introduction to quantum computing with aws braket

brettkoonce.com/talks

january 27th, 2021

overview

- quantum computing theory overview
- quantum hardware/simulators, software
- applications: entanglement, chemistry, grover, optimization, graphs, neural networks
- quantum supremacy, future

quantum mechanics

- schrödinger: $E\Psi = \hat{H}\Psi$
- wavefunctions: state (Ψ) + conjugate (Ψ^*)
- <bra|ket>: $< \Psi \mid H \mid \Psi >$
- linear operators: $\hat{H} == A \bullet (\Psi \bullet \Psi^* == [I])$

self study

- Quantum Physics 130 (ucsd)
- Quantum Mechanics in Simple Matrix Form / Linear Operators for Quantum Mechanics (Jordan)
- Differential Forms with Applications to the Physical Sciences (Flanders) / Linear and Geometric Algebra / Vector and Geometric Calculus (Macdonald)

quantum computing

- scott aaronson: arXiv:1607.05256 (chapter 1+2)
- np complete problems
- sat solvers vs oracles
- linear vs exponential complexity

aws braket

- rigetti: 30 qubits
- ion trap: 11 qubits
- simulator: sv1: 34 qubits, tn1: 50 'qubits'
- dwave: quantum annealing, 2k/5k 'qubits'

cirq + qsim

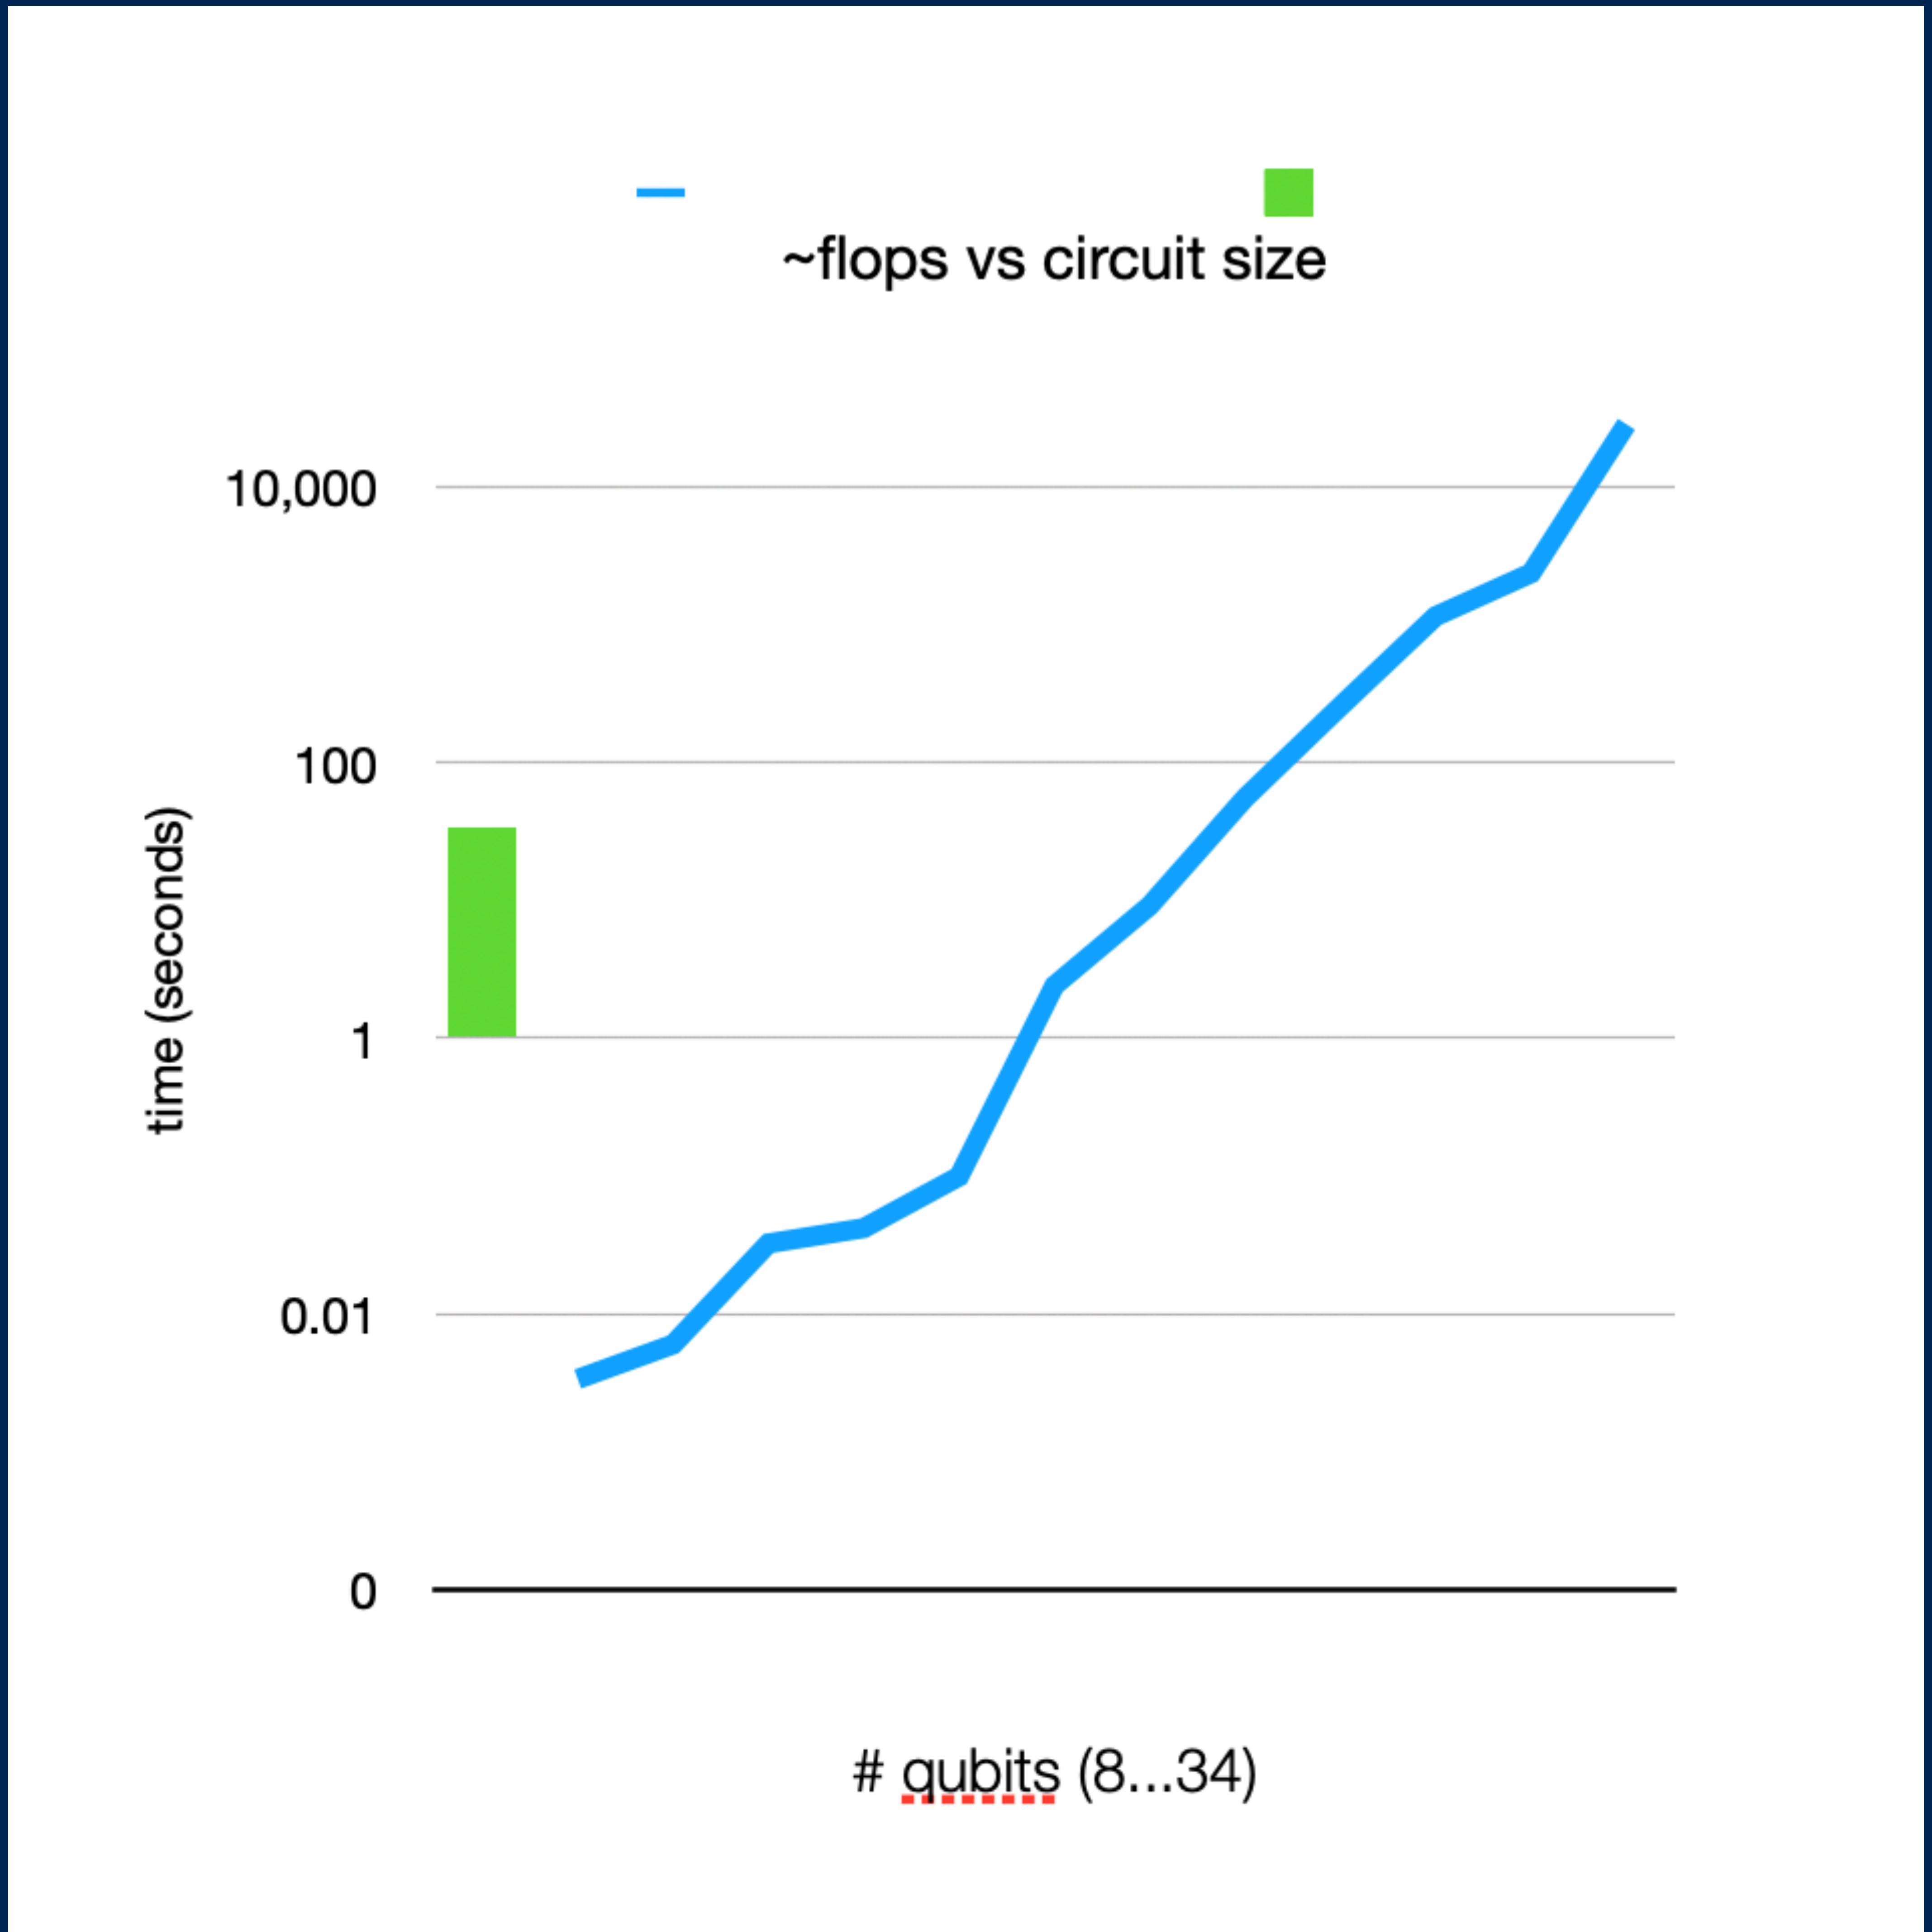
~flops vs circuit size
10,000
100
1
0.01
0
time (seconds)
qubits (8...34)

entanglement

- bell's theorem
- entanglement
- rigetti demo

quantum chemistry

- physical chemistry
- simulate hydrogen atoms
- measure hartree-fock energy directly
- pennylane vqe demo: h2

