Sandia National Laboratories

Exceptional service in the national interest

QuAAC

PRIORITY RESEARCH DIRECTIONS FOR QUANTUM COMPUTING AND NETWORKING

Ojas Parekh (Quantum Algorithms and Applications Collaboratory)

U.S. DEPARTMENT OF ENERGY | Office of Science

Advanced Scientific Computing Advisory Committee Sep 27, 2024

U.S. DEPARTMENT OF ENERGY

NNSA

WORKSHOP-ORGANIZING/REPORT-WRITING TEAM

Report Authors

- Pavel Lougovski (co-chair), Amazon Web Services
- Ojas Parekh (co-chair), Sandia National Laboratories

- Joe Broz, IBM
- Mark Byrd, Southern Illinois University
- Joseph C. Chapman, Oak Ridge National Laboratory
- Yanne Chembo, University of Maryland
- Wibe A. de Jong, Lawrence Berkeley National Laboratory
- Eden Figueroa, Stony Brook University
- Travis S. Humble, Oak Ridge National Laboratory
- Jeffrey Larson, Argonne National Laboratory
- Gregory Quiroz, Johns Hopkins University Applied Physics Laboratory
- Gokul Ravi, University of Michigan
- Nathan Shammah, Unitary Fund
- Krysta M. Svore, Microsoft
- Wenji Wu, Lawrence Berkeley National Laboratory
- William J. Zeng, Unitary Fund

DOE Points of Contact

- Kalyan Perumalla, DOE, Advanced Scientific Computing Research
- Thomas Wong, DOE, Advanced Scientific Computing Research
- Margaret Lentz, DOE, Advanced Scientific Computing Research
- Steven Lee, DOE, Advanced Scientific Computing Research
- William Spotz, DOE, Advanced Scientific Computing Research
- Marco Fornari, DOE, Advanced Scientific Computing Research

DOI: https://doi.org/10.2172/2001045

ASCR COMPUTING AND NETWORKING WORKSHOP


BASIC RESEARCH NEEDS IN

Quantum Computing and Networking

JULY 11–13, 2023

U.S. DEPARTMENT OF ENERGY | Office of Science

PREPARED FOR THE U.S. DEPARTMENT OF ENERGY ADVANCED SCIENTIFIC COMPUTING

Grand Challenge

Demonstrate an end-to-end rigorously quantifiable quantum performance improvement over classical analogs, especially for problems of practical value

5 Priority Research Directions across a general stack

- Applications
- Programming Models
- Algorithms
- Compilation
- Error Resilience
- Hardware Architectures

