QUANTINUUM

Quantum Computing: Progress Towards Real World Applications

MulticoreWorld 2025

Russell Stutz, Sr Director of Product Technologies at Quantinuum

Quantum computing: Superposition & Entanglement

0

FALSE/OFF

1

TRUE/ON

0	1	0	0	1

2^N possibilities:
holds only 1 value at a time

BITS

N-BIT WORD

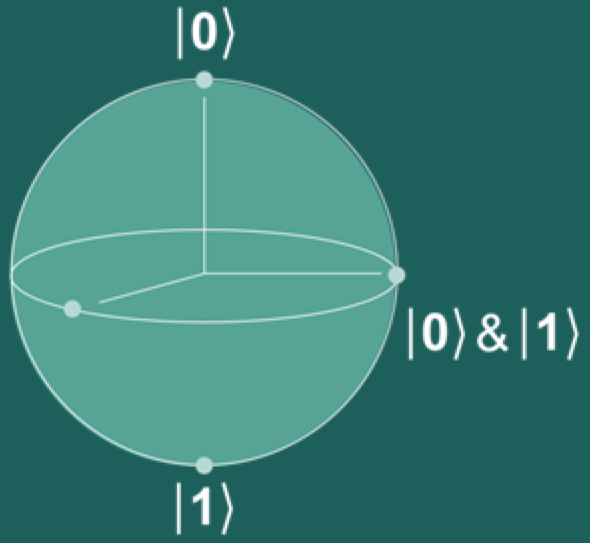

Superposition:

State = $\propto |0\rangle + \beta|1\rangle$

There is an amplitude and phase relationship between 0 and 1 states

QUBITS

Classical States

Quantum States

Quantum computing: Superposition & Entanglement

0

FALSE/OFF

1

TRUE/ON

BITS

0	1	0	0	1

2^N possibilities:
holds only 1 value at a time

N-BIT WORD

Classical States

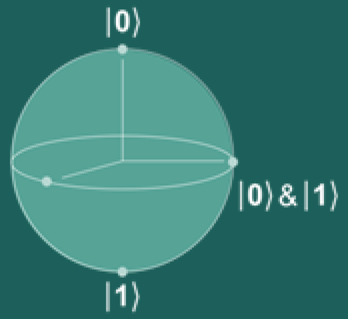

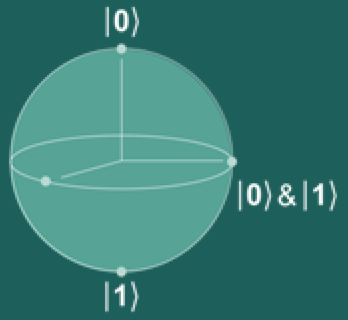

QUBITS

Entanglement:

State = $\propto |00\rangle + \beta|01\rangle + \gamma|10\rangle + \delta|11\rangle$

Bell State = $\propto |00\rangle + \beta|11\rangle$

Quantum States

QUANTINUUM

Quantum computing: Superposition & Entanglement

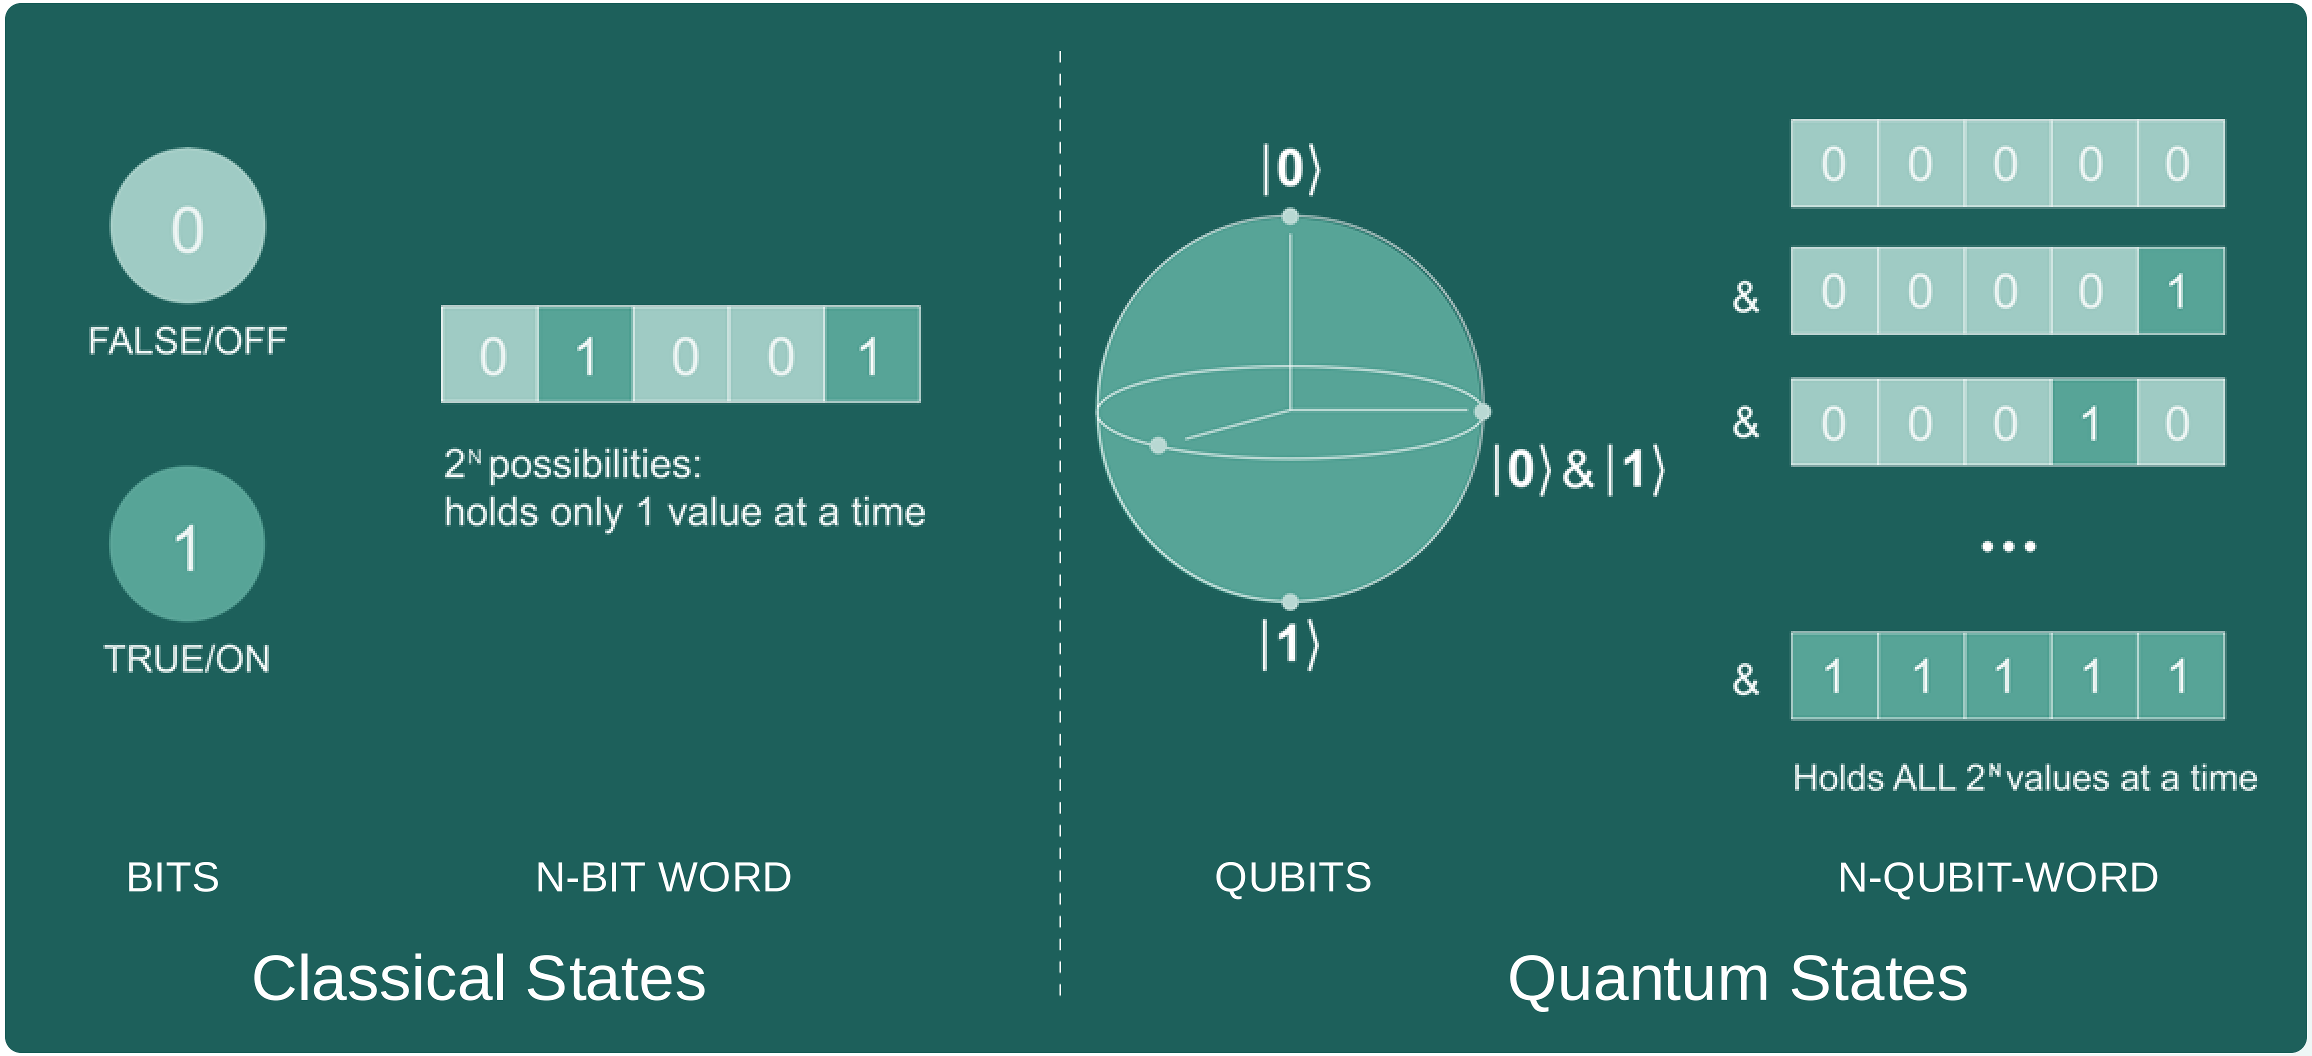

QUANTINUUM

Quantum computing: Superposition & Entanglement

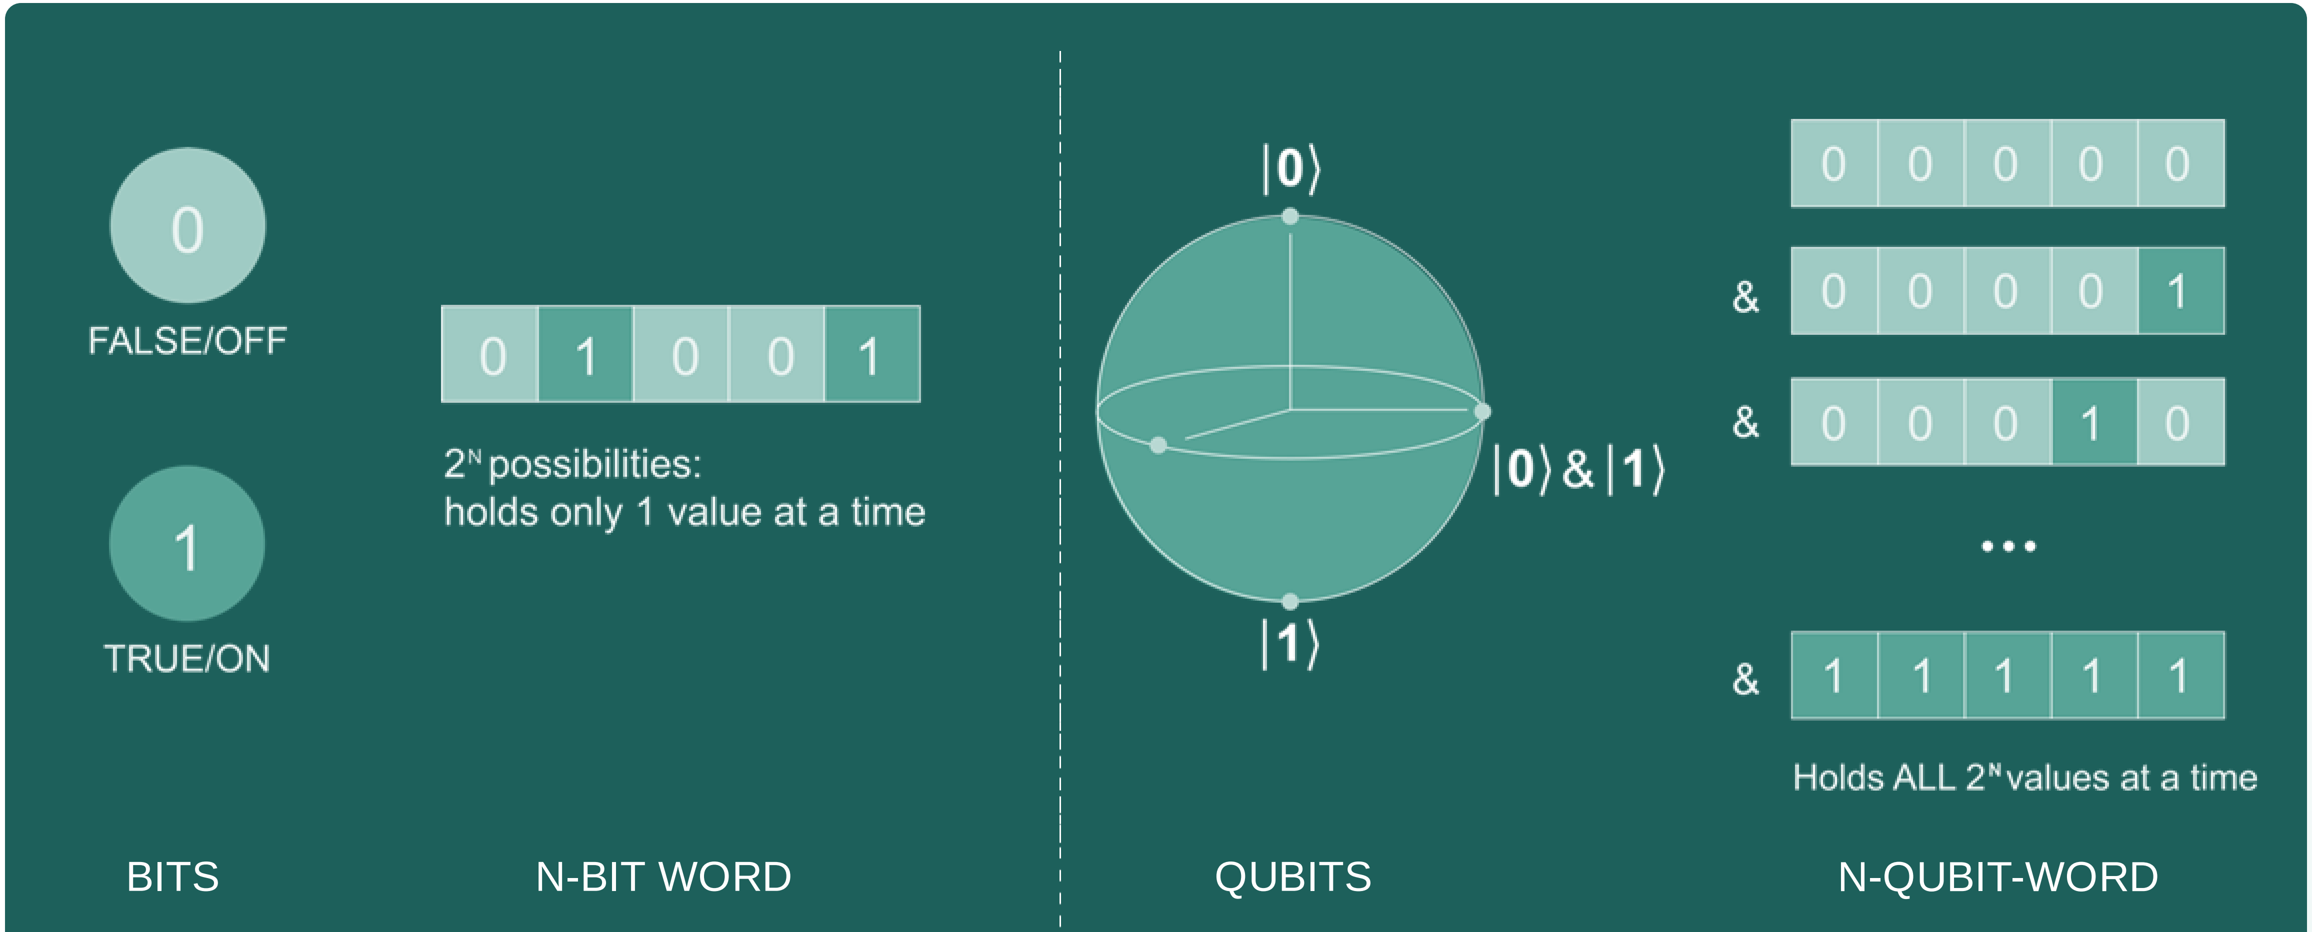

100 qubits = 10^{30} complex amplitudes to describe

QUBITs are fragile…

Transistor in a processor (2008)

