Quantum Computing and Fundamental Physics

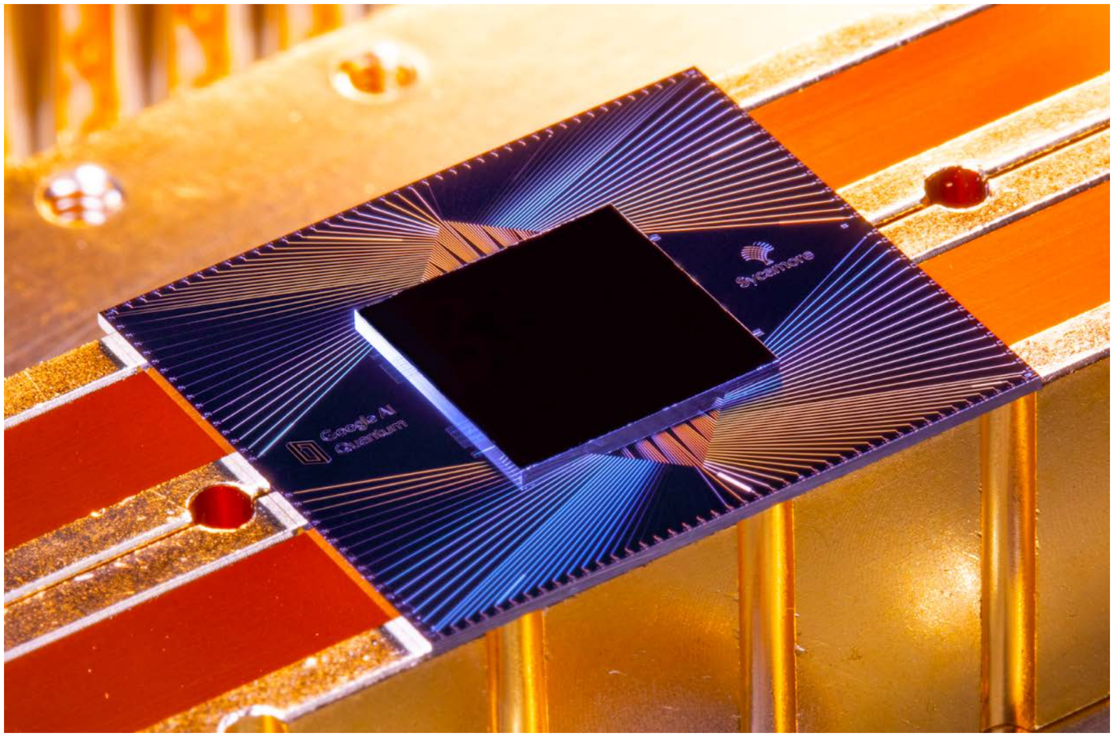

Credit: Erik Lucero/Google

IQIM
INSTITUTE FOR QUANTUM INFORMATION AND MATTER

John Preskill, Caltech

QuantHEP Seminar

7 October 2020

CALIFORNIA INSTITUTE OF TECHNOLOGY 1891

This talk has three parts

(1) *Near-term prospects for quantum computing.*

(2) *Opportunities in quantum simulation of quantum field theory.*

(3) *Recent and ongoing work on quantum and classical algorithms for simulating quantum field theory.*

Collaborators: Stephen Jordan, Keith Lee, Hari Krovi
arXiv: 1111.3633, 1112.4833, 1404.7115, 1703.00454, 1811.10085

Work in progress with:
Junyu Liu, Ashley Milsted, Burak Şahinoğlu, Guifre Vidal

Opportunities in quantum simulation of quantum field theories

Exascale classical computers will advance our knowledge of quantum chromodynamics (QCD), but some challenges will remain, especially concerning real-time evolution (e.g. scattering) and properties of nuclear matter and quark-gluon plasma as a function of both temperature and chemical potential.

Classical computers may never be able to address these (and other) problems; quantum computers will solve them eventually, though we're not sure when. The big physics payoff may still be far away, but today's research can hasten the arrival of a new era in which quantum simulation fuels progress in fundamental physics. Even in the near term, studies of dynamics in strongly-coupled many-particle systems can provide revealing insights.

Frontiers of Physics

short distance	long distance	complexity
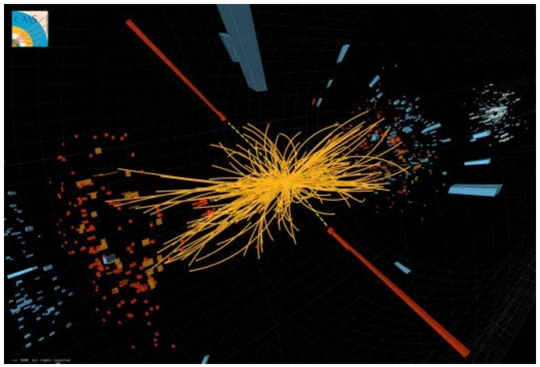		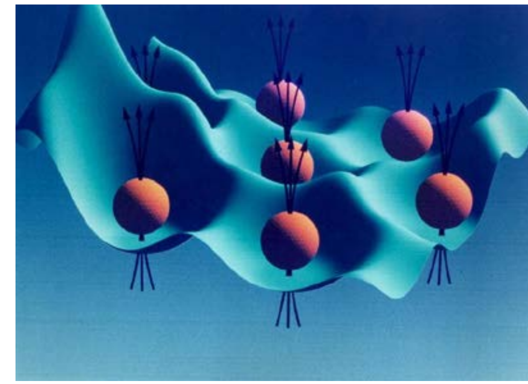
Higgs boson	Large scale structure	“More is different”
Neutrino masses	Cosmic microwave background	Many-body entanglement
Supersymmetry	Dark matter	Phases of quantum matter
Quantum gravity	Dark energy	Quantum computing
String theory	Gravitational waves	Quantum spacetime

Two fundamental ideas

(1) *Quantum complexity*

Why we think quantum computing is powerful.

(2) *Quantum error correction*

Why we think quantum computing is scalable.

Quantum entanglement

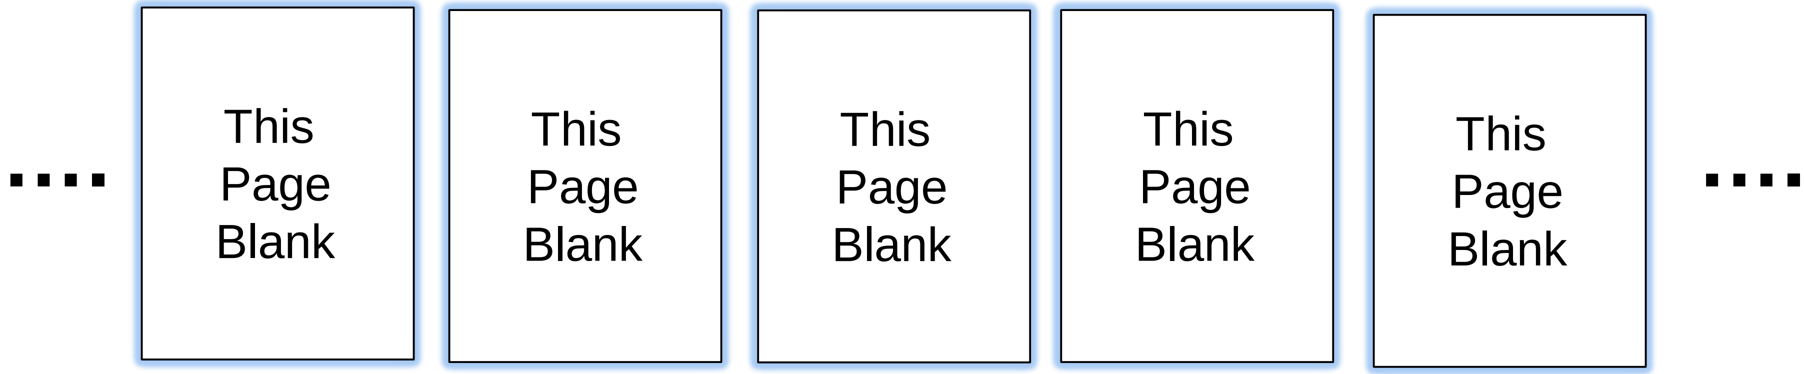

Nearly all the information in a typical entangled "quantum book" is encoded in the correlations among the "pages".

You can't access the information if you read the book one page at a time.

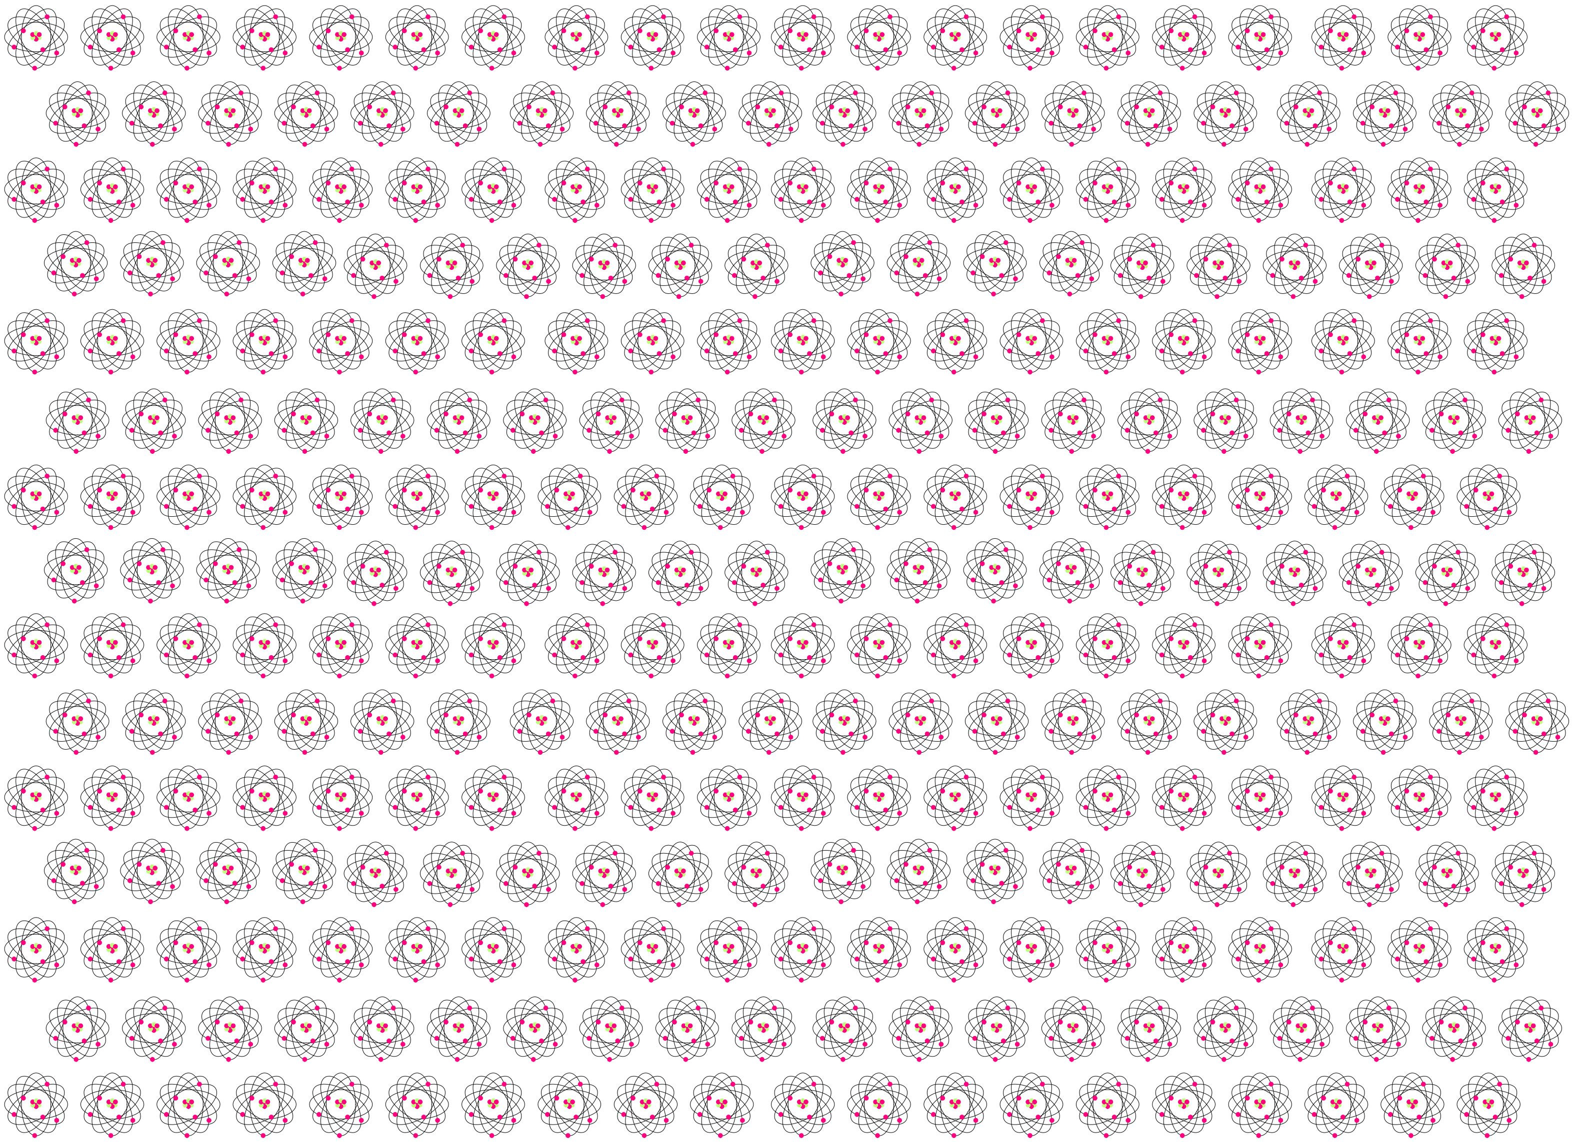

A complete description of a typical quantum state of just 300 qubits requires more bits than the number of atoms in the visible universe.