SOME CONTEXT FOR THE GRAND CHALLENGE

EXAMPLE OF A STACK FOCUSED ON QUANTUM ADVANTAGES

Novel applications of quantum technologies

New algorithms and analysis techniques

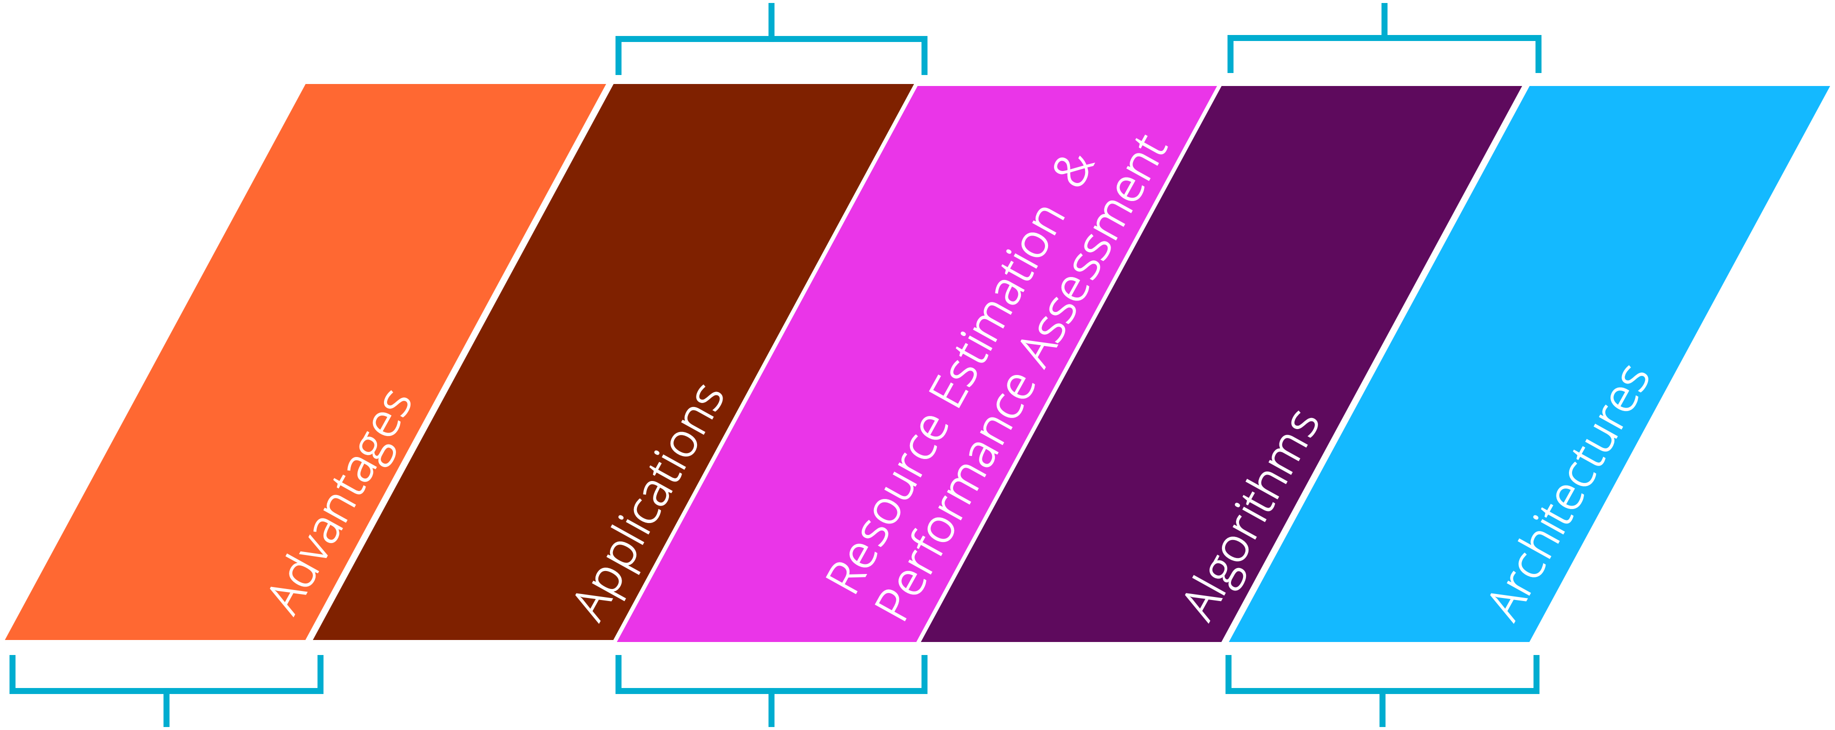

Fair and rigorous evaluation of quantum advantages

Realistic estimates of hardware resources, evaluation of hardware execution

Quantum computing system design informed by insights from higher levels of the stack

DIFFERENT QUANTUM ALGORITHMS FOR DIFFERENT FOLKS

Quantum Algorithms for *Ideal Abstract* Quantum Computers

Models: based on abstract models of quantum computation
Goal: identification of rigorous asymptotic quantum advantages
Challenge: potentially difficult to practically realize advantages

Quantum Algorithms for *Physically-inspired Abstract* Quantum Computers

Models: abstract imbued with physically-inspired features
(e.g. using few ancilla, restricted gate sets or topologies, fidelity/noise-based limitations)
Goal: rigorous quantum advantages under resource restrictions
Challenge: models and results should help bridge ideal-physical gap

Quantum Algorithms for *Physical* Quantum Computers

Models: implementation on current- and near-term quantum computers
(e.g. "quantum software engineering" on commercial systems)
Goal: empirical demonstration of quantum "wins"
Challenge: wins may be platform-specific, not sustainable asymptotically as problems grow,
or have no immediate practical applications

QUANTUM ADVANTAGES NEED MANY THINGS TO GO RIGHT

Find efficient quantum algorithm

Find a problem admitting (exponential) quantum advantage

Find proof or evidence *no* efficient classical algorithm exists

Find applications where discovered advantages translate to impact

See Aaronson [arXiv:2209.06930] for recent high-level survey of exponential quantum advantages

LOOKING BEYOND QUANTUM SPEEDUPS

Known quantum speedups are limited, so we look for quantum algorithms that are better (for some resource), not necessarily faster

	Against best-known classical algorithm	Against best-possible classical algorithm
Polynomial speedup	e.g. checking if an integer is prime	e.g. Grover's algorithm for unstructured search
Exponential speedup	e.g. Shor's algorithm for integer factorization	?

