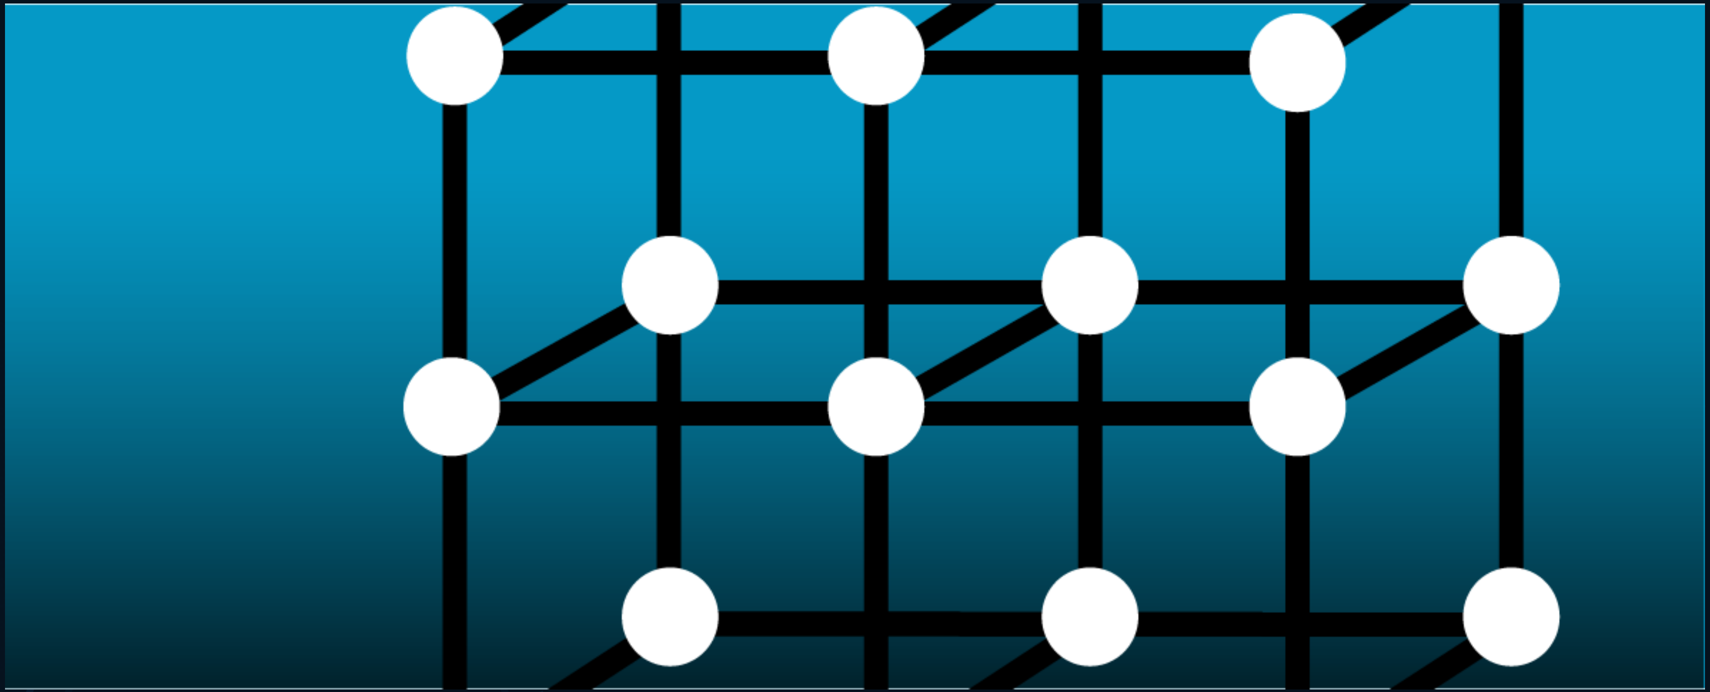

Quantum Simulation

Two styles: digital vs analog

- Digital (gate-based): write the evolution as a circuit of discrete gates; universal, programmable, and compatible with error correction
  - *Almost every nuclear-physics result so far is digital, on superconducting or trapped-ion hardware*
- Analog: engineer a physical device whose natural Hamiltonian is the one you want (cold atoms, ion crystals, Rydberg arrays)
  - *Fewer knobs, but far less overhead: a strong match for lattice gauge theories and spin models*
- Hybrid digital-analog approaches are an active middle ground

Daley et al., Nature 607, 667 (2022): practical quantum advantage via analog simulation.

