What's Next After Quantum Supremacy?

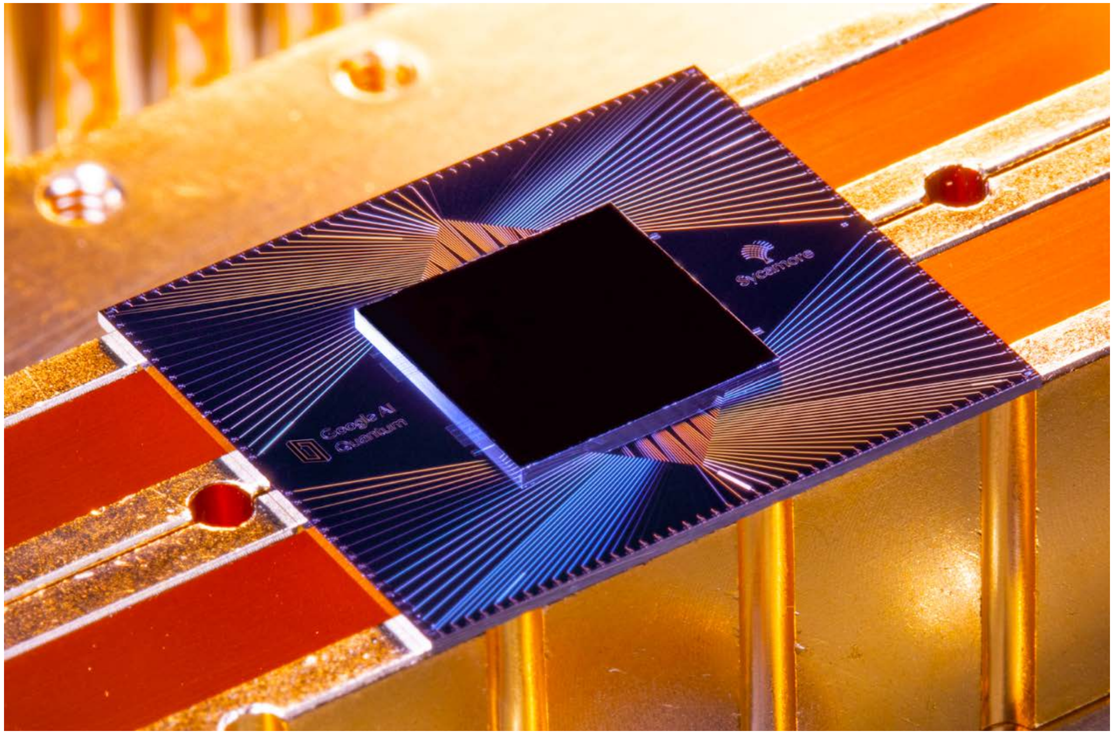

Credit: Erik Lucero/Google

INSTITUTE FOR QUANTUM INFORMATION AND MATTER

John Preskill

Q2B 2019

11 December 2019

CALIFORNIA INSTITUTE OF TECHNOLOGY 1891

Quantum Computing in the NISQ era and beyond

John Preskill

Institute for Quantum Information and Matter and Walter Burke Institute for Theoretical Physics, California Institute of Technology, Pasadena CA 91125, USA

30 July 2018

Noisy Intermediate-Scale Quantum (NISQ) technology will be available in the near future. Quantum computers with 50-100 qubits may be able to perform tasks which surpass the capabilities of today's classical digital computers, but noise in quantum gates will limit the size of quantum circuits that can be executed reliably. NISQ devices will be useful tools for exploring many-body quantum physics, and may have other useful applications, but the 100-qubit quantum computer will not change the world right away — we should regard it as a significant step toward the more powerful quantum technologies of the future. Quantum technologists should continue to strive for more accurate quantum gates and, eventually, fully fault-tolerant quantum computing.

Quantum 2, 79 (2018), arXiv:1801.00862

Based on a Keynote Address delivered at Quantum Computing for Business, 5 December 2017

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], Zhang Jiang[1], Dvir Kafri[1], Kostyantyn Kechedzhi[1], Julian Kelly[1], Paul V. Klimov[1], Sergey Knysh[1], Alexander Korotkov[1,8], Fedor Kostritsa[1], David Landhuis[1], Mike Lindmark[1], Erik Lucero[1], Dmitry Lyakh[9], Salvatore Mandrà[3,10], Jarrod R. McClean[1], Matthew McEwen[5], Anthony Megrant[1], Xiao Mi[1], Kristel Michielsen[11,12], Masoud Mohseni[1], Josh Mutus[1], Ofer Naaman[1], Matthew Neeley[1], Charles Neill[1], Murphy Yuezhen Niu[1], Eric Ostby[1], Andre Petukhov[1], John C. Platt[1], Chris Quintana[1], Eleanor G. Rieffel[3], Pedram Roushan[1], Nicholas C. Rubin[1], Daniel Sank[1], Kevin J. Satzinger[1], Vadim Smelyanskiy[1], Kevin J. Sung[1,13], Matthew D. Trevithick[1], Amit Vainsencher[1], Benjamin Villalonga[1,14], Theodore White[1], Z. Jamie Yao[1], Ping Yeh[1], Adam Zalcman[1], Hartmut Neven[1] & John M. Martinis[1,5]*

The promise of quantum computers is that certain computational tasks might be executed exponentially faster on a quantum processor than on a classical processor. A fundamental challenge is to build a high-fidelity processor capable of running quantum algorithms in an exponentially large computational space. Here we report the use of a processor with programmable superconducting qubits to create quantum states on 53 qubits, corresponding to a computational state-space of dimension 253 (about 1016). Measurements from repeated experiments sample the resulting probability distribution, which we verify using classical simulations. Our Sycamore processor takes about 200 seconds to sample one instance of a quantum circuit a million times—our benchmarks currently indicate that the equivalent task for a state-of-the-art classical supercomputer would take approximately 10,000 years. This dramatic increase in speed compared to all known classical algorithms is an experimental realization of quantum supremacy for this specific computational task, heralding a much anticipated computing paradigm.