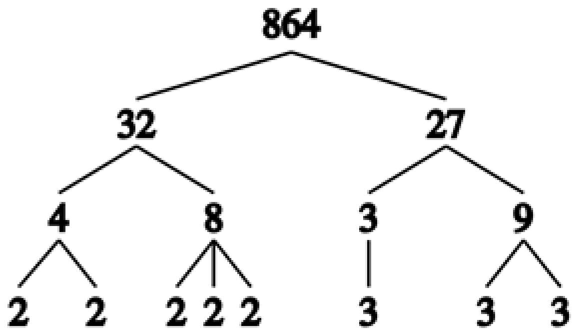

See https://quantumalgorithmzoo.org for state of quantum speedups

QUANTUM ADVANTAGES BY MAKING PROBLEMS HARDER

Conventional Input

$$A = \begin{bmatrix} 2 & -1 & 0 & 0 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ 0 & 0 & -1 & 2 \end{bmatrix}$$

$$\lambda_{max}(A)$$

Input: direct representation of matrix *A*
(size of input depends directly on *A*)

Output: numerical value
(precision depends on *A*)

Compressed Input

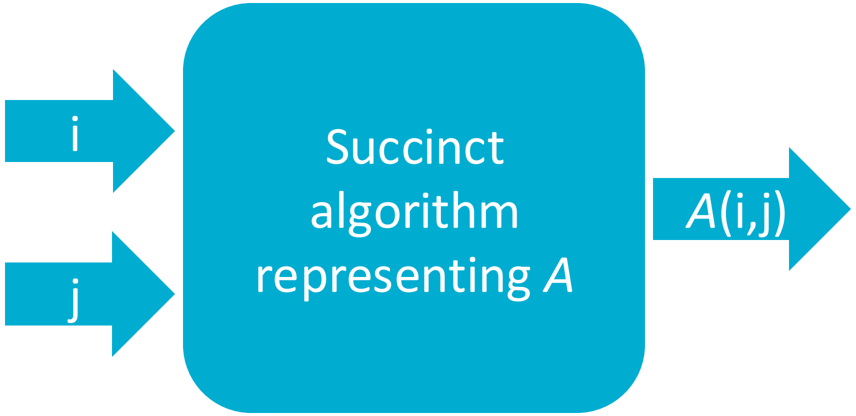

Exponential quantum advantages possible!

$$\lambda_{max}(A)$$

Input: algorithm indirectly representing matrix *A*
(*A* could be exponentially large in size of input)

Output: numerical value
(precision depends on *A*)

HOW ABOUT SPACE INSTEAD OF TIME?

Qubits are expensive, so we want algorithms using few of them as possible (minimize space)

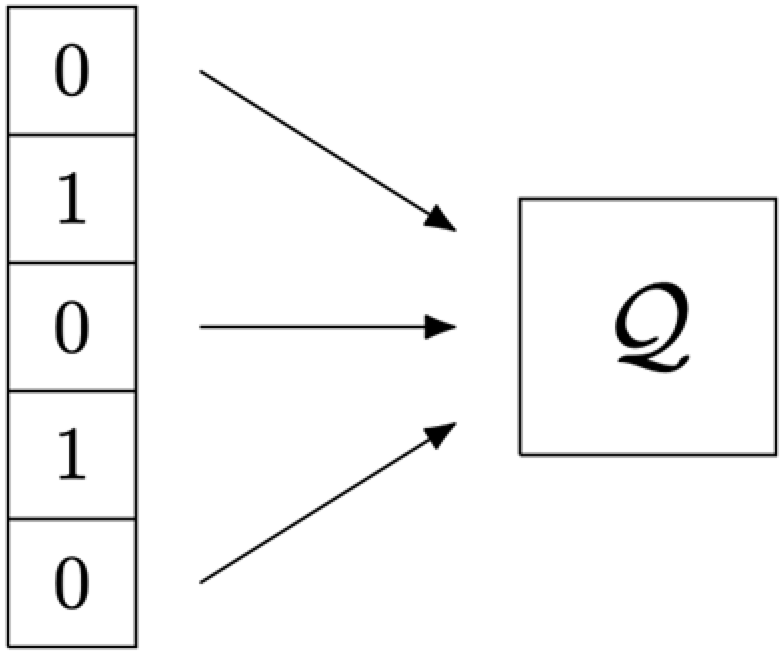

Ideally a number *sublinear* in the size of the input, e.g. $O(\sqrt{n})$ or $O(\log(n))$ for a size-n input