Four questions you can ask a Hamiltonian

- **Static:** ground-state energy, spectrum, transition matrix elements $\langle \psi_n | O | \psi_m \rangle$ (structure)
- **Dynamic:** real-time evolution, response functions $S(\omega, q)$, transport, scattering (reactions)
- **Thermal:** $\rho = e^{-\beta H}/Z$ at finite temperature and density (hot/dense matter)
- **Open coupling to a continuum**: decay widths, three-body breakup, transport coefficients

Classical methods are excellent at the first and progressively fail down the list, quantum advantage grows in the same order

Carlson et al., Rev. Mod. Phys. 87, 1067 (2015).

Step 1: encoding fermions as qubits

Chemistry example of fermion-to-qubit mapping:

- A qubit register stores occupation numbers: each spin-orbital is filled ($|1\rangle$) or empty ($|0\rangle$)
- The fermion minus-sign (swap two particles → minus sign) must be built into the operators
- Jordan–Wigner: exact, but a single fermion operator becomes a long string of Z gates
- Qubit count ≈ number of orbitals, not the dimension of the many-body space

Jordan & Wigner, Z. Phys. 47, 631 (1928); Whitfield et al., Mol. Phys. 109, 735 (2011).
https://qiskit-community.github.io/qiskit-nature/tutorials/06_qubit_mappers.html

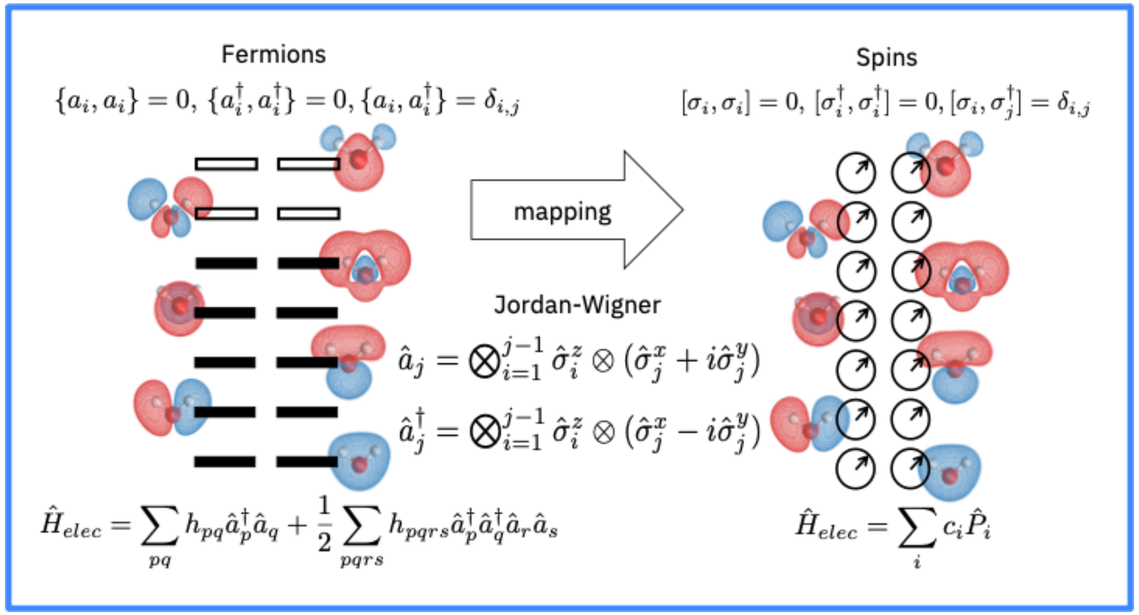

$$\hat{H} = \sum_{pq} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s .$$

where:

- $a^\dagger$ → adds an electron in an orbital
- a → removes an electron in an orbital
- h_{pq} and h_{pqrs} come from molecular integrals computed classically

Step 2: preparing the starting state

- Most algorithms need an initial state with good overlap onto the target (e.g. the true ground state)
- A Hartree–Fock / mean-field determinant is cheap and is the usual starting point
- Adiabatic preparation: start in an easy ground state, deform the Hamiltonian slowly toward the hard one
- For phase estimation, the success probability is the overlap $|\langle \psi_{trial} | \psi_0 \rangle|^2$ (poor overlap is a real bottleneck)

Tubman et al., arXiv:1809.05523 (2018); Lee et al., PRX Quantum 4, 020331 (2023).