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.

Andrew Yang 🧢
@AndrewYang

Google achieving quantum computing is a huge deal. It means, among many other things, that no code is uncrackable.

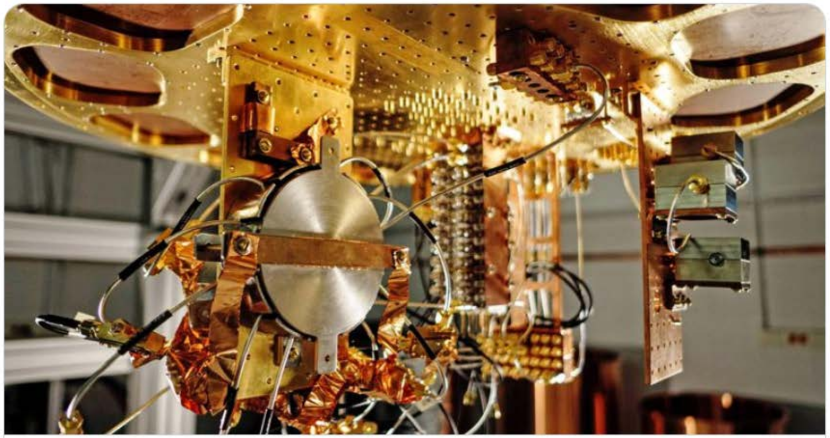

Google reportedly attains 'quantum supremacy'
Its quantum computer can solve tasks that are otherwise unsolvable, a report says.
cnet.com

5:11 PM · Sep 20, 2019 · Twitter for iPhone

1.2K Retweets **5.5K** Likes

Andrew Yang 🧢 @AndrewYang · Sep 20
Replying to @AndrewYang
We need to catch up with our approach to encryption

Ivanka Trump
@IvankaTrump

It's official! 💥 The US has achieved quantum supremacy!
In a collaboration between the Trump Admin, @Google and UC Santa Barbara, quantum computer Sycamore has completed a calculation in 3 min 20 sec that would take about 10,000 years for a classical comp.

5:47 AM · Oct 23, 2019 · Twitter for iPhone

2.1K Retweets **8.9K** Likes

Ivanka Trump @IvankaTrump · Oct 23
Replying to @IvankaTrump
#QIS is a critical industry of the future. That's why @POTUS signed the National Quantum Initiative Act into law, supporting robust quantum R&D. We're proud to have contributed to this major milestone, ushering in the next gen of quantum tech in the USA! 🇺🇸

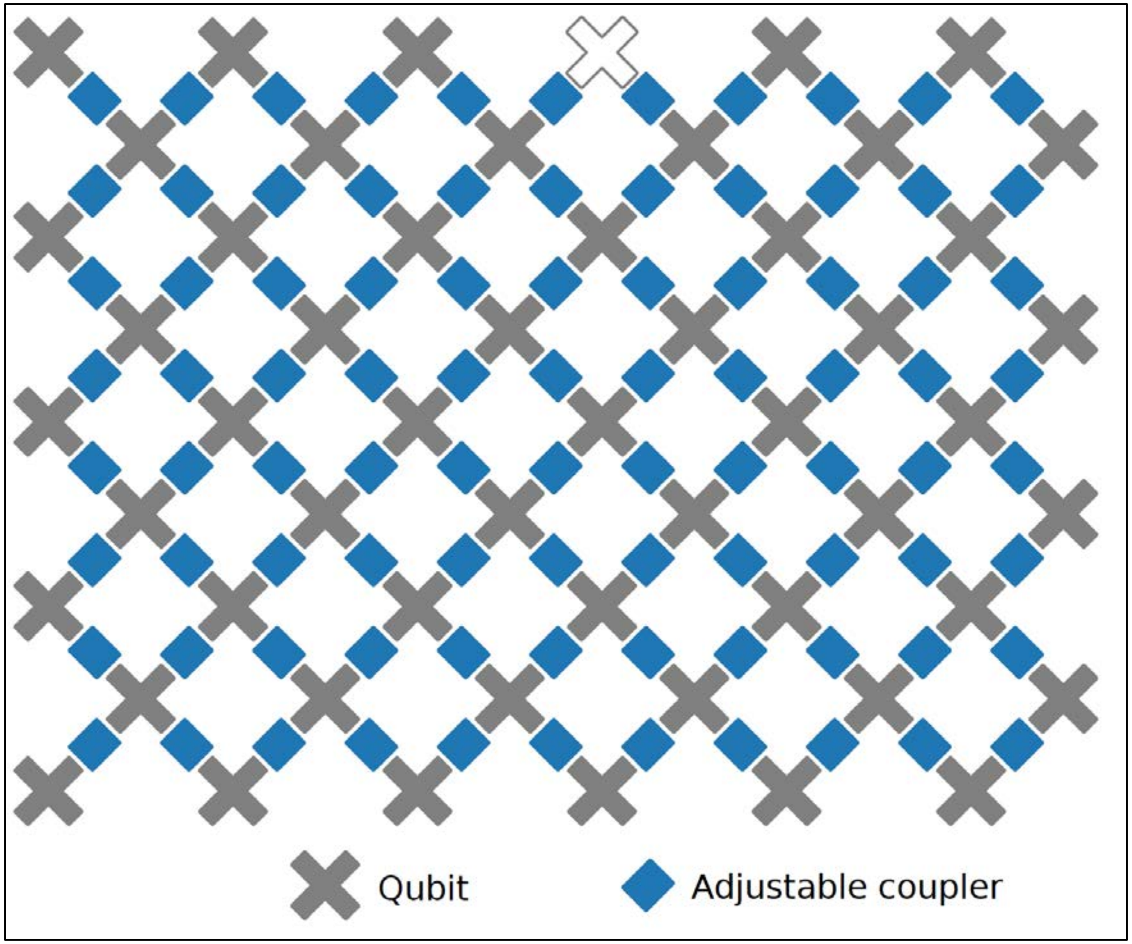

Each qubit is also connected to its neighboring qubits using a new adjustable coupler [31, 32]. Our coupler design allows us to quickly tune the qubit-qubit coupling from completely off to 40 MHz. Since one qubit did not function properly the device uses 53 qubits and 86 couplers.