variational quantum eigensolver

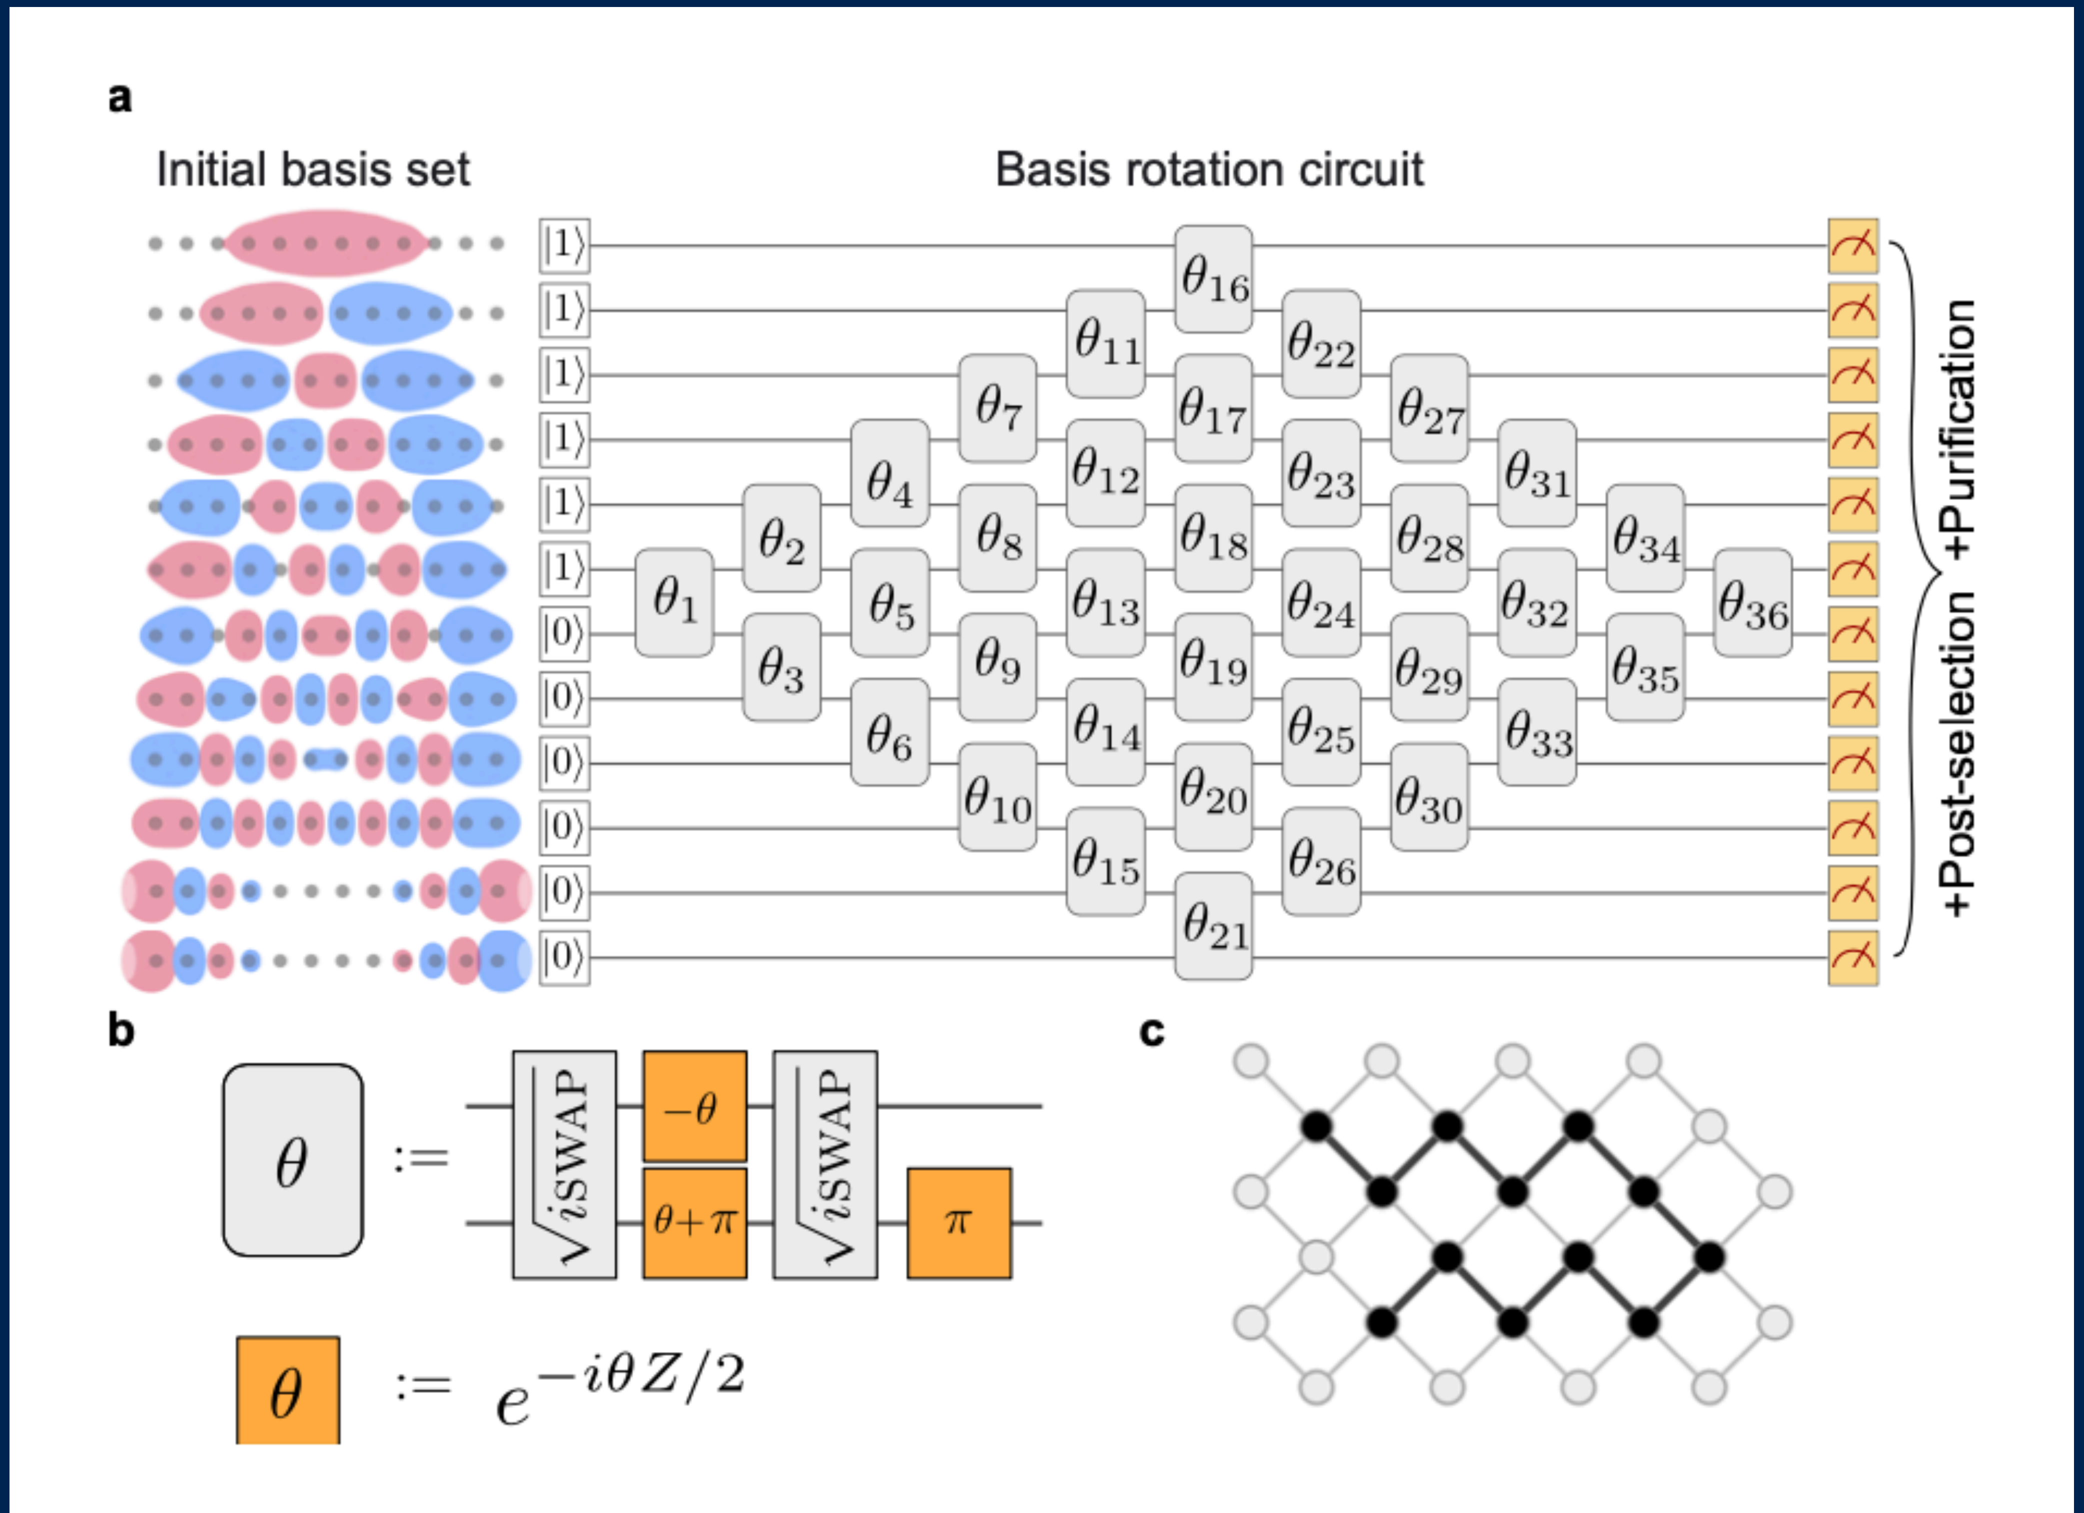

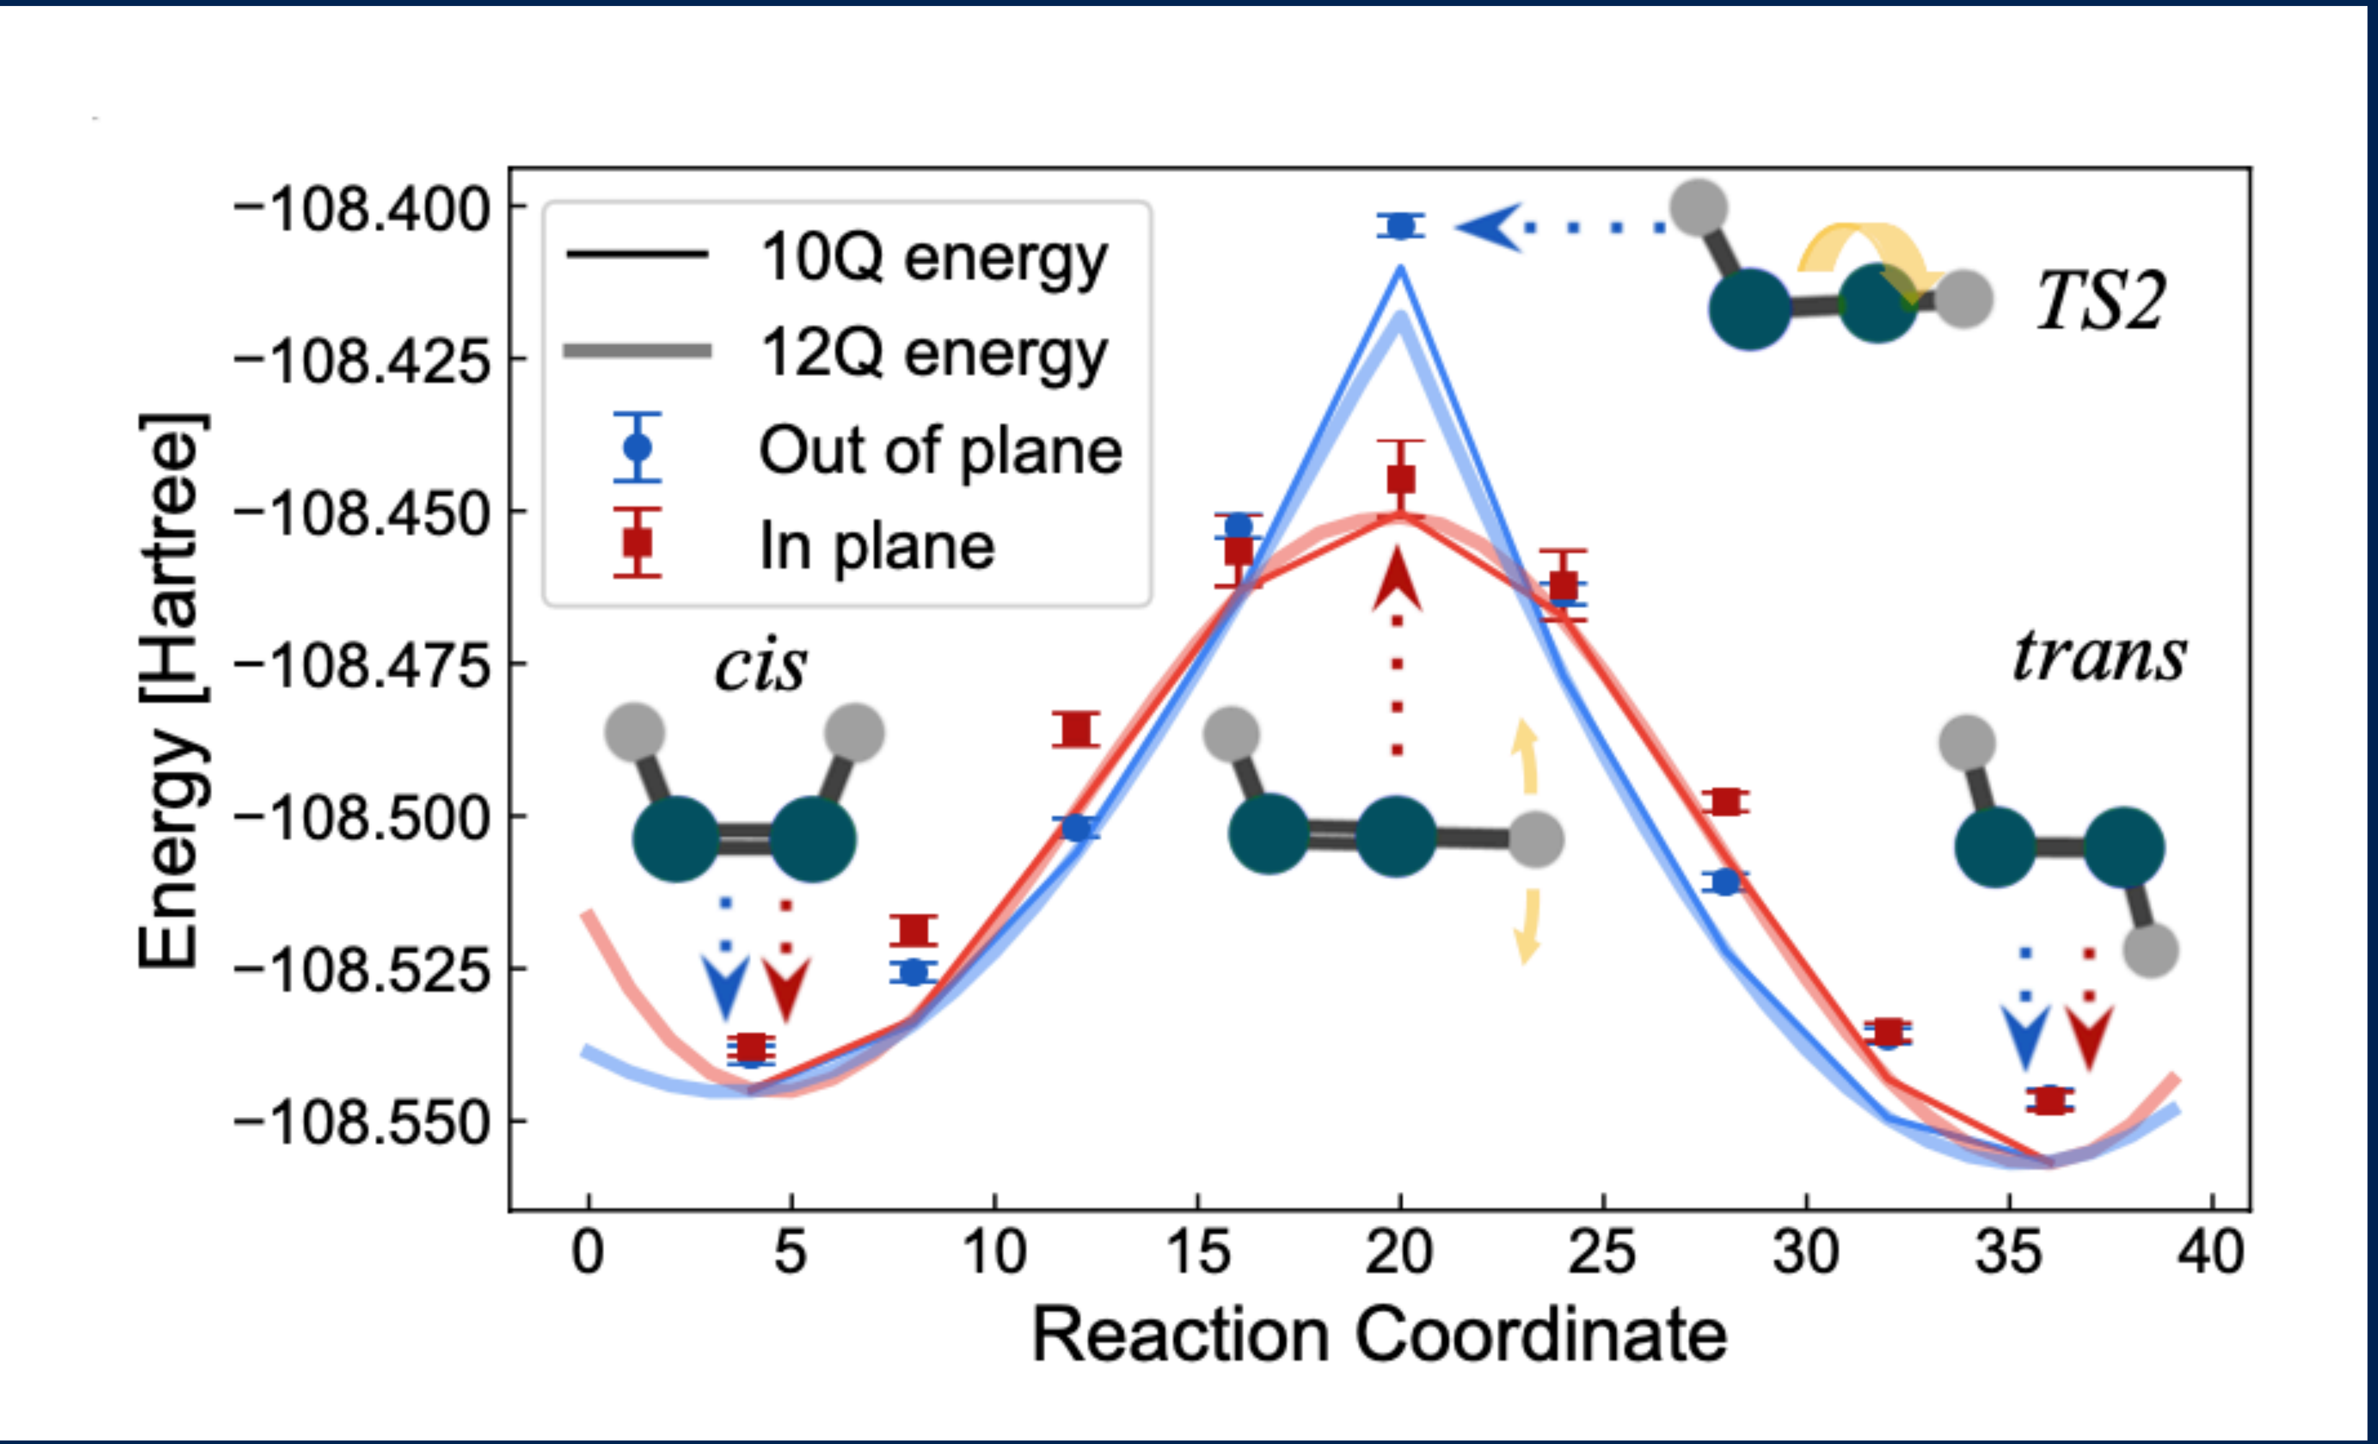

- **arxiv:2004.04174**

grover's algorithm