Why we think quantum computing is powerful

(1) Problems believed to be hard classically, which are easy for quantum computers. Factoring is the best known example.

(2) Complexity theory arguments indicating that quantum computers are hard to simulate classically.

(3) We don't know how to simulate a quantum computer efficiently using a digital ("classical") computer. The cost of the best known simulation algorithm rises exponentially with the number of qubits.

But … the power of quantum computing is limited. For example, we don't believe that quantum computers can efficiently solve worst-case instances of NP-hard optimization problems (e.g., the traveling salesman problem).

Problems

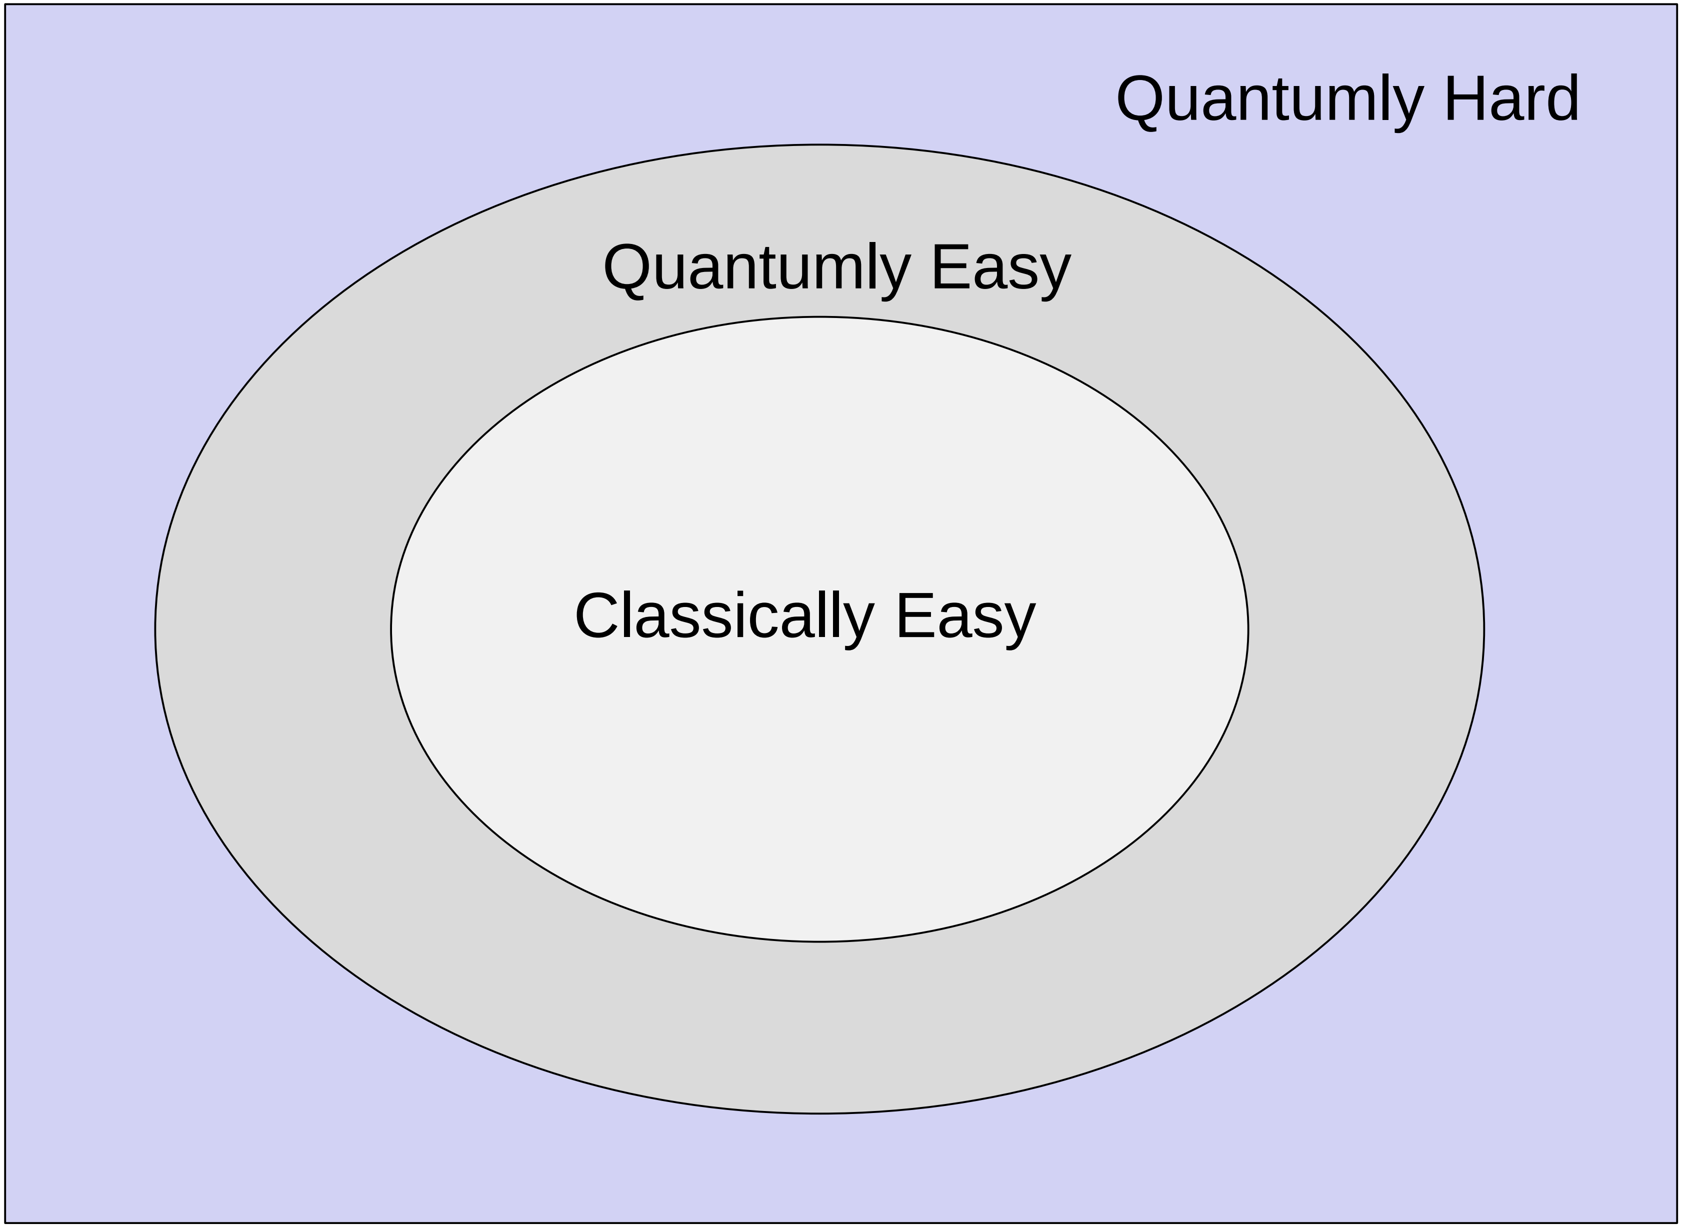

Problems

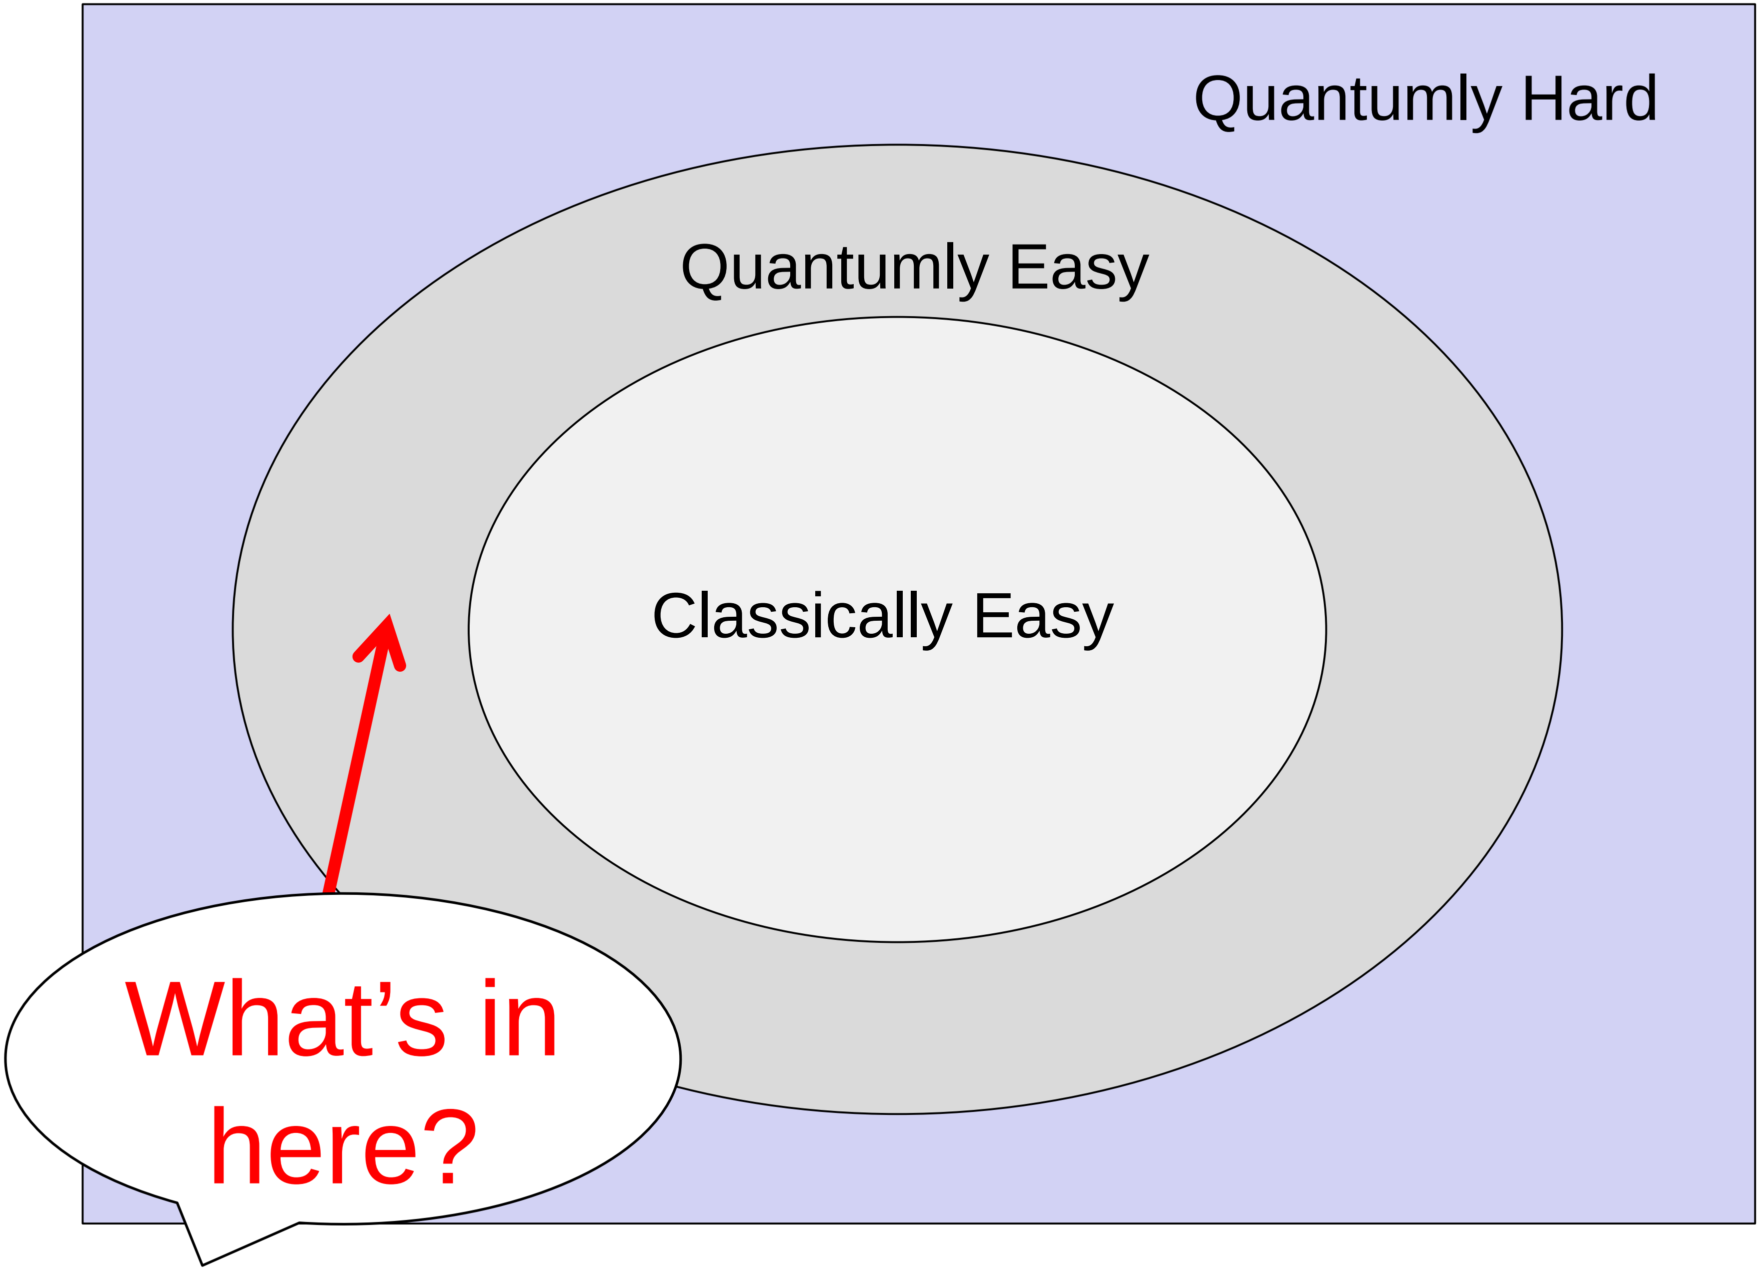
Quantumly Hard
Quantumly Easy
Classically Easy
What's in here?