[31] Martinis group, UCSB, 2014 (inductor coupled).
[32] Oliver group, MIT Lincoln Laboratory, 2018 (capacitor coupled)

About Sycamore

Greg Kuperberg: "Quantum David vs. Classical Goliath"

A fully programmable circuit-based quantum computer. n= 53 working qubits in a 2D array with coupling of nearest neighbors.

Entangling 2-qubit gates with error rate .6% (in parallel), executed in 12 ns.

Estimated global circuit fidelity F = .2% for circuit with 20 "cycles" of 2-qubit gates: 430 2-qubit gates and 1113 1-qubit gates.

A circuit with fixed 2-qubit gates and randomly-chosen 1-qubit gates is chosen and executed millions of times; Each time, all qubits are measured, generating a 53-bit string.

The collected sample of 53-bit strings is not uniformly distributed. Comparing with classical simulations one can verify "heavy output generation" --- that the average probability of strings in the sample is greater than 2^{-n}.

Because a random circuit has no structure, and the Hilbert space is exponentially large in n, simulation using a classical supercomputer is hard. (At least days, while the Sycamore generates a large sample in minutes.)

Experiment verifies that 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.

Other takes:
John Martinis and Sergio Boixo on Google QI Blog, 23 October 2019.
Scott Aaronson's "Quantum Supremacy FAQ" on Shtetl Optimized.
Scott's New York Times Op-Ed, 30 October 2019.
My column in Quanta Magazine, 2 October 2019.

What's Next?

Real-world applications.

Dramatically extended qubit lifetimes using quantum error correction.

Significant improvements in 2-qubit gate fidelity.

In various platforms: more qubits, better gates.

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

How to find applications?

Scott Aaronson: "Instead of thinking of a hard problem and asking how to speed it up, ask what quantum computers are good at and build an application from that."

For example, certified randomness. If you receive putatively random numbers from the cloud, how can you be sure they are really uniformly random?

Idea (Aaronson/Google):
Random quantum circuits generate a lot of entropy, which can be distilled using classical methods to obtain very nearly uniform randomness.

Client generates a random quantum circuit, (untrusted) server executes it to generate a sample, returning the result so quickly that only a quantum computer could have done it (if the server is honest). Client does an (exponential time) classical simulation to validate the sample, testing for heavy output generation.

Questions:
Can we have both poly time classical verification and NISQ implementation?
What natural complexity assumptions suffice to ensure security?
More certified randomness by running the same circuit over and over again?

Simulation of quantum dynamics is another application in a similar spirit. What else?

Hybrid quantum/classical optimizers

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

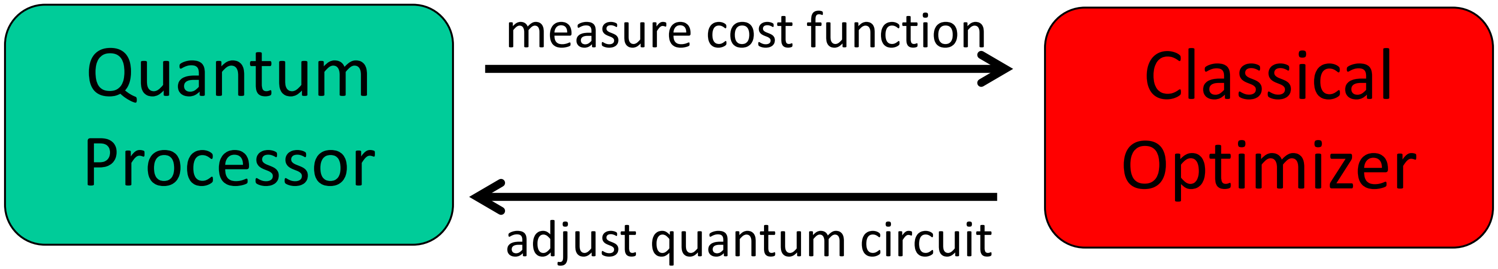

We don't expect a quantum computer to solve worst case instances of NP-hard problems, but it might find better approximate solutions, or find them faster.

Combine quantum evaluation of a cost function with a classical feedback loop for seeking a quantum state with a lower value.

Quantum approximate optimization algorithm (QAOA).
In effect, seek low-energy states of a classical spin glass.

Variational quantum eigensolvers (VQE).
Seek low energy states of a quantum many-body system with a local Hamiltonian.

Classical optimization algorithms (for both classical and quantum problems) are sophisticated and well-honed after decades of hard work. Will NISQ be able to do better?

Hybrid quantum/classical optimizers

Eddie Farhi: “Try it and see if it works!”