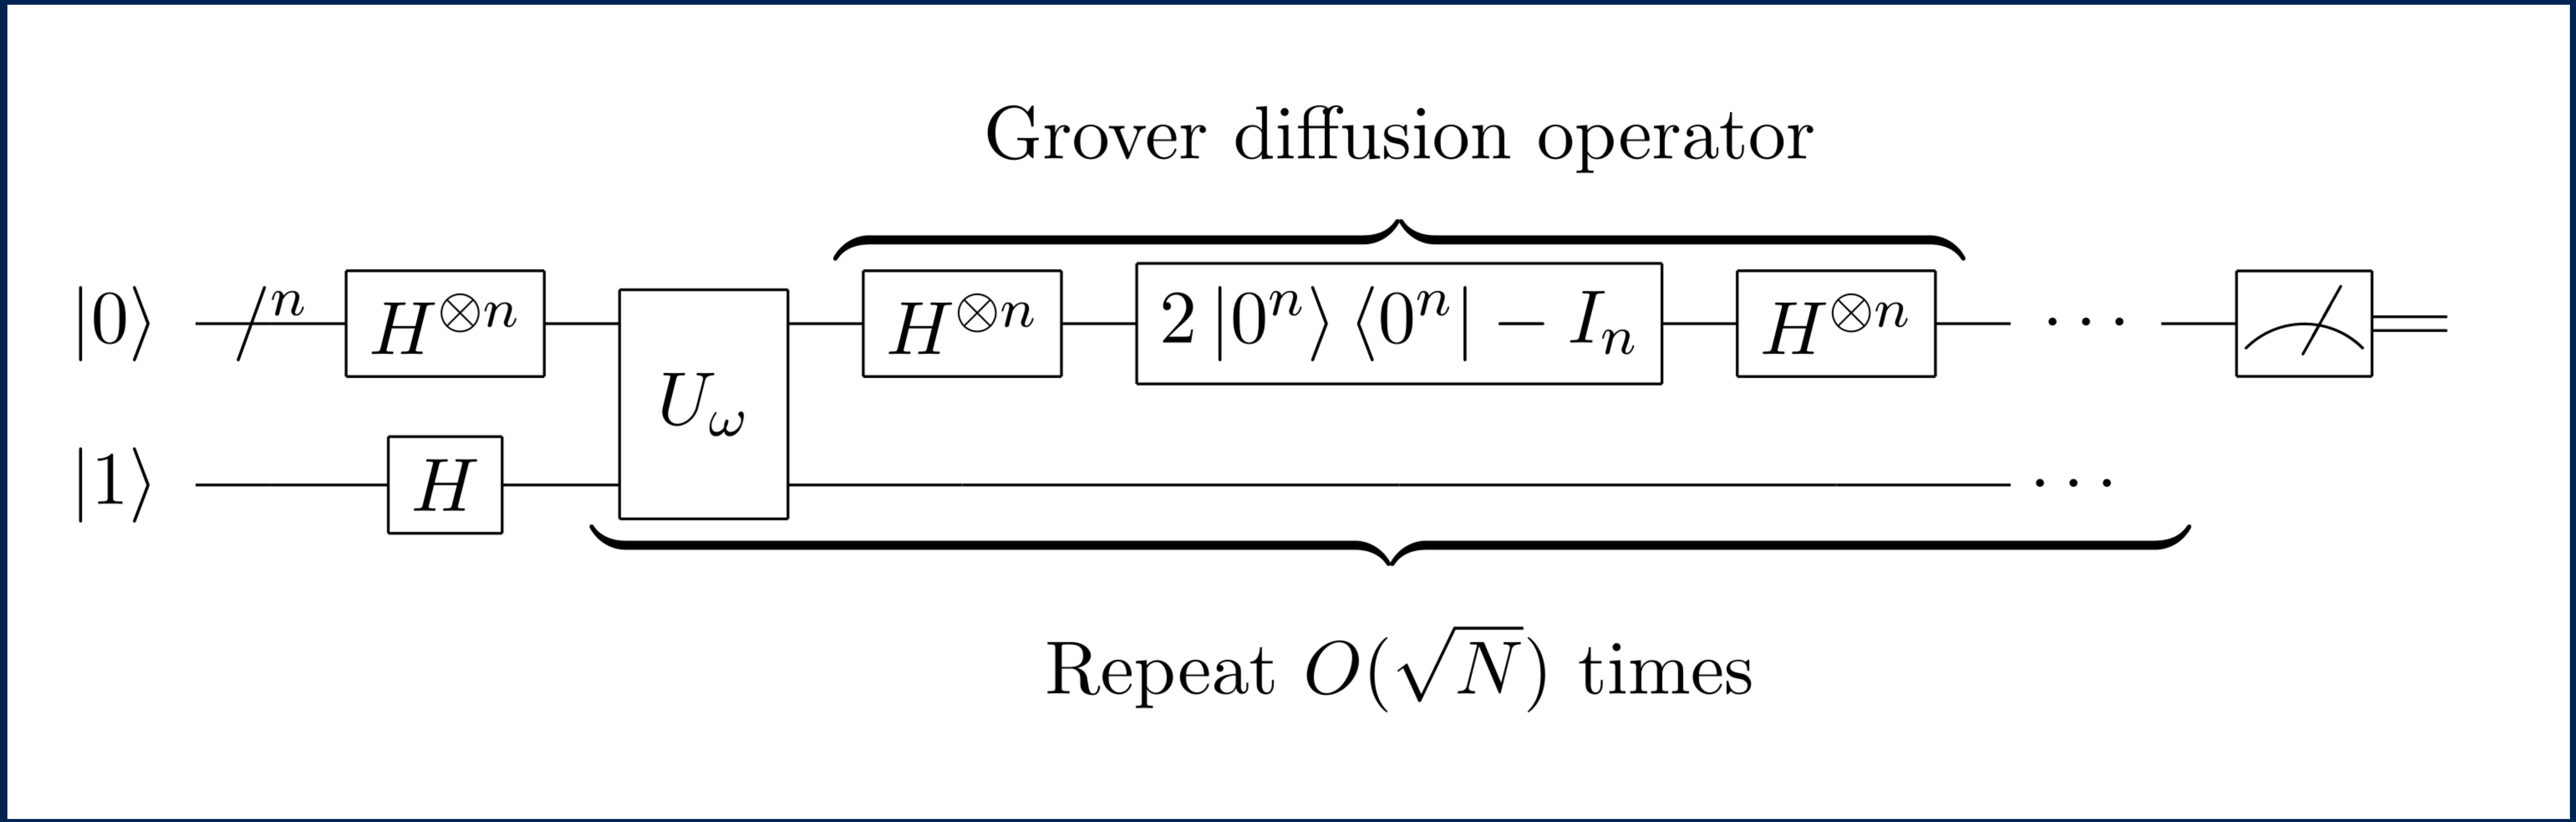

- **ion trap demo**

qubo/ising

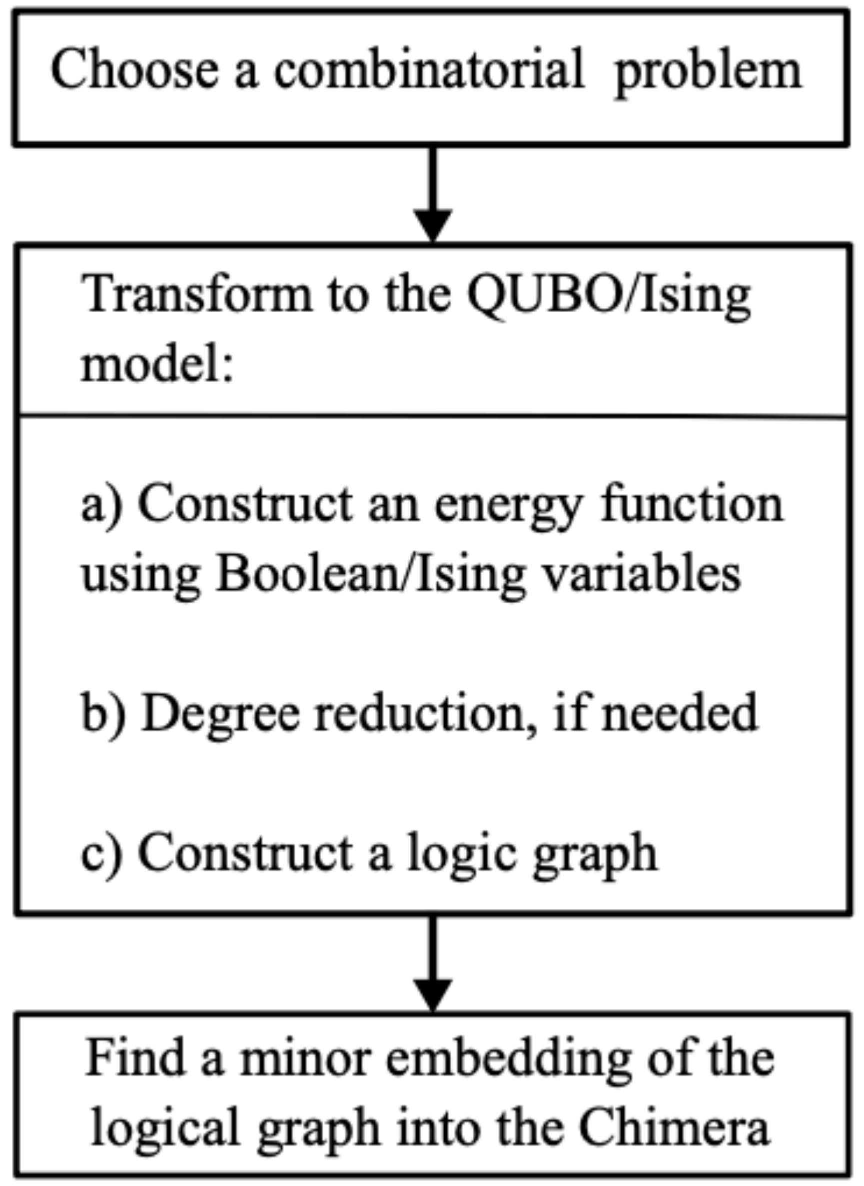

Figure 8. Main steps for programming a D-Wave system.