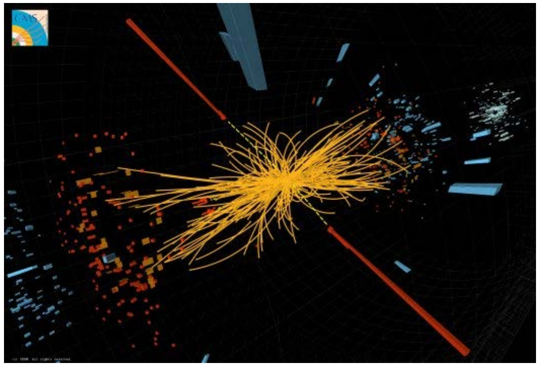

particle collision

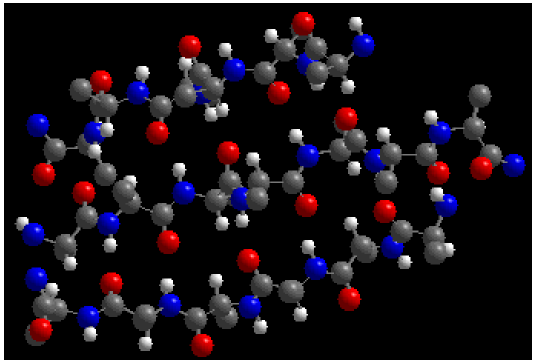

molecular chemistry

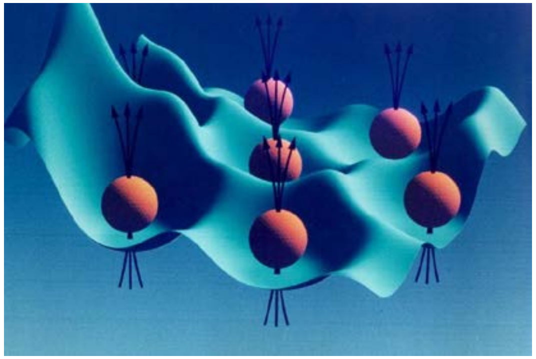

entangled electrons

A quantum computer can simulate efficiently any physical process that occurs in Nature.

(Maybe. We don’t actually know for sure.)

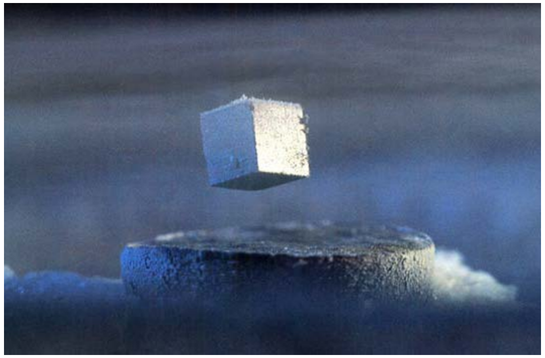

superconductor

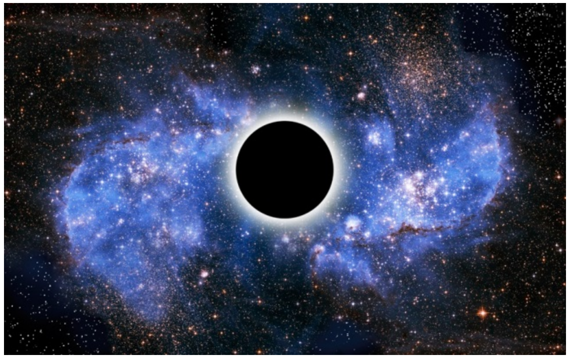

black hole

early universe

Why quantum computing is hard

We want qubits to interact strongly with one another.

We don't want qubits to interact with the environment.

Except when we control or measure them.

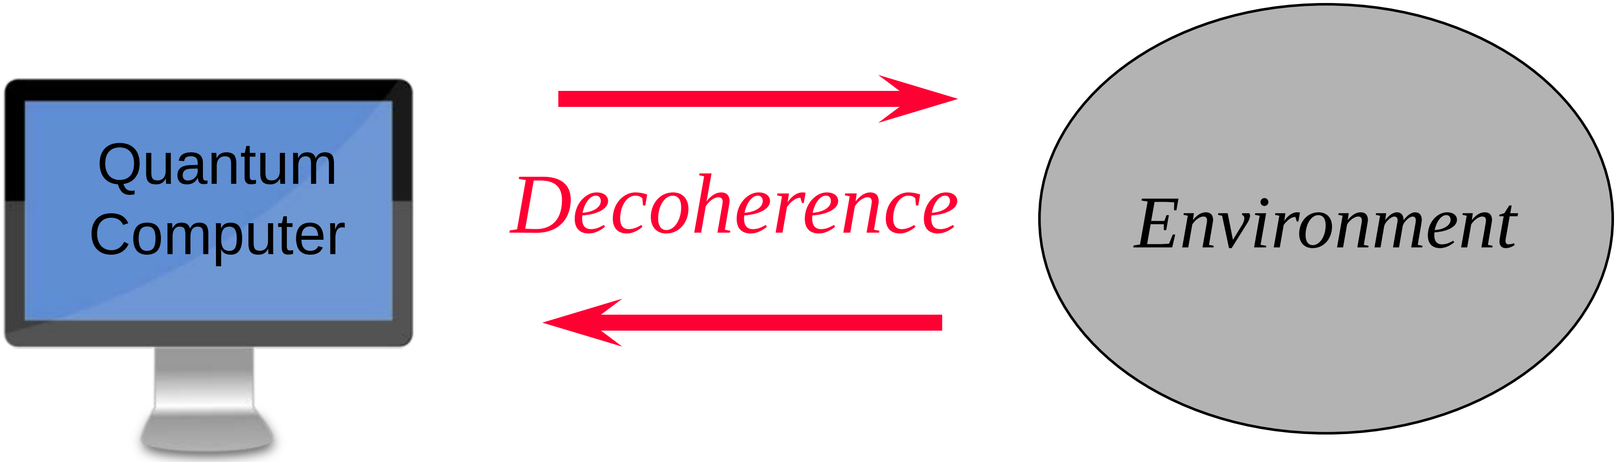

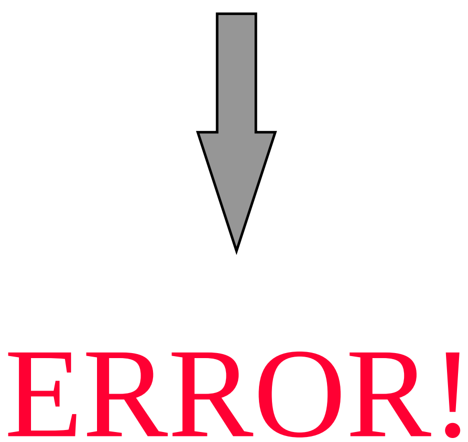

How can we protect a quantum computer from decoherence and other sources of error?

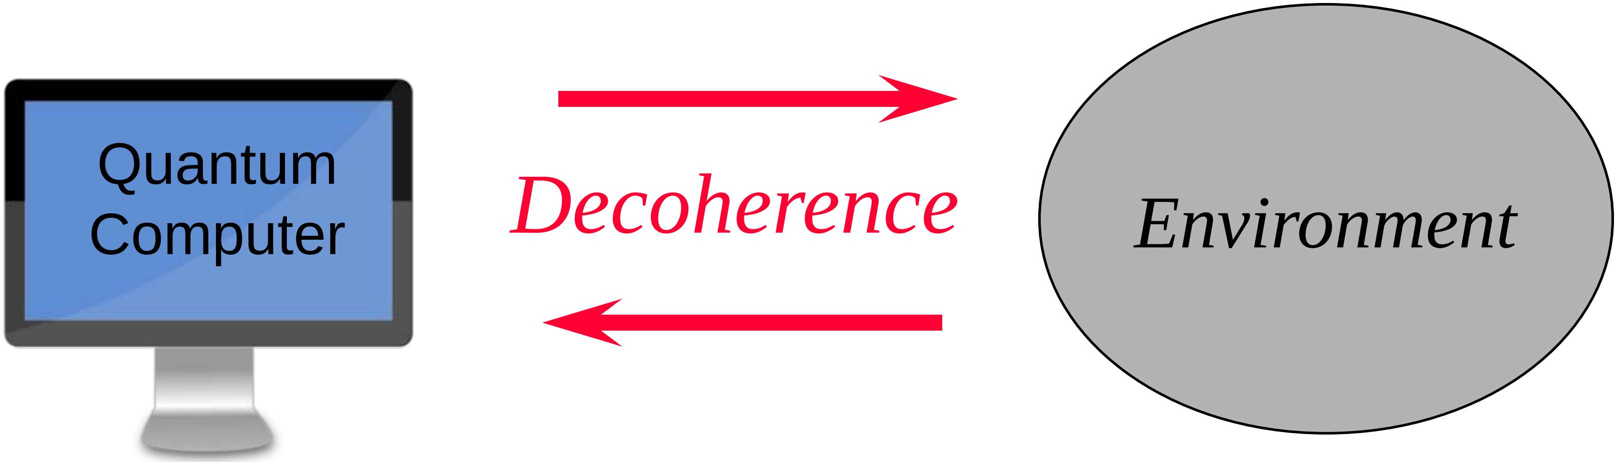

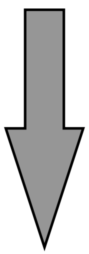

ERROR!

To resist decoherence, we must prevent the environment from "learning" about the state of the quantum computer during the computation.