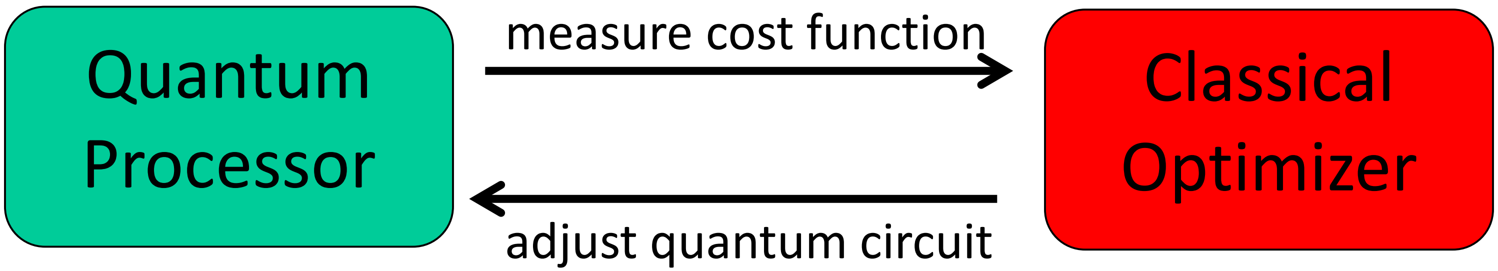

Can we streamline the classical loop? Naively, the number of variational parameters needed scales with the size of the instance. That’s a problem.

Optimized classical parameters for one instance of a problem provide a good starting point for other instances, including larger ones. (*Brandão, Broughton, Farhi, Gutmann, Neven 2018*). Even if we don’t need to adjust the classical parameters for each instance, we still need to run the quantum computer to find an input string that (approximately) solves the problem for that instance.

Symmetries might help reduce the classical load (e.g., translation invariant quantum optimization).

A concern (*Hastings 2019*): Why should quantum bounded-depth approximations be better than classical ones?

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, … 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 …

Quantum machine learning

Jordan Kerenidis: "Overhyped but underestimated"

Perhaps a quantum deep learning network can be trained more efficiently, e.g. using a smaller training set. We don't know. We'll have to try it to see how well it works.

High-dimensional classical data can be encoded very succinctly in a quantum state. In principle log N qubits suffice to represent a N-dimensional vector. Such "quantum Random Access Memory" (= QRAM) *might* have advantages for machine learning applications.

However, many proposed quantum machine learning applications are hampered by input/output bottlenecks.

Loading classical data into QRAM is slow, nullifying the potential advantage, and the output is a quantum state, and only a limited amount of information can be accessed by measuring the state.

Perhaps it's more natural to consider quantum inputs / outputs; e.g. better ways to characterize or control quantum systems. Quantum networks might have advantages for learning about *quantum* correlations, rather than classical ones.

From my talk in 2017:

Applications of quantum linear algebra

Given classical input A (N X N matrix, sparsity s and condition number κ) and N-qubit quantum input $|b\rangle$, algorithm outputs
$|y\rangle = |A^{-1} b\rangle$ with error ε.

It is more promising if the input b is computed rather than entered from a database.

Example: Solving (monochromatic) Maxwell's equations in a complex 3D geometry; e.g., for antenna design (Clader et al. 2013). Discretization and preconditioner needed.

Alternative method for solving classical scattering problems: quantum simulation of N X N Laplacian using O(log N) qubits (Jordan et al. 2017). Need efficient preparation of initial state (e.g. input Gaussian wavepacket).

Recommendation systems (e.g. Netflix/Amazon with $m=10^8$ users and $n=10^6$ products). Sample rapidly from preference matrix with *low-rank* $k \cong 100$ (Kerenidis & Prakash 2016). Quantum queries to a classical data structure: Linear-time offline preprocessing, online processing of quantum queries in time poly(k) polylog(mn).

Dequantization!