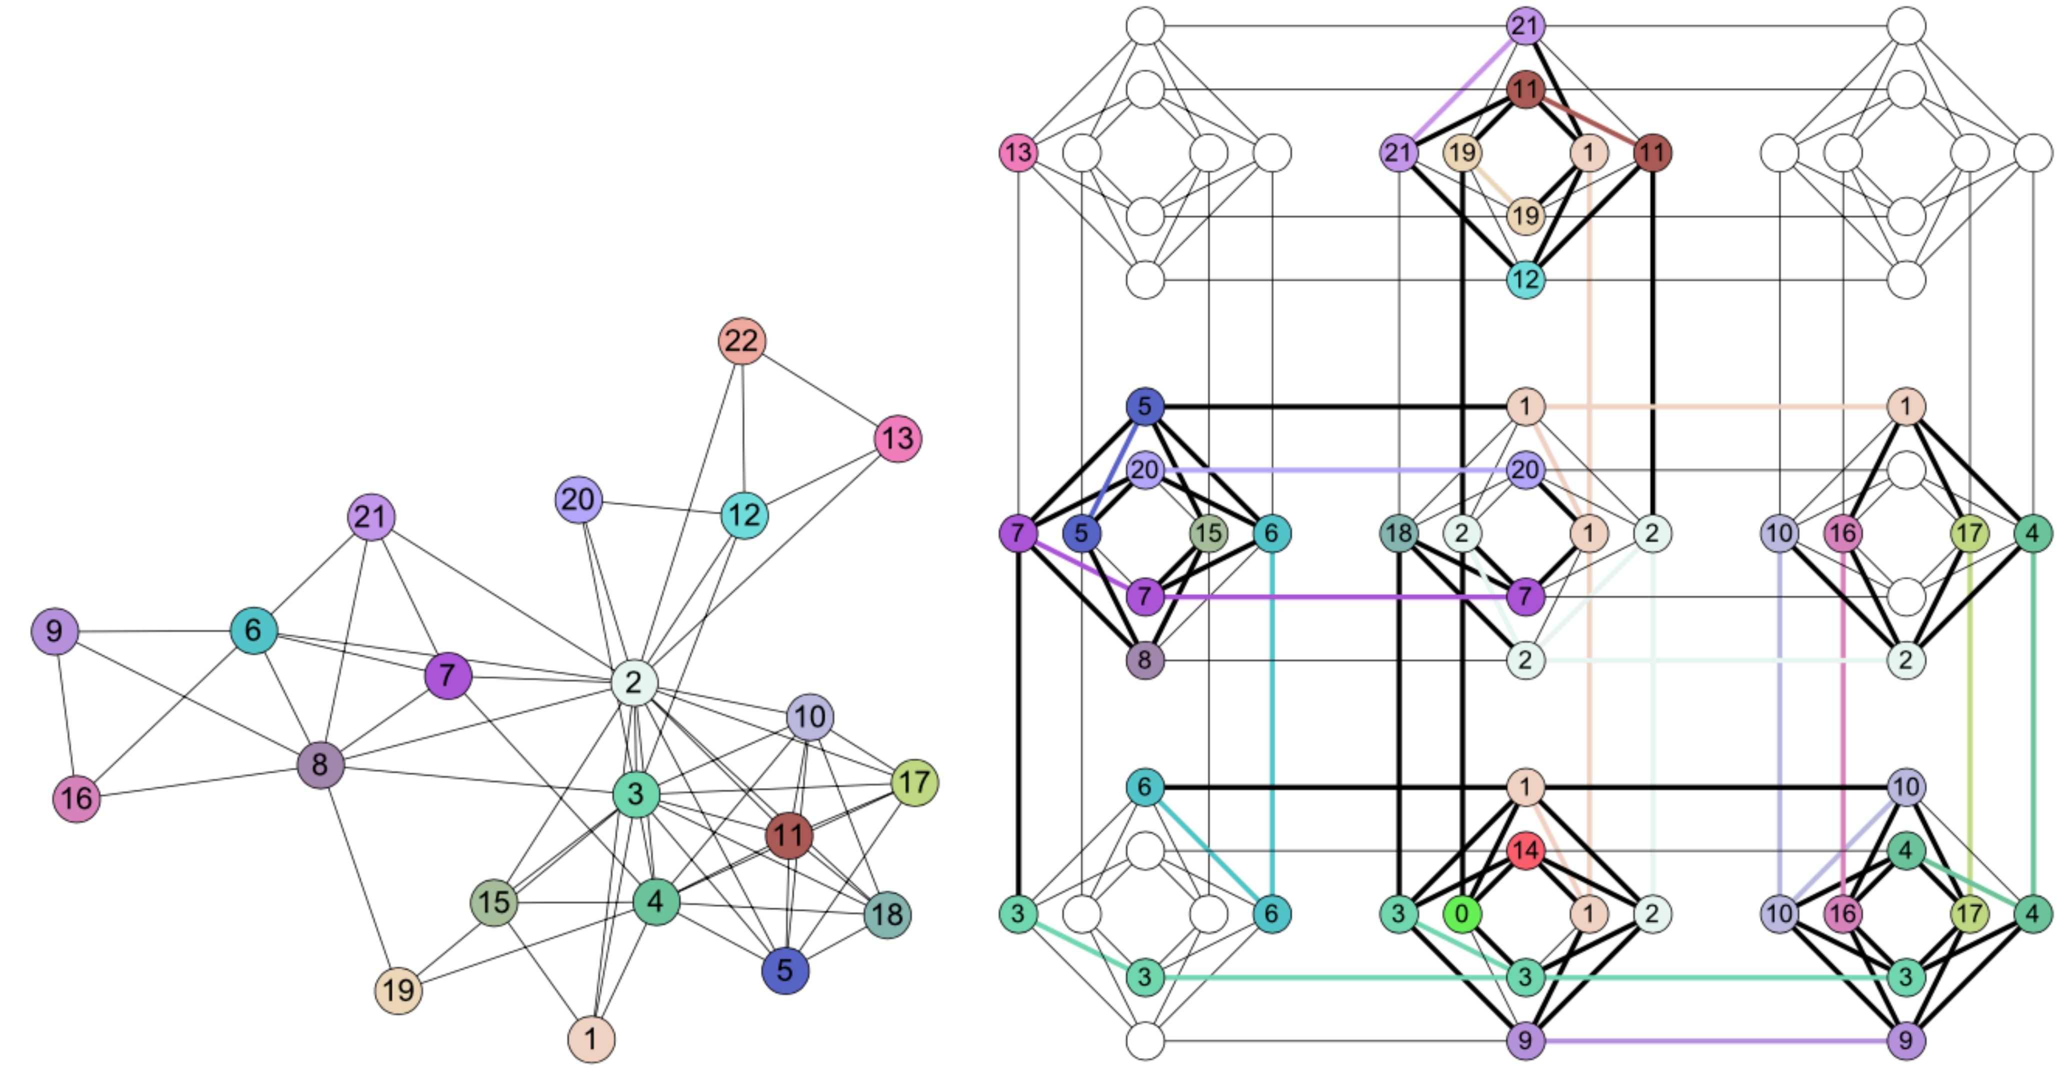

Figure 12. (left) Logical Ising graph for the QUBO function h_G in Eq. (17) and (right) a minor embedding of this graph. The logical graph has 23 variables and its minor embedding requires 51 physical qubits. The numbers and colors of the vertices in the logical graph are the same as in the minor embedding. Bold black lines correspond to the mapped edges and bold color lines correspond to chain of qubits.

- **arxiv:1803.03372**

arxiv:1604.05796

factoring

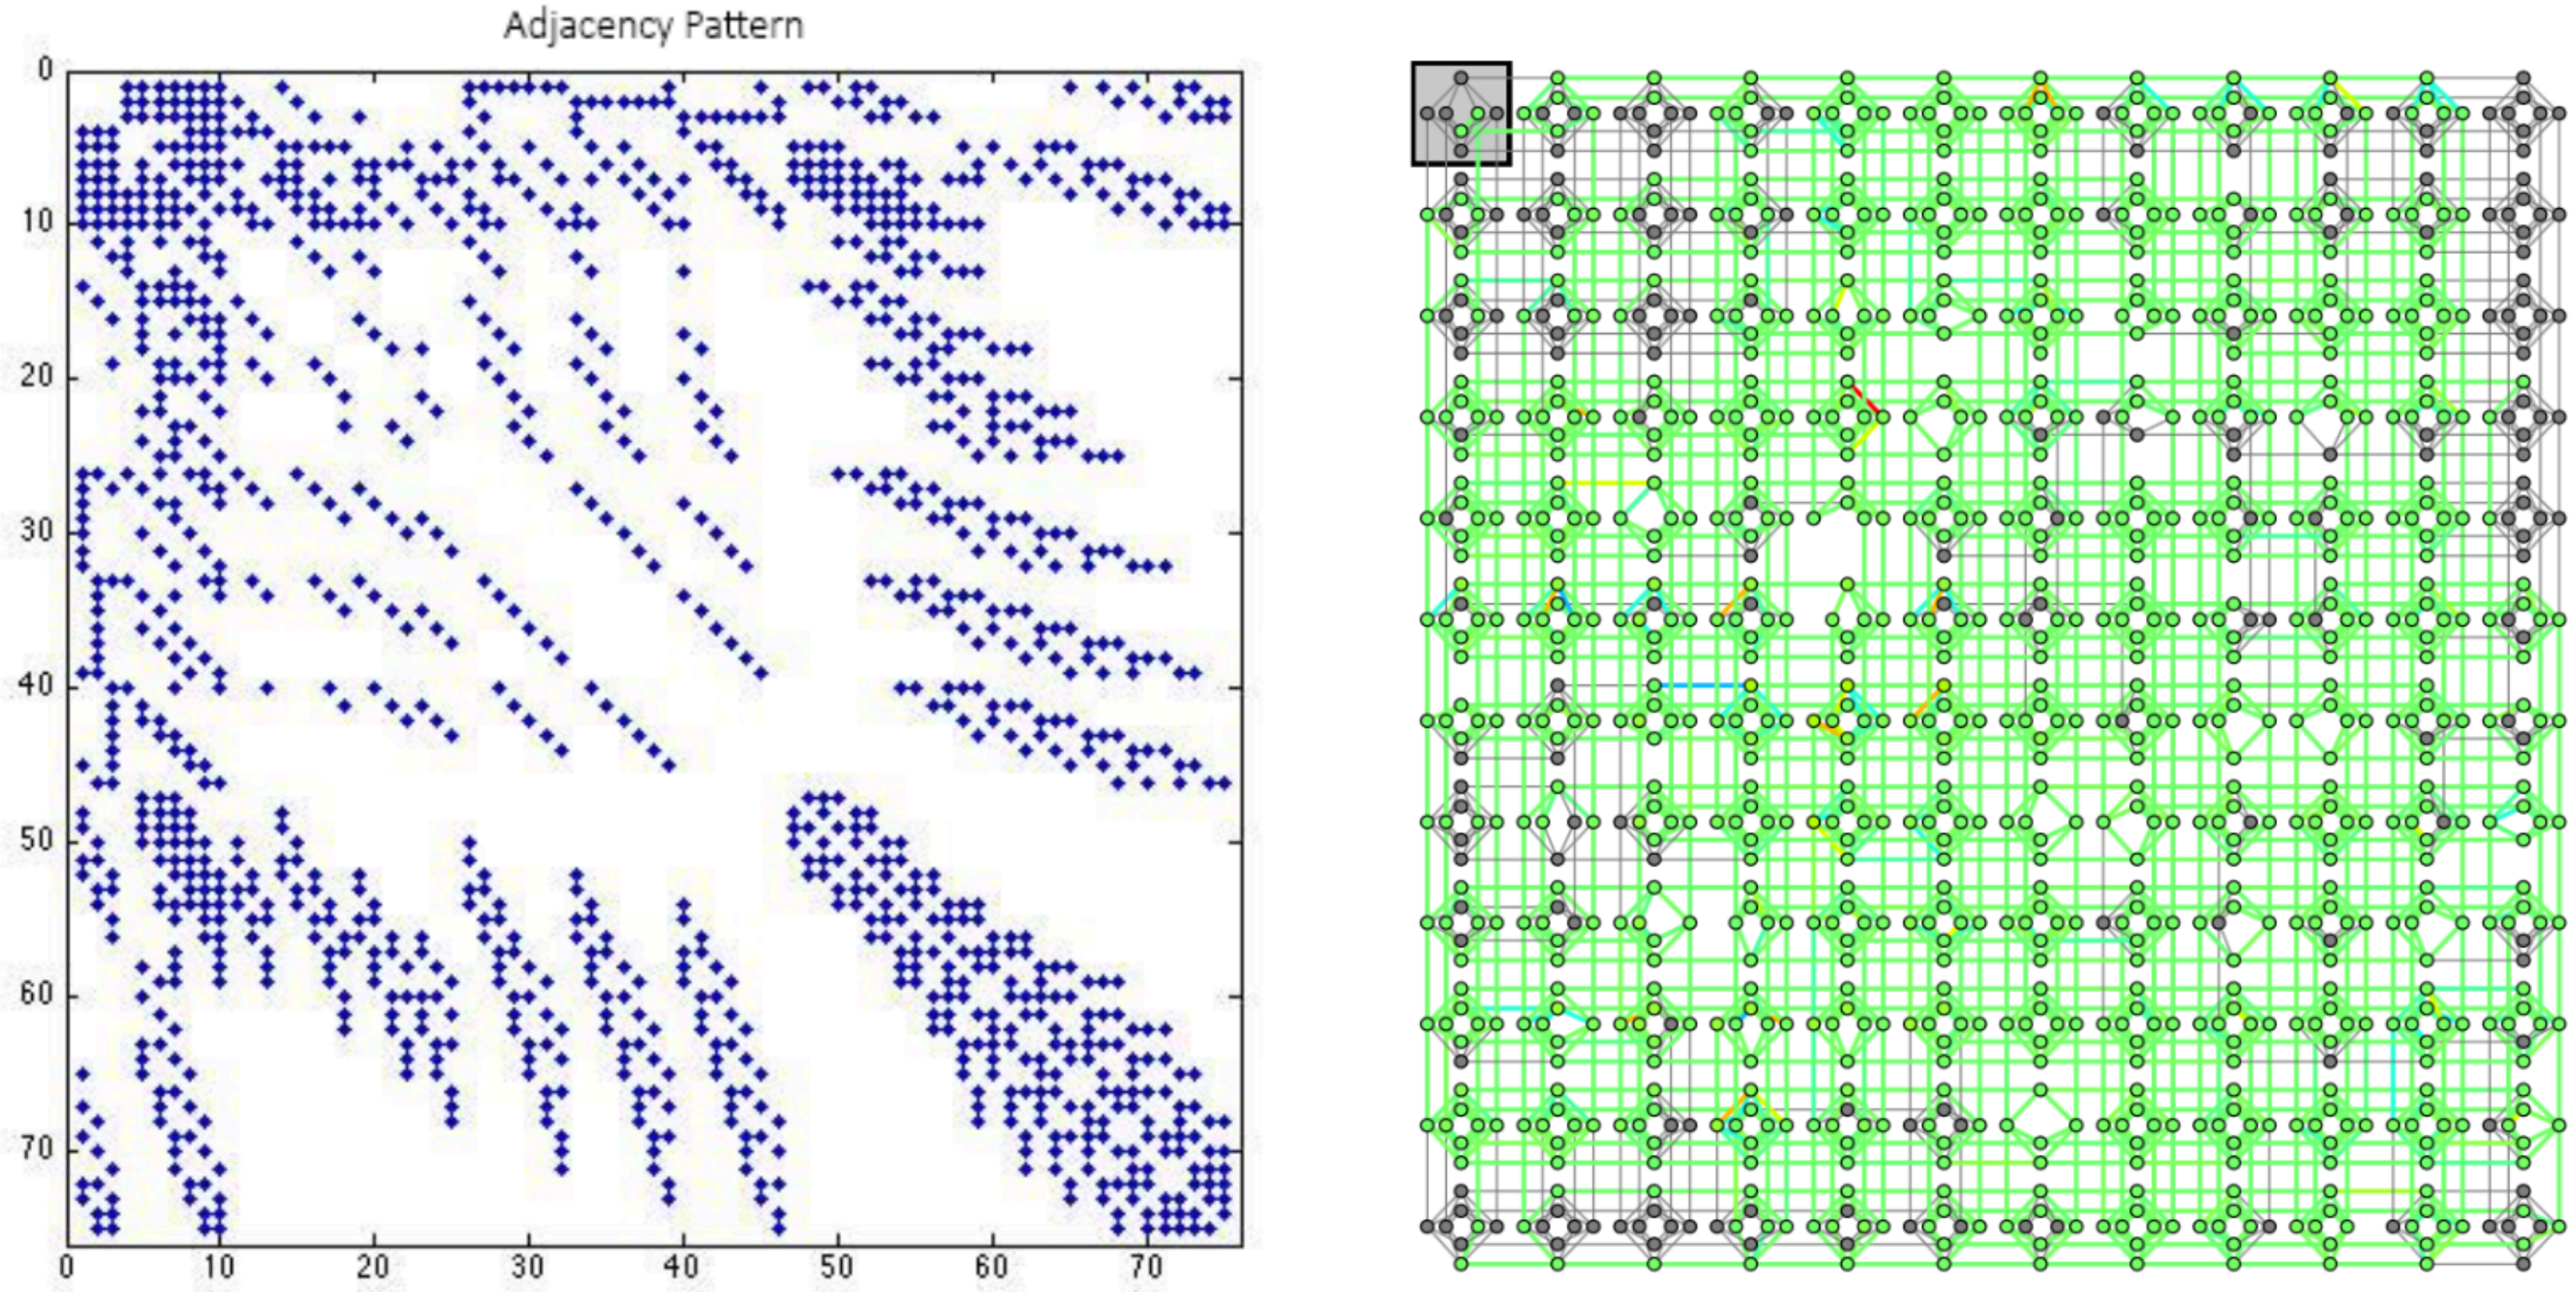

Figure 1: The column algorithm: the adjacency matrix pattern (left) and embedding into the the D-Wave 2X quantum processor (right) of the quadratic binary polynomial for $M = 200\,099$.