Quantum error correction

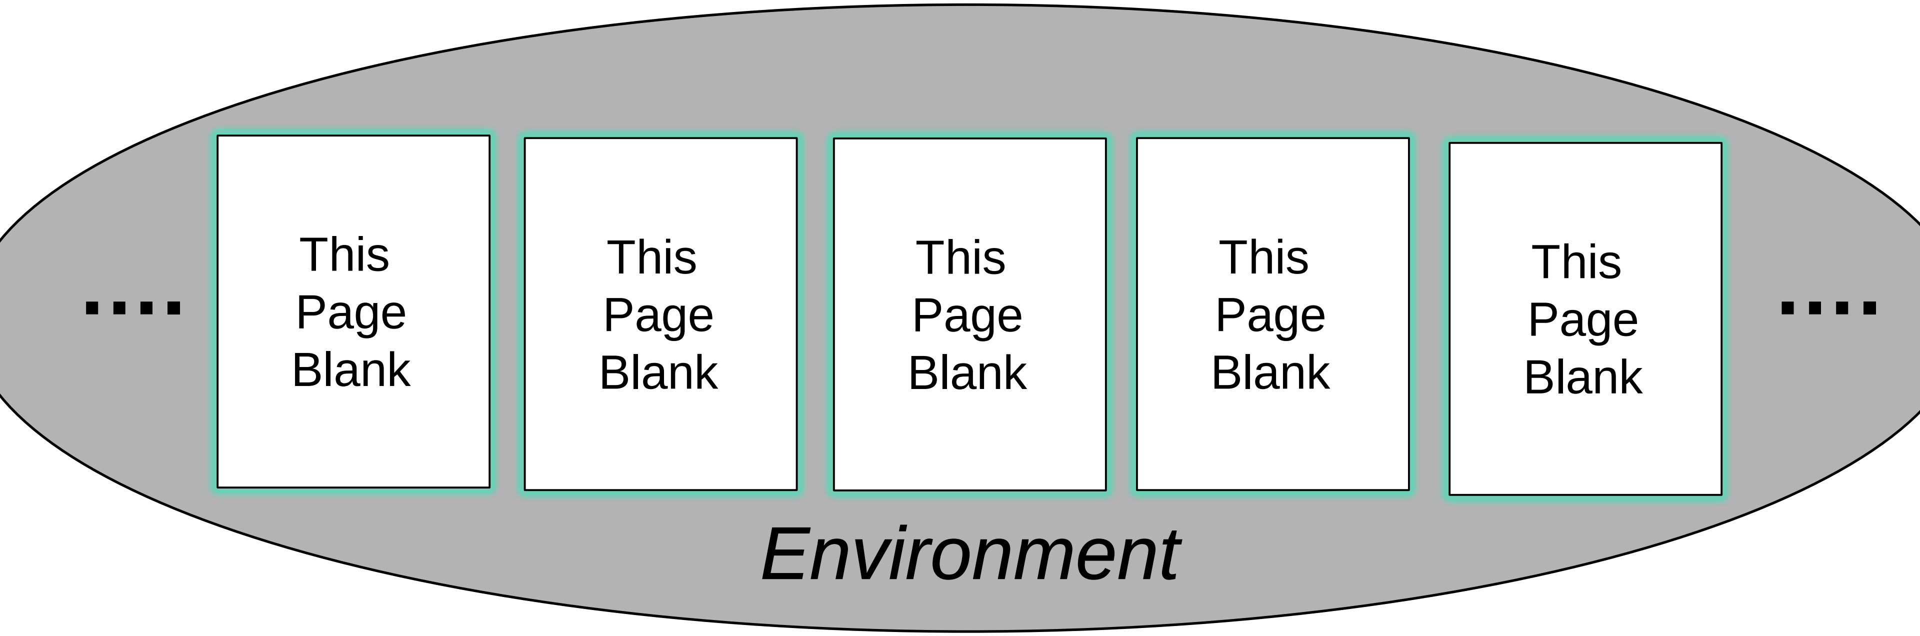

The protected "logical" quantum information is encoded in a highly entangled state of many physical qubits.

The environment can't access this information if it interacts locally with the protected system.

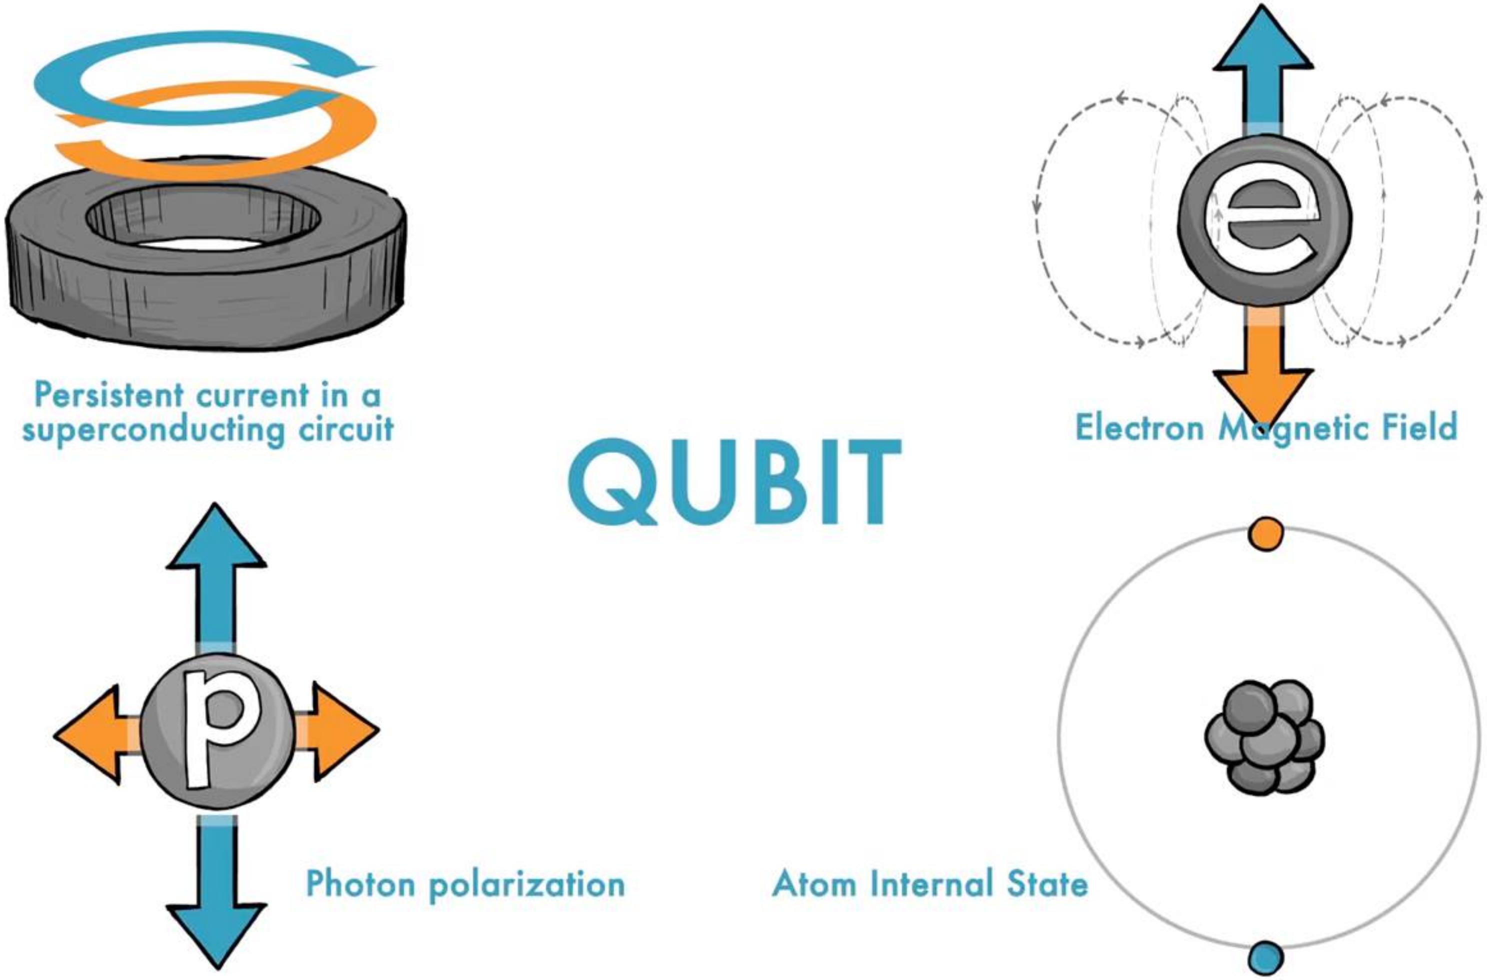
QUBIT
Persistent current in a superconducting circuit
Electron Magnetic Field
Photon polarization
Atom Internal State