Step 3: time evolution by Trotterization

$$e^{-i(A+B)t} \approx \left(e^{\frac{-iAt}{n}}\, e^{\frac{-iBt}{n}}\right)^n$$

- Split H into easy pieces and evolve under each in turn for a short slice t/n
- This is the split-operator idea: but each factor is now a quantum circuit
- First-order error is $O(t^2/n)$; higher-order (Suzuki) and randomized (qDRIFT) formulas cut the cost
- Empirical errors run far below worst-case bounds

Childs et al., PNAS 115, 9456 (2018); Campbell, PRL 123, 070503 (2019).

Reading out energies: Quantum Phase Estimation

$$U|\psi_k\rangle = e^{2\pi i \varphi_k}|\psi_k\rangle$$

- Feed in $U = e^{-iHt}$ and an approximate eigenstate, QPE writes the phase φ_k (hence the energy) into ancilla qubits
- QPE can achieve $\sim 1/\varepsilon$ in coherent evolution time, versus $\sim 1/\varepsilon^2$ shot scaling for direct expectation-value sampling
- The catch: long coherent circuits → a fault-tolerant-era tool
- **Lighter variants trade depth for more repetitions**

Kitaev, arXiv:quant-ph/9511026 (1995); Abrams & Lloyd, PRL 83, 5162 (1999).

The current era workhorse: VQE

$$E(\theta) = \left\langle \psi(\theta) \middle| H \middle| \psi(\theta) \right\rangle \geq E_0$$

- Prepare a parameterized trial state $|\psi(\theta)\rangle$, measure its energy, let a classical optimizer adjust θ
- The variational principle guarantees an upper bound on the true ground-state energy
- Ansatz design is everything: physics-inspired (UCC, ADAPT-VQE) beats generic circuits
- Honest limits: barren plateaus (vanishing gradients), hard optimization, no convergence guarantee