optimization

- quantum annealing linear optimization
- quadratic unconstrained binary optimization
- quantum approximate optimization algorithm
- traveling salesman problem dwave demo
- quantum algorithms, arxiv: 1804.03719

arxiv:1904.04767

quanvolutional networks

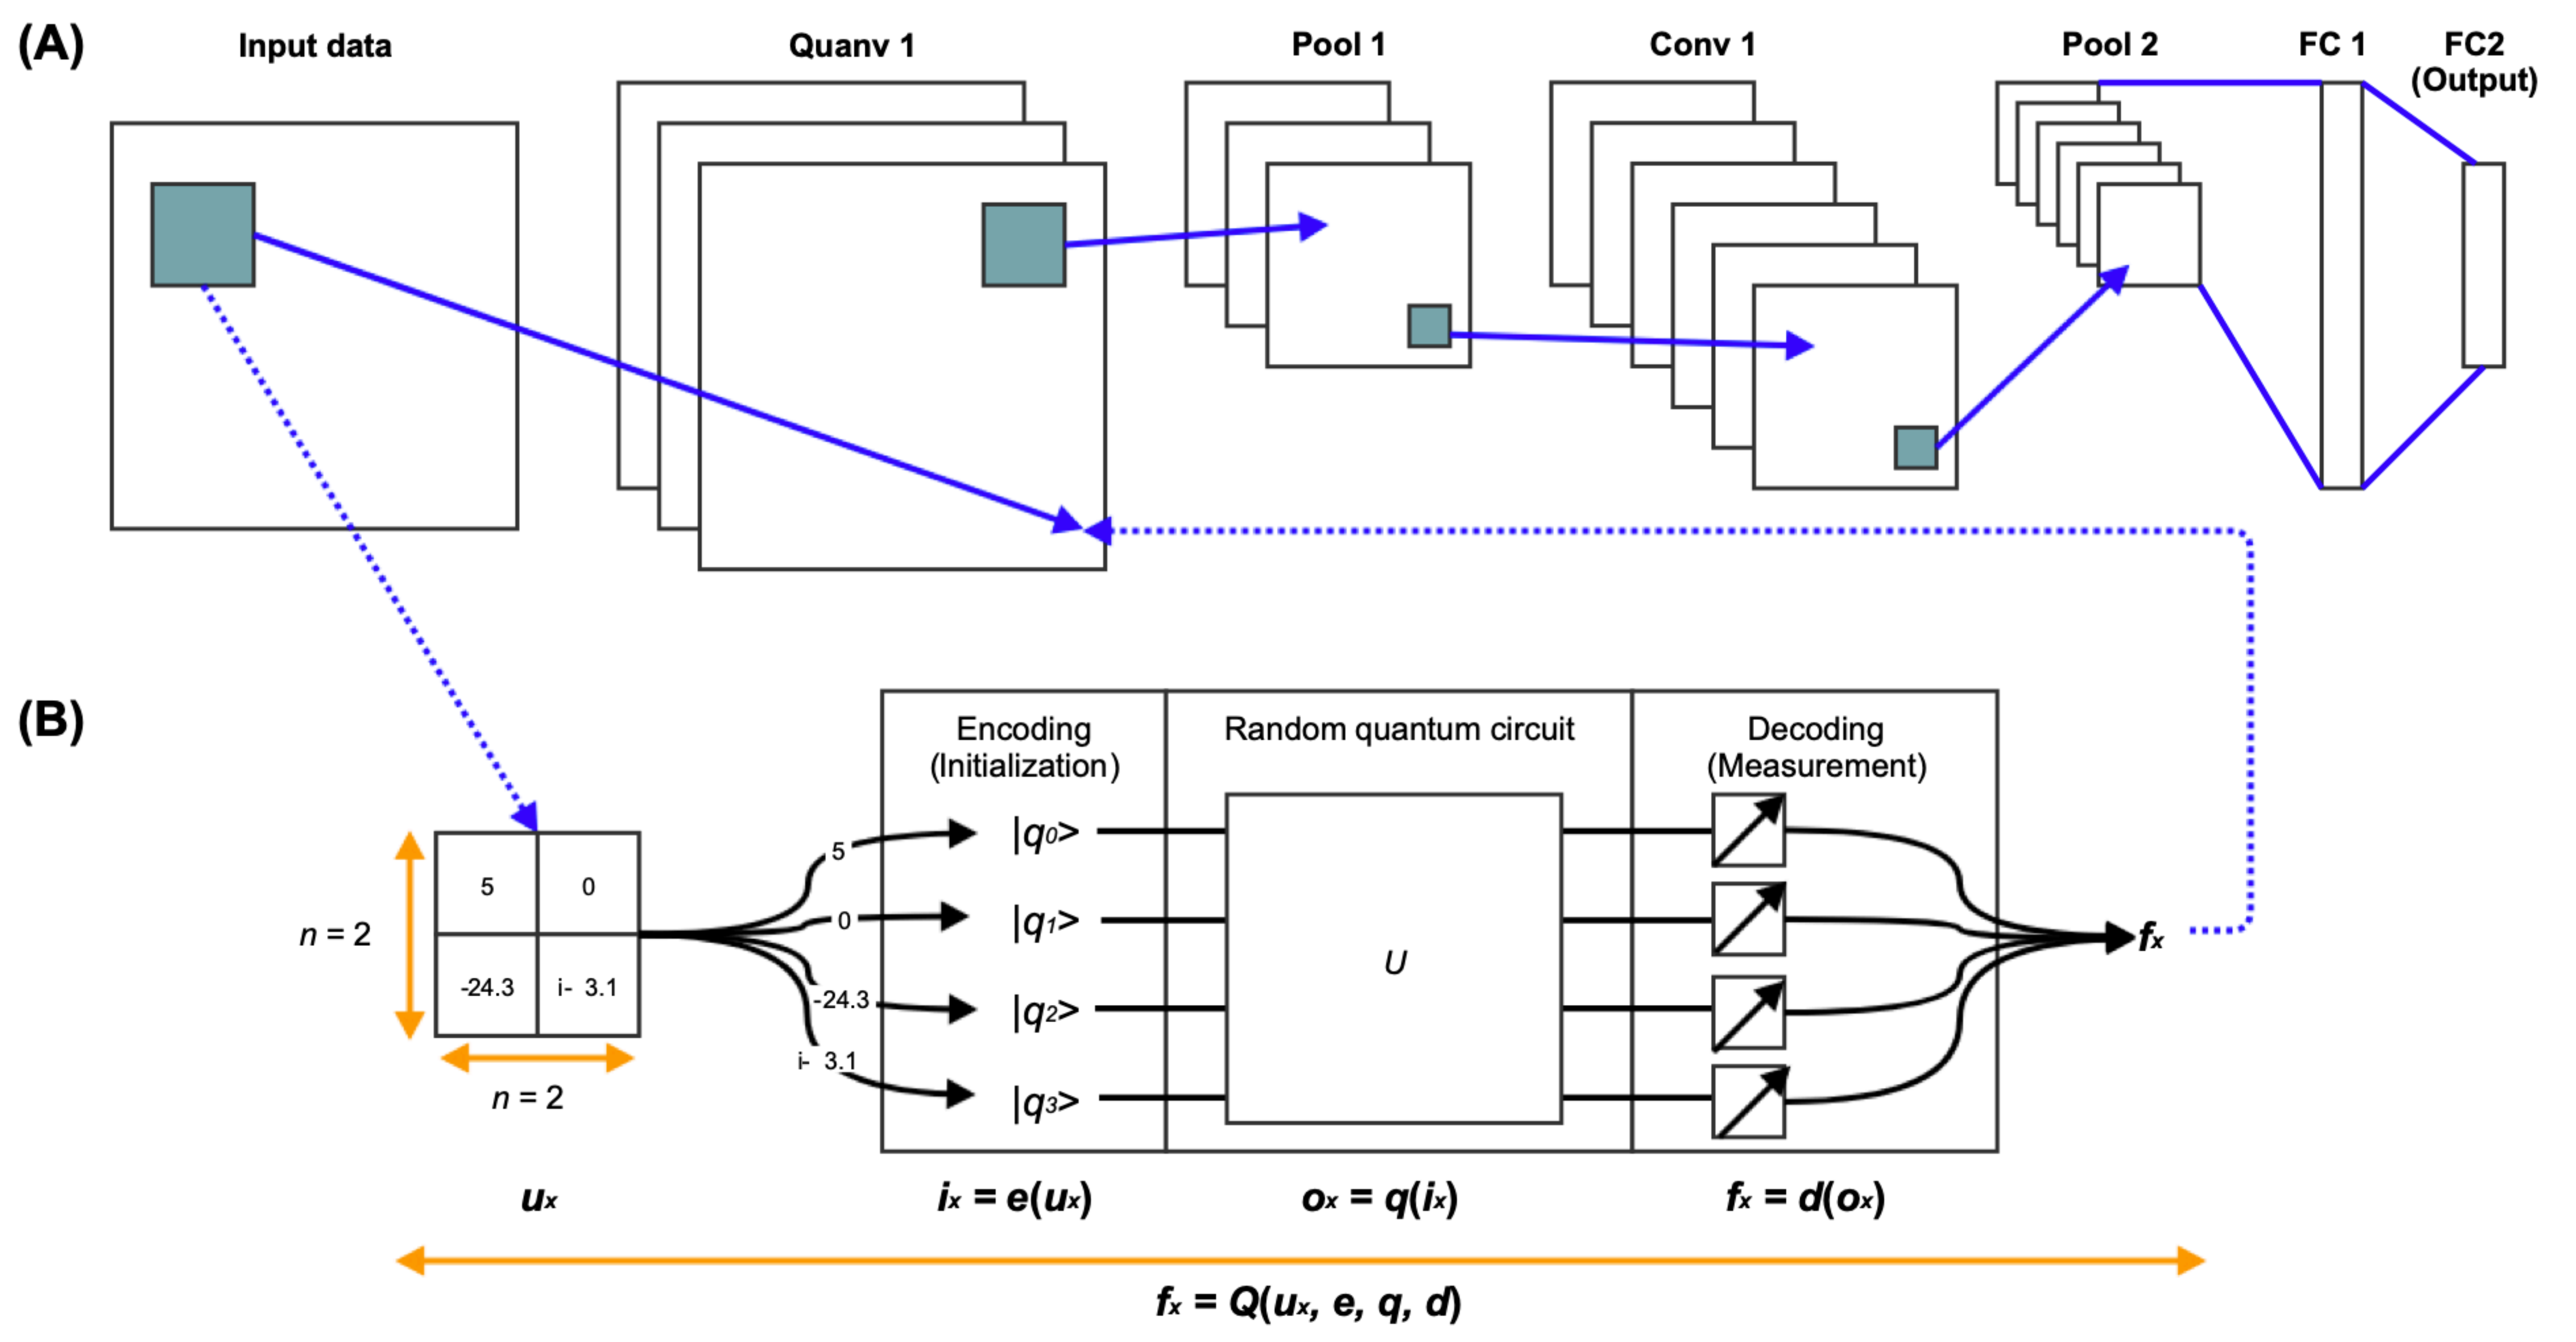

Fig. 1.: A. Simple example of a quanvolutional layer in a full network stack. The quanvolutional layer contains several quanvolutional filters (three in this example) that transform the input data into different output feature maps. B. An in-depth look at the processing of classical data into and out of the random quantum circuit in the quanvolutional filter.