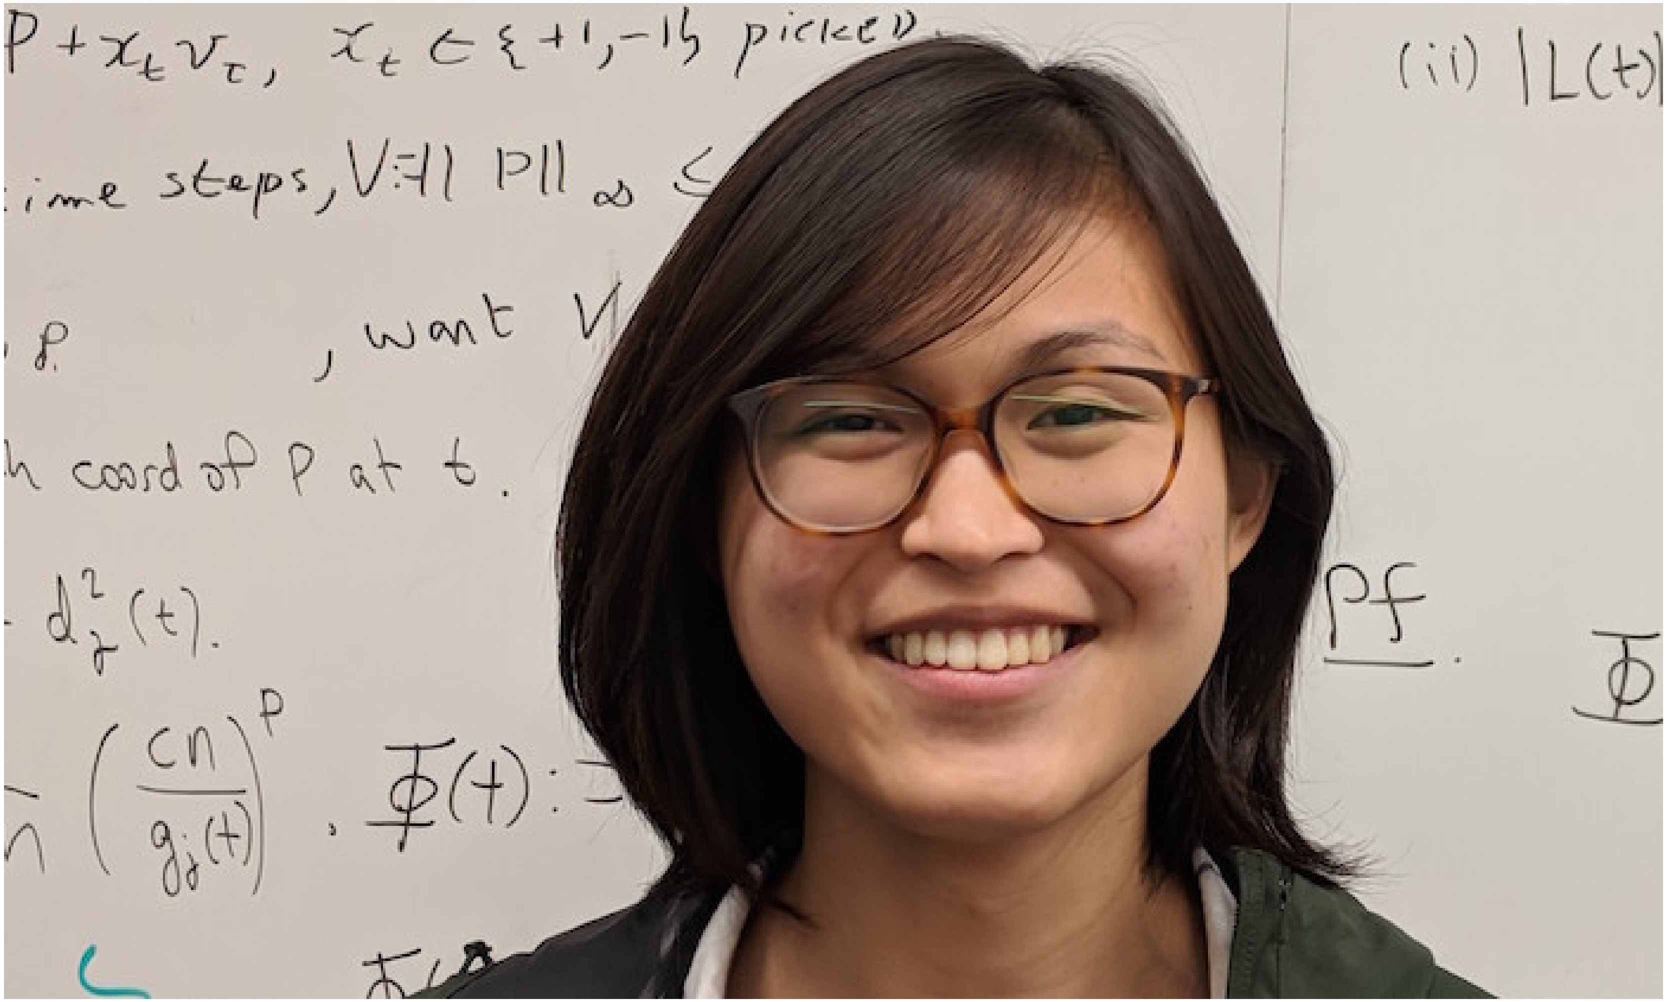

Ewin Tang (U. Washington)

Dequantization!

Recommendation systems: preference matrix has low rank.

Quantum-inspired classical sampling algorithms also have polylog(mn) complexity for low-rank $m \times n$ matrix.

Other tasks: recommendation systems, principal component analysis, supervised clustering, low-rank matrix inversion, low-rank semidefinite programming, support vector machines, etc.

A significant theoretical advance in classical algorithms:
Classical sampling methods with runtime polylog in matrix size. This applies to many tasks, for low-rank matrices.

How practical? Classical methods scale badly with rank and error.

Superpolynomial quantum speedups for sparse high-rank matrices, still hold, but the applications remain unclear.

About quantum error correction

Quantum error correction (QEC) will be essential for solving some hard problems.

It works against noise that is sufficiently weak and also sufficiently weakly correlated. Parallel operations and entropy extraction are necessary.

In a universal gate set acting on encoded quantum information, some gates are "easy" and some are "hard." The hard gates dominate the overhead cost.

Ideal circuit with T gates can be simulated using (asymptotically)
T [polylog (T) / polylog (ε_{th} / ε)] noisy gates, where ε is gate error rate, and ε_{th} is "accuracy threshold." Overhead cost improves as gate error rate declines.

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. Mainstream users may need to patient!

Quantum error correction: long term and near term

Surface code is leading contender for scalability. That could change: e.g. better error rates, higher connectivity. Or a better idea.

Opportunities to explore fault tolerance for qubits with all-to-all connectivity.

Customize fault tolerance to algorithm design and properties of noise.

In the near term, noise mitigation without full-blown quantum error correction.

Neven/Dowling Law?

(1) Gate error rates for two-qubit quantum gates are improving exponentially with time.* (Debatable, and can't go on for long. But by some estimates the error rate is decreasing by a factor of two every two to three years.)

(2) Therefore, the volume of a quantum circuit that can be executed with fixed circuit fidelity is increasing exponentially with time. (Not exactly, but close enough to make a point.)

(3) Furthermore, the classical cost of simulating the quantum circuit increases exponentially with the circuit volume. (Maybe not exactly, but definitely superpolynomial.)

(4) Therefore (Neven/Dowling): for the largest quantum circuit that can be executed with fixed fidelity, the classical cost of the simulation is increasing doubly exponentially with time.

(5) That's really fast. (Even if you don't believe the details.)

(*) Caveat: *Actually, further improvements in gate error rates are difficult to achieve, and progress may be stalling.*

(Much) better gate error rates?

Topological quantum computing. Quantum error correction at the physical level in a highly correlated material. Visionary idea and beautiful physics. Challenging!

Zero-pi qubit. Superconducting qubit designed for robustness against noise. Demonstration by Houck group 2019.