[1]: avg. one error per $\mathbf{10^{28}}$ operations

Quantum technologies (nowish):

avg. one error per $\mathbf{10^{3}}$ operations (99.9% fidelity)

Operations for large-scale quantum algorithms:

RSA/Shor4096 ≈ $\mathbf{10^{12}}$ CNOTs

[1] Pradip Bose, Designing Reliable Systems with Unreliable Components; IEEE Computer Society

Trapped-Ion QCCD Architecture

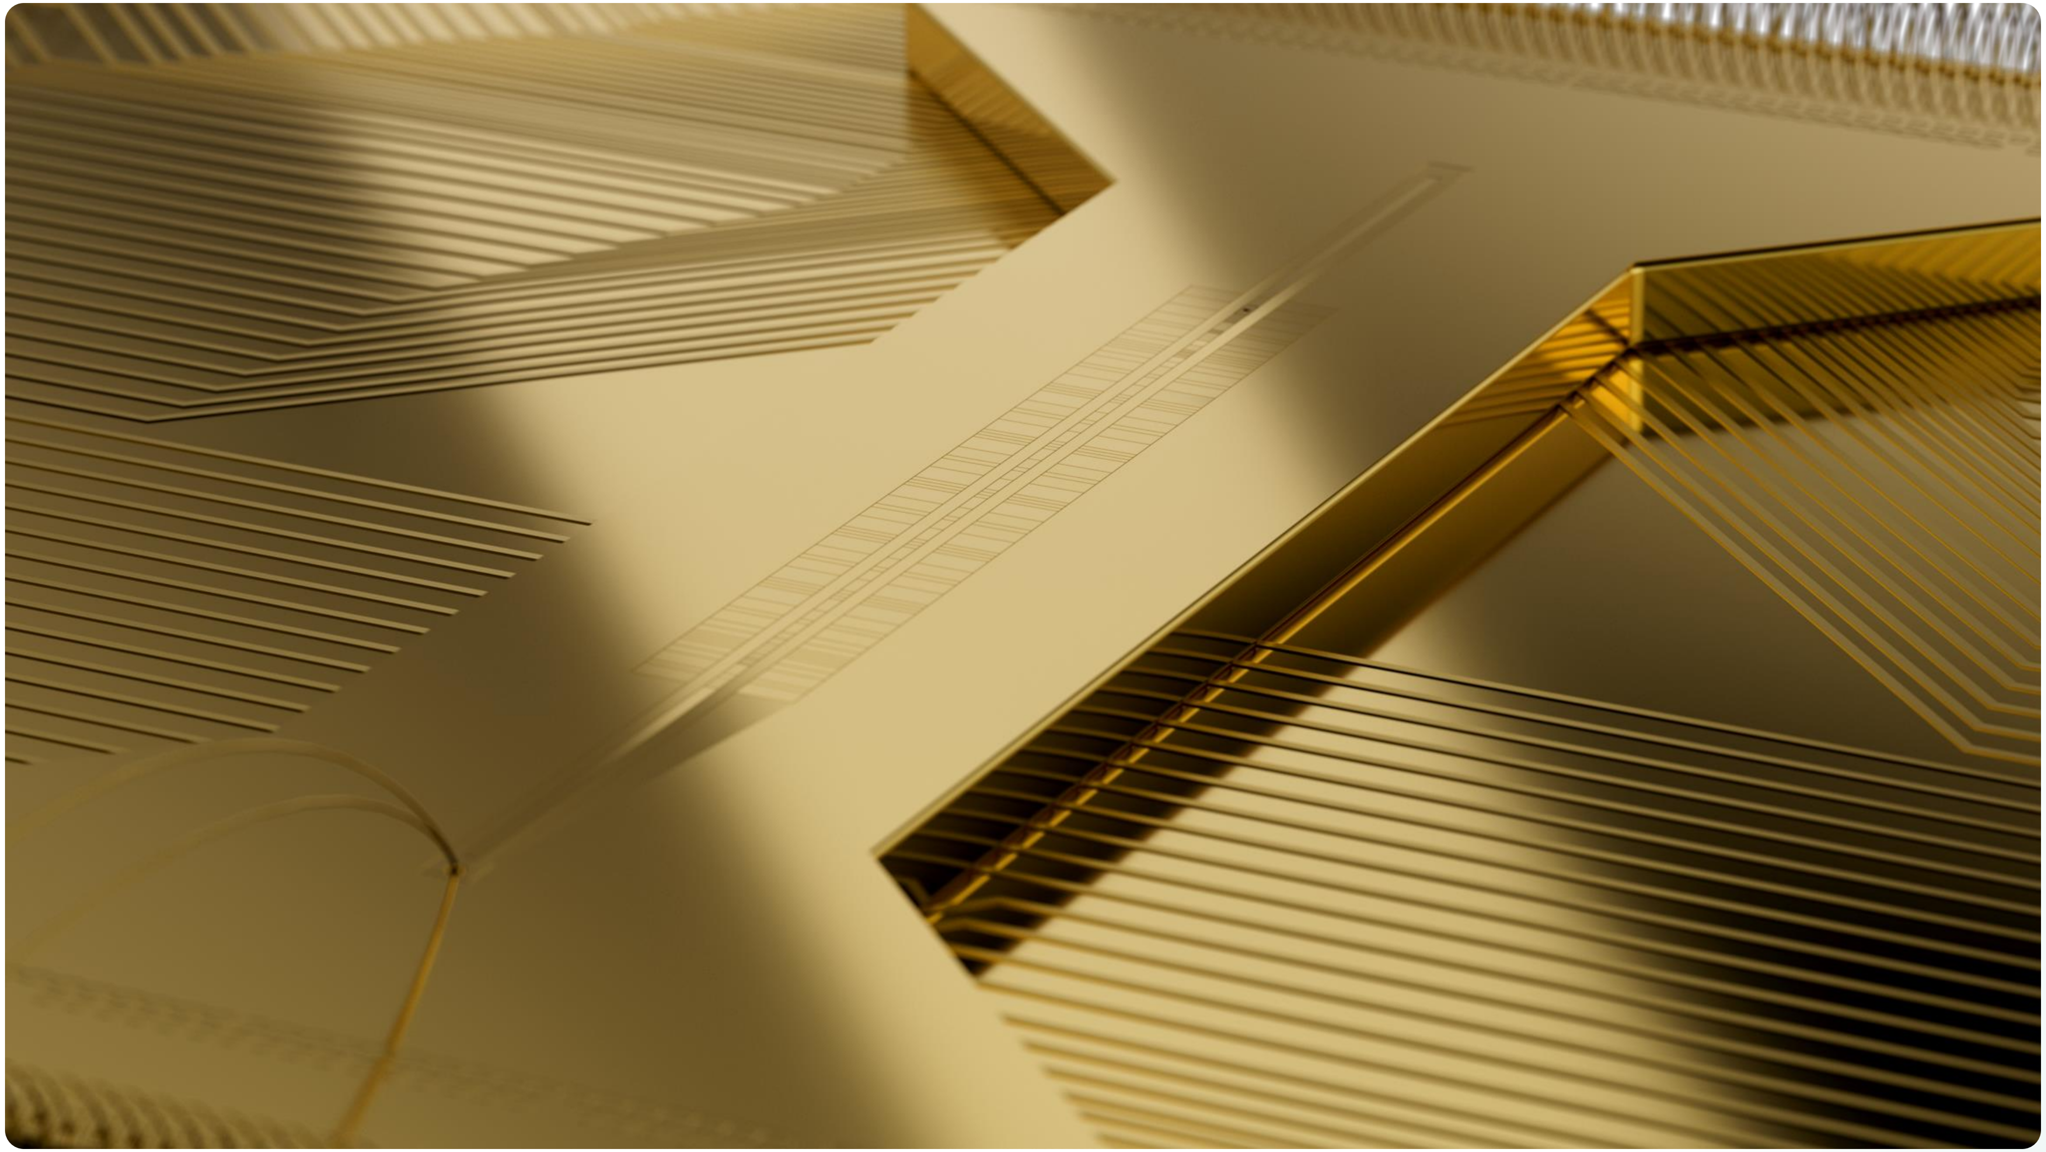

QUANTINUUM

ION TRAP ARCHITECTURE

System Model H1 with N=20 qubits

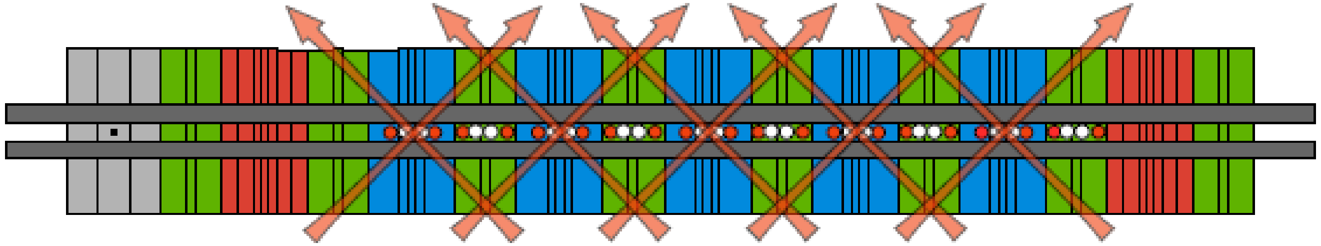

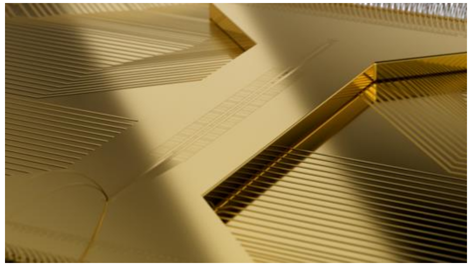

ARCHITECTURE FEATURES

- Identical, high-quality qubits
- Dedicated interaction zones
- Short ion chains
- High fidelity quantum gates
- Ions transport from zone to zone