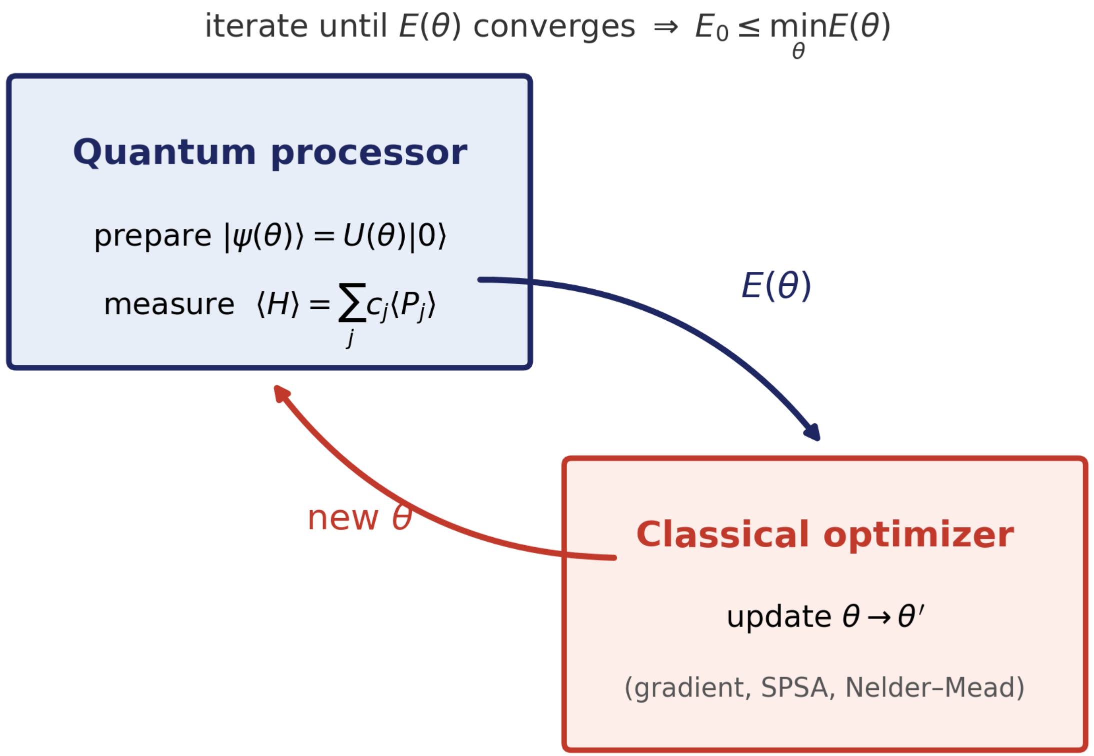

Peruzzo et al., Nat. Commun. 5, 4213 (2014); Tilly et al., Phys. Rep. 986, 1 (2022).

Everything is sampling

- Each run yields one bitstring; expectation values are estimated from many shots
- Statistical error falls as $1/\sqrt{N_{shots}}$, each extra digit costs $\sim 100\times$ more shots
- Nuclear energies are wanted to keV accuracy, so smart measurement is essential
- Tools: Pauli grouping, classical shadows, and (in the FT era) phase estimation to beat the $1/\varepsilon^2$ wall

Living with noise

- Zero-noise extrapolation: amplify the noise, then extrapolate back to the zero-noise limit
- Probabilistic error cancellation: characterize the noise channel and invert it statistically
- Symmetry verification: discard any shot that violates a conserved quantity (number, J, parity)

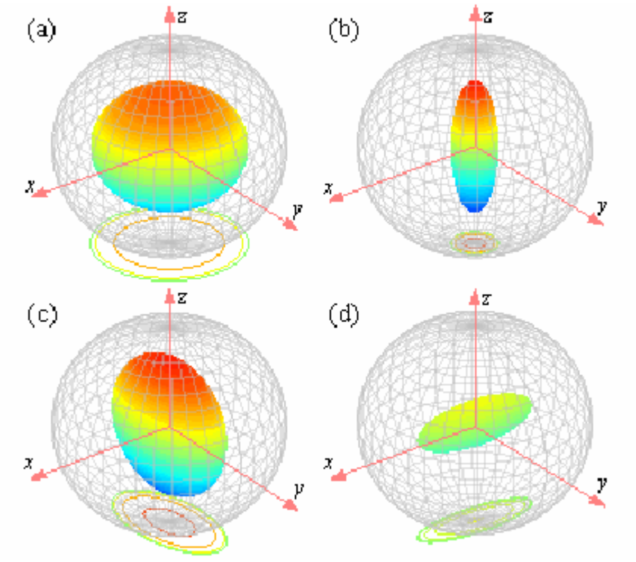

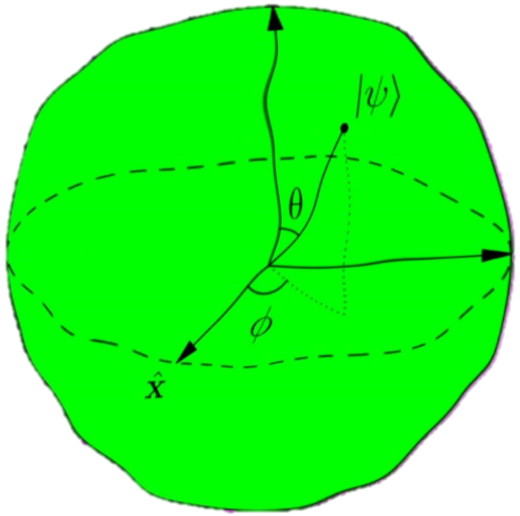

Teleportation of the one-qubit state in decoherence environments November 2011, Journal of Physics B Atomic Molecular and Optical Physics 44(2)