STREAMING ALGORITHMS

When dealing with very small space algorithms, it matters *how* you receive the input dataset

Streaming

- Dataset is built up by a "stream" of small updates
- Answer is expected at the end of the stream

192.168.0.1
10.0.0.2
172.16.20.5
192.0.2.123
203.0.113.5
172.16.0.3
203.0.113.4

STREAMING ALGORITHMS

When dealing with very small space algorithms, it matters *how* you receive the input dataset

Streaming

Dataset is built up by a “stream” of small updates

Answer is expected at the end of the stream

10.0.0.2
172.16.20.5
192.0.2.123
203.0.113.5
172.16.0.3
203.0.113.4
127.0.0.1

STREAMING ALGORITHMS

When dealing with very small space algorithms, it matters *how* you receive the input dataset

Streaming

Dataset is built up by a “stream” of small updates

Answer is expected at the end of the stream

172.16.20.5
192.0.2.123
203.0.113.5
172.16.0.3
203.0.113.4
127.0.0.1
8.8.8.8

STREAMING ALGORITHMS

When dealing with very small space algorithms, it matters *how* you receive the input dataset

Streaming

- Dataset is built up by a "stream" of small updates
- Answer is expected at the end of the stream

172.16.20.5
192.0.2.123
203.0.113.5
172.16.0.3
203.0.113.4
127.0.0.1
8.8.8.8

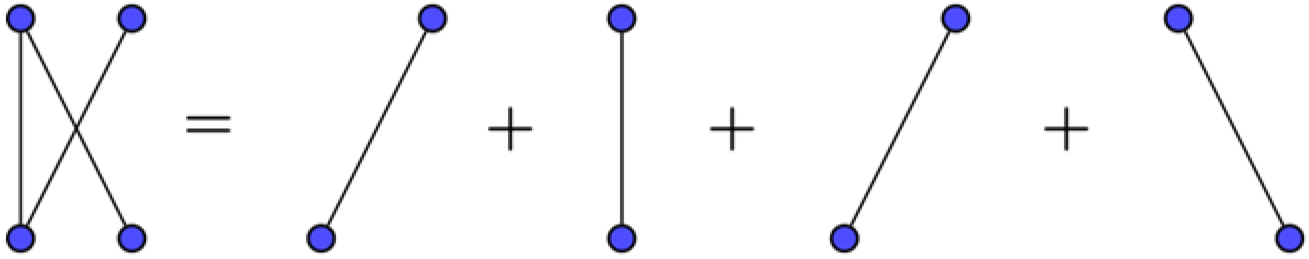

Large social network arising from real-time frenemy-ships

QUANTUM STREAMING ADVANTAGES FOR GRAPH PROBLEMS

Exponential advantage for Boolean Hidden Matching
[Gavinsky, Kempe, Kerenidis, Raz, and de Wolf 2008]

First natural problem: polynomial advantage for triangle counting
[Kallaugher 2021]

No quantum advantage possible: Max Cut graph partitioning problem
[Kallaugher, P 2022] ASCR ARQC FAR-QC Project

Exponential advantage for natural problem: *Directed* Max Cut problem
[Kallaugher, P, Voronova 2023] ASCR ARQC FAR-QC Project

A first quantum approximation advantage for approximating discrete optimization (albeit in streaming model)