Quantum bits (qubits) are stored in the electronic states of Yb^{+} ions

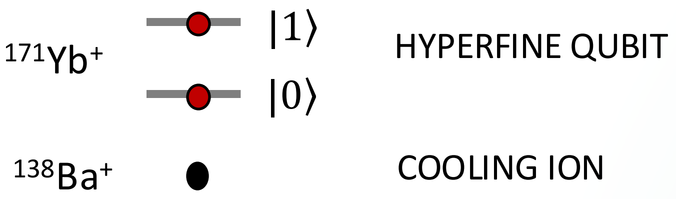

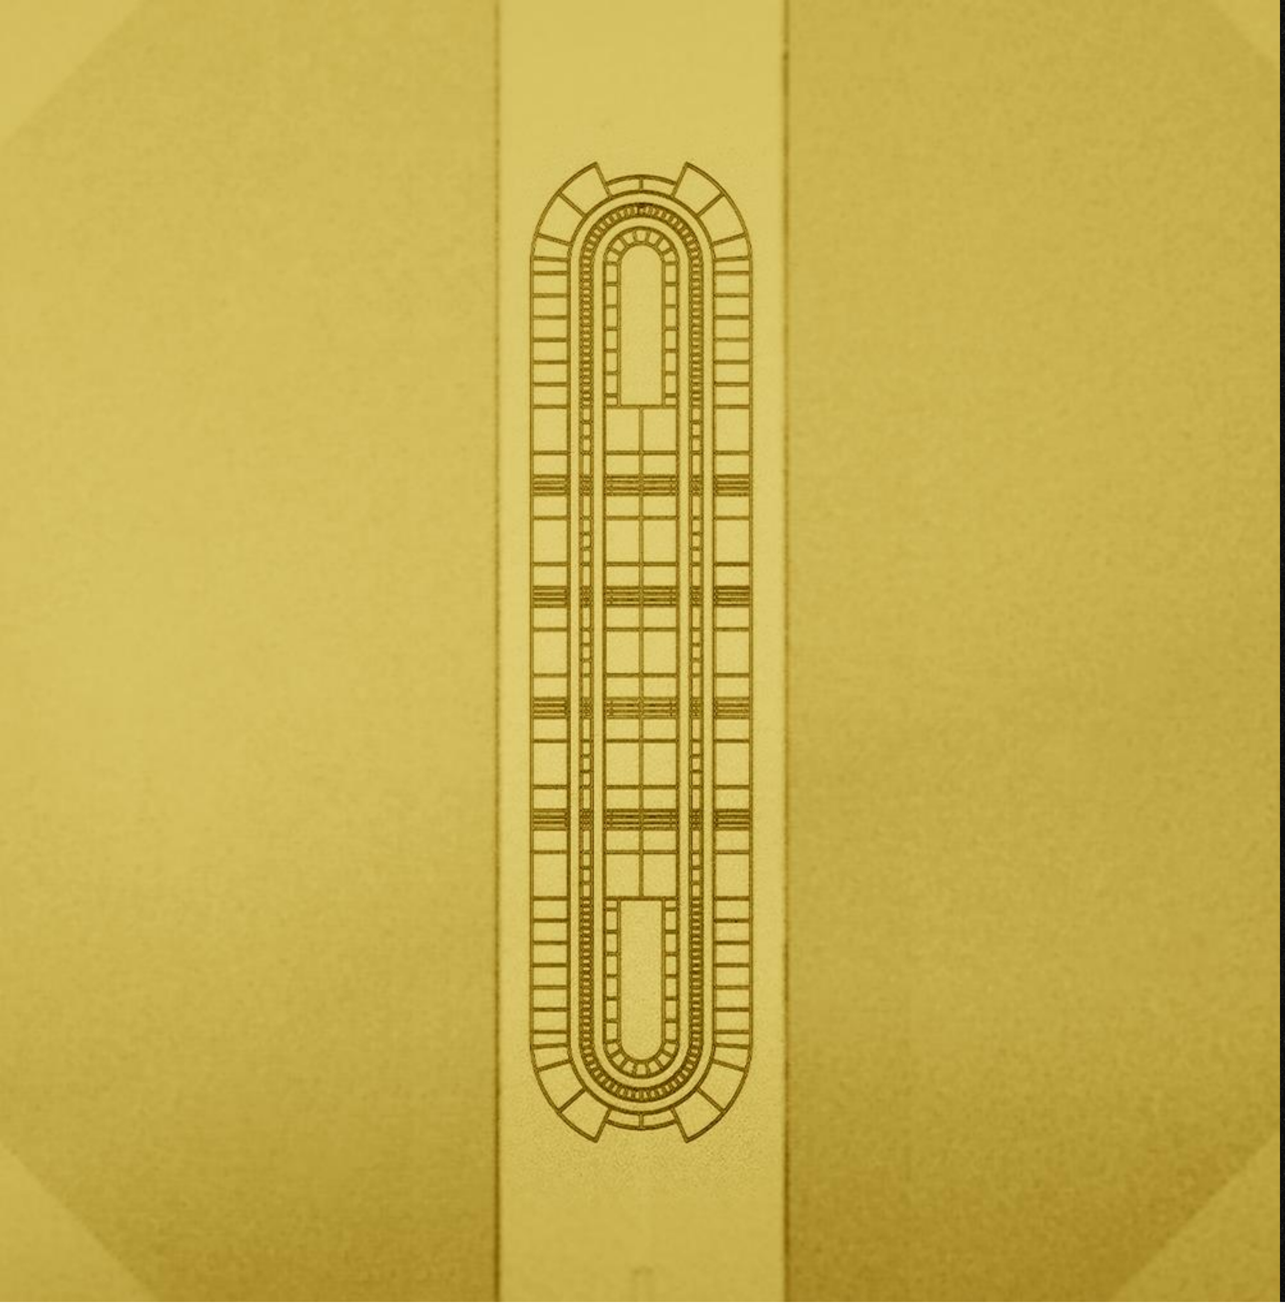

QUANTINUUM

Now at 56 Qubits

QUANTINUUM

Quantinuum High Quality Qubit Features

Arbitrary connectivity

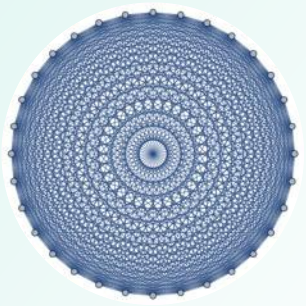

Direct entangling operations between any qubit pairing

Parameterized 2Q gate

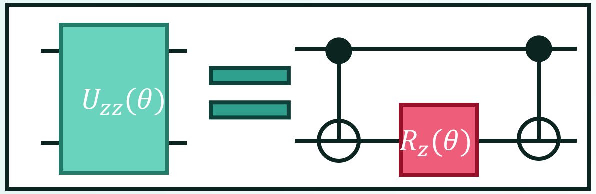

$$U_{ZZ}(\theta) = e^{-iZZ\theta/2}$$

Mid-circuit measurement and reset

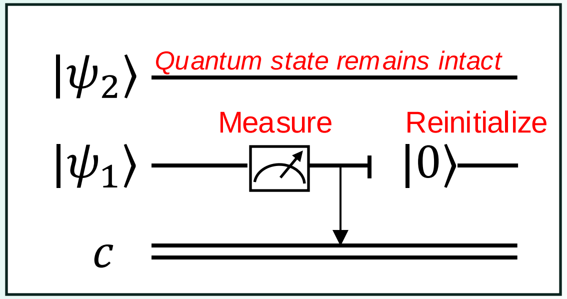

Allows qubit reuse on physical layer

Quantum computing benchmarks

Component

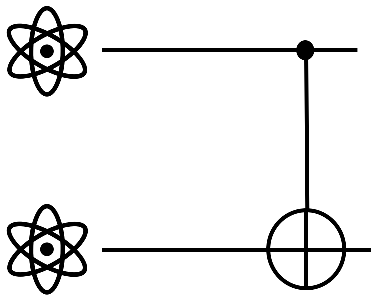

Failure rate of quantum operations

e.g. 1Q randomized benchmarking

System-level

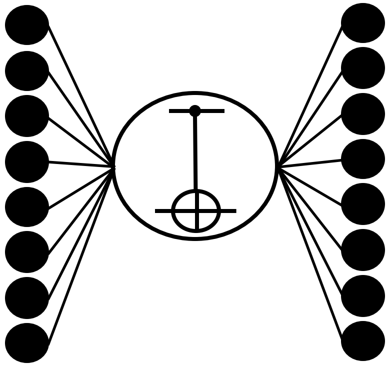

Complex but usually random operations

e.g. quantum volume

Applications

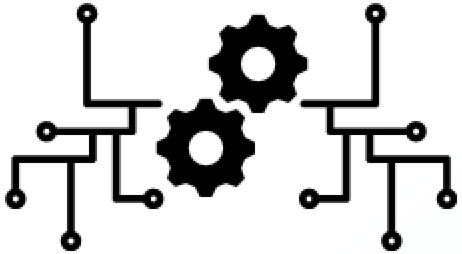

Testable versions of algorithms for potential applications

e.g. Hamiltonian simulation

Quantinuum Systems
The Most Benchmarked. The Highest Performing.

Component Level
- Single-qubit randomized benchmarking
- Two-qubit randomized benchmarking
- Two-qubit SU(4) randomized benchmarking
- Two-qubit parameterized gate randomized benchmarking
- Measurement crosstalk bright state depumping
- Reset crosstalk bright state depumping
- SPAM test
- Two-qubit cycle benchmarking

System Level
- Mirror benchmarking
- Quantum volume
- Random circuit sampling
- GHZ state fidelity

Algorithmic Level
- 1D transverse field Ising model simulation
- QAOA
- Repetition code
- HoloQUADS