Fragment-Based quantum simulation

Radical-fragment MBE2

- Linear alkane C_nH_{2n+2} needs q = 14n + 4 qubits (JW, STO-3G): 368 qubits for $C_{26}H_{54}$, beyond any device
- Homolytic C–C cleavage → open-shell radical fragments (CH_3•, CH_2••): ROHF, no caps, no embedding
- MBE2 assembly: monomer energies + bonded-pair corrections; translational symmetry → only 4 unique fragments for any chain length
- Constant 30-qubit ceiling: 12.3× qubit reduction and 12.8× fewer unique calculations at C_{26}
- O(n) → O(1): the fragment, not the molecule, sets the quantum hardware budget

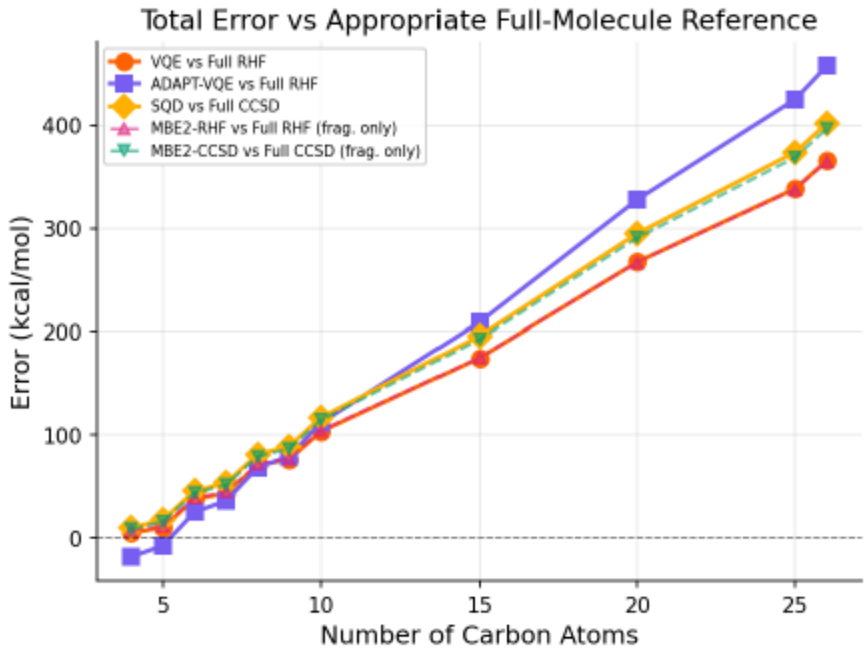

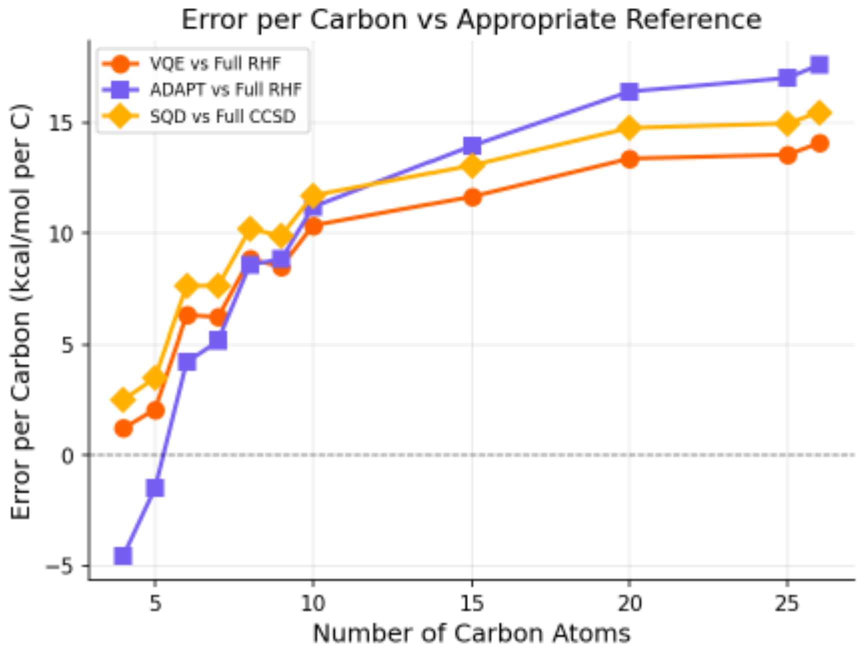

D. Sierra-Sosa, J. Saavedra, S. Solares, G. Toscano-Pulido — The Catholic University of America; run on IBM Quantum hardware.