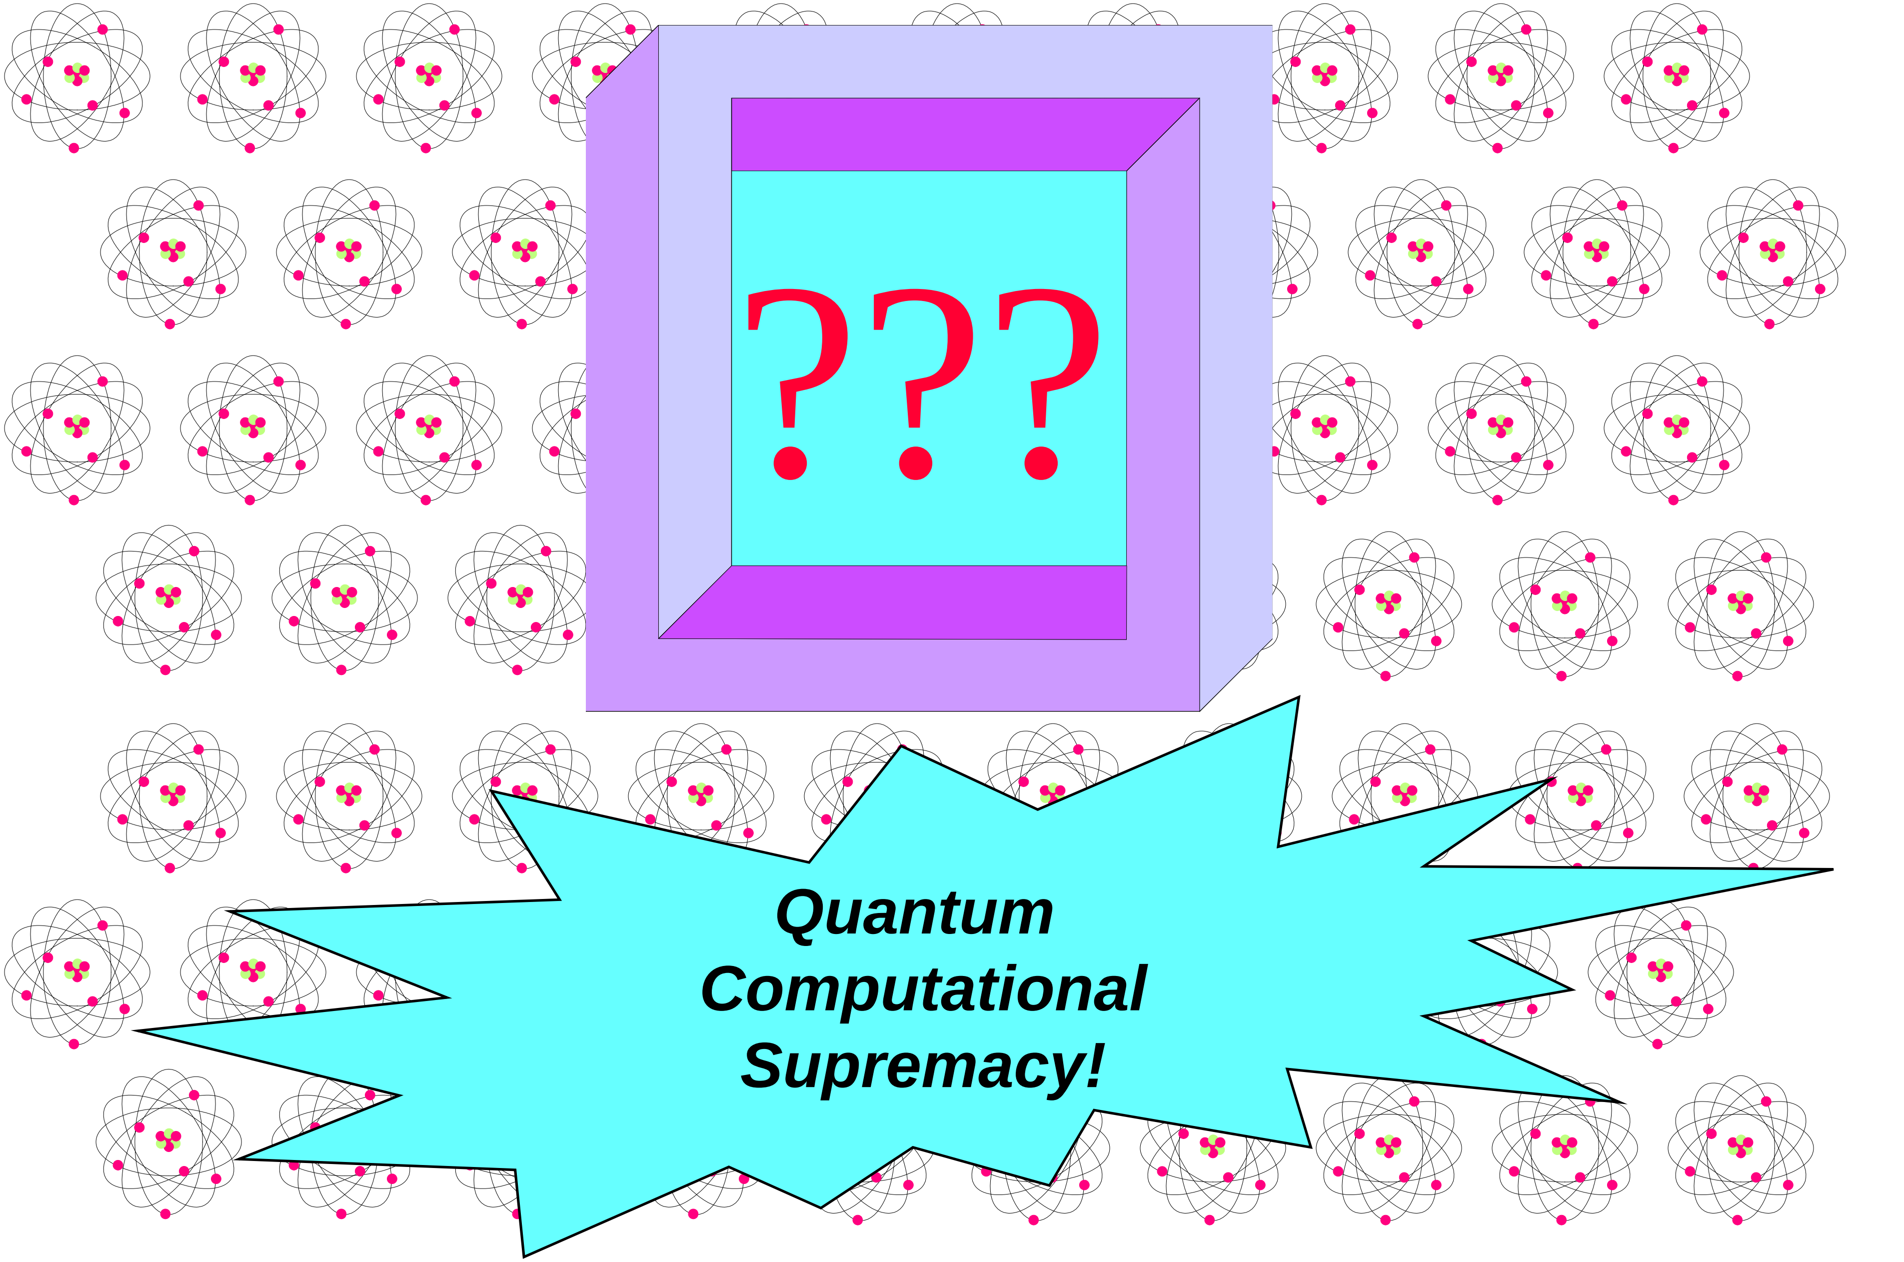
???
Quantum
Computational
Supremacy!

Nature **574**, pages 505–510 (2019), 23 October 2019

Article

Quantum supremacy using a programmable superconducting processor

https://doi.org/10.1038/s41586-019-1666-5

Received: 22 July 2019

Accepted: 20 September 2019

Published online: 23 October 2019

Frank Arute[1], Kunal Arya[1], Ryan Babbush[1], Dave Bacon[1], Joseph C. Bardin[1,2], Rami Barends[1], Rupak Biswas[3], Sergio Boixo[1], Fernando G. S. L. Brandao[1,4], David A. Buell[1], Brian Burkett[1], Yu Chen[1], Zijun Chen[1], Ben Chiaro[5], Roberto Collins[1], William Courtney[1], Andrew Dunsworth[1], Edward Farhi[1], Brooks Foxen[1,5], Austin Fowler[1], Craig Gidney[1], Marissa Giustina[1], Rob Graff[1], Keith Guerin[1], Steve Habegger[1], Matthew P. Harrigan[1], Michael J. Hartmann[1,6], Alan Ho[1], Markus Hoffmann[1], Trent Huang[1], Travis S. Humble[7], Sergei V. Isakov[1], Evan Jeffrey[1],

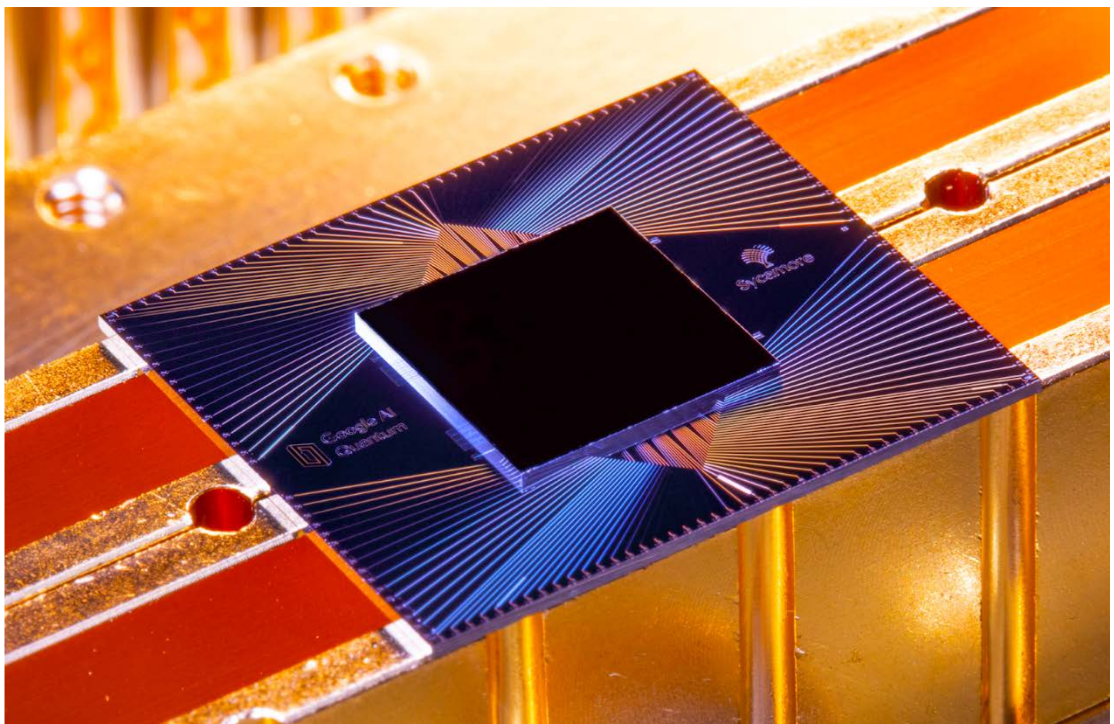

Credit: Erik Lucero/Google

Classical systems cannot simulate quantum systems efficiently (a widely believed but unproven conjecture).

Arguably the most interesting thing we know about the difference between quantum and classical.

About Sycamore

“Quantum David vs. Classical Goliath”

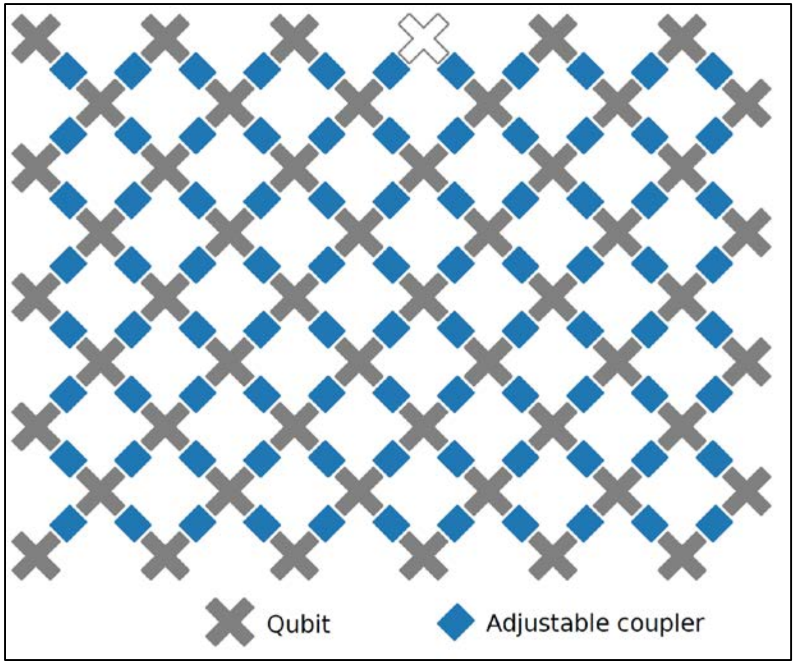

A fully programmable circuit-based quantum computer. *n*= 53 working qubits in a two-dimensional array with coupling of nearest neighbors.

A circuit with 20 layers of 2-qubit gates *can be executed millions of times in a few minutes*, yielding verifiable results.

Simulating this quantum circuit using a classical supercomputer is hard. *It would take at least days*, possibly much longer.

Furthermore, the cost of the classical simulation grows exponentially with the number of qubits.

Conclusion: the hardware is working well enough to produce meaningful results in a regime where classical simulation is very difficult.

What quantum computational supremacy means

"Quantum David vs. Classical Goliath"

It's a programmable circuit-based quantum computer.

An impressive achievement in experimental physics and a testament to ongoing progress in building quantum computing hardware.

We have arguably entered the regime where the extravagant exponential resources of the quantum world can be validated.

This confirmation does not surprise (most) physicists, but it's a milestone for technology on planet earth.

Building a quantum computer is merely *really, really hard,* not *ridiculously hard*. The hardware is working; we can begin a serious search for useful applications.

But the specific task performed by Sycamore to demonstrate quantum computational supremacy is not particularly useful.

Quantum computing in the NISQ Era

The (noisy) 50-100 qubit quantum computer has arrived.
(*NISQ* = noisy intermediate-scale quantum.)

NISQ devices *cannot be simulated* by brute force using the most powerful currently existing supercomputers.

Noise limits the computational power of NISQ-era technology.

NISQ will be an interesting tool for exploring physics. It *might* also have other useful applications. But we're not sure about that.

NISQ will not change the world by itself. Rather it is a step toward more powerful quantum technologies of the future.

Potentially transformative scalable quantum computers may still be decades away. *We're not sure how long it will take.*

Quantum 2, 79 (2018), arXiv:1801.00862

Hybrid quantum/classical optimizers

Eddie Farhi: "Try it and see if it works!"

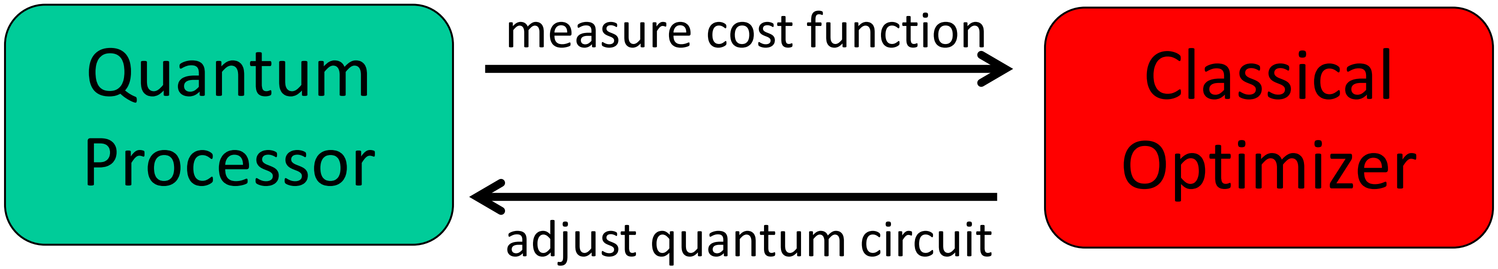

We don't expect a quantum computer to solve worst case instances of NP-hard problems, but it *might* find better approximate solutions, or find them faster.