PRX **13**, 041052 (2023)
github.com/CQCL/quantinuum-hardware-h2-benchmark

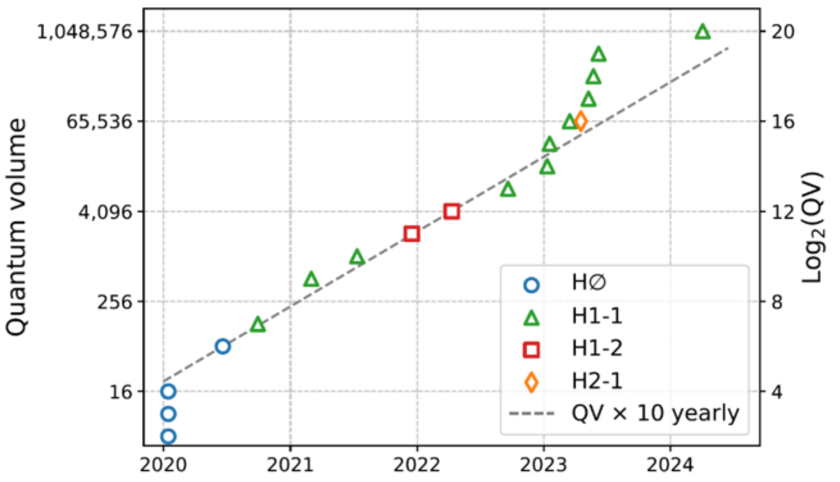

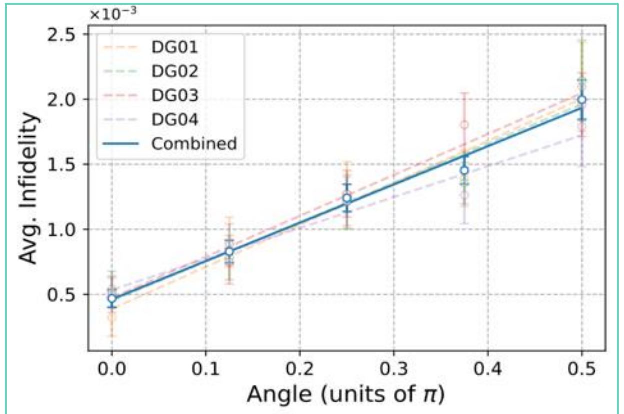

$$\epsilon_{2Q}(\theta)\ 10^{-3}[2.9(2)\,\theta/\pi + 0.46(6)]$$

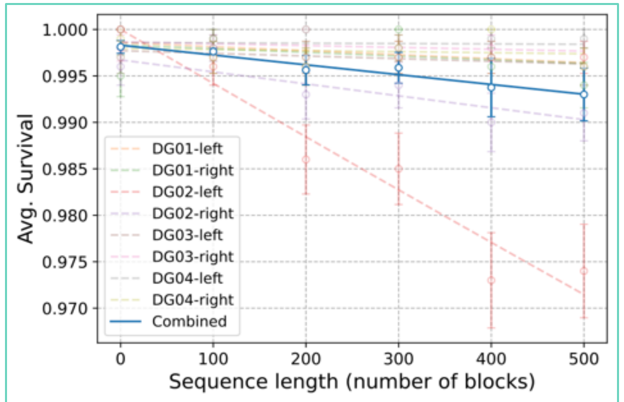

$$\epsilon_{MCM} = 4.5(3) \times 10^{-6}$$

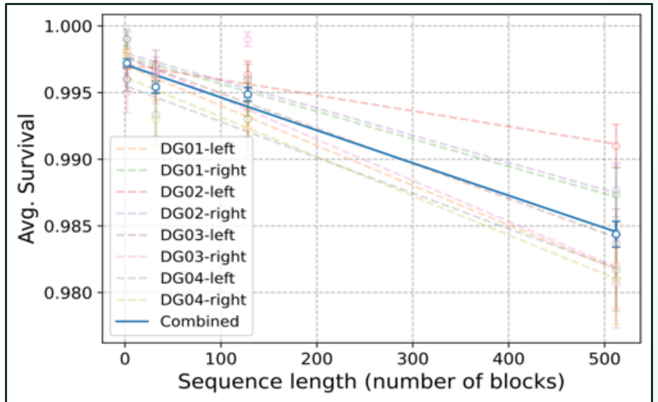

$$\epsilon_{1Q} = 2.5(3) \times 10^{-5}$$

QUANTINUUM

Random circuit sampling

Random quantum circuit

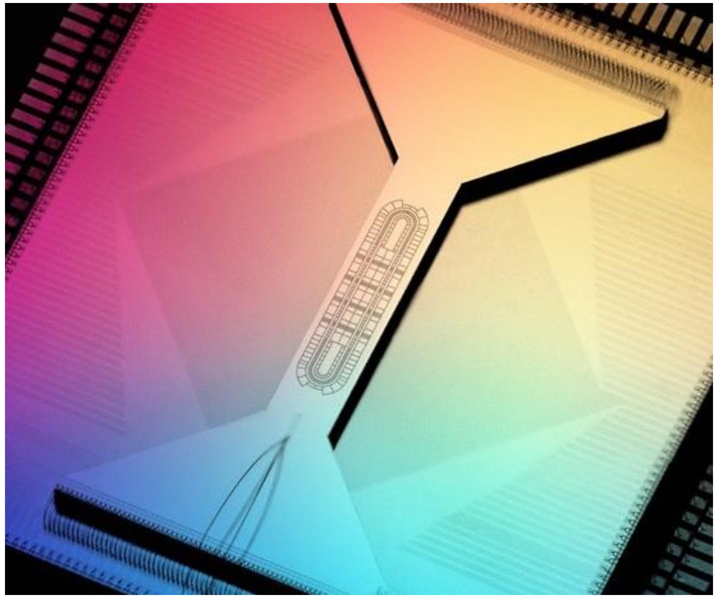

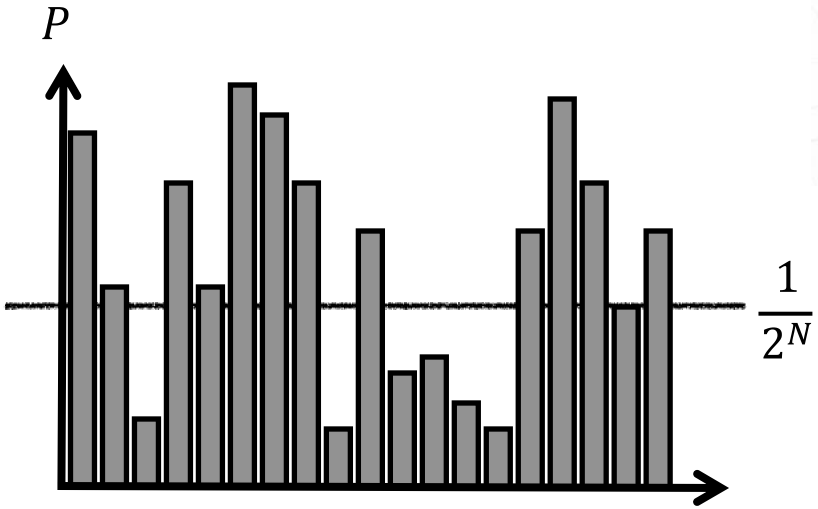

010101…
100110…
110010…

111001…
101110…
010011…

Random circuit sampling

Random quantum circuit

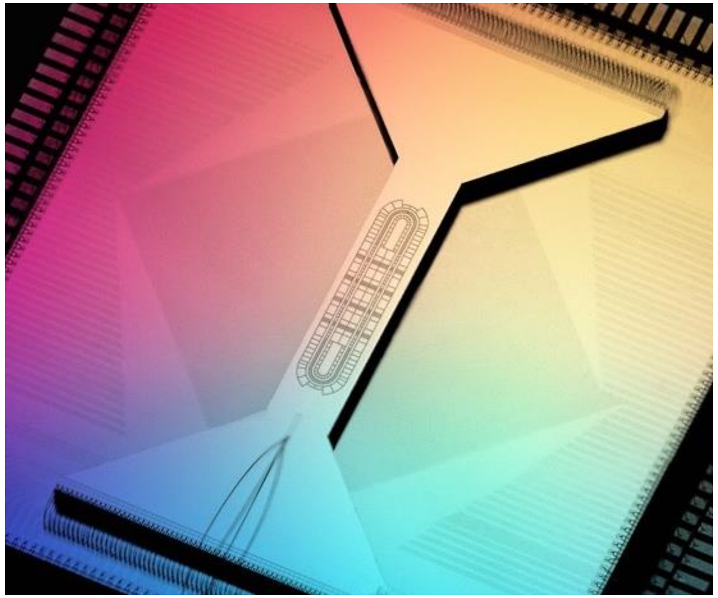

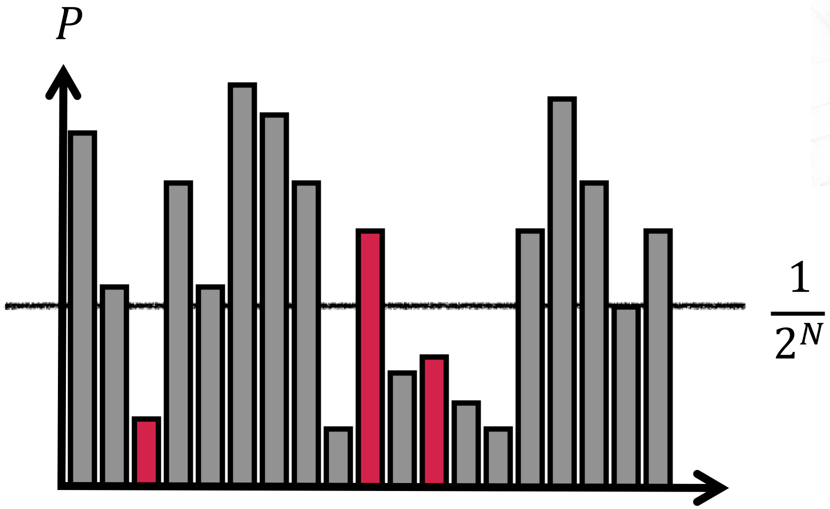

010101…
100110…
110010…

Bad random circuit sampling

111001…
101110…
010011…

QUANTINUUM

Random circuit sampling

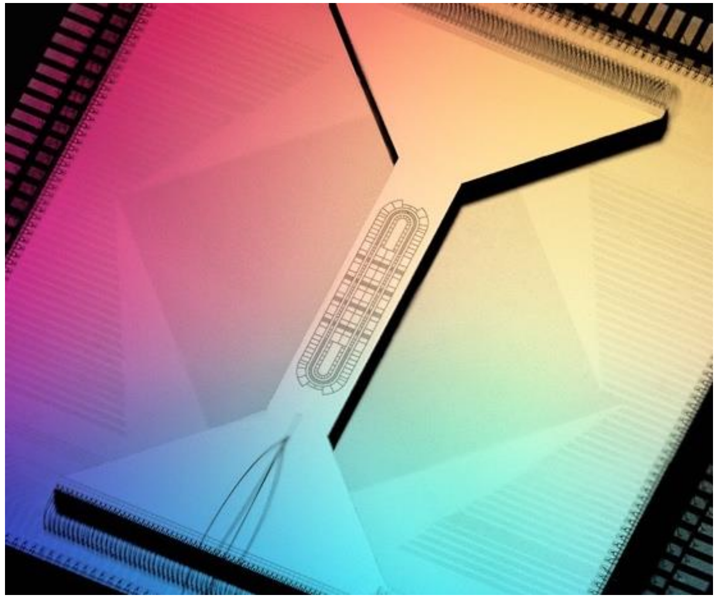

Random quantum circuit

010101…
100110…
110010…

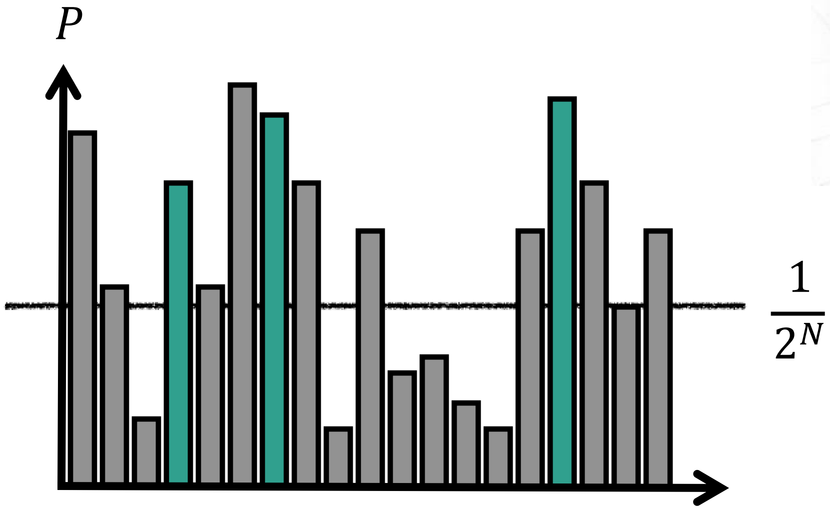

Good random circuit sampling

111001…
101110…
010011…

Mirrored circuits estimate fidelity

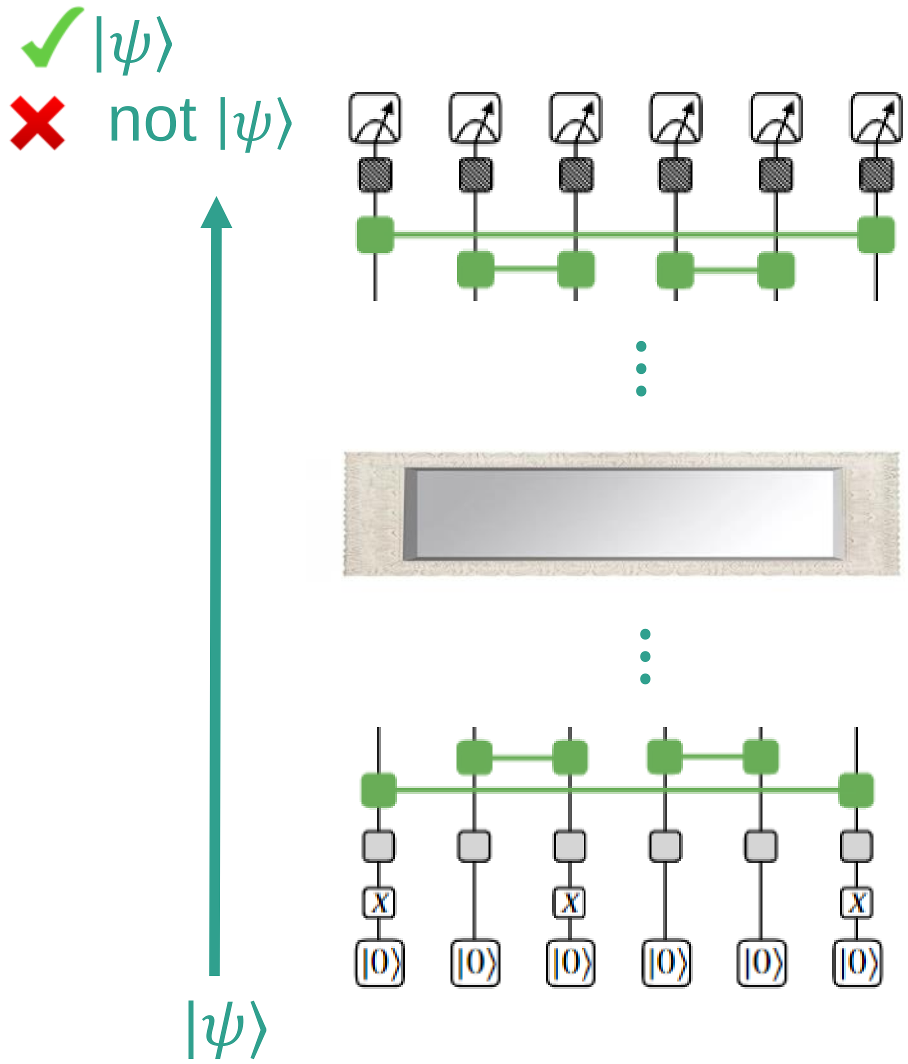

How hard is this classically?

Cost estimates based on state-of-the-art TN contraction heuristics:

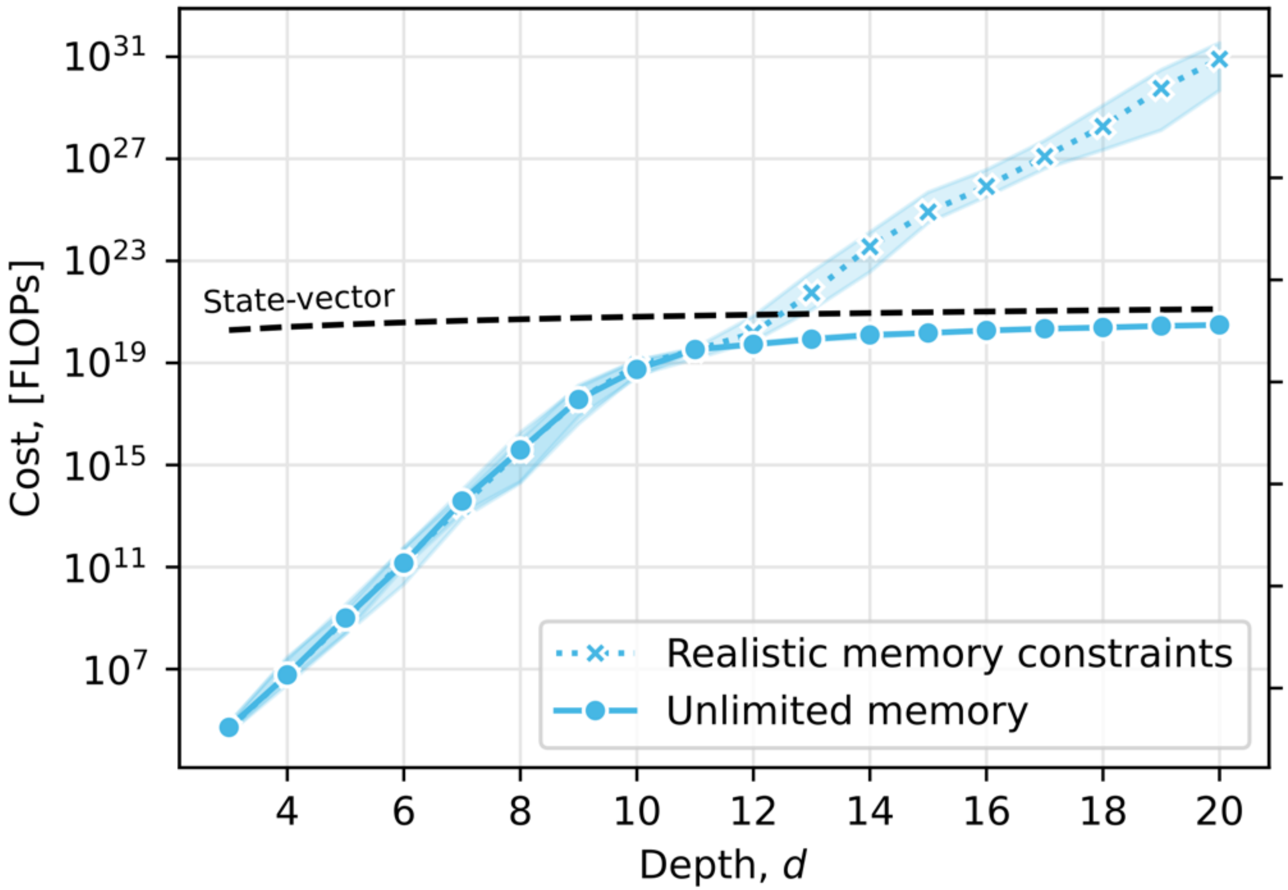

How hard is this classically?

Cost estimates based on state-of-the-art TN contraction heuristics:

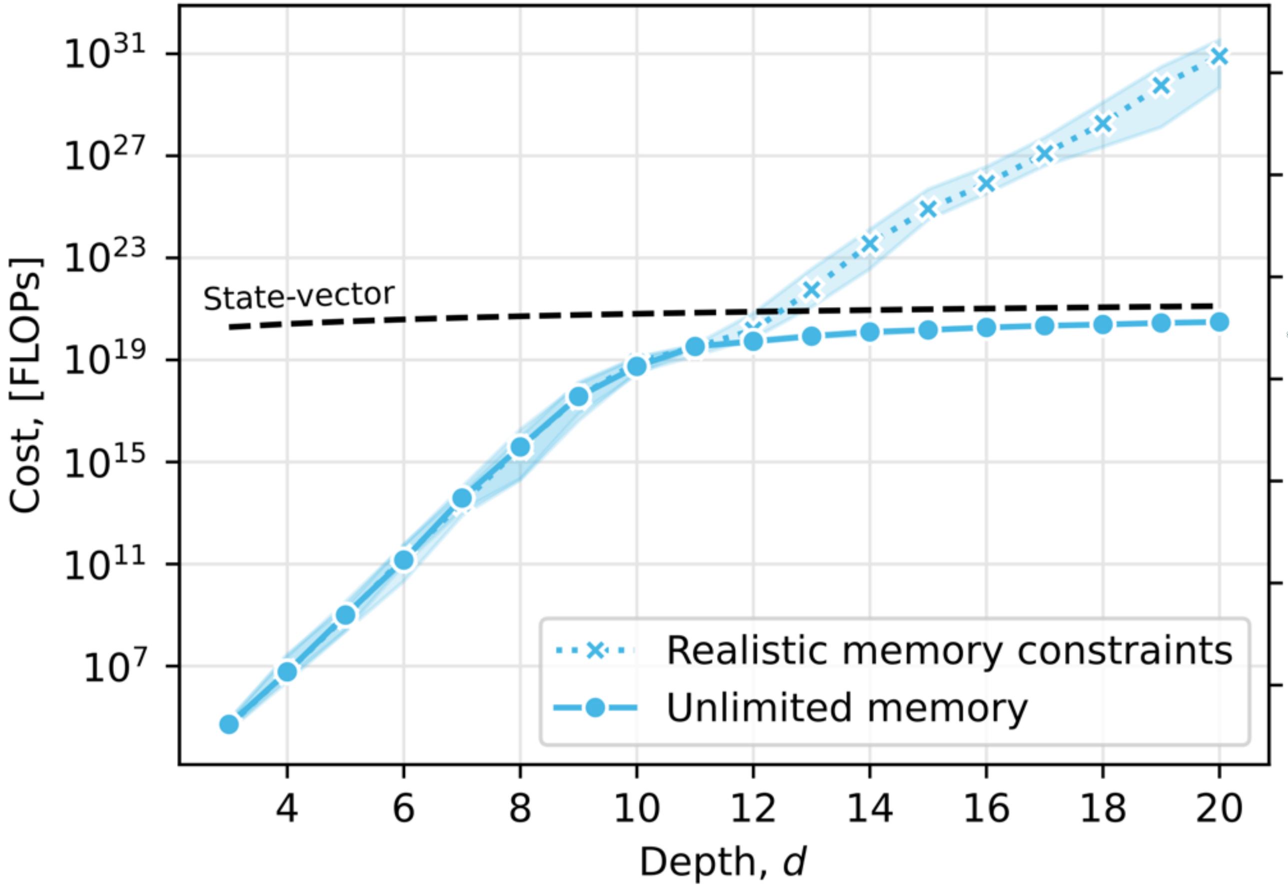

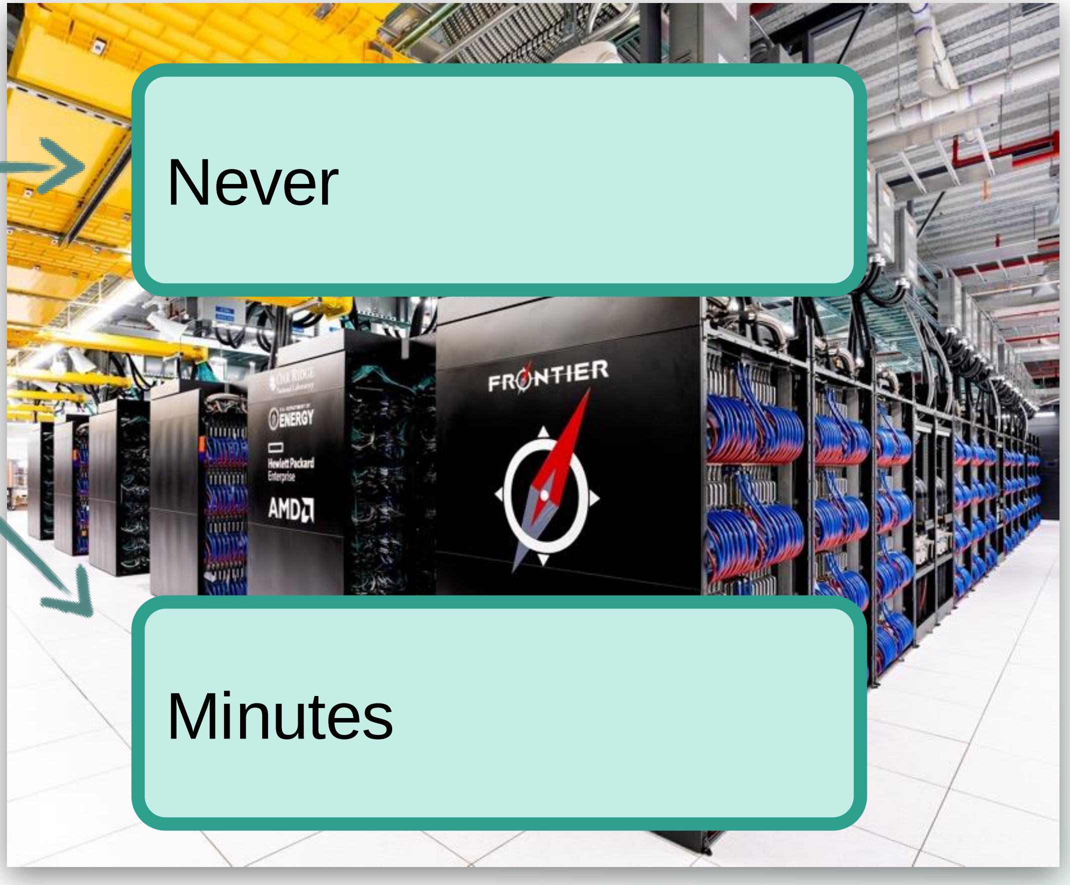

A brief history of quantum advantage claims

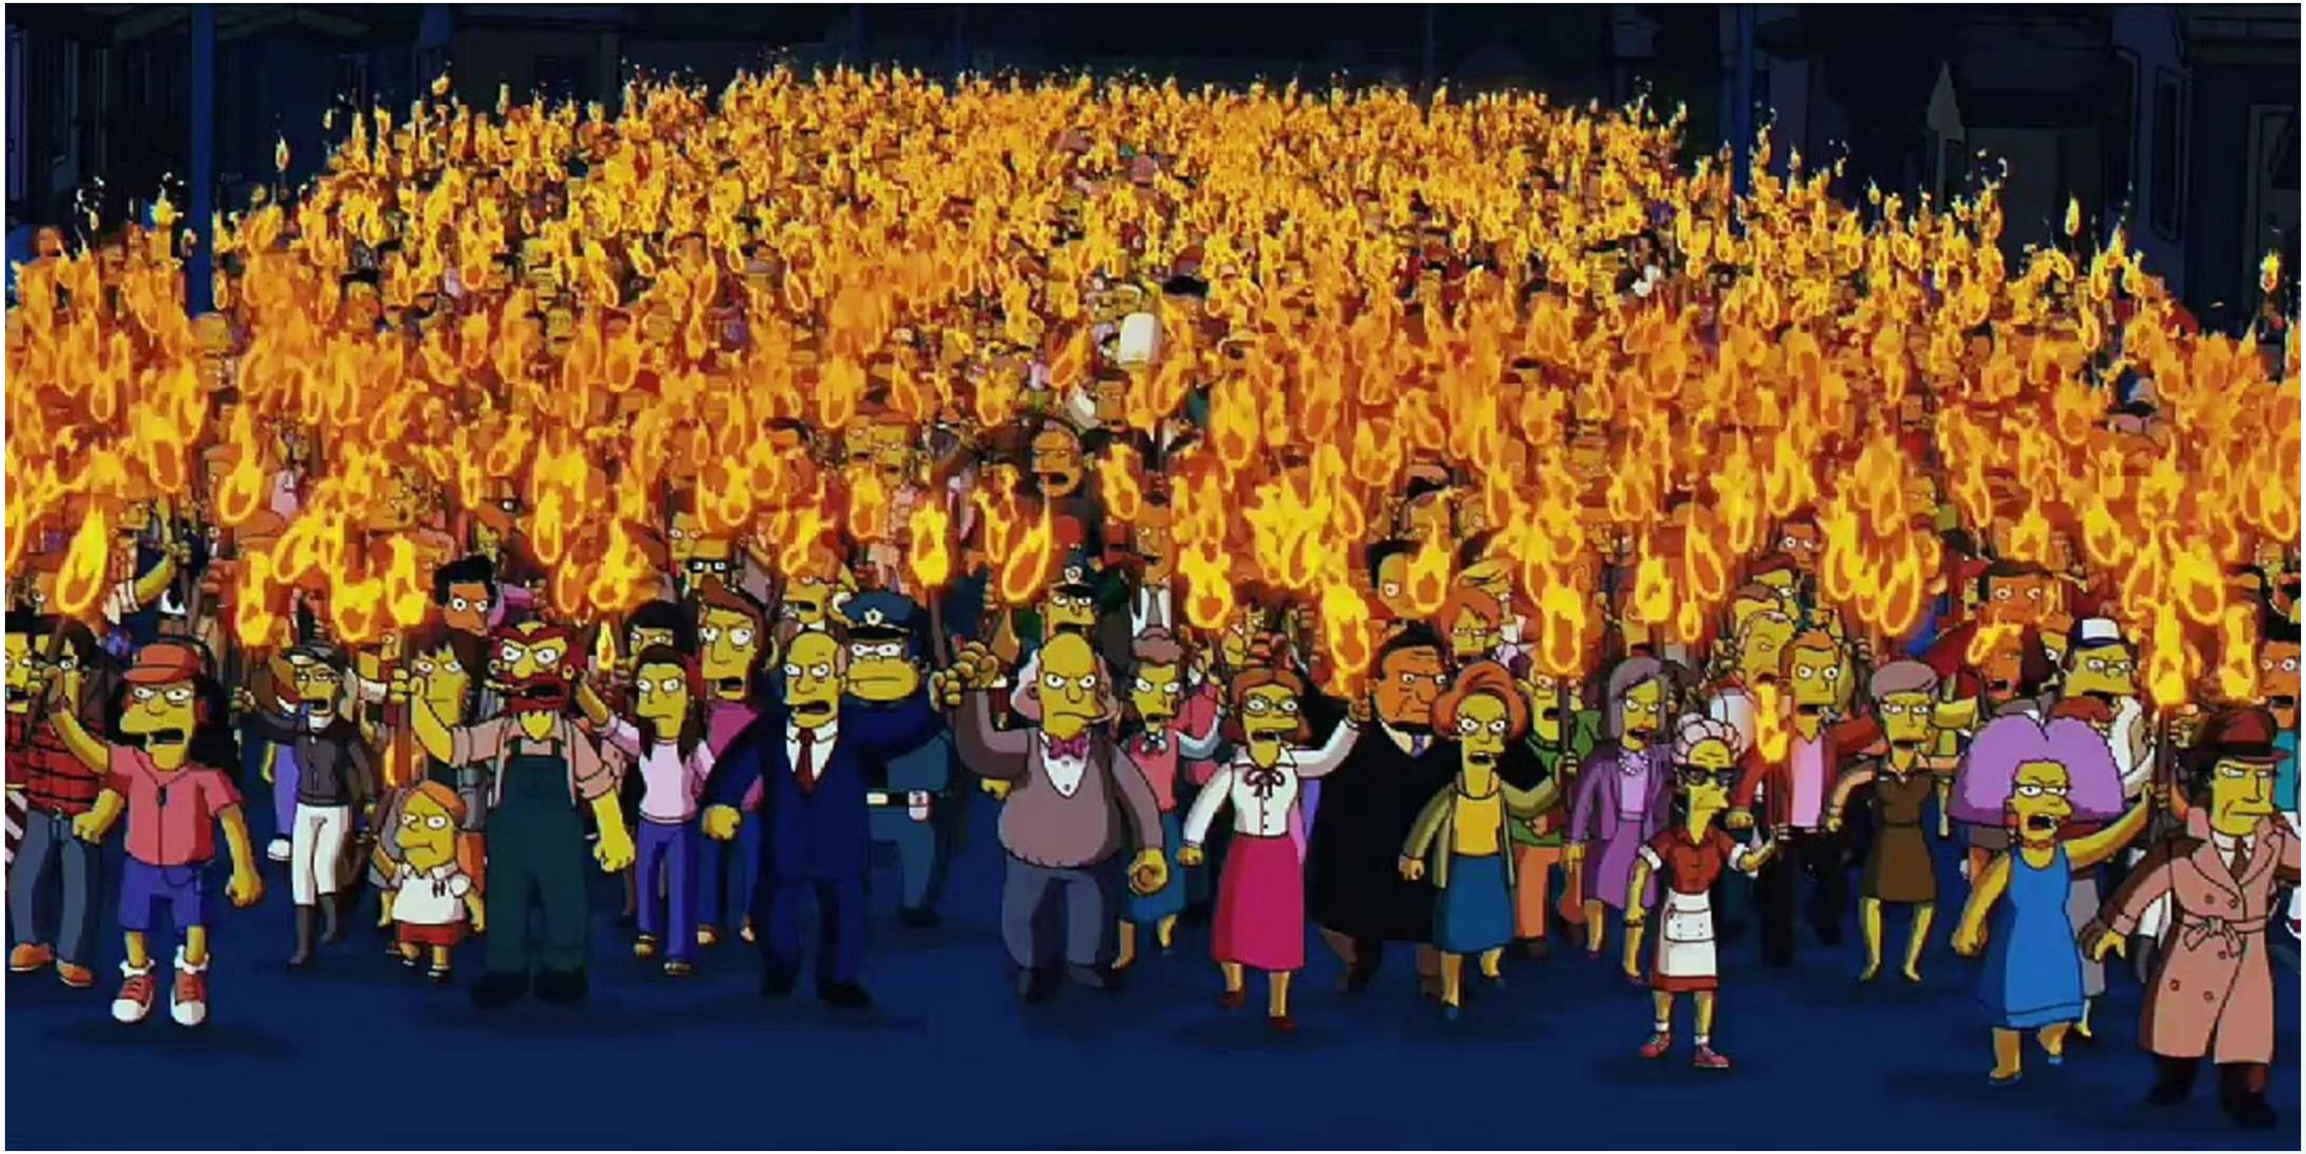

A brief history of quantum advantage claims

What limits the simulation of quantum computers?

Yiqing Zhou, E. Miles Stoudenmire, Xavier Waintal

It is imperative that useful quantum computers be very difficult to simulate classically; otherwise classical computers could be used for the applications envisioned for the quantum ones. Perfect quantum computers are unarguably exponentially difficult to simulate: the classical resources required grow exponentially with the number of qubits N or the depth D of the circuit. Real quantum computing devices, however, are characterized by an exponentially decaying fidelity $\mathcal{F} \sim (1-\epsilon)^{ND}$ with an error rate ϵ per operation as small as $\approx 1\%$ for current devices. In this work, we demonstrate that real quantum computers can be simulated at a tiny fraction of the cost that would be needed for a perfect quantum computer. Our algorithms compress the representations of quantum wavefunctions using matrix product states (MPS), which capture states with low to moderate entanglement very accurately. This compression introduces a finite error rate ϵ so that the algorithms closely mimic the behavior of real quantum computing devices. The computing time of our algorithm increases only linearly with N and D. We illustrate our algorithms with simulations of random circuits for qubits connected in both one and two dimensional lattices. We find that ϵ can be decreased at a polynomial cost in computing power down to a minimum error ϵ_∞. Getting below ϵ_∞ requires computing resources that increase exponentially with ϵ_∞/ϵ. For a two dimensional array of $N = 54$ qubits and a circuit with Control-Z gates, error rates better than state-of-the-art devices can be obtained on a laptop in a few hours. For more complex gates such as a swap gate followed by a controlled rotation, the error rate increases by a factor three for similar computing time.

Hyper-optimized tensor network contraction

Johnnie Gray, Stefanos Kourtis

Tensor networks represent the state-of-the-art in computational methods across many disciplines, including the classical simulation of quantum many-body systems and quantum circuits. Several applications of current interest give rise to tensor networks with irregular geometries. Finding the best possible contraction path for such networks is a central problem, with an exponential effect on computation time and memory footprint. In this work, we implement new randomized protocols that find very high quality contraction paths for arbitrary and large tensor networks. We test our methods on a variety of benchmarks, including the random quantum circuit instances recently implemented on Google quantum chips. We find that the paths obtained can be very close to optimal, and often many orders or magnitude better than the most established approaches. As different underlying geometries suit different methods, we also introduce a hyper-optimization approach, where both the method applied and its algorithmic parameters are tuned during the path finding. The increase in quality of contraction schemes found has significant practical implications for the simulation of quantum many-body systems and particularly for the benchmarking of new quantum chips. Concretely, we estimate a speed-up of over 10,000× compared to the original expectation for the classical simulation of the Sycamore `supremacy' circuits.

Solving the sampling problem of the Sycamore quantum circuits

Feng Pan, Keyang Chen, Pan Zhang

We study the problem of generating independent samples from the output distribution of Google's Sycamore quantum circuits with a target fidelity, which is believed to be beyond the reach of classical supercomputers and has been used to demonstrate quantum supremacy. We propose a new method to classically solve this problem by contracting the corresponding tensor network just once, and is massively more efficient than existing methods in obtaining a large number of uncorrelated samples with a target fidelity. For the Sycamore quantum supremacy circuit with 53 qubits and 20 cycles, we have generated one million uncorrelated bitstrings $\{\mathbf{s}\}$ which are sampled from a distribution $\hat{P}(\mathbf{s}) = |\hat{\psi}(\mathbf{s})|^2$, where the approximate state $\hat{\psi}$ has fidelity $F \approx 0.0037$. The whole computation has cost about 15 hours on a computational cluster with 512 GPUs. The obtained one million samples, the contraction code and contraction order is made public. If our algorithm could be implemented with high efficiency on a modern supercomputer with ExaFLOPS performance, we estimate that ideally, the simulation would cost a few dozens of seconds, which is faster than Google's quantum hardware.

Limitations of Linear Cross-Entropy as a Measure for Quantum Advantage

Xun Gao, Marcin Kalinowski, Chi-Ning Chou, Mikhail D. Lukin, Boaz Barak, Soonwon Choi

Demonstrating quantum advantage requires experimental implementation of a computational task that is hard to achieve using state-of-the-art classical systems. One approach is to perform sampling from a probability distribution associated with a class of highly entangled many-body wavefunctions. It has been suggested that this approach can be certified with the Linear Cross-Entropy Benchmark (XEB). We critically examine this notion. First, in a "benign" setting where an honest implementation of noisy quantum circuits is assumed, we characterize the conditions under which the XEB approximates the fidelity. Second, in an "adversarial" setting where all possible classical algorithms are considered for comparison, we show that achieving relatively high XEB values does not imply faithful simulation of quantum dynamics. We present an efficient classical algorithm that, with 1 GPU within 2s, yields high XEB values, namely 2-12% of those obtained in experiments. By identifying and exploiting several vulnerabilities of the XEB, we achieve high XEB values without full simulation of quantum circuits. Remarkably, our algorithm features better scaling with the system size than noisy quantum devices for commonly studied random circuit ensembles. To quantitatively explain the success of our algorithm and the limitations of the XEB, we use a theoretical framework in which the average XEB and fidelity are mapped to statistical models. We illustrate the relation between the XEB and the fidelity for quantum circuits in various architectures, with different gate choices, and in the presence of noise. Our results show that XEB's utility as a proxy for fidelity hinges on several conditions, which must be checked in the benign setting but cannot be assumed in the adversarial setting. Thus, the XEB alone has limited utility as a benchmark for quantum advantage. We discuss ways to overcome these limitations.