VQE, ADAPT-VQE, and SQD inside MBE2

- Five solvers across 11 alkanes ($C_4 \rightarrow C_{26}$): RHF and CCSD references vs VQE, ADAPT-VQE, SQD
- MBE2-SQD on IBM Heron (ibm_pittsburgh): full 14–30-qubit fragment Hamiltonians track MBE2-CCSD within +1.4 to +5.2 kcal/mol
- Fragmentation, not the quantum solver, dominates the error budget; quantum results faithfully preserve the classical baseline
- Cautionary tale: ADAPT-VQE converged monomers tighter than dimers; the MBE2 formula amplifies the imbalance ×(n−3) → +457 kcal/mol by C_{26}
- Lesson: balanced accuracy across fragments beats maximal accuracy per fragment

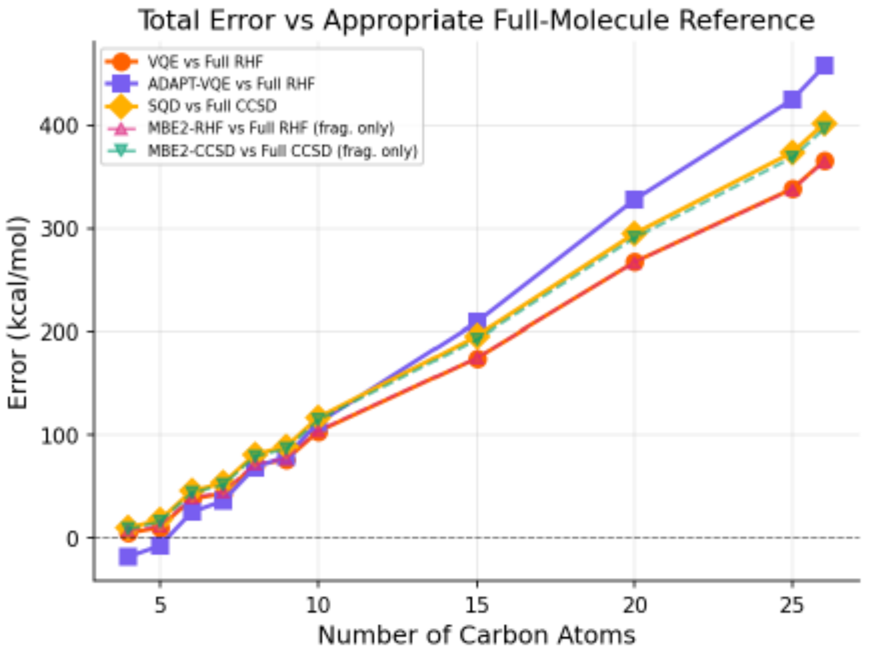

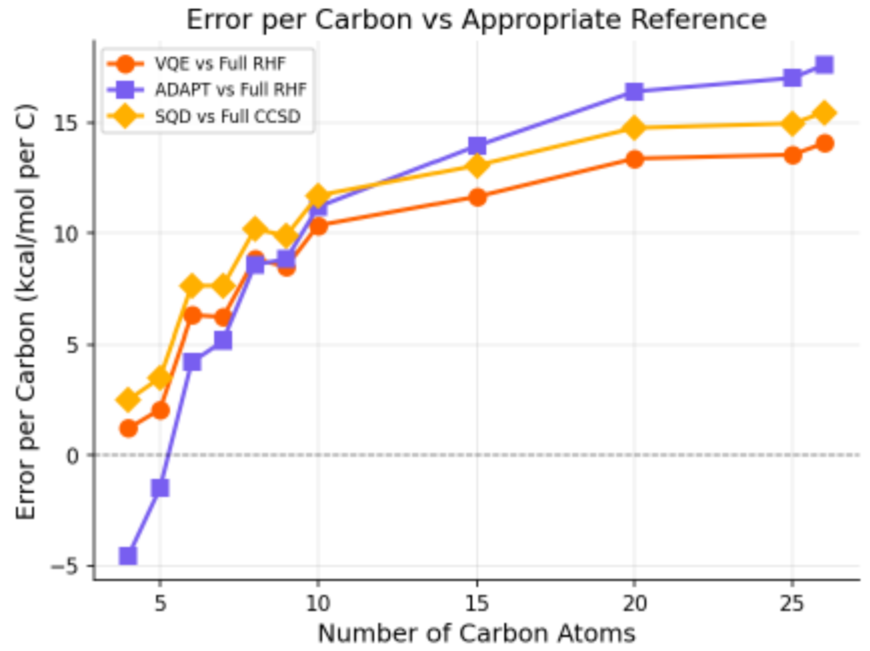

What this suggests for nuclear physics

- Same goal as valence-space downfolding: keep the quantum register fixed while the physical system grows
- SQD is the chemistry twin of the QSCI methods in shell-model demos (cf. slide 34): sample configurations on hardware, diagonalize classically
- The balanced-convergence lesson transfers: any assembly formula (MBE, cluster expansion) amplifies uneven solver errors
- Open-shell fragments with correct spin, no artificial caps: a natural mindset for pairing and odd-A subsystems
- An invitation: port the radical-fragment playbook to nucleon clusters and pairing Hamiltonians

Our ongoing directions

- Template-driven hybrid algorithms (IBM Quantum): pre-optimized fragment wavefunctions assembled into long chains, then globally refined — from ~20 toward 50–100 qubits
- From σ to π: alkanes → polyacetylene; delocalized, strongly correlated electrons as the stress test (conducting polymers)
- Pipelines: VQE, VarQITE, QAOA — with machine-learning-guided encodings and ansatz selection
- Quantum optimization taxonomy (IEEE Access survey): hardware / pipeline / kernel / order / continuity / metaheuristics — a map for choosing the right tool
- The nuclear analog of all of the above is unwritten — collaboration welcome

Future Directions

Nuclear Problem Map

Target	Classical bottleneck	Quantum-computing role	Likely era
Bound states / spectra	Hilbert-space growth; correlation; sign structure	Small benchmarks now; QPE later for sharp energies	NISQ → FT
Response / scattering	Real-time dynamics and continuum coupling	Natural unitary time evolution; phase estimation / LCU / qubitization in FT era	Mostly FT