quantum convolutional neural networks

arxiv:1810.03787

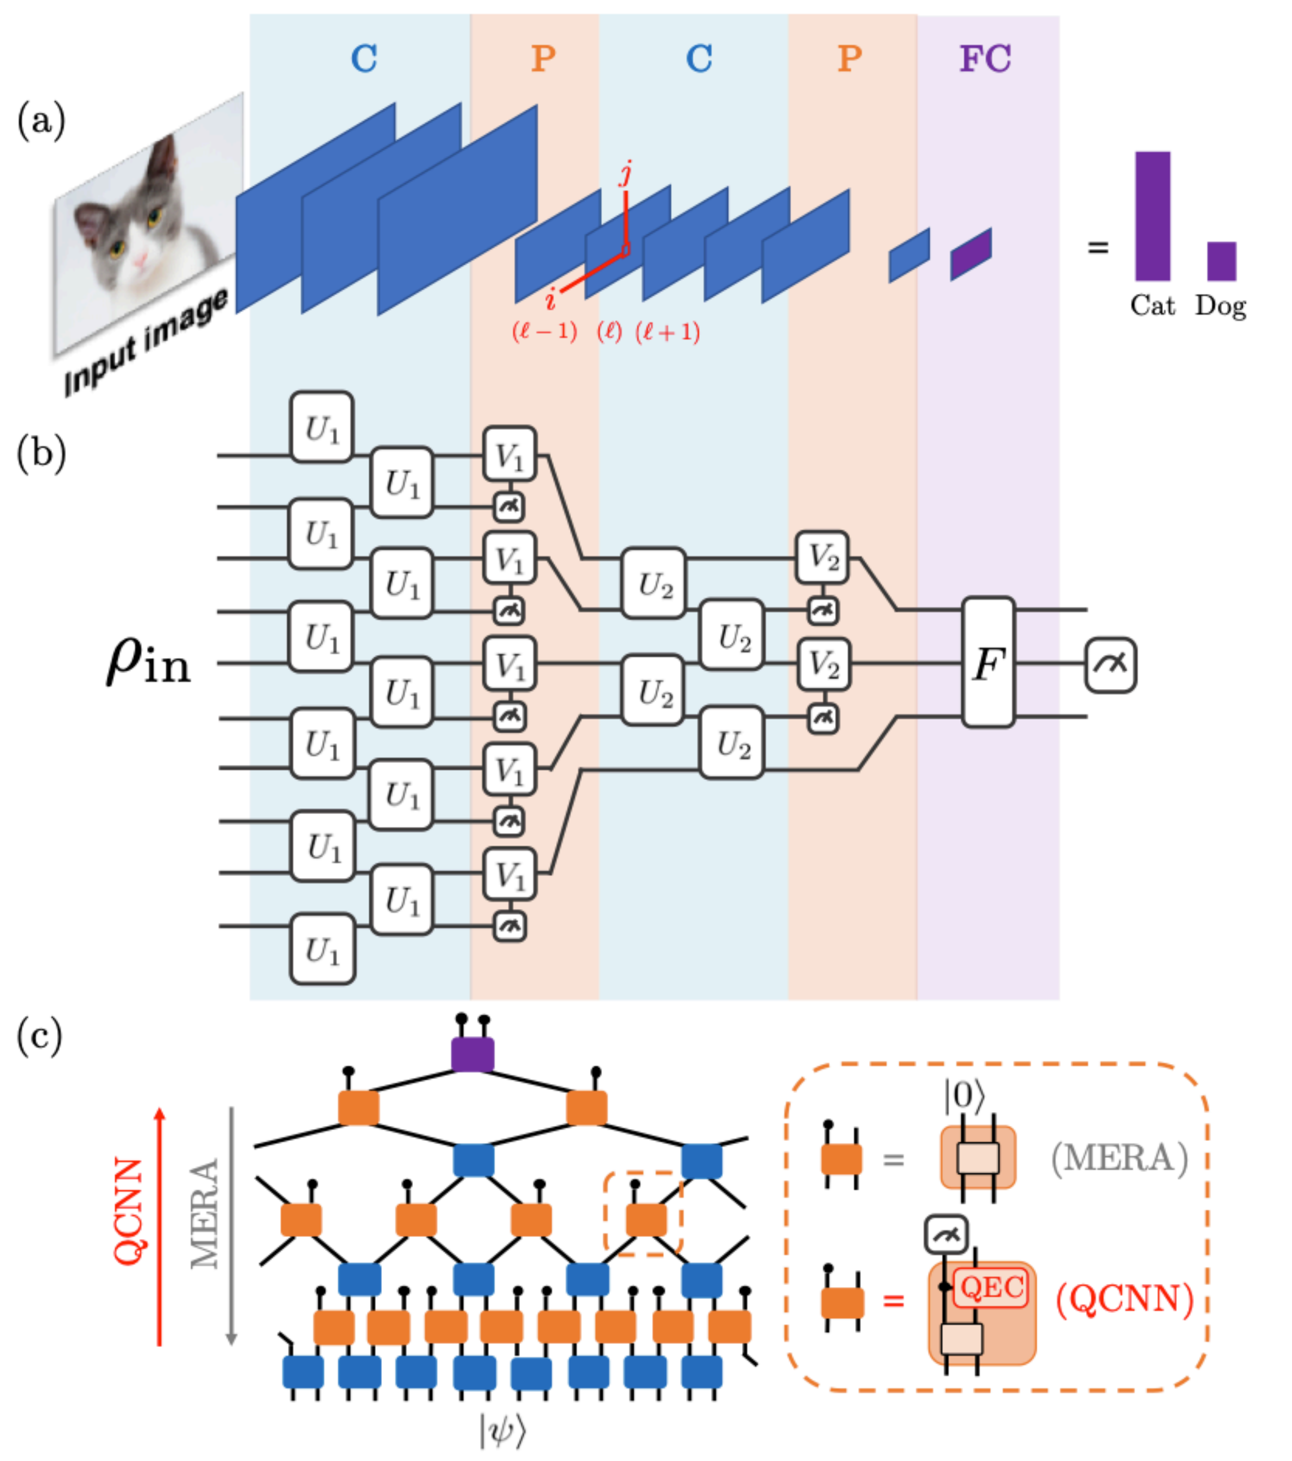

Figure 1: (a) Simplified illustration of CNNs. A sequence of image processing layers—convolution (C), pooling (P), and fully connected (FC)—transforms an input image into a series of feature maps (blue rectangles), and finally into an output probability distribution (purple bars). (b) QCNNs inherit a similar layered structure. (c) QCNN and MERA share the same circuit structure, but run in reverse directions.

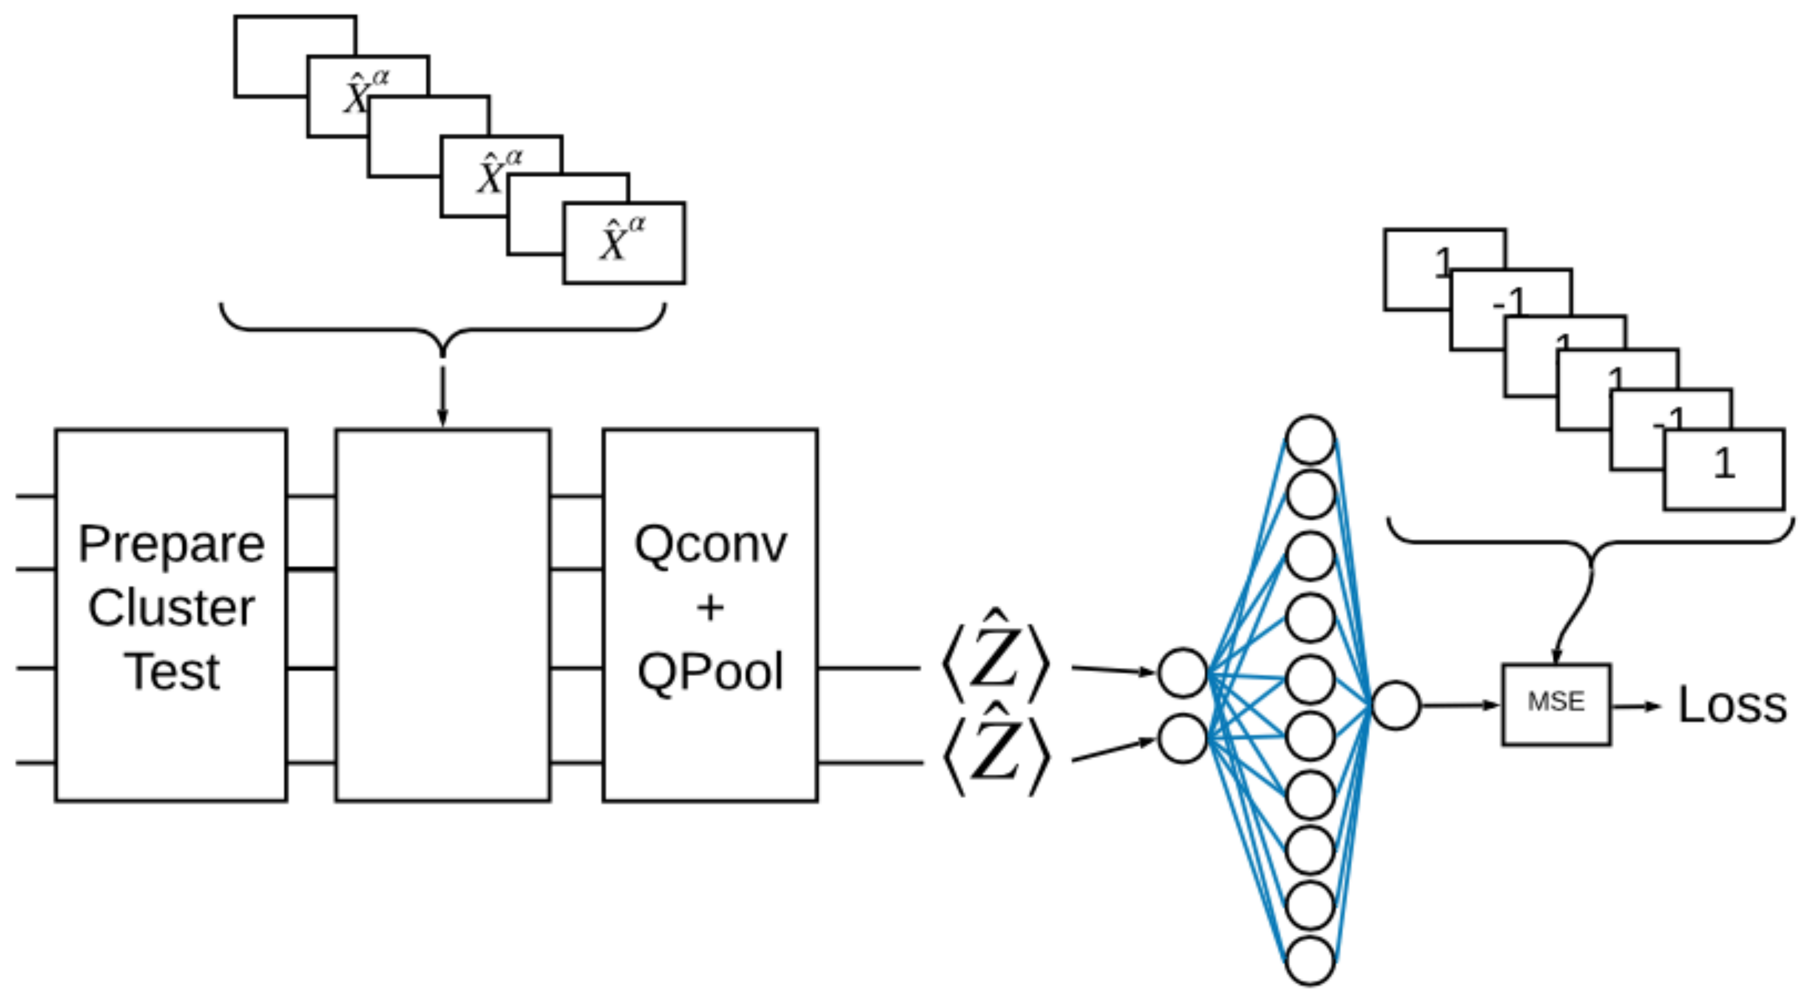

Figure 15. A simple hybrid architecture in which the outputs of a truncated QCNN are fed into a classical neural network.