A brief history of quantum advantage claims

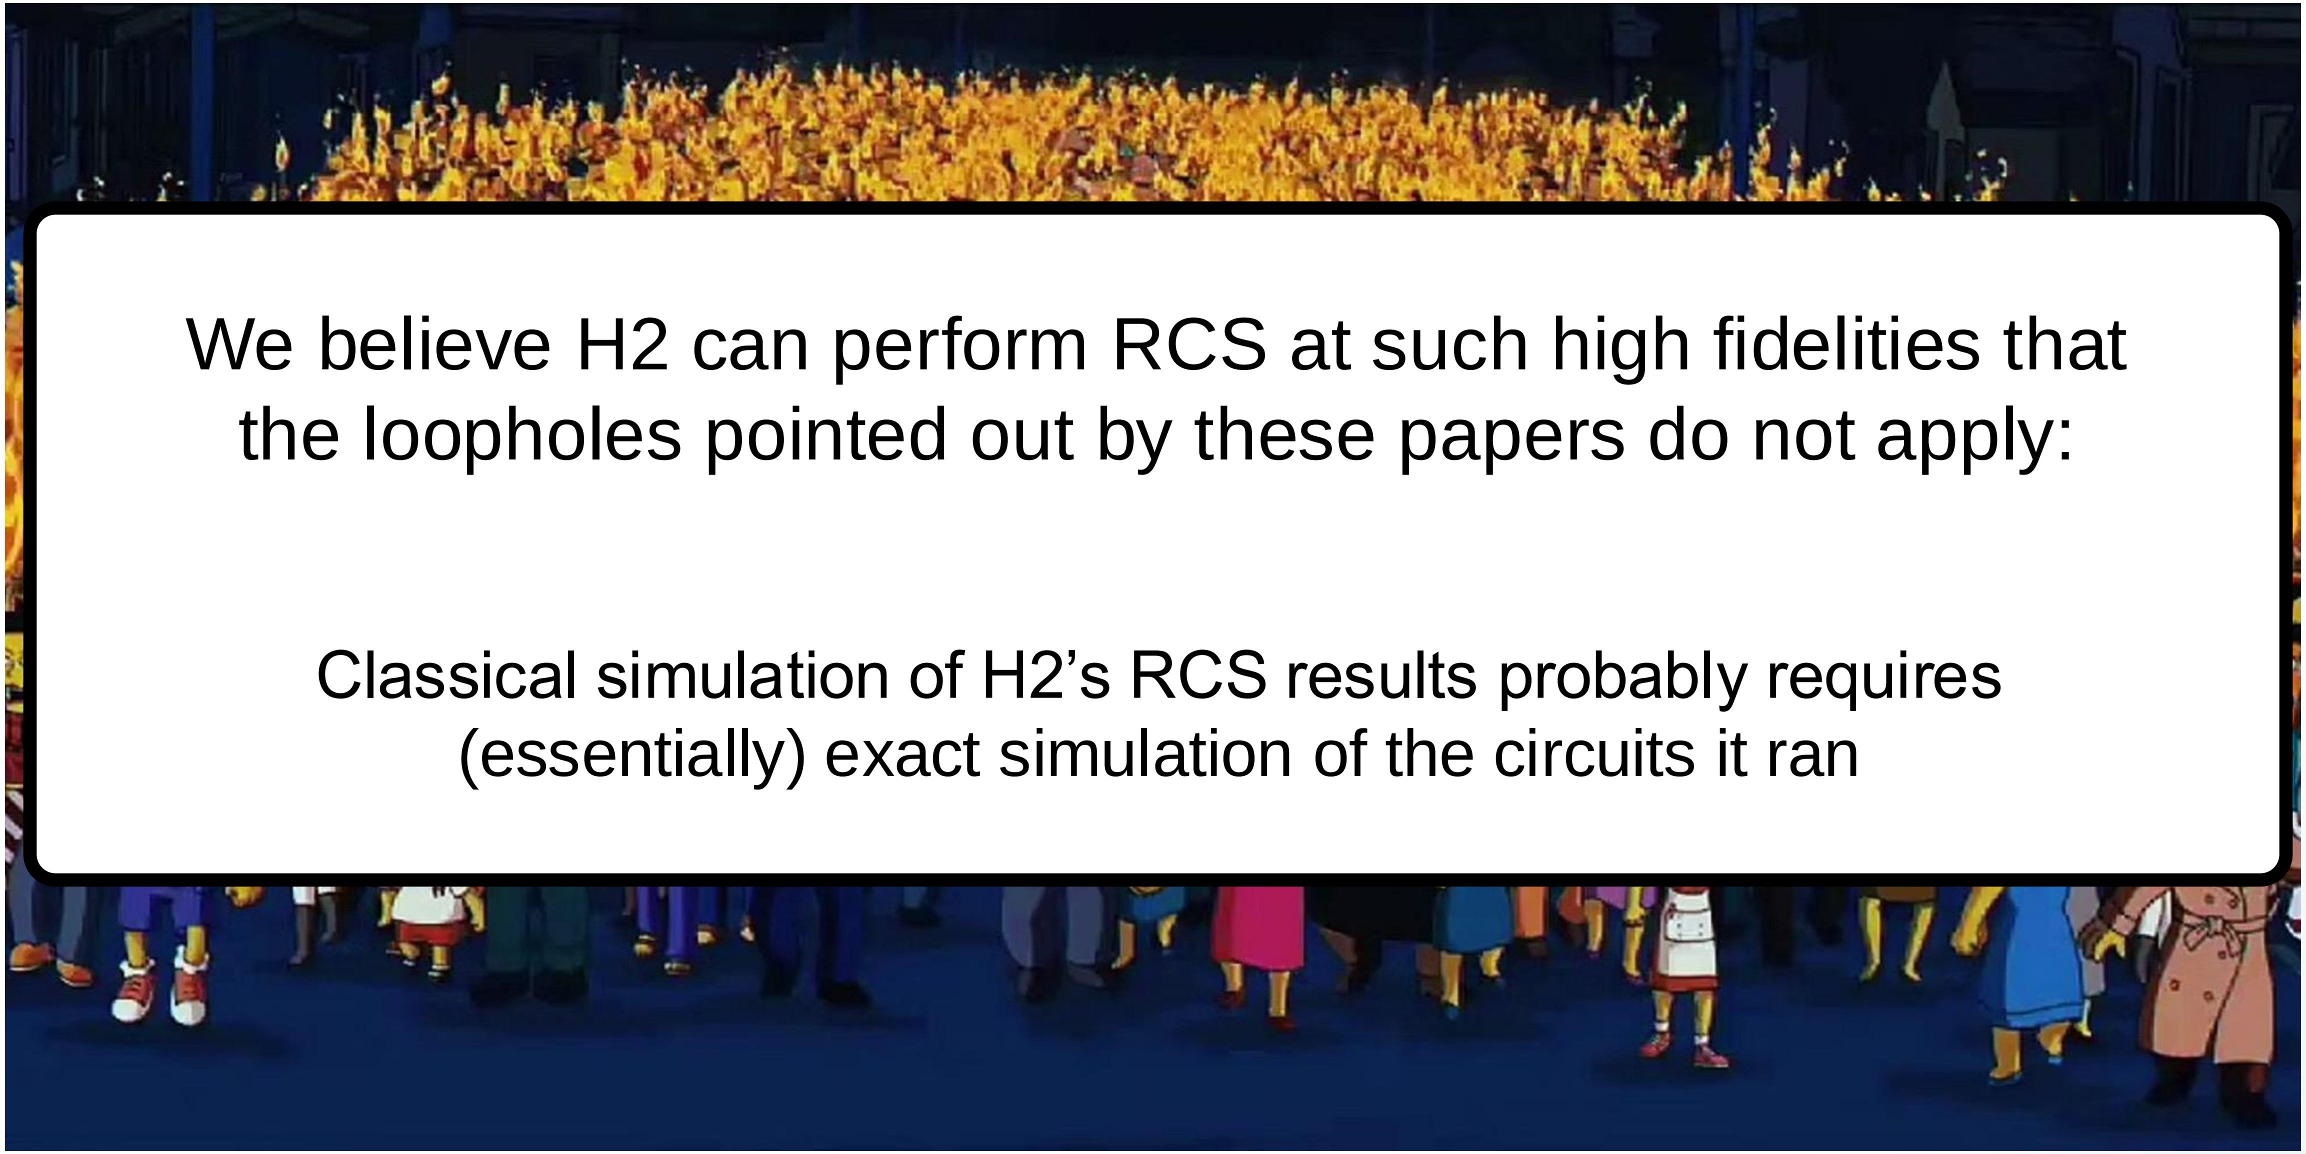

We believe H2 can perform RCS at such high fidelities that the loopholes pointed out by these papers do not apply:

Classical simulation of H2's RCS results probably requires (essentially) exact simulation of the circuits it ran

Pushing Quantum Advantage Into New Regime

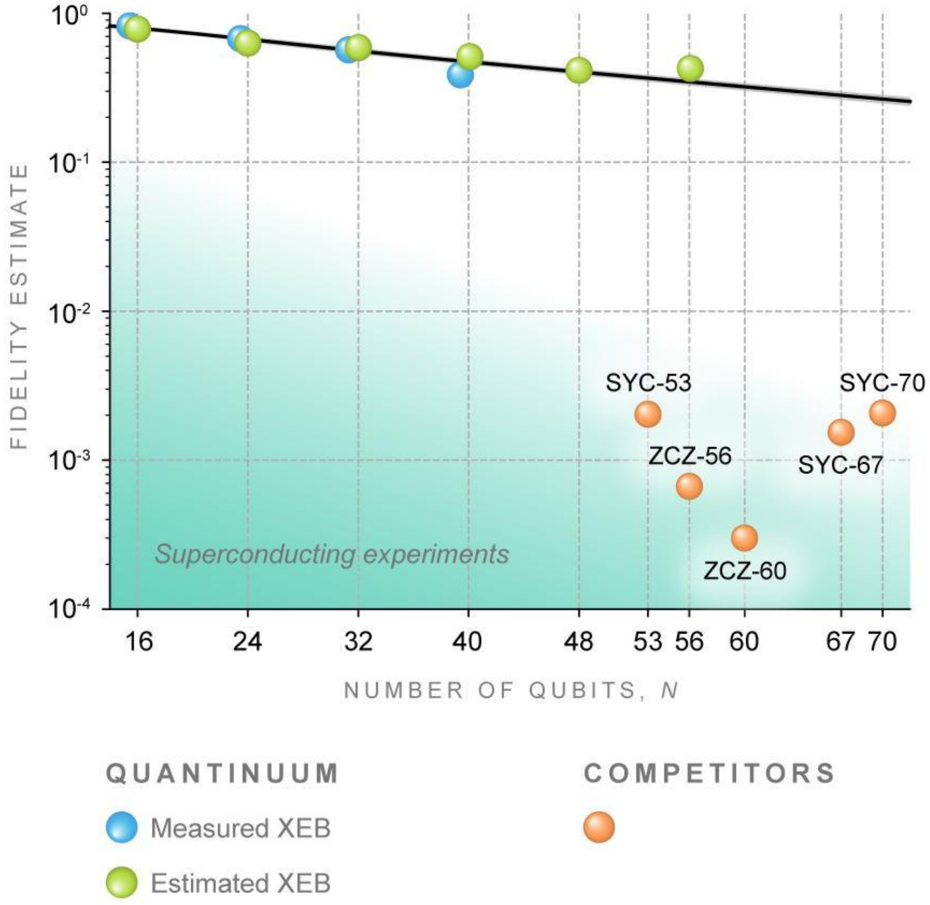

High fidelity of classically hard circuits gives room for classically hard results on less contrived problems

From random circuits to Hamiltonian simulation

Easy — Hard

Random circuit sampling

- Threshold for interesting problems
- Benchmark for hardware errors

Perfect entanglers

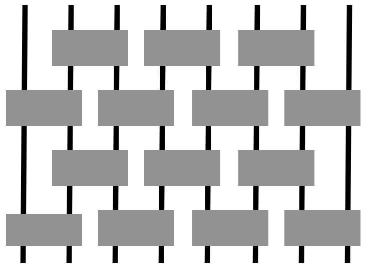

$e^{iA}\ e^{iB}\ e^{iA}\ e^{iB}\ e^{iA}\ e^{iB}$

Floquet circuits

- Prethermalization
- OTOCS / scrambling / teleportation
- Dynamics on random graphs

Partial entanglers

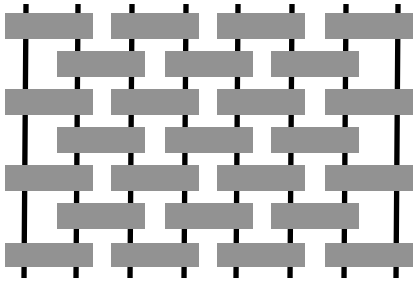

$e^{iA/f}\, e^{iB/f}\, e^{iA/f}\, e^{iB/f}\, e^{iA/f}\, e^{iB/f}$

Hamiltonian simulation

- Non-equilibrium quenches
- Thermal physics

Weak entanglers

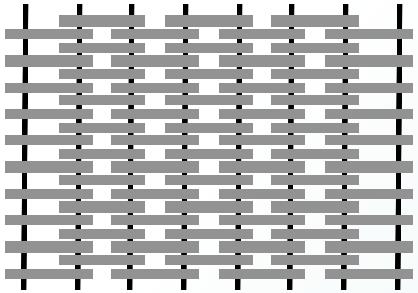

$\approx e^{i(A+B)}$

Our greatest aid: angle-dependent TQ gate errors

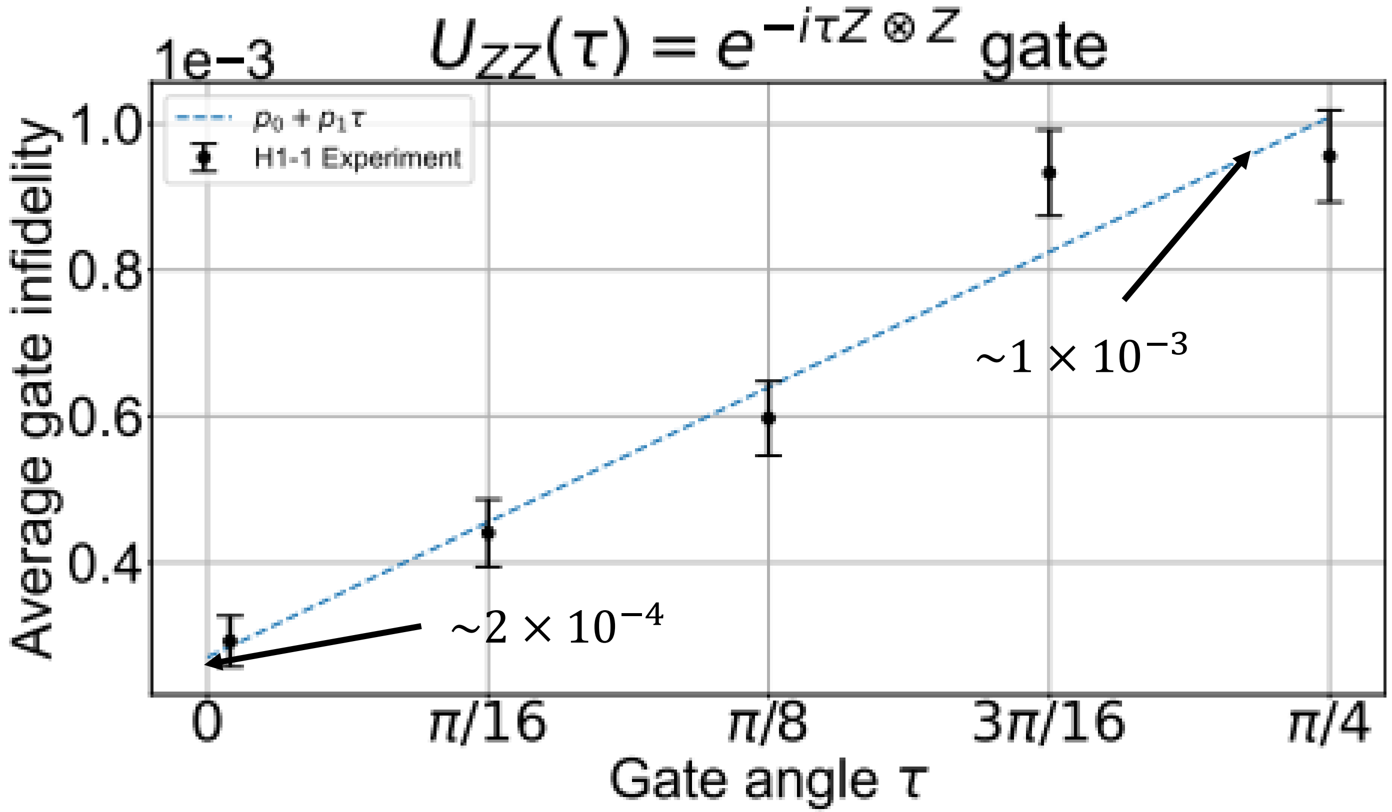

Prethermalization

Assessing the stability of digital quantum matter

Teeny time steps
↓
Continuous time translation invariance
↓
Energy Conservation
↓
Thermalization

Huge time steps (e.g. perfect entanglers)
↓
Discrete time translation invariance
↓
No energy Conservation
↓
Infinite temperature / random states