Classical optimization algorithms (for both classical and quantum problems) are sophisticated and well-honed after decades of hard work.

We don't know whether NISQ devices can do better, but we can try it and see how well it works.

The era of quantum heuristics

Peter Shor: "You don't need them [testbeds] to be big enough to solve useful problems, just big enough to tell whether you can solve useful problems."

Sometimes algorithms are effective in practice even though theorists are not able to validate their performance in advance.

Example: Deep learning. Mostly tinkering so far, without much theory input.

Possible quantum examples:
Quantum annealers, approximate optimizers, variational eigensolvers, quantum machine learning … playing around may give us new ideas.

What can we do with, say, < 100 qubits, depth < 100? We need a dialog between quantum algorithm experts and application users.

Maybe we'll get lucky …

The steep climb to scalability

NISQ-era quantum devices will not be protected by quantum error correction. Noise will limit the scale of computations that can be executed accurately.

Quantum error correction (QEC) will be essential for solving some hard problems. But QEC carries a high overhead cost in number of qubits & gates.

This cost depends on both the hardware quality and algorithm complexity. With today's hardware, solving (say) useful chemistry problems may require hundreds to thousands of physical qubits for each protected logical qubit.

Recent estimate: 20 million physical qubits to break RSA 2048 (*Gidney, Ekerå 2019*), for gate error rate 10^{-3}.

To reach scalability, we must cross the daunting "quantum chasm" from hundreds to millions of physical qubits. This may take a while.

Advances in quantum gate fidelity, systems engineering, algorithm design, and error correction protocols can hasten the arrival of the fully fault-tolerant quantum computer.

Quantum simulation

Classical computers are especially bad at simulating quantum *dynamics* --- predicting how highly entangled quantum states change with time. Quantum computers will have a big advantage in this arena. Physicists hope for noteworthy advances in quantum dynamics during the NISQ era.

For example: Classical *chaos theory* advanced rapidly with onset of numerical simulation of classical dynamical systems in the 1960s and 1970s. Quantum simulation experiments may advance the theory of *quantum* chaos. Simulations with ~ 100 qubits could be revealing, if not too noisy.

Near-term quantum simulators can be either *digital* (circuit based) or *analog* (tunable Hamiltonians).

Digital provides more flexible Hamiltonian and initial state preparation. We can use hybrid quantum/classical methods. But gate based simulations of time evolution are expensive.

Experience with near-term digital simulators will lay foundations for fault-tolerant simulations in the future (applies to NISQ more broadly).

Quantum simulation of quantum field theories.

Beyond Euclidean Monte Carlo on classical computers

-- Improved predictions for QCD backgrounds in collider experiments

-- Equation of state for nuclear matter, quark gluon plasma, early universe

-- Electroweak response of hadronic matter, e.g. intensity frontier

-- Simulation of nuclear reactions, e.g. astrophysical modeling

-- Exploration of other strongly-coupled theories, beyond-standard-model physics

-- Stepping stone to quantum gravity, e.g. through holographic duality

-- New insights!

What quantum simulators can do

-- Sample accurately from outgoing states in simulation of scattering event.

-- Real-time correlation functions, including at nonzero temp and chem potential.

-- Transport properties, far from equilibrium phenomena.

Prototypical quantum simulation task

(1) State preparation. E.g., incoming scattering state.

(2) Hamiltonian evolution. E.g. Trotter approximation.

(3) Measure an observable. E.g., a simulated detector.

Goal: sample accurately from probability distribution of outcomes.

Determine how computational resources scale with: error, system size, particle number, total energy of process, energy gap, …

Resources include: number of qubits, number of gates, …

Hope for polynomial scaling! Or even better: polylog scaling.

Need an efficient preparation of initial state.

Approximating a continuous system incurs discretization cost (smaller lattice spacing improves accuracy).

What should we simulate, and what do we stand to learn?

Entanglement in high-energy scattering

Two incoming high-energy particles, many soft outgoing particles. Crudely model the outgoing particles as a thermal ensemble with temperature $T = O(1)$.

The overall state is pure – the thermodynamic entropy of left-movers and right-movers is really entanglement entropy.

If the particles are emitted from the interaction region in a time $t \sim L$, they occupy a region of size L. Entropy S, particle number N, energy E, are related by

$$S \sim N \sim TL \sim E / T$$

The bond dimension is *exponential in the entropy*, therefore a classical simulation of the scattering would be very difficult for about 20 particles.

We could measure time-dependence of the outgoing particle flux, particle-particle correlations, etc.

Broken symmetry phase

$$H = \int dx \left(\frac{1}{2}\Pi^2 + \frac{1}{2}(\nabla\phi)^2 + U(\phi) \right)$$

$$U(\phi) = \frac{\lambda}{8}(\phi^2 - v^2)^2$$

weak coupling (large v^2):

$$\langle \text{vac} | \phi | \text{vac} \rangle = \pm v$$

$$m_{\text{scalar}}^2 = \lambda v^2, \quad \frac{m_{\text{kink}}}{m_{\text{scalar}}} = \frac{2}{3} v^2 \gg 1$$

strong coupling (small v^2):

$$\frac{m_{\text{kink}}}{m_{\text{scalar}}} \approx \frac{1}{2}$$

kink-antikink scattering: nonperturbative, and a toy model for colliding bubble walls in the early universe.

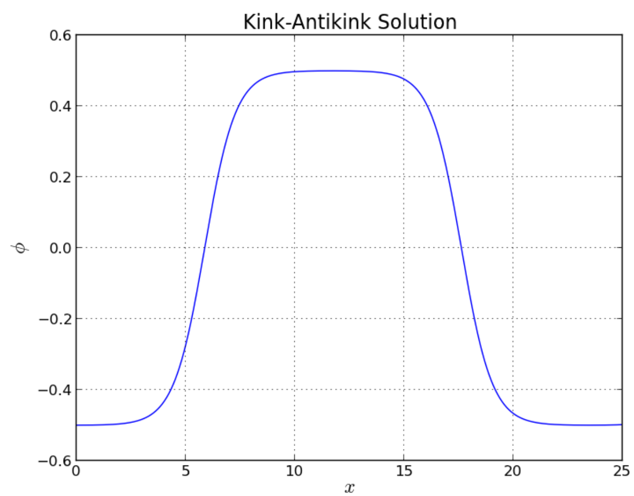

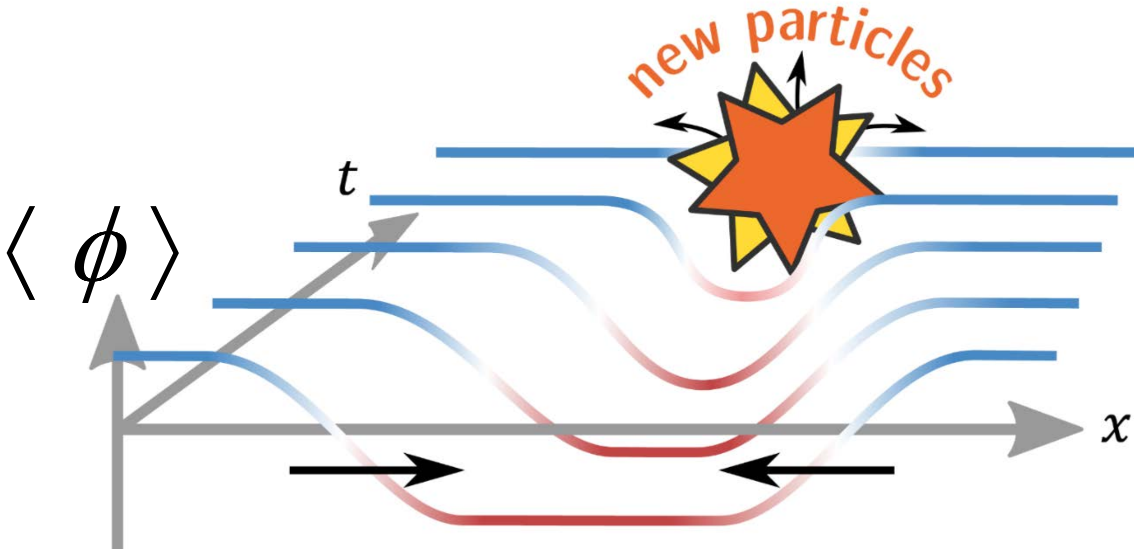

Lots of entropy generated, so hard to simulate classically.

A case where adiabatic state preparation is difficult. And the excitations are nonperturbative (cf. proton in large-*N* QCD). A baby version of what we'll need in QCD.

And it's not obvious what happens at strong coupling.

1) Brute force simulation of the field theory (as in our earlier work). But using a hybrid quantum/classical method (classical computer guides the state prep).

2) Emergent field theory in a spin system. More heuristic, but more likely to be feasible using near-term platforms. And interesting in its own right.

Junyu Liu, Burak Şahinoğlu, John Preskill

From matrix-product state to quantum circuit

MPS:

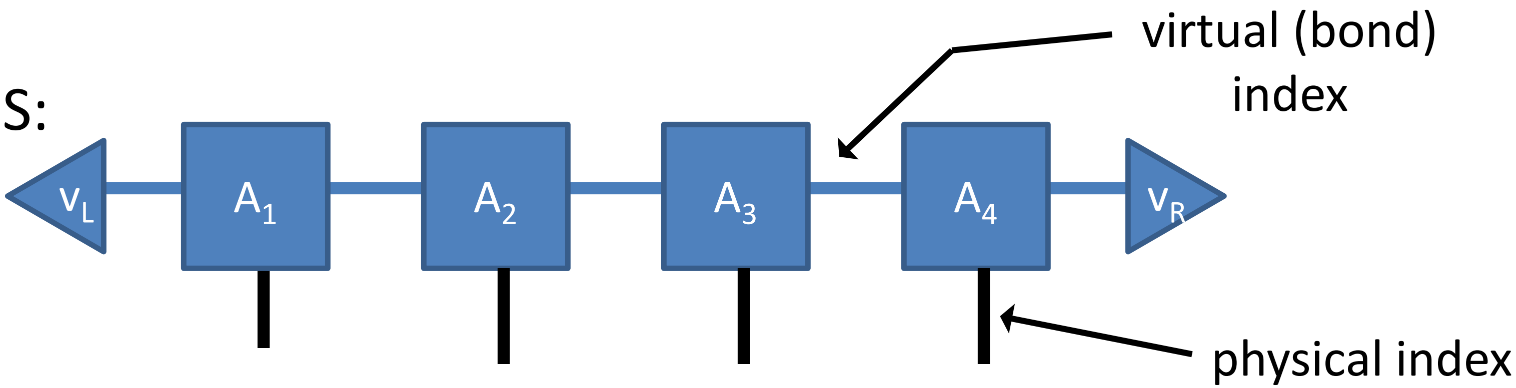

Circuit:

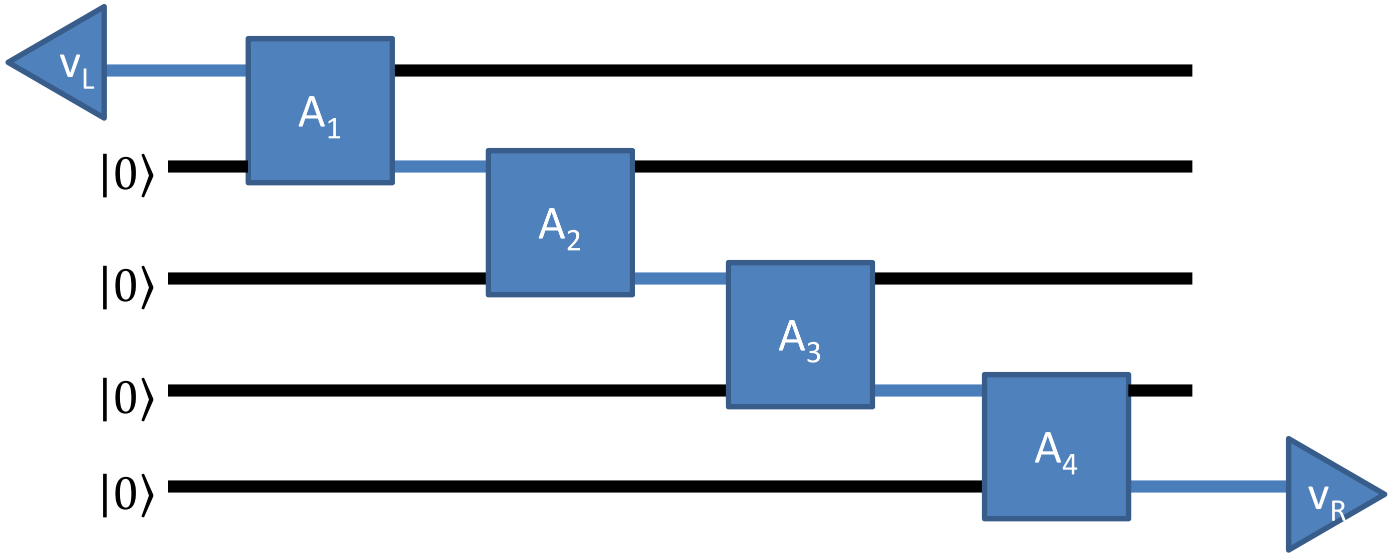

Constructing vacua, kinks, and kink-antikink pairs

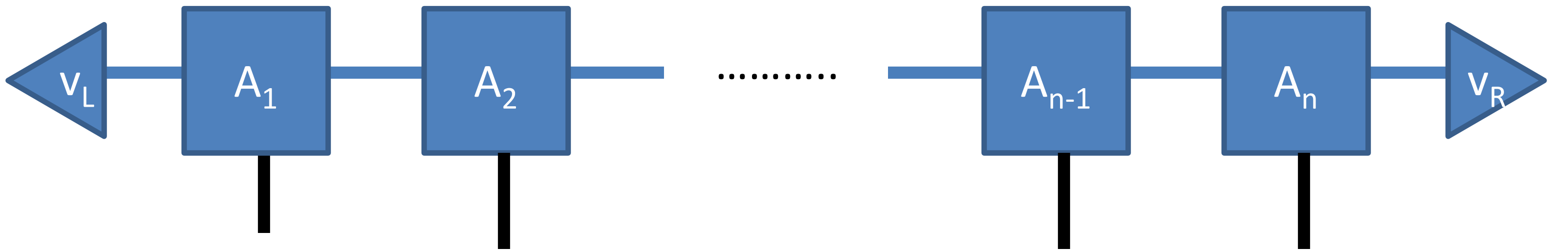

Vacuum:

$$|\mathrm{vac},+\rangle = v_L^\dagger A_+ A_+ \cdots A_+ A_+ v_R, \quad |\mathrm{vac},-\rangle = v_L^\dagger A_- A_- \cdots A_- A_- v_R,$$

Zero-momentum kink:

$$|\mathrm{kink}, p=0\rangle = \sum_x v_L^\dagger A_- A_- \cdots B(x) \cdots A_+ A_+ v_R$$

Distantly separated propagating kink-antikink wave packets:

$$|\mathrm{kink}, f; \mathrm{antikink}, g\rangle = \sum_{x,y} f(x) g(y) v_L^\dagger A_- \cdots A_- B(x) A_+ \cdots A_+ \overline{B}(y) A_- \cdots A_- v_R.$$

Alternative: construct separated static kink and antikink, then adiabatically break the symmetry to accelerate them.

Efficient classical algorithm for constructing the initial state

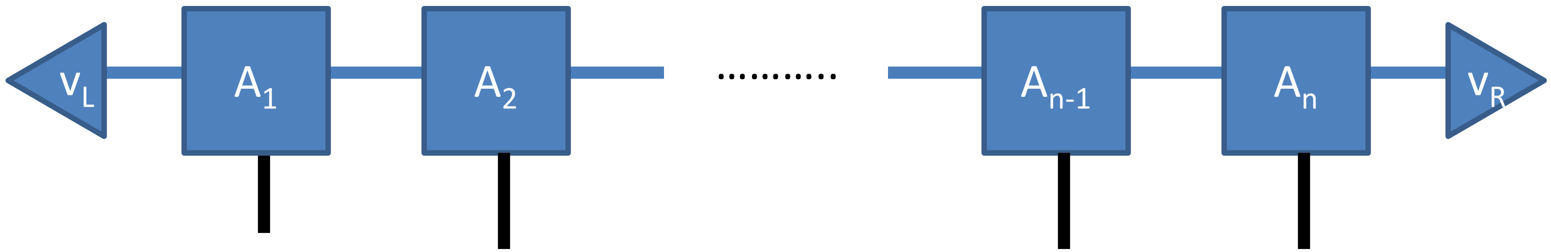

It helps a lot that the theory is in one spatial dimension and is superrenormalizable! (We can bound field fluctuations.)

A rigorous version of DMRG can find an MPS approximation to the ground state in poly(n) time. Landau, Vazirani, Vidick 2015

The "rigorous renormalization group" finds the kink state in quasi-poly(n) time. Arad, Landau, Vazirani, Vidick 2017

Applying a filtering matrix-product-operator to the "bare" kink-antikink state finds an MPS approximation to initial state with $n^{O(\log \log n)}$ bond dimension.

Though formally "efficient" (runtime exp[polylog(n)]), the classical algorithm might not be practical.

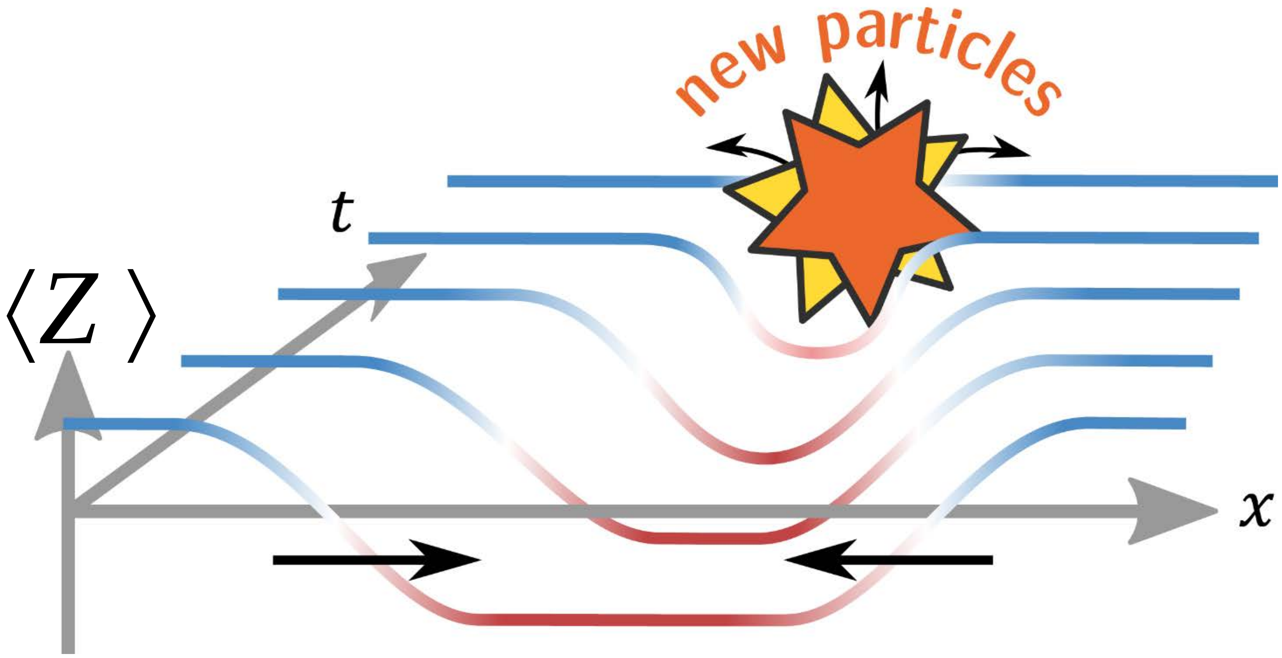

Explicit symmetry breaking (false vacuum)

$$H = \sum_{j=1}^{N} \left[-Z_j Z_{j+1} - gX_j - hZ_j + \lambda \left(X_j Z_{j+1} Z_{j+2} + Z_j Z_{j+1} X_{j+2} \right) \right]$$

We want a spin chain with these features:

-- Emergent Lorentz-invariant field theory
-- Spontaneously broken discrete symmetry
-- Not integrable (or too close to integrable)

Ashley Milsted, Junyu Liu, John Preskill, Guifre Vidal

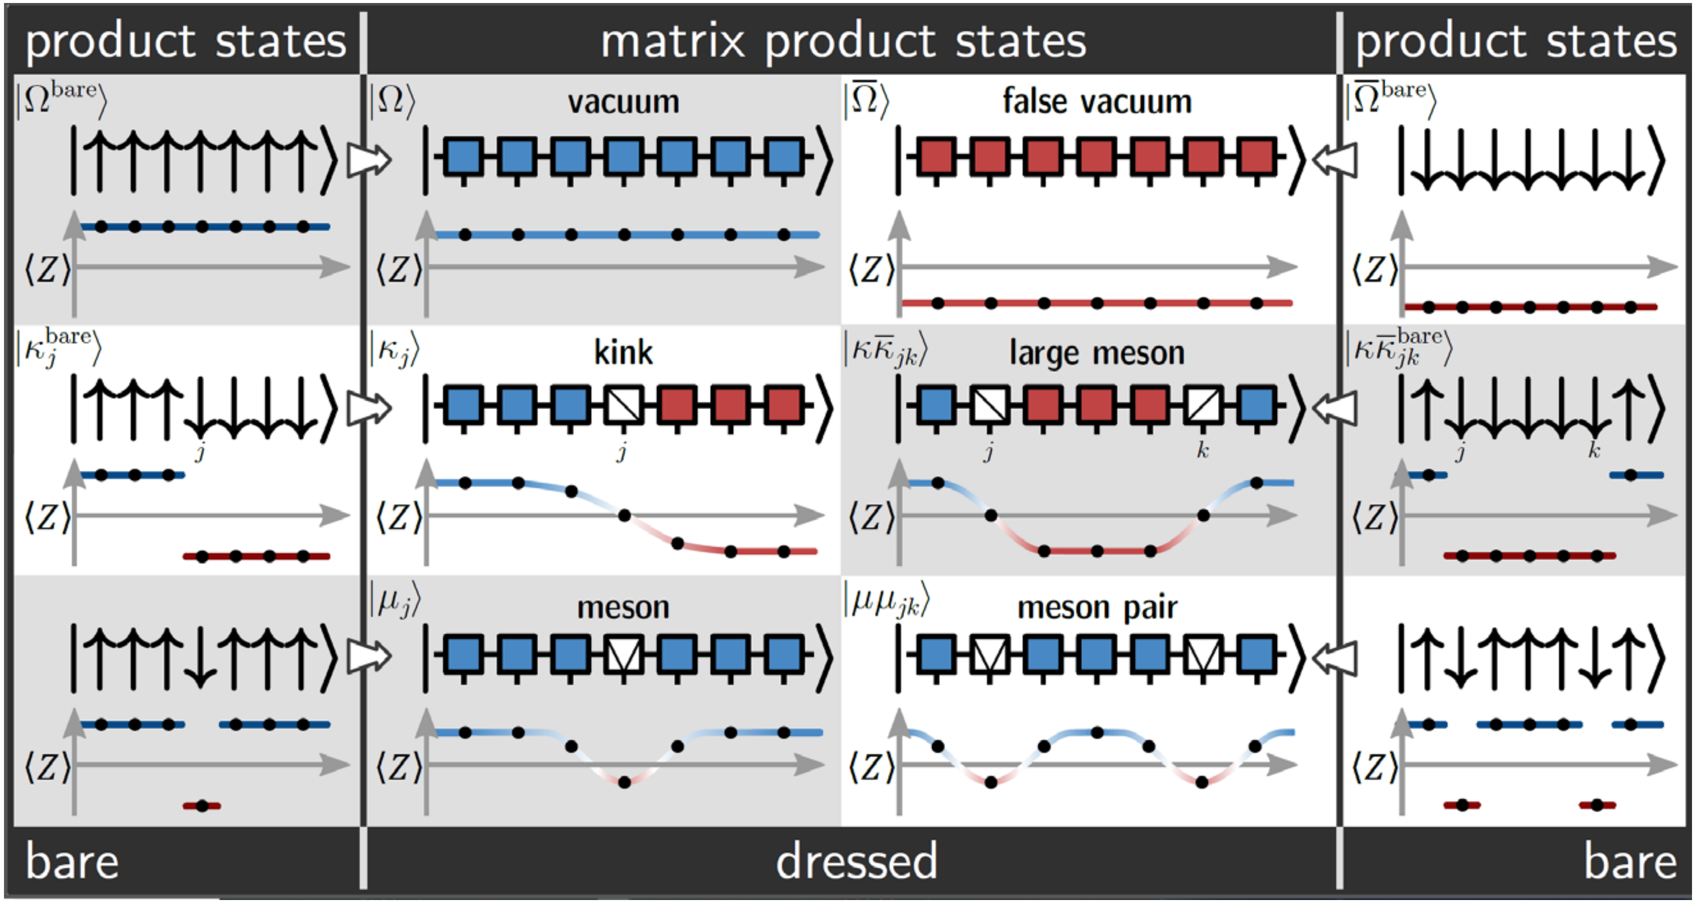

Break the symmetry slightly – there are true and false vacuum states.
Find approximations to the true and false vacuum.
Find MPS approximations to the kink and antikink momentum eigenstates.
Prepare kink-antikink wave packets, with false vacuum in between.
Kink and antikink accelerate toward one another and collide.
Update the MPS as the state evolves.
Measure energy density, spin expectation value, entanglement entropy, etc.
Cost is O(D^3), where D is bond dimension (in our simulations $D \leq 128$).