- **tensorflow quantum mnist demo**

sycamore: 53 qubit

nature.com/articles/
s41586-019-1666-5

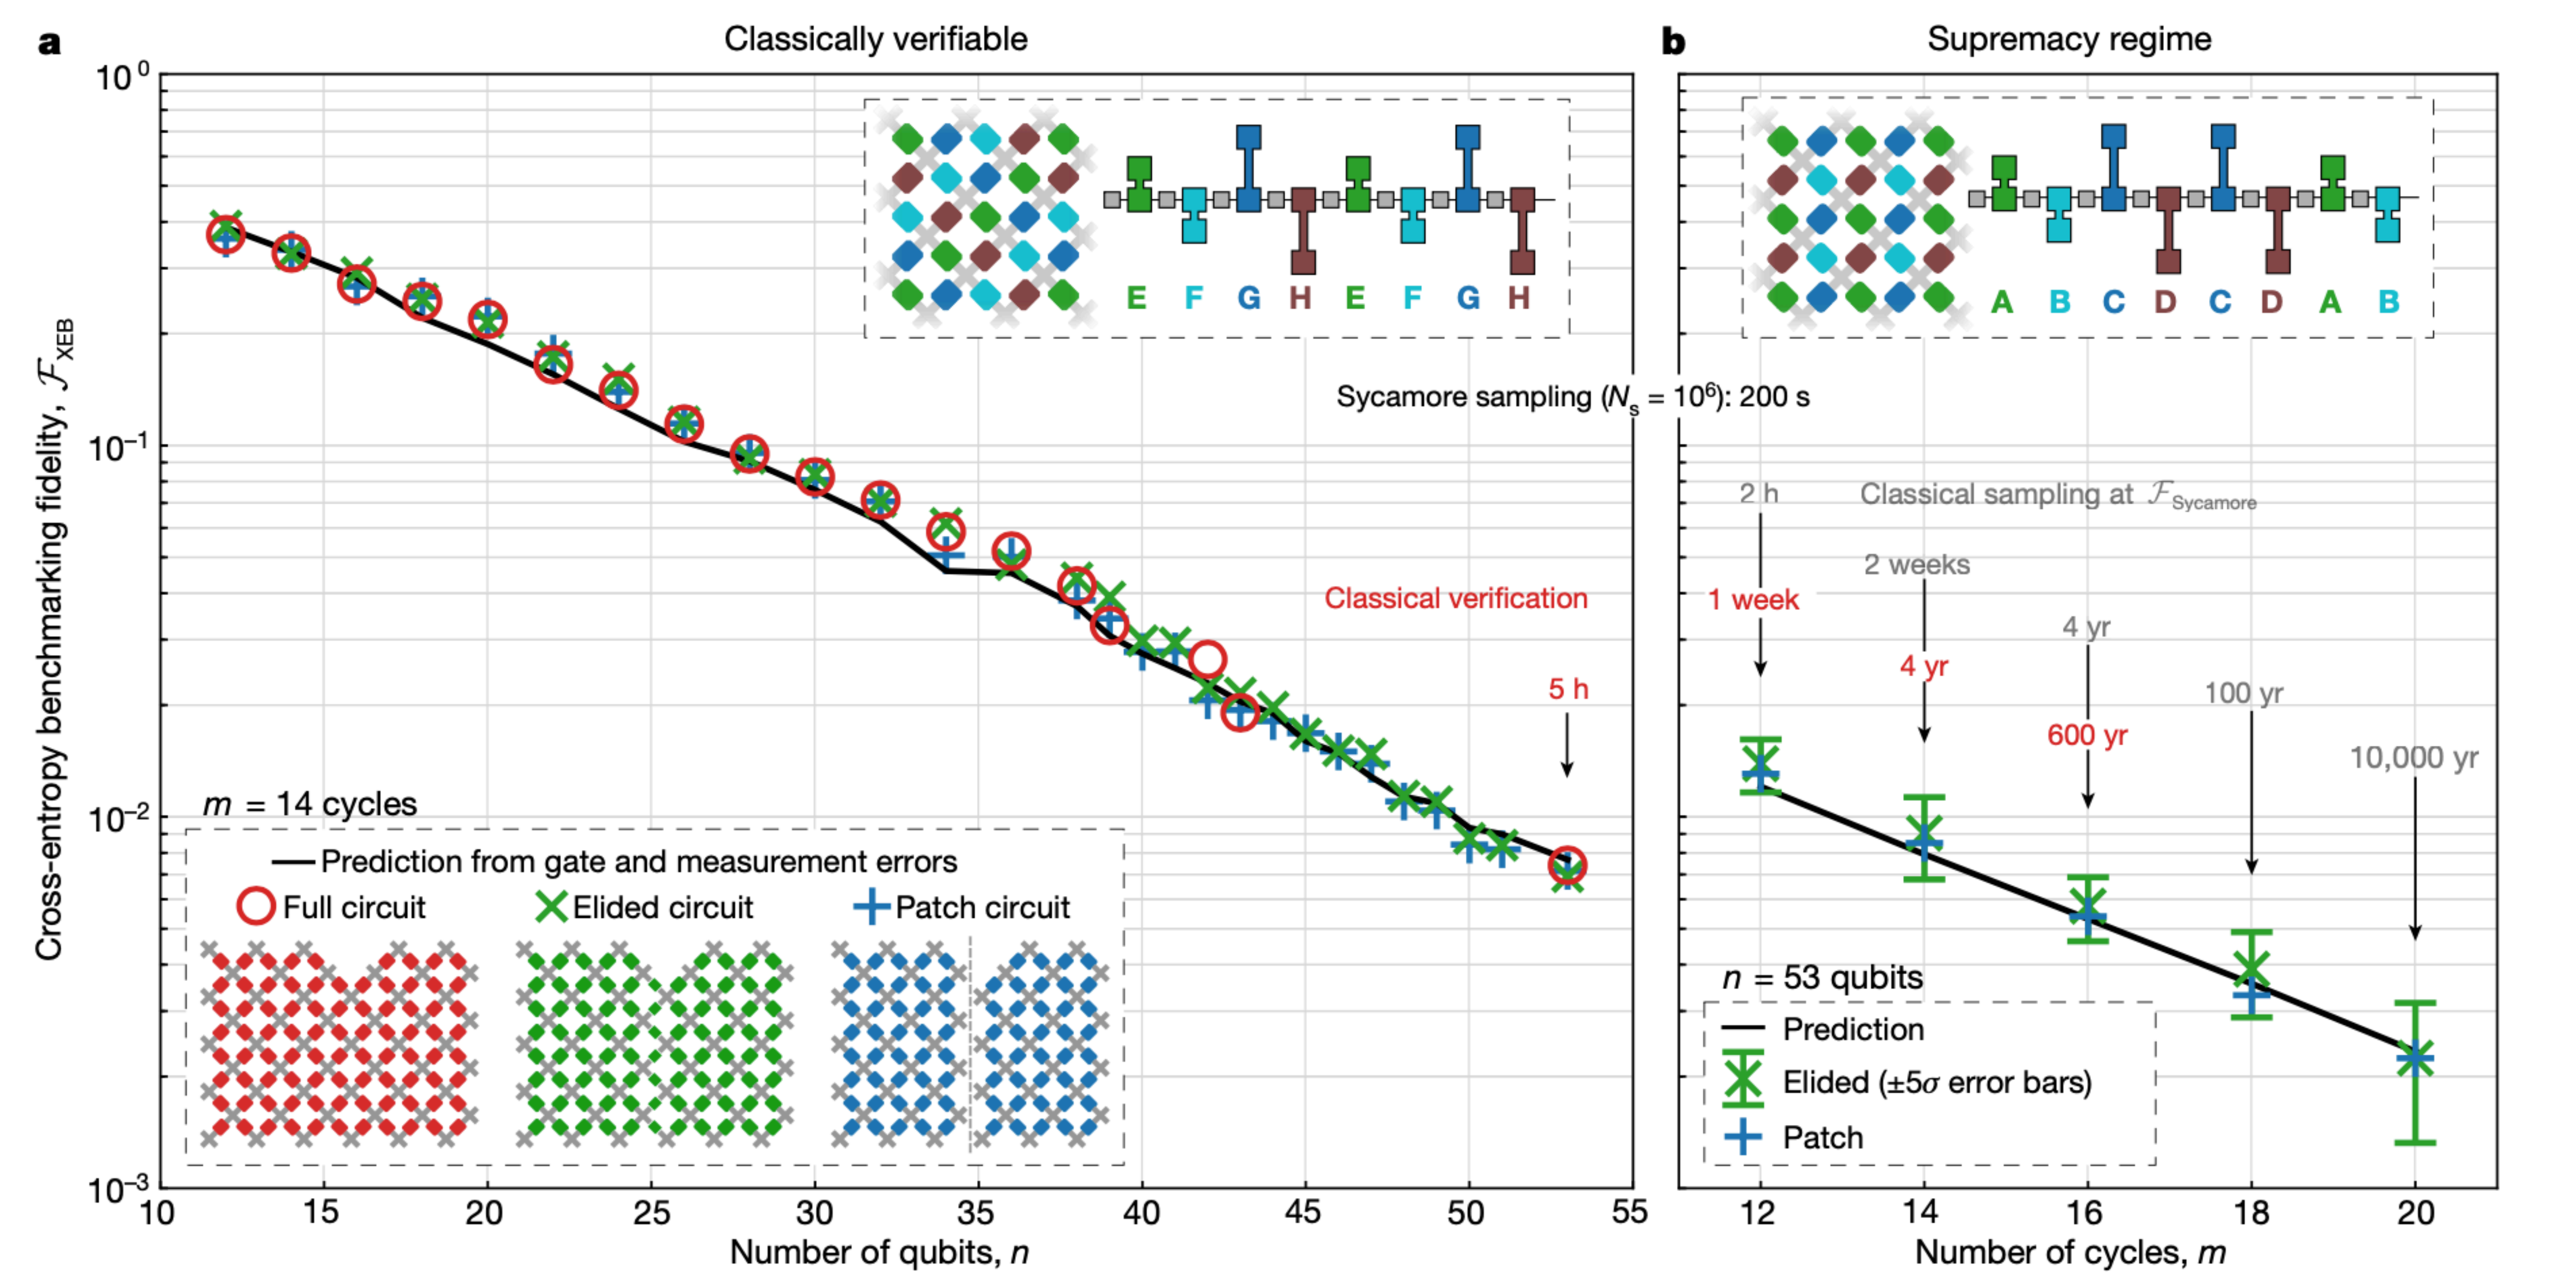

Fig. 4 | Demonstrating quantum supremacy. a, Verification of benchmarking methods. $\mathcal{F}_{XEB}$ values for patch, elided and full verification circuits are calculated from measured bitstrings and the corresponding probabilities predicted by classical simulation. Here, the two-qubit gates are applied in a simplifiable tiling and sequence such that the full circuits can be simulated out to $n = 53$, $m = 14$ in a reasonable amount of time. Each data point is an average over ten distinct quantum circuit instances that differ in their single-qubit gates (for $n = 39$, 42 and 43 only two instances were simulated). For each n, each instance is sampled with N_s of 0.5–2.5 million. The black line shows the predicted $\mathcal{F}_{XEB}$ based on single- and two-qubit gate and measurement errors. The close correspondence between all four curves, despite their vast differences in complexity, justifies the use of elided circuits to estimate fidelity in the supremacy regime. **b**, Estimating $\mathcal{F}_{XEB}$ in the quantum supremacy regime. Here, the two-qubit gates are applied in a non-simplifiable tiling and sequence for which it is much harder to simulate. For the largest elided data ($n = 53$, $m = 20$, total $N_s = 30$ million), we find an average $\mathcal{F}_{XEB} > 0.1\%$ with 5σ confidence, where σ includes both systematic and statistical uncertainties. The corresponding full circuit data, not simulated but archived, is expected to show similarly statistically significant fidelity. For $m = 20$, obtaining a million samples on the quantum processor takes 200 seconds, whereas an equal-fidelity classical sampling would take 10,000 years on a million cores, and verifying the fidelity would take millions of years.