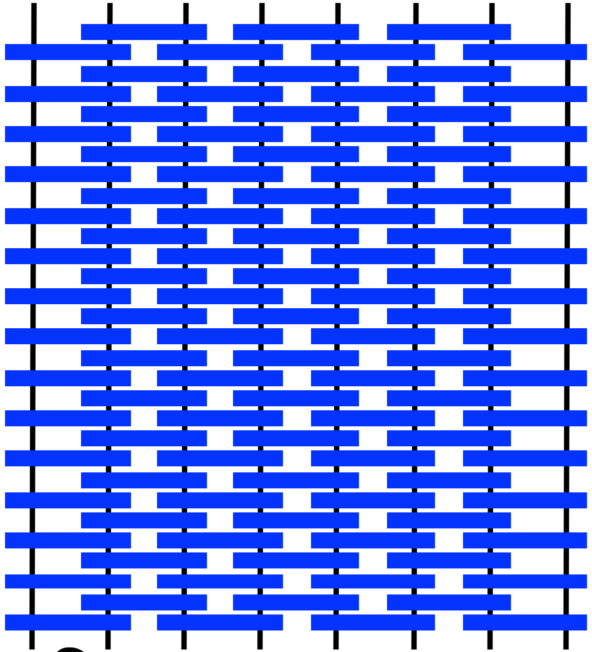

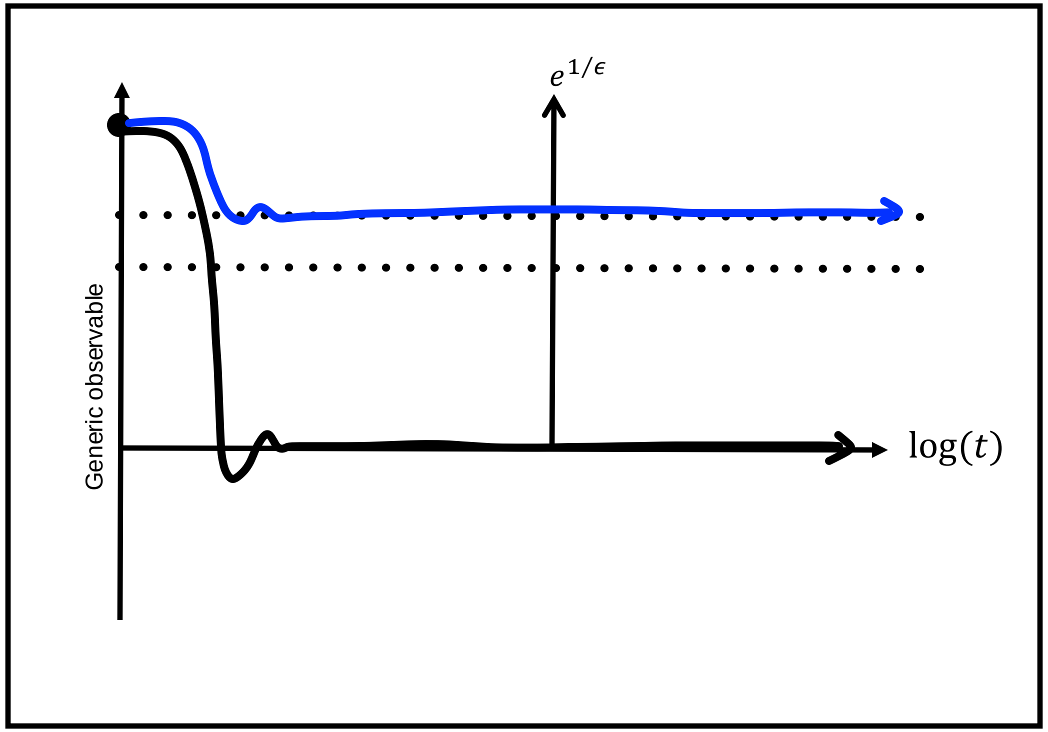

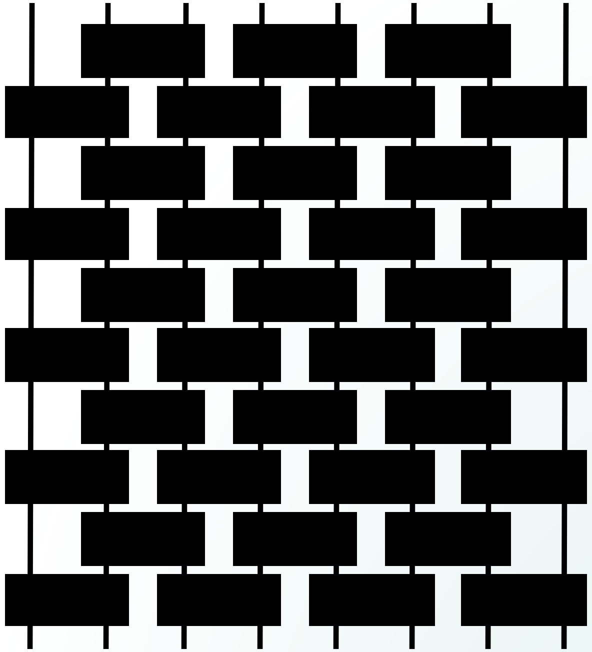

QUANTINUUM

Prethermalization

Assessing the stability of digital quantum matter

Teeny time steps

Continuous time translation invariance

Energy Conservation

Thermalization

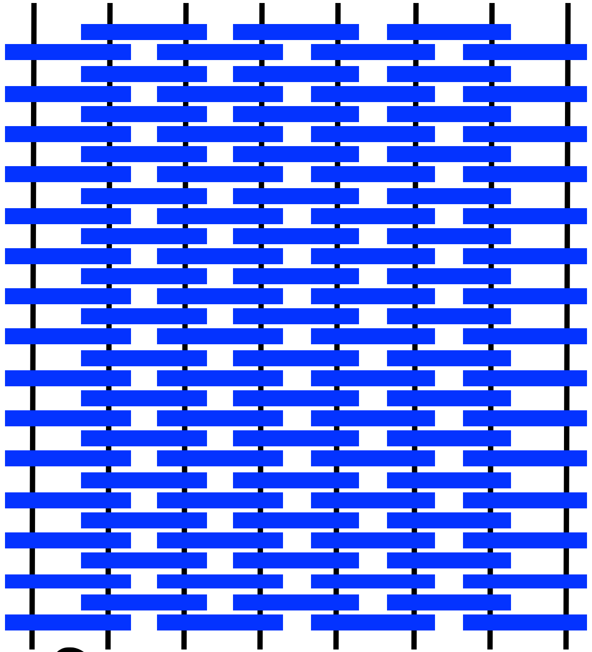

Middle ground: Emergent energy conservation

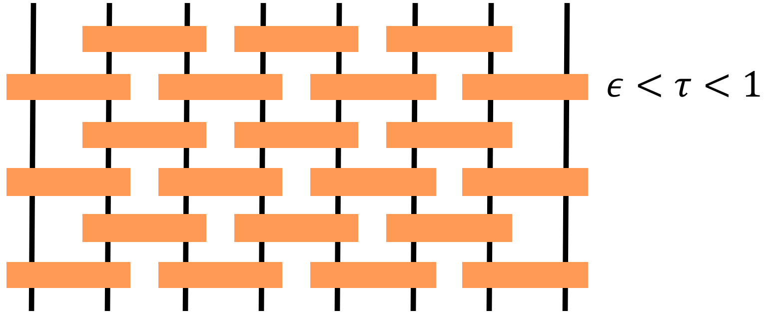

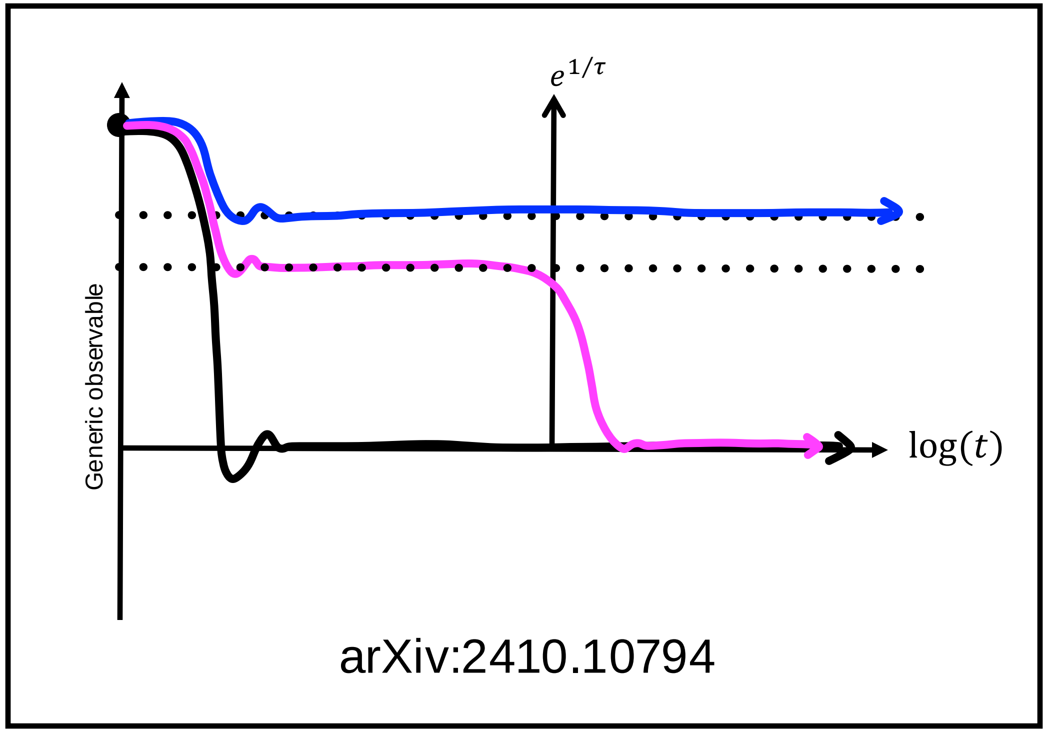

Huge time steps (e.g. perfect entanglers)

Discrete time translation invariance

No energy Conservation

Infinite temperature / random states

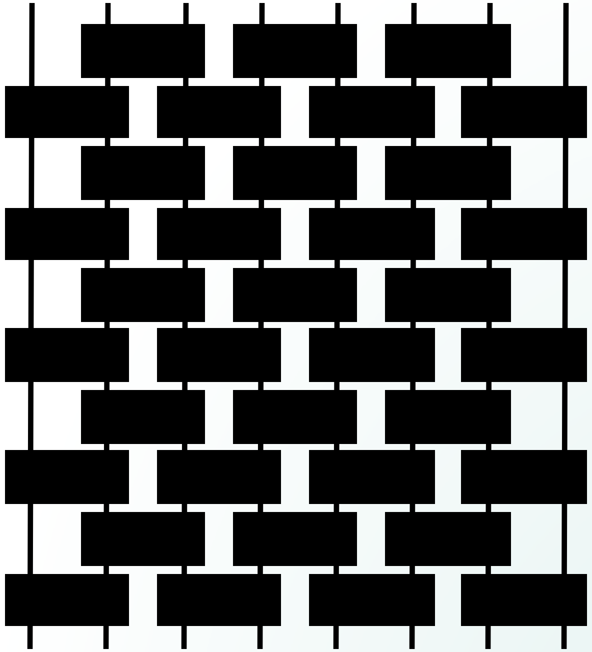

Introducing Helios

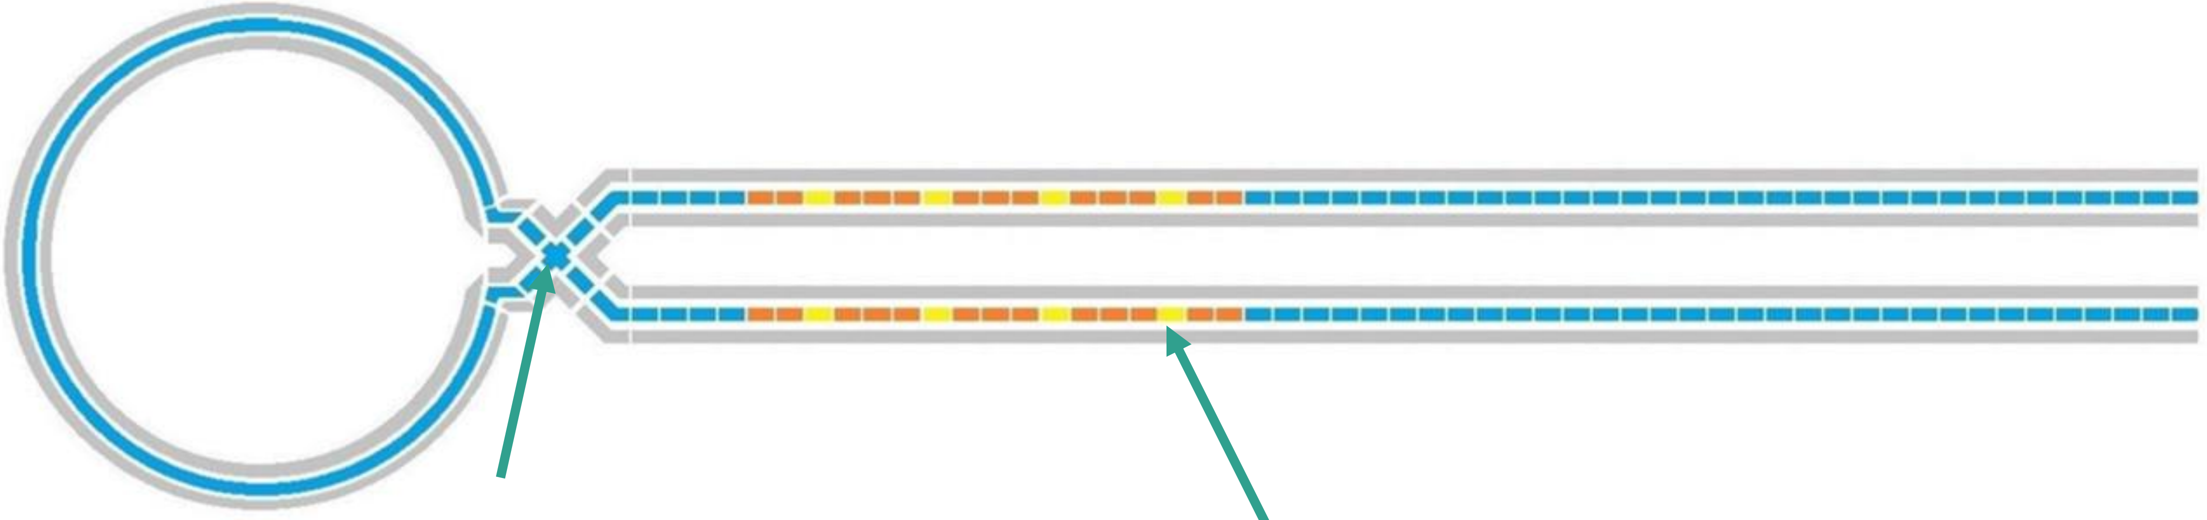

Junction transport allows for higher qubit counts and full connectivity at higher speeds

Barium-137 qubits will provide lower error rates

~100 fully connected qubits, lower errors, faster operations

From here to Apollo

	2020	2023	2025	2027	2029
SYSTEMS:	H1	H2	QUANTINUUM HELIOS	QUANTINUUM SOL	QUANTINUUM APOLLO
PHYSICAL QUBITS:	20	56	96	192	1000's
PHYSICAL 2-QUBIT GATE ERROR:	1×10^{-3}	1×10^{-3}	$< 5 \times 10^{-4}$	$< 2 \times 10^{-4}$	1×10^{-4}
LOGICAL QUBITS:		> 12	~ 50	~ 100	100's
LOGICAL ERROR RATES:		1×10^{-3}	$< 10^{-4}$	$\sim 10^{-5}$	1×10^{-5} to 1×10^{-10}*

*analysis based on recent literature in new, novel error correcting codes predict that error could be as low as 1E-10 in Apollo (ref: arXiv:2403.16054, arXiv:2308.07915)

Enabling Large-Scale Quantum Computation

QEC: Quantum error correction

Suppresses noise in a quantum computer → more operations

How? Encode qubits into many other qubits → logical qubits

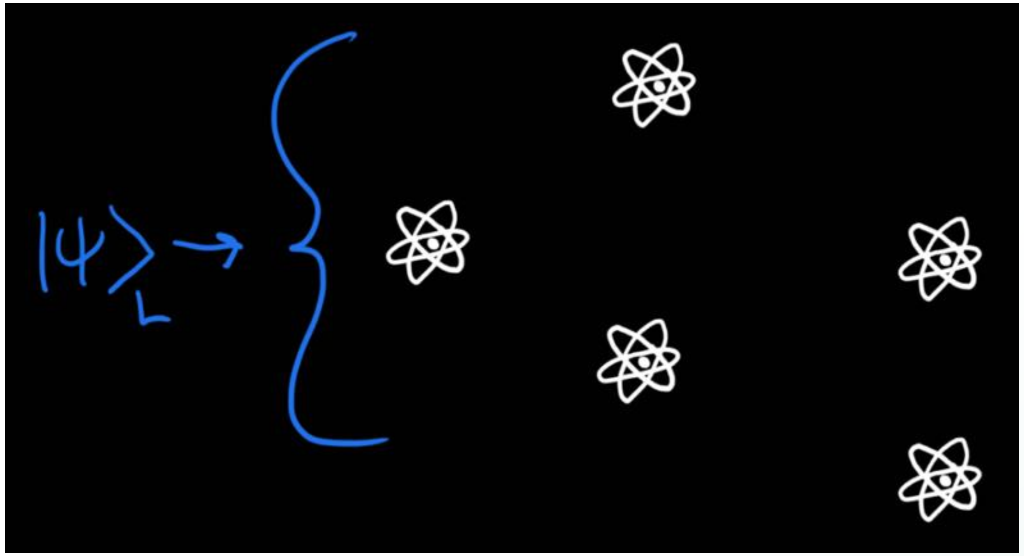

Error Correction

- Bit-flip error:

$$\mathbf{0} \leftrightarrow \mathbf{1}$$

- A key idea in both classical and quantum error correction is redundancy.