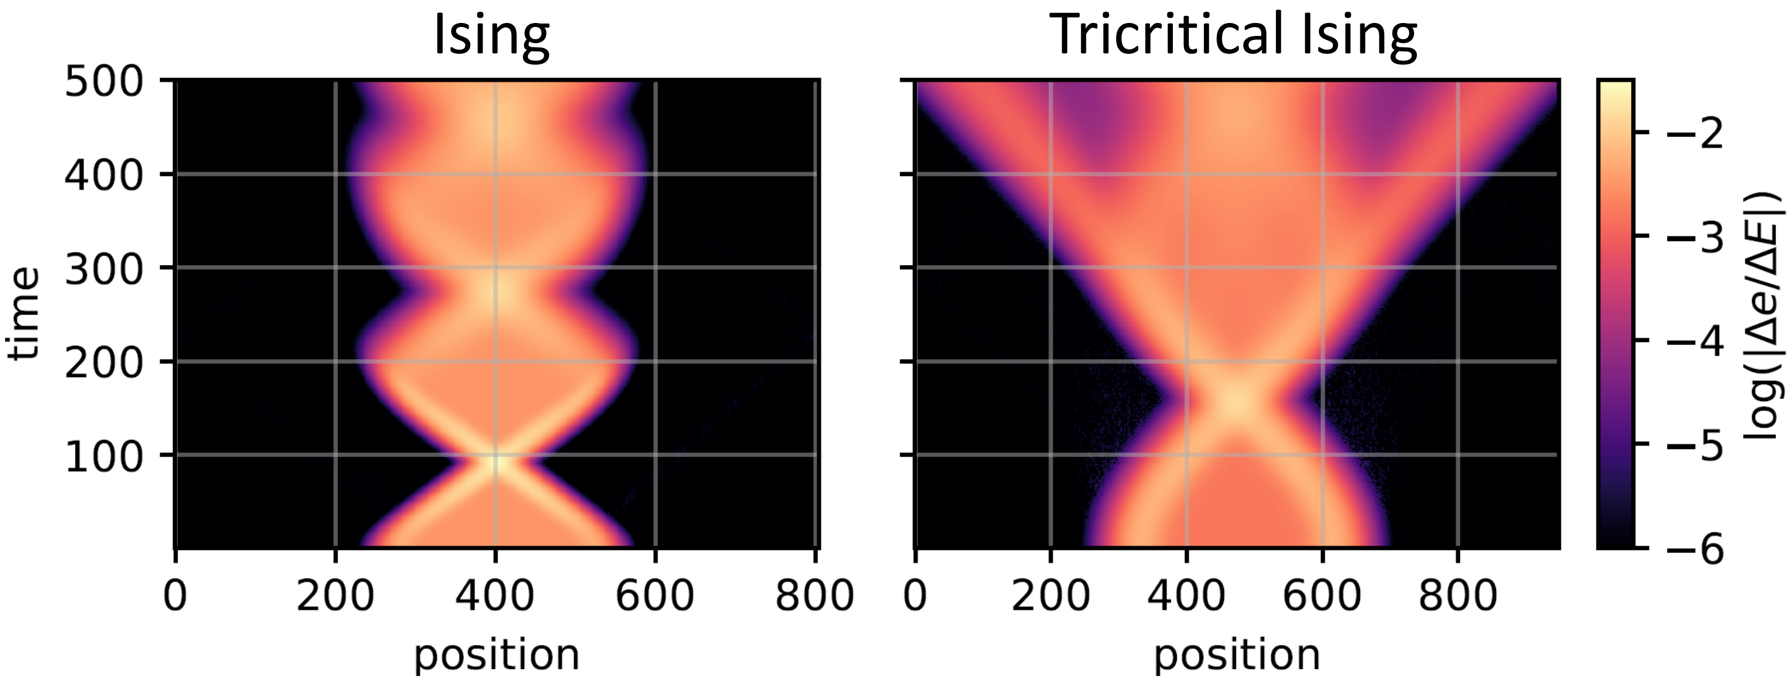

Here excess energy density (relative to the true vacuum) is plotted as a function of position and time.

A kink-antikink pair with false vacuum in between is confined. They collide repeatedly.

The IR limit of the Ising model with $Z2$ symmetry broken ($\lambda = 0$) is an integrable theory. Kinks “bounce” --- no new particles are produced.

The tricritical Ising model ($\lambda \neq 0$) is a nonintegrable field theory. Unconfined “mesons” are produced which propagate ballistically.

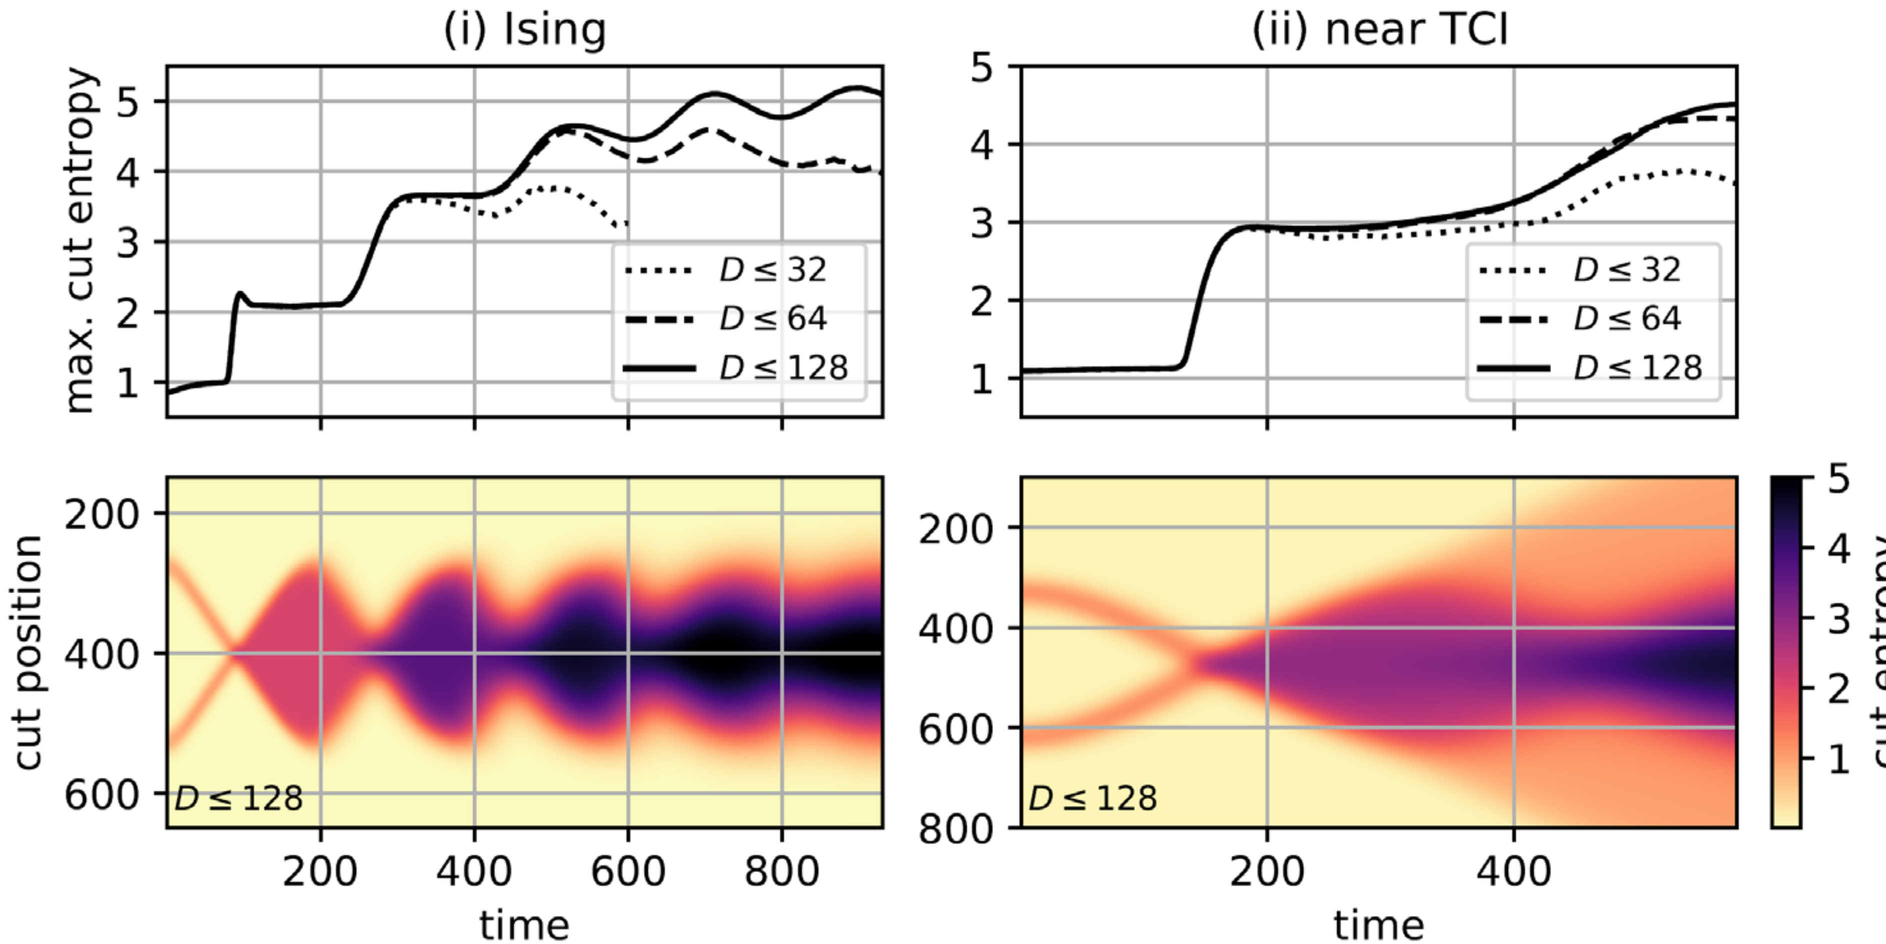

Here entanglement entropy across a cut is plotted as a function of time.

Entropy can increase due to either elastic scattering of wave packets (momentum dependence of scattering phase shift), or due to inelastic particle production.

As entropy increases, a larger bond dimension *D* is needed to approximate the state accurately.

Quantum simulation of quantum field theories.

Where are we now?

-- Resource scaling estimates (number of qubits and gates) for scattering simulations in scalar and Yukawa theories.
-- *Classical* tensor-network simulation of massive 1D QED.
Static and dynamic studies of strings and string breaking.
-- Few-site *quantum* simulations of 1D QED with trapped ions and superconducting circuits.
-- Binding energies of deuteron, ^{3}He, ^{4}He in (pionless) effective field theory.
-- Proposals for analog simulation using ultracold atoms, etc.
-- In progress: Classical and quantum simulations of nonabelian gauge symmetry, higher dimensions.

Where to seek quantum advantage?

-- How to outperform classical tensor network calculations?
-- Classical simulation methods fail for highly entangled states.
-- High-energy scattering with multiple particle production.
-- Dynamics after a quench, or many successive scattering events.

Quantum simulation of quantum field theories: What next?

-- More qubits, better precision, greater programmability

-- Access to a variety of platforms, for exploration and benchmarking

-- Stepping stones toward QFT simulators, for both analog and digital approaches

-- Hybrid quantum / classical methods (focusing quantum resources where they are most needed)

-- Protocols for state preparation, evolution, readout, classical post-processing

-- Hamiltonian simulation theory: gauge invariance, errors, renormalization, scaling

-- Clarify the hardware / software requirements for a special-purpose QFT/QCD quantum machine

-- Exploit quantum advantage in sampling, matrix inversion, semidefinite programs

-- Elucidate the path forward, both near term and long term