gaussian boson sampling

doi: 10.1126/
science.abe8770

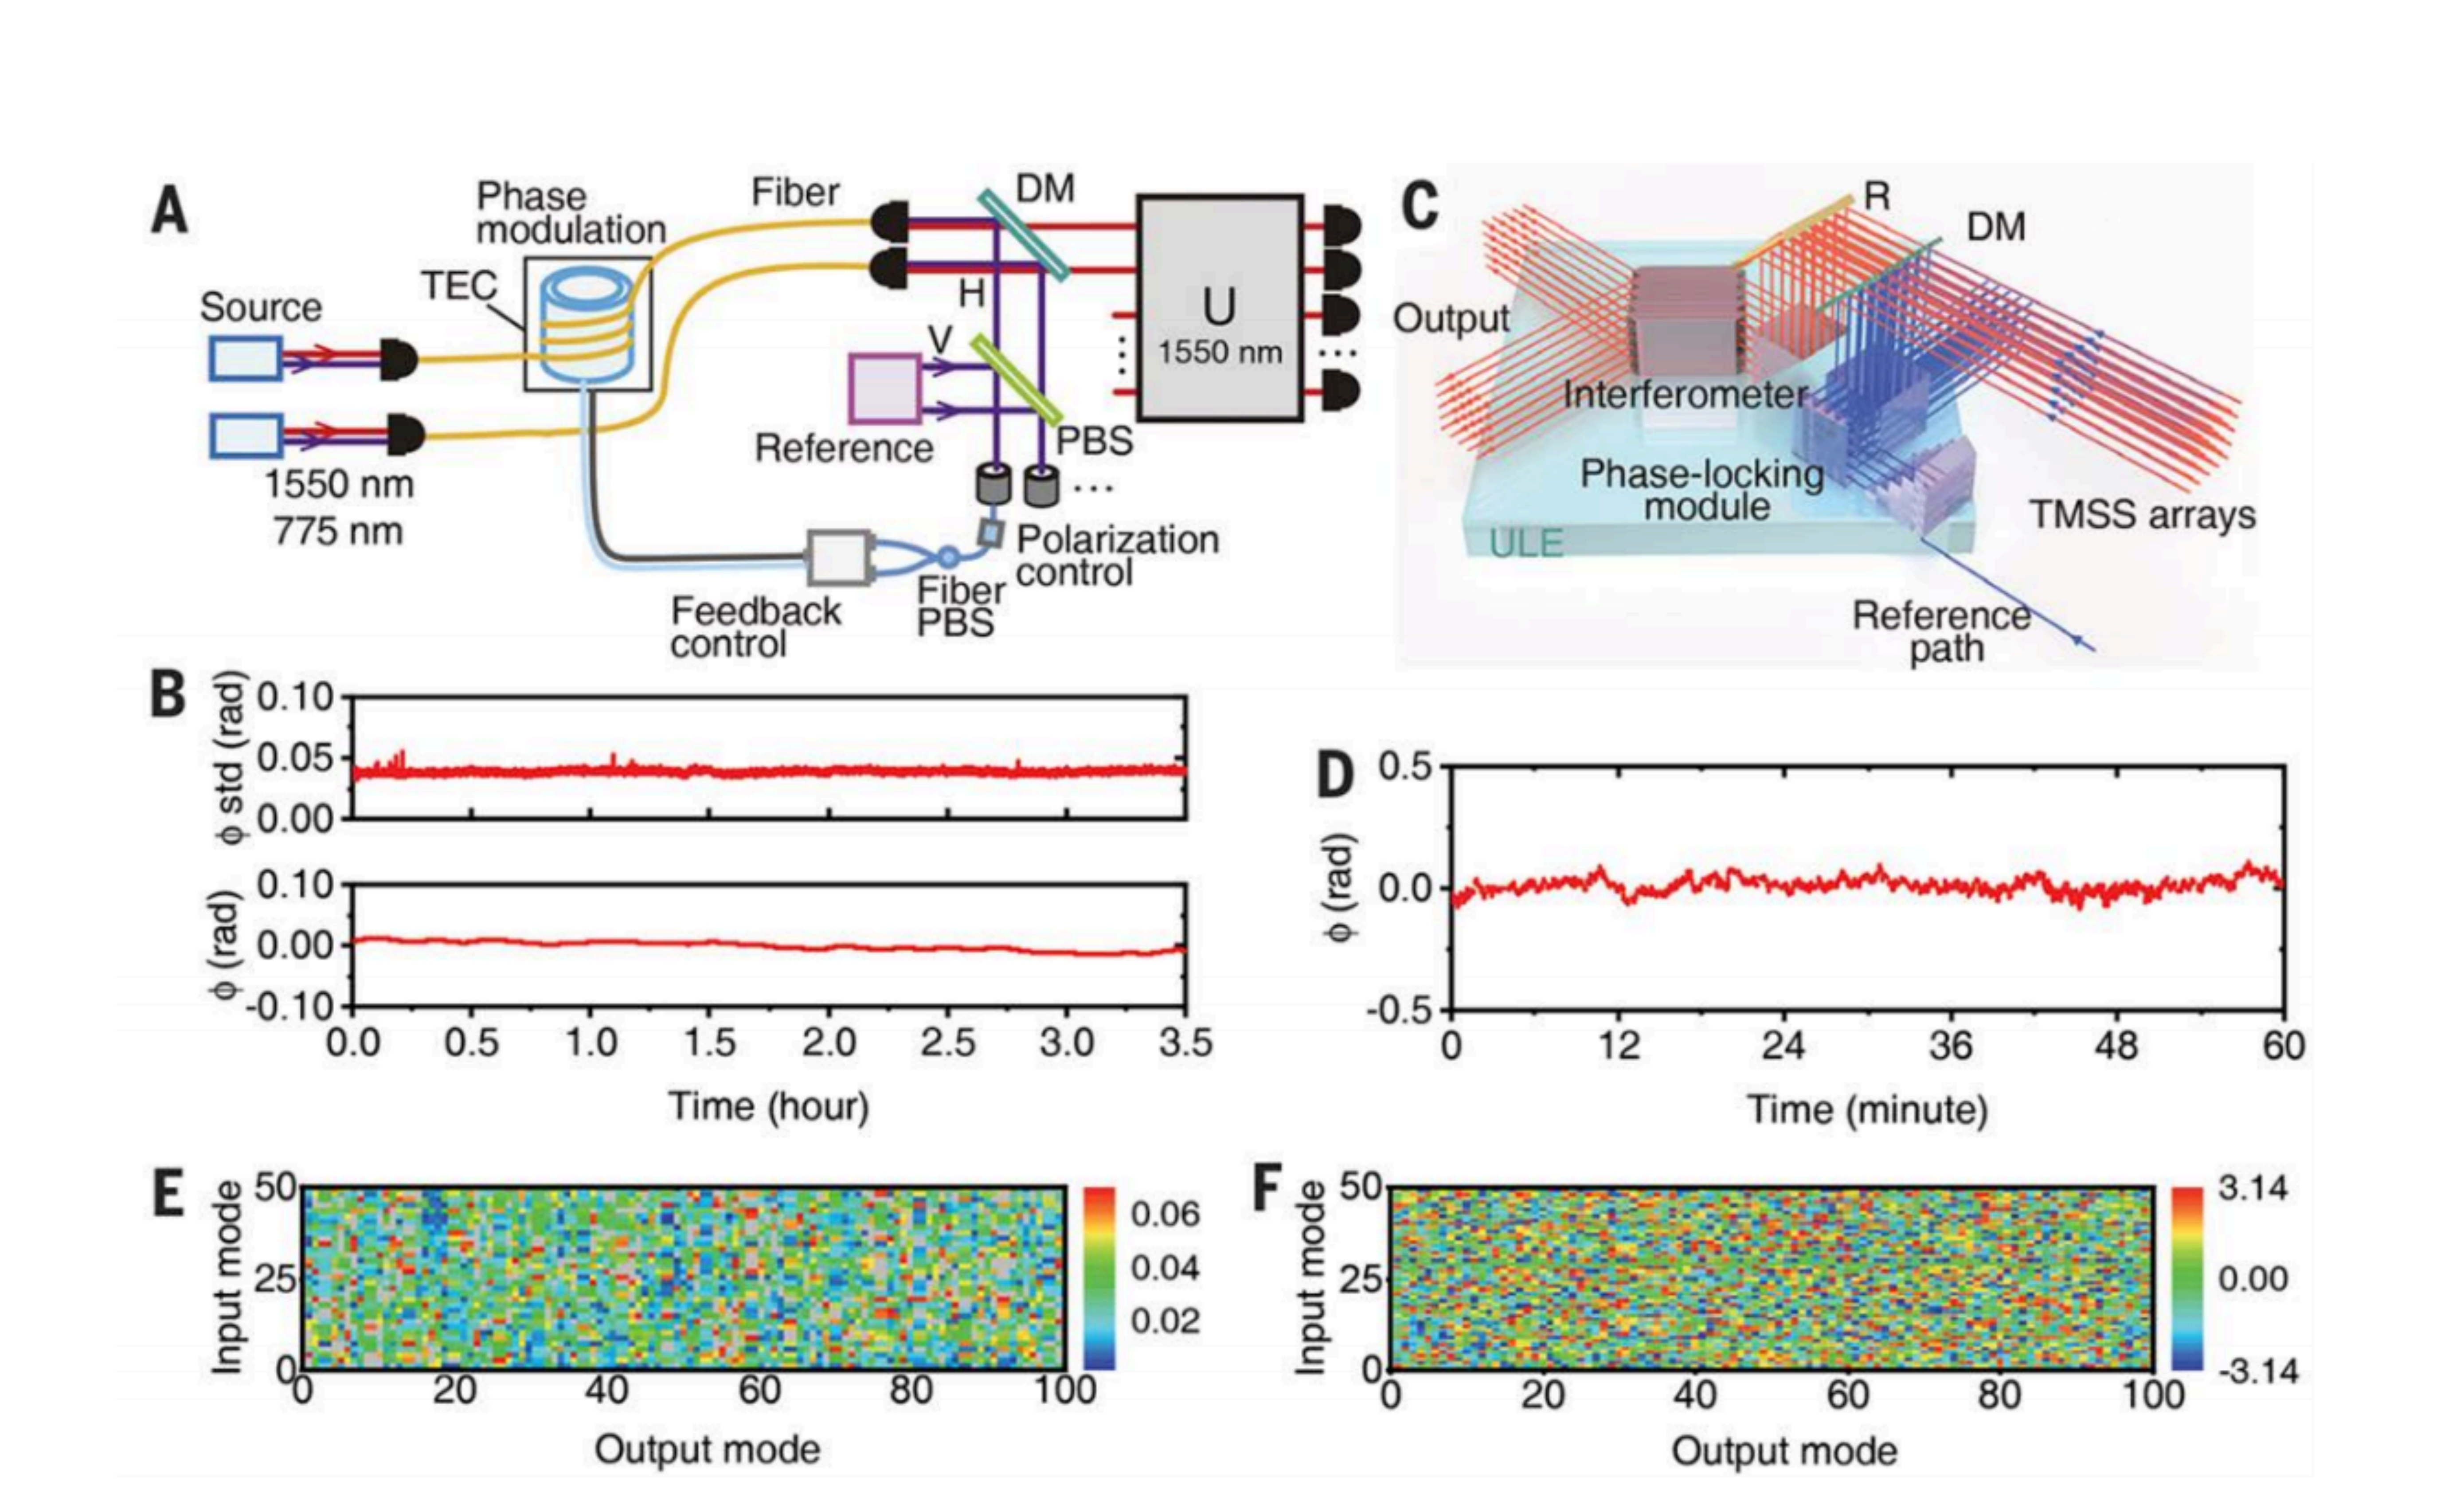

- 43 qubit average, demoed 76 qubit

frontier

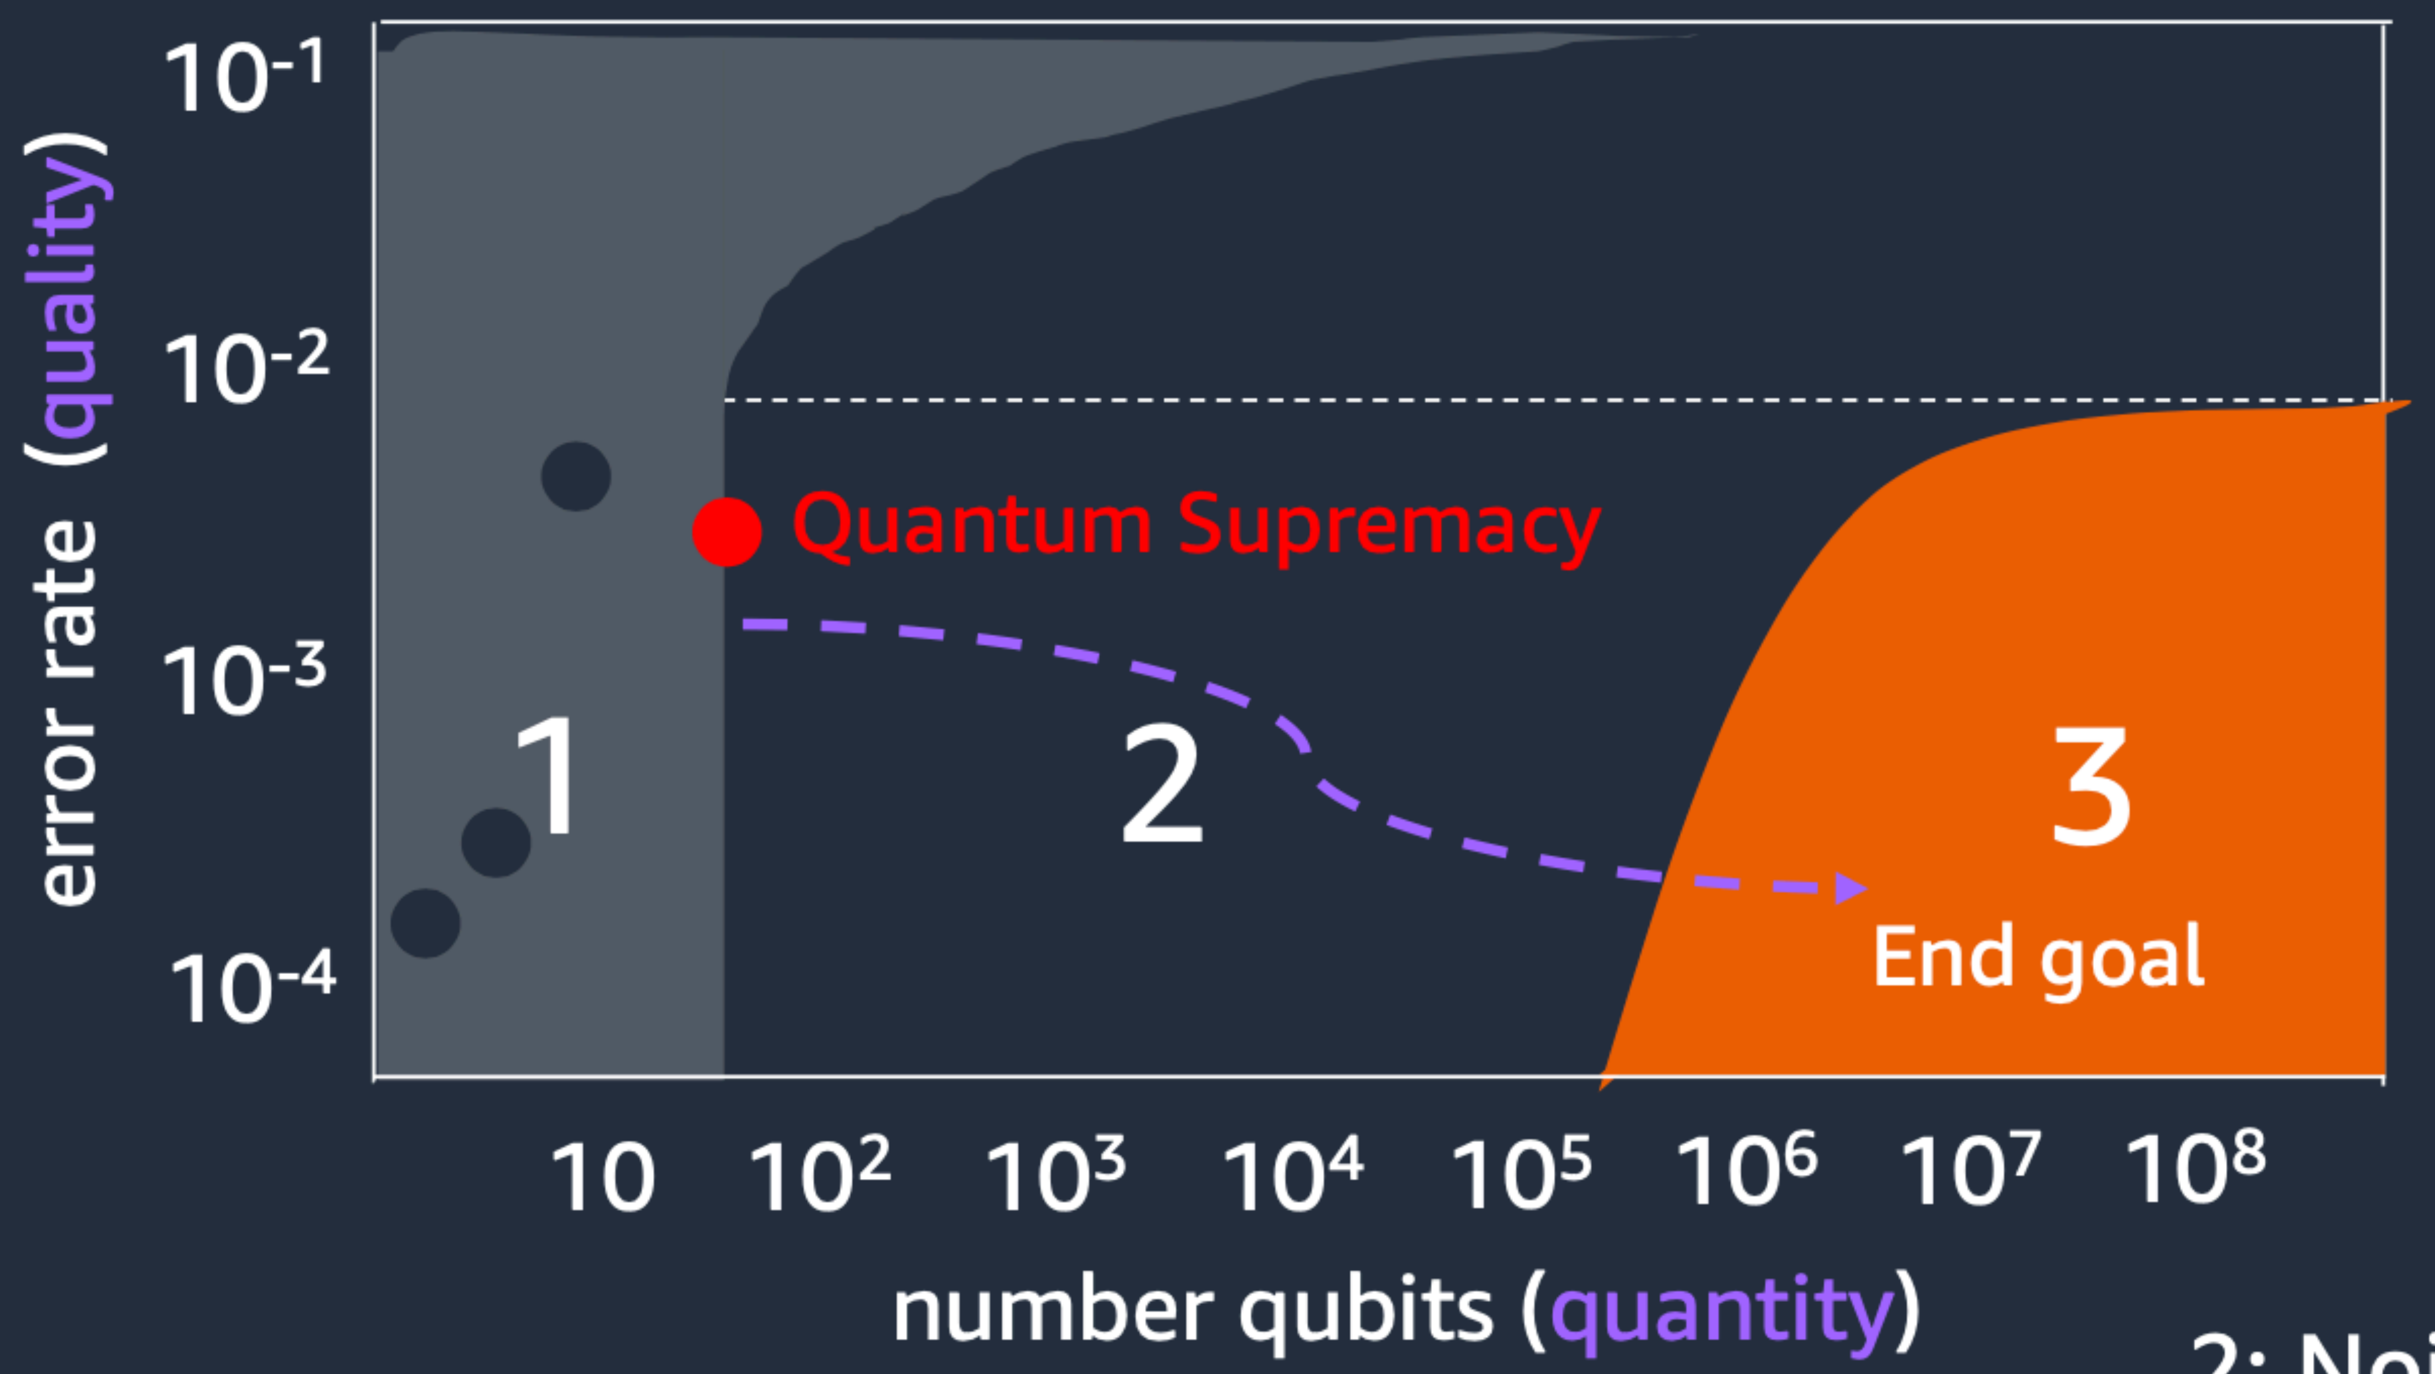

1: Classically simulatable

2: Noisy Intermediate-Scale Quantum (NISQ)

3: Quantum Computing with error correction

recap

- theory of quantum computing
- different applications
- demos of different hardware and software approaches
- future is you