- The repetition code:

$$\mathbf{0} \rightarrow \mathbf{000}$$

$$\mathbf{1} \rightarrow \mathbf{111}$$

Quantum is A BIT different…

1. Not just bit flip errors
2. Measurement collapse
3. No cloning (need encoding circuits)
4. Universal gate sets are not easy

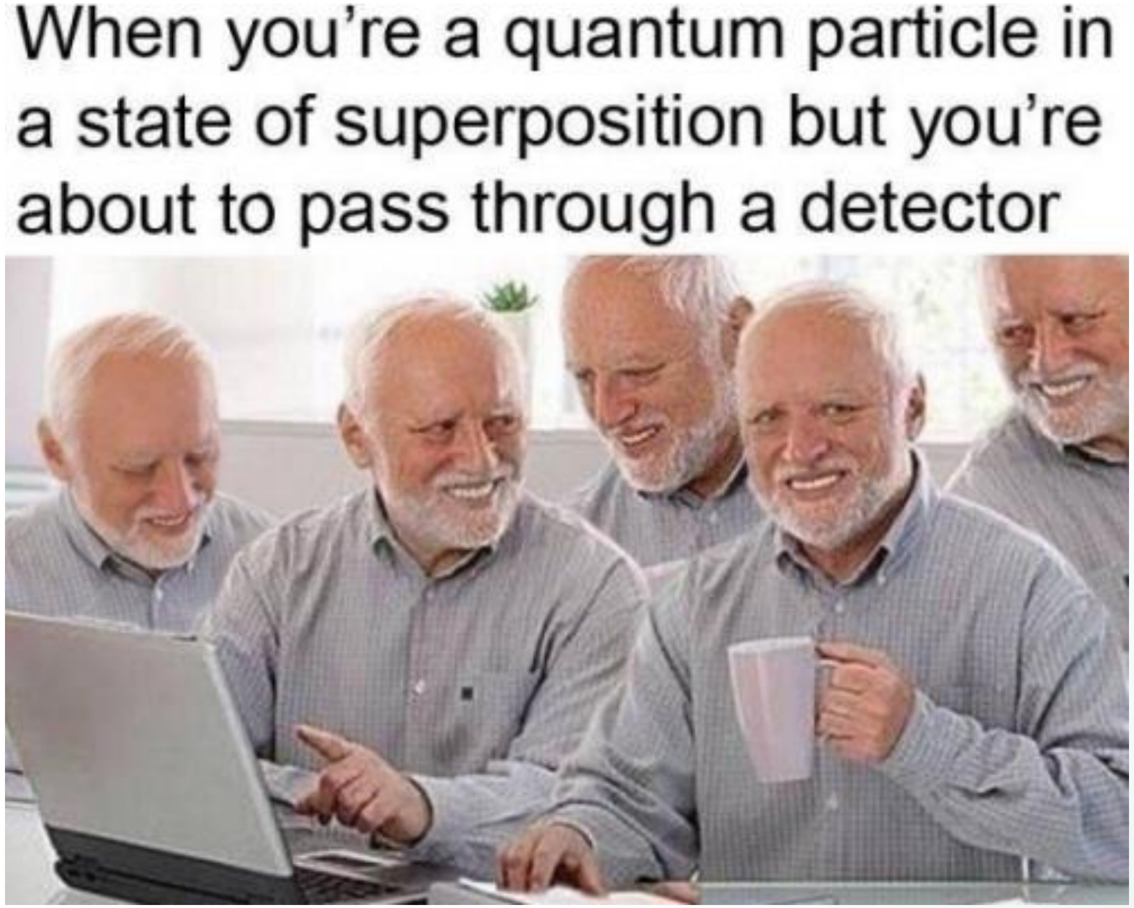

https://meme.xyz/m/meme/15316/when-youre-a-quantum-particle-in-a-state-of-superposition-but-youre-about-to-pass-through-a-detector.html

QUANTINUUM

MSFT x Quantinuum Announcement

- We demonstrated beyond break even error rates using both the Steane code and a code developed by Microsoft, the Carbon Code [[12,2,4]]

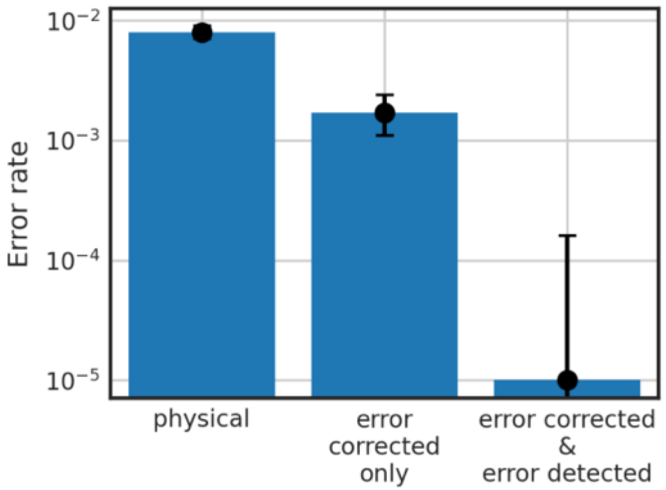

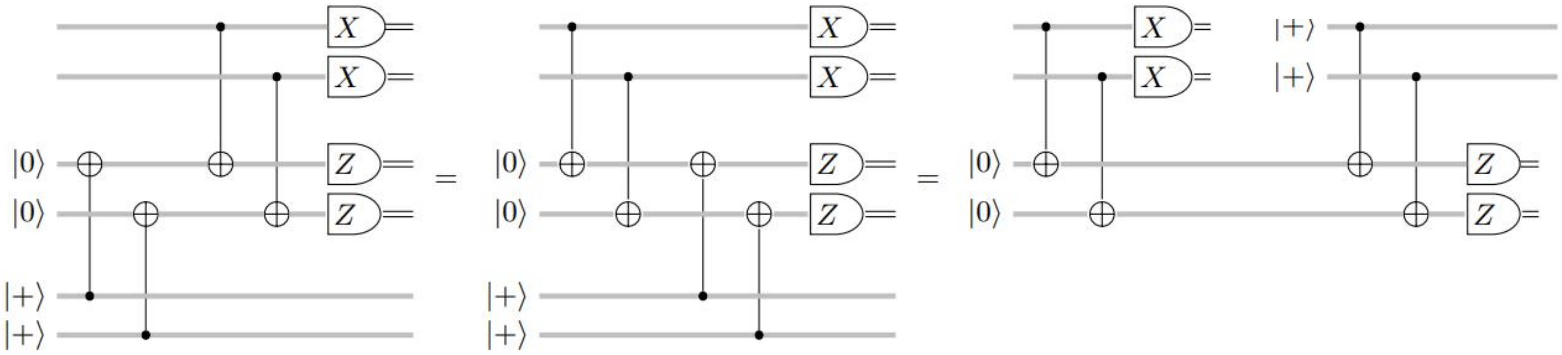

FIG. 5. Syndrome information can be obtained by performing teleportation at the logical level, as described by Knill [50]. Taking the realization of the original teleportation circuit for two logical qubits encoded in a Carbon block, which required 3 encoded blocks (left), it is possible to rearrange commuting circuit components to arrive at a circuit that uses a sequence of two 1-bit teleportations [52] to extract syndrome information requiring only 2 encoded blocks at any given time (right) [49].

QUANTINUUM

MSFT x Quantinuum Announcement

- We demonstrated beyor
both the Steane code ar
the Carbon Code [[12,2,

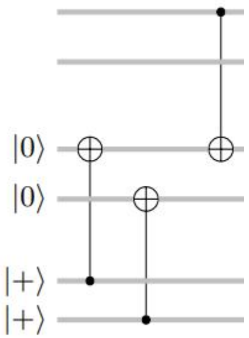

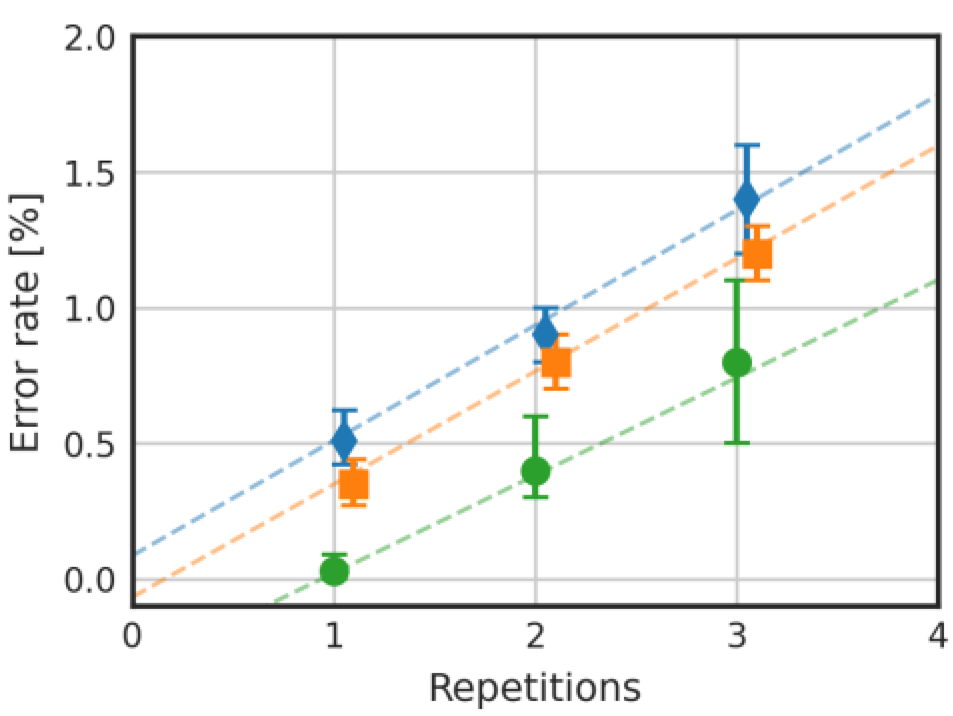

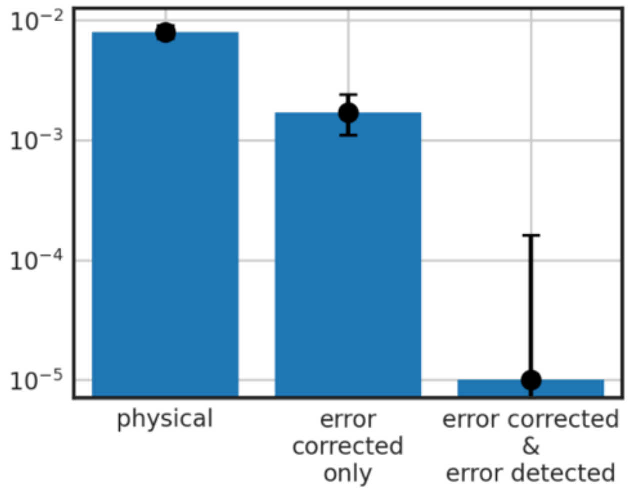

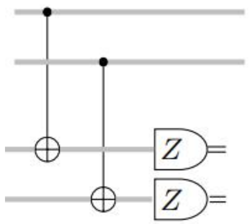

FIG. 5. Syndrome info
Taking the realization
encoded blocks (left), it
1-bit teleportations [52

escribed by Knill [50].
lock, which required 3
uses a sequence of two
ime (right) [49].

FIG. 7. Observed error rate for circuits with 1 to 3 rounds of error correction with the [[12, 2, 4]] Carbon code (green circles) and physical baselines (blue diamond for pairs of 1-bit teleportations, and orange squares for pairs of CNOTs). Results are offset along the x-axis for clarity. Linear fits are obtained by maximum-likelihood estimation (see Appendix A for details).

TRY H1/H2!

- **Quantum Computing User Program** from Oak Ridge National Laboratory

  https://www.olcf.ornl.gov/olcf-resources/compute-systems/quantum-computing-user-program/

- **Azure Quantum Credits Program** from Microsoft

  https://aka.ms/aq/credits

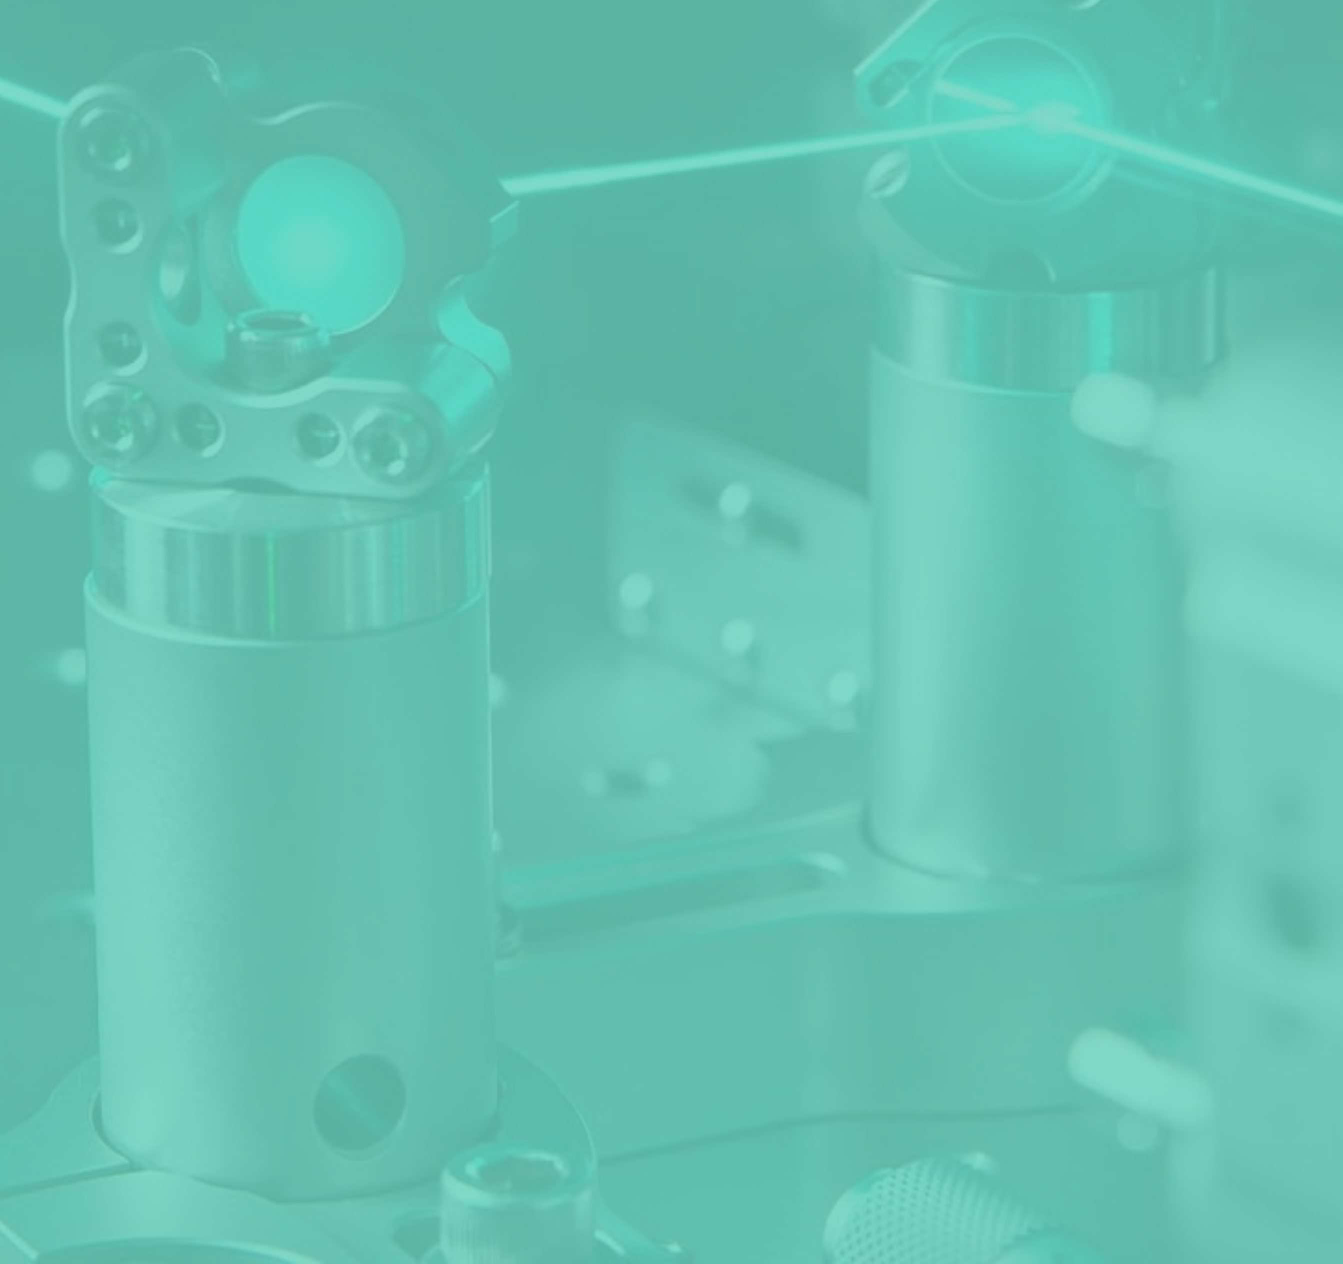

QUANTINUUM

We deliver full-stack value

INQUANTO™

Next generation of molecular and materials discovery

Algorithm Libraries

Quantum Machine Learning
Quantum Monte Carlo Integration
Quantum Natural Language Processing

Third party software

Enables other partners to leverage the power of quantum

QUANTINUUM NEXUS

Quantum workflow orchestration platform

TKET

High-performance quantum SDK | Open-source

Quantum Error Correction: Quantinuum and partners

H-SERIES

Powered by Honeywell

The world's highest-performing quantum hardware

Other quantum computers

Thank you

QUANTINUUM