GKP codes. A grid state of a continuous-variable system. Demonstrated by Home group 2018 (ion trap), Devoret group 2019 (superconducting). Suited for photonic computing, too.

Important to continue to develop alternative platforms that could potentially enable big improvements in gate fidelity! (Notable recent improvements for trapped neutral atoms and spin qubits, etc.)

(Much) better gate error rates?

GKP codes

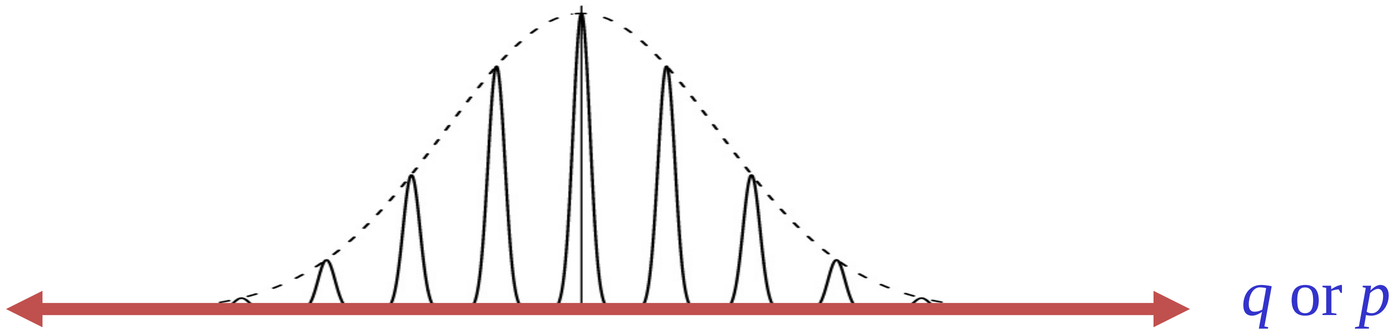

Zero-pi qubit

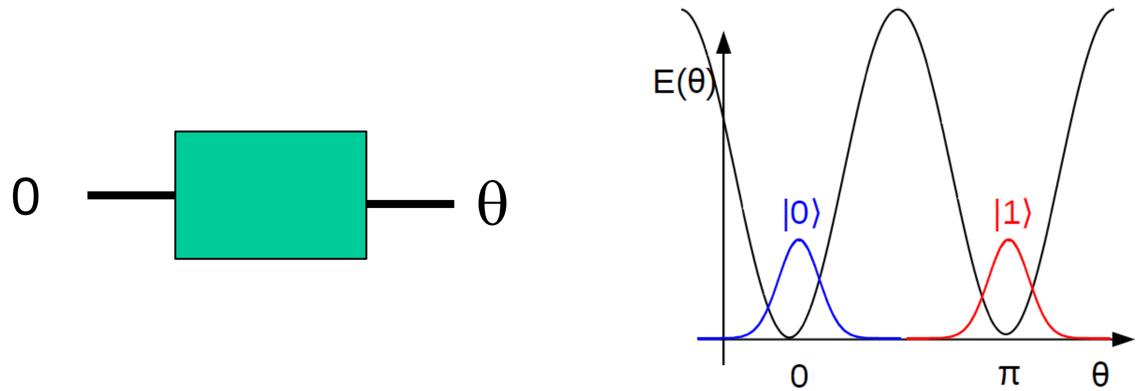

Topological quantum computing

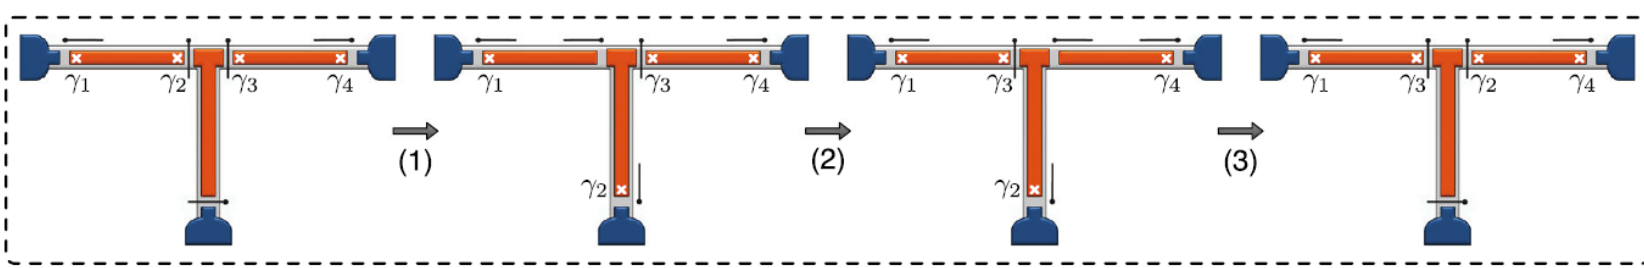

"We are living in a materials world
And I am a materials guy!"