Max Cut

Following slides courtesy of John Wright

Max Cut

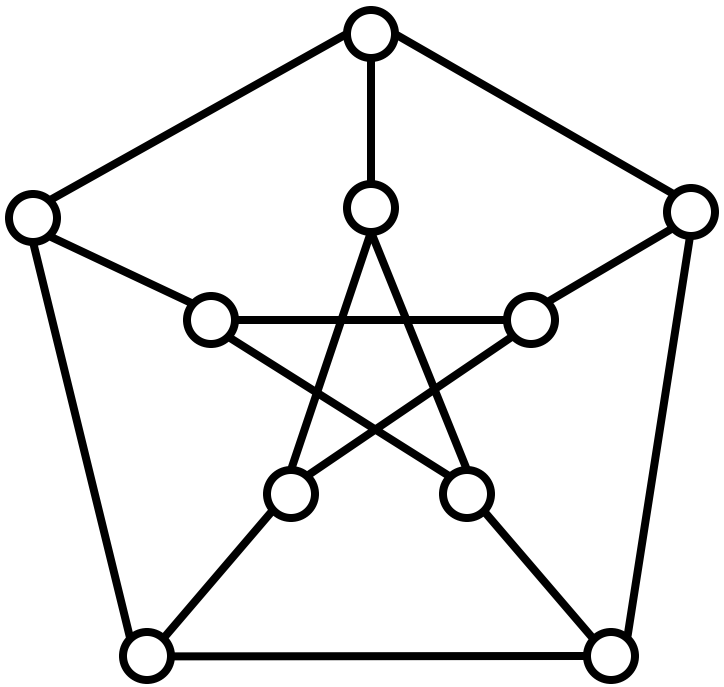

$$G = (V, E)$$

Max Cut

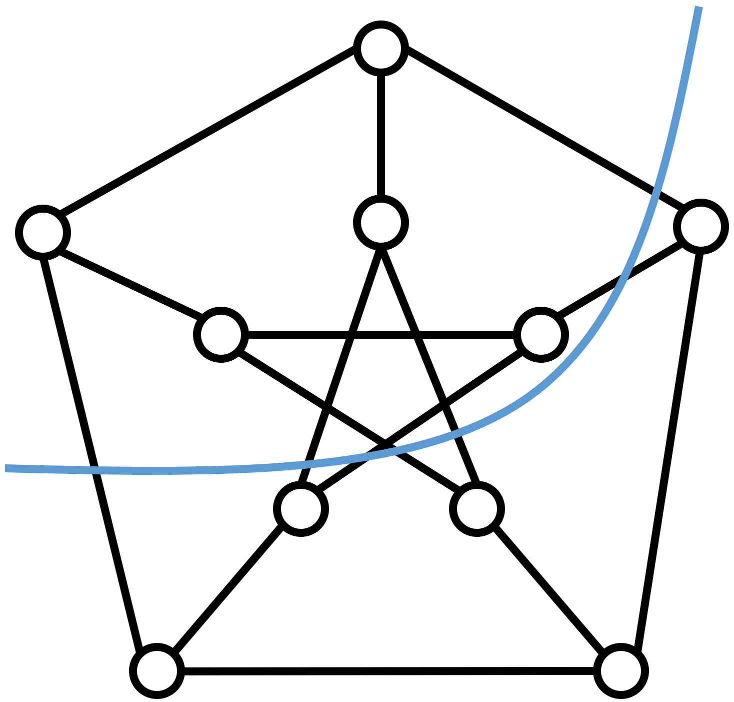

$G = (V, E)$

Max Cut

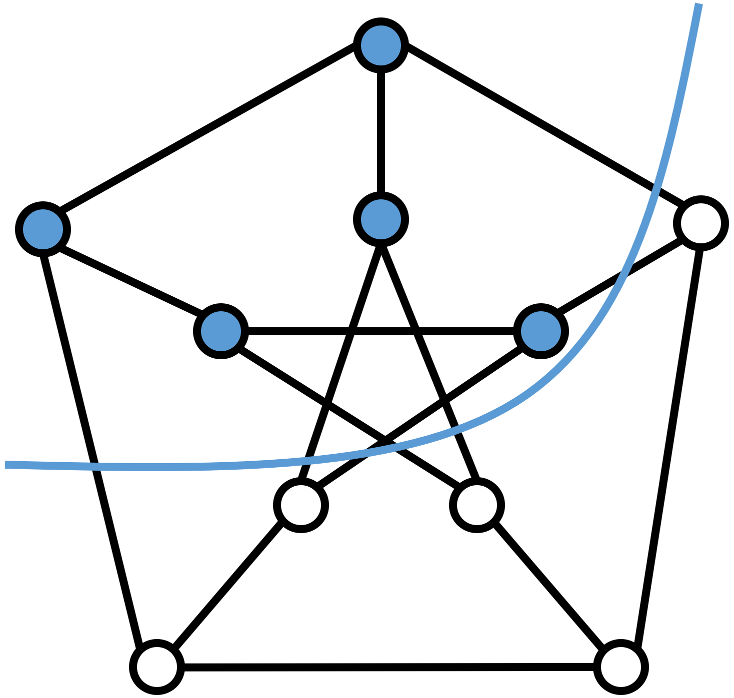

$$G = (V, E)$$

Max Cut