Finite density / hot matter	Monte Carlo sign problem and analytic continuation	Thermal-state preparation and real-time correlators	FT / analog
Gauge dynamics / QCD	Large local Hilbert spaces; gauge constraints; non-equilibrium dynamics	Digital or analog lattice-gauge simulation with symmetry-preserving encodings	Co-design now, FT later
Neutrinos / spin models	Collective many-body dynamics can grow rapidly with modes	Analog/digital simulators can directly implement spin Hamiltonians	NISQ demos → FT

Beyond simulation

- Hybrid quantum ML replaces a targeted neural-network layer with a parameterized quantum circuit (PQC), exploiting Hilbert-space feature maps and entanglement-based correlations
- The open question is a confounder: does the quantum layer genuinely help, or was a single circuit simply hand-picked well? Most studies cannot separate the two
- This bites hardest in nuclear / hadronic physics, where detector-level classification (particle ID) directly limits the precision of extracted QCD observables
- Any advantage claim must beat a strong classical baseline and a fair sample of circuit designs, not one hand-designed ansatz

Removing the human from circuit design

- Quantum Neural Architecture Search: a reinforcement-learning agent edits a candidate PQC step by step: encoding, ansatz depth, gate insert/delete, qubit rewiring, entanglement pattern
- A graph-neural-network surrogate predicts a candidate's performance from its circuit structure, scoring ~50× faster than training
- Train classical, validate quantum: every candidate is trained on statevector simulation; scarce QPU time (IBM Heron via Qiskit Runtime) is reserved for final validation of top models

Summary

- Nuclear physics is a quantum many-body problem; classical methods face exponential Hilbert-space growth.
- The key hard regimes also include sign problems, real-time dynamics, continuum/open-system physics, and dense matter.
- Quantum computers represent amplitudes natively and can simulate few-body/local Hamiltonians with polynomial overhead.
- Near-term value: small demonstrations, hybrid solvers, SQD/QSCI-style subspace methods, and error mitigation.
- Fault-tolerant value: QPE and long real-time evolution for response, scattering, gauge dynamics, and QCD.
- First real advantage will likely come from real-time or finite-density problems, not easy small ground-state calculations.

Questions?