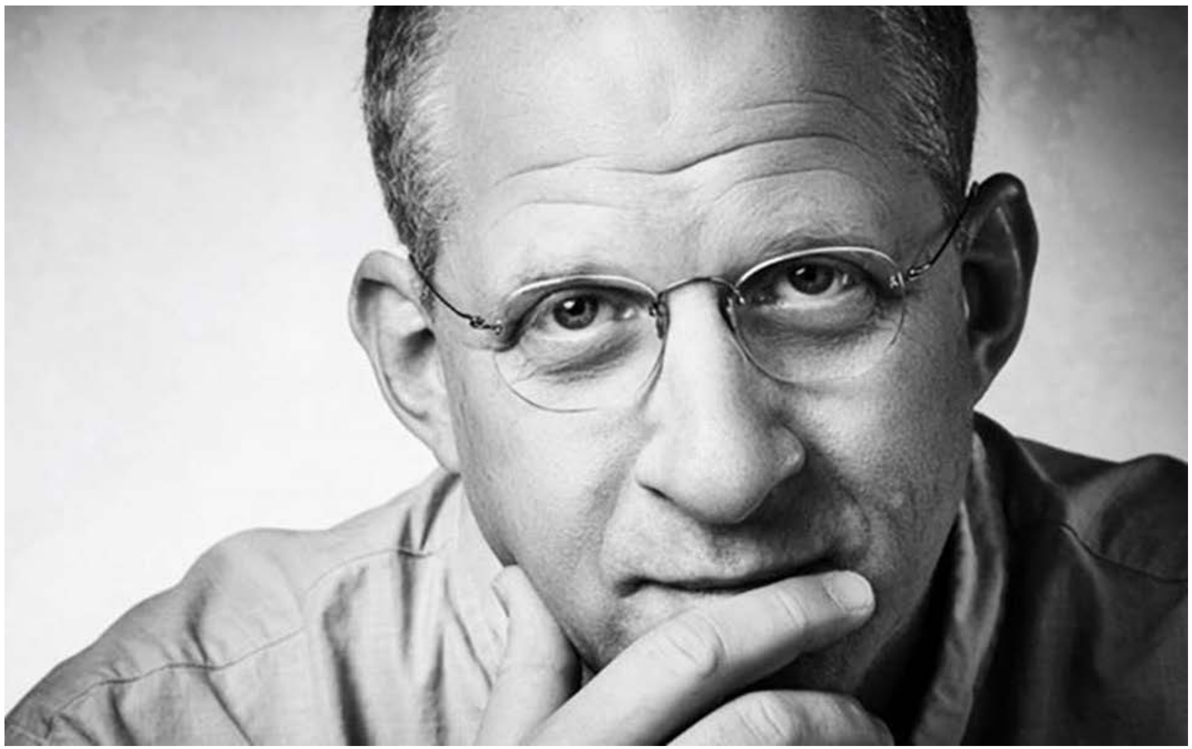

Charlie Marcus (Microsoft/U. Copenhagen)

“We are living in a materials world”

Partially true, but not the whole story.

We also need new ideas.

Frontiers of Physics

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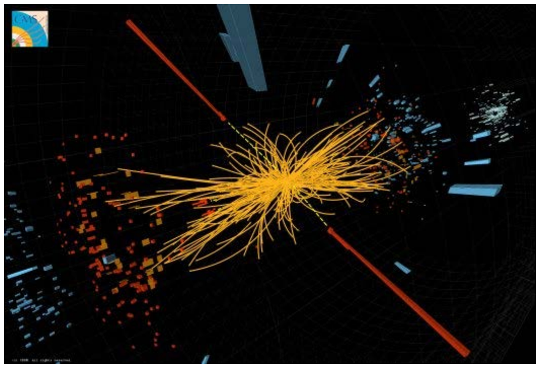		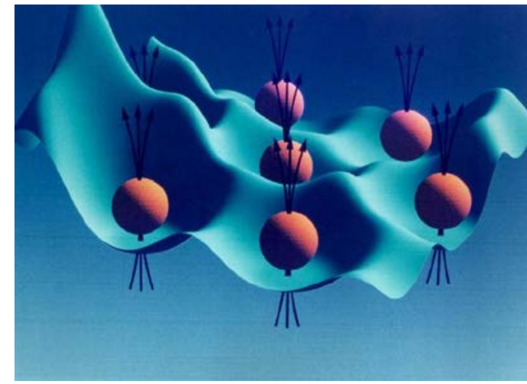
Higgs boson	Large scale structure	“More is different”
Neutrino masses	Cosmic microwave background	Many-body entanglement
Supersymmetry	Dark matter	Phases of quantum matter
Quantum gravity	Dark energy	Quantum computing
String theory	Gravitational waves	Quantum spacetime

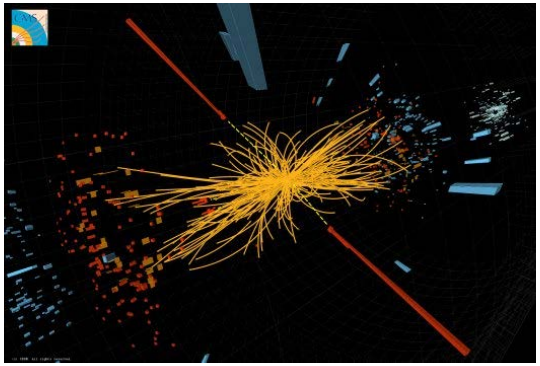
particle collision

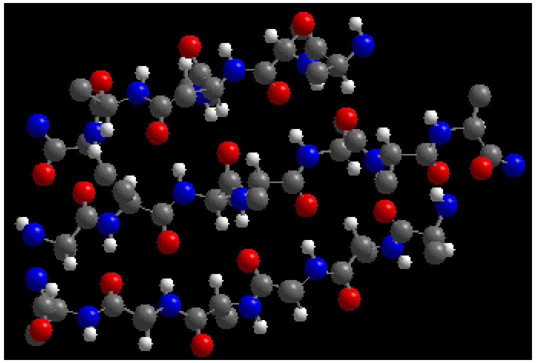
molecular chemistry

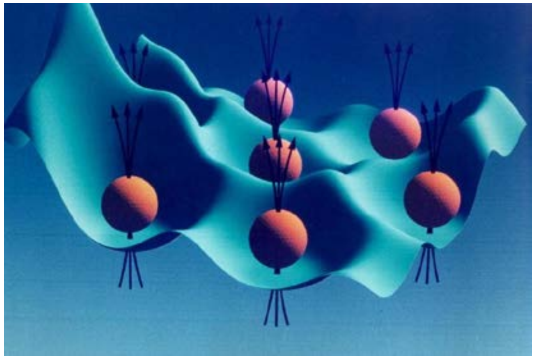
entangled electrons

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

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

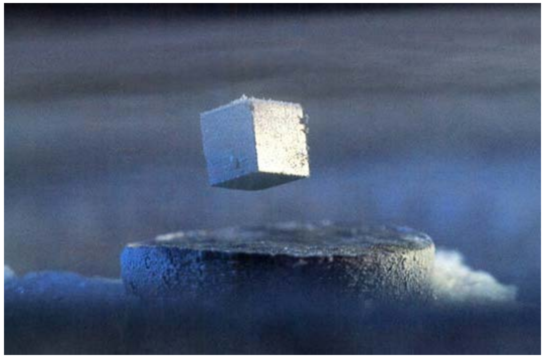
superconductor

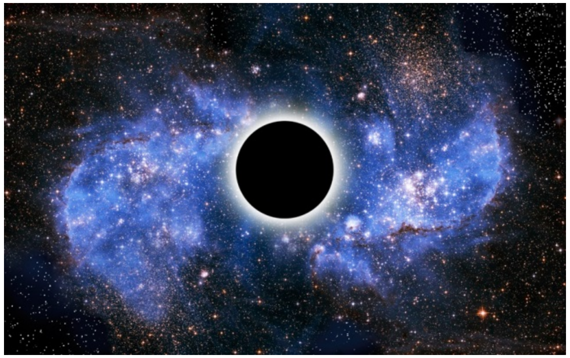
black hole


early universe

"Understanding strongly interacting quantum systems will be like the transition from alchemy to chemistry."

Piers Coleman (Rutgers)

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).

Department of unlikely headlines (Gizmodo 22 Nov. 2019)

PHYSICS

Google Scientists Are Using Quantum Computers to Study Wormholes

Ryan F. Mandelbaum
Yesterday 1:15PM • Filed to: GOOGLE

19.7K 8 Save

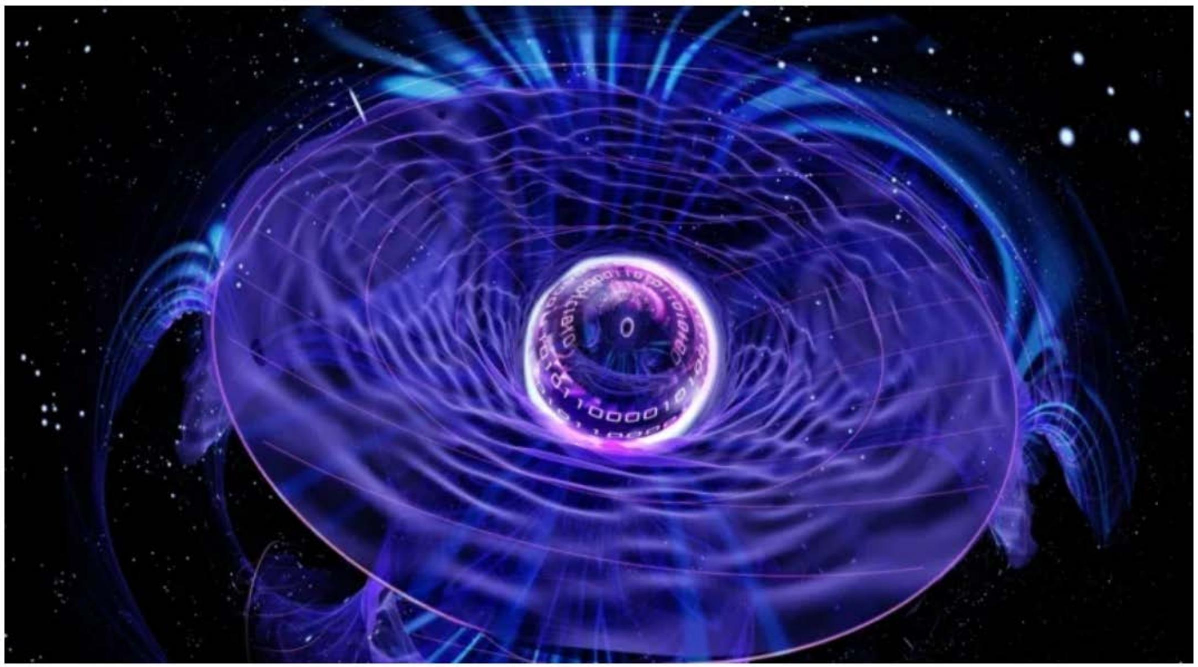

Artist's concept of information falling into a black hole
Illustration: E. Edwards/JQI

Google researchers are figuring out how to study some of the weirdest theorized physics phenomena, like wormholes that link pairs of black holes, using experiments in a lab.

“Quantum computing is a marathon not a sprint”

Chris Monroe (UMD/IonQ)

Quantum speedups in the NISQ era and beyond

Quantum supremacy demonstrations confirm the extravagant computation resources provided by the quantum world.

In the NISQ era we can explore *heuristic* quantum algorithms. Near-term quantum advantage for useful applications is possible, but not guaranteed.

Near-term algorithms should be designed with *noise resilience* in mind.

Lower quantum gate error rates will lower the overhead cost of quantum error correction, and also extend the reach of quantum algorithms which do not use error correction.

Dequantization: Practical uses of quantum linear algebra and of quantum-inspired classical algorithms are still unclear.

Quantum dynamics of highly entangled systems is especially hard to simulate, and is therefore an especially promising arena for quantum advantage.

NISQ will not change the world by itself. Realistically, the goal for near-term quantum platforms should be to *pave the way for bigger payoffs using future devices*. Progress toward fault-tolerant QC must continue to be a high priority for quantum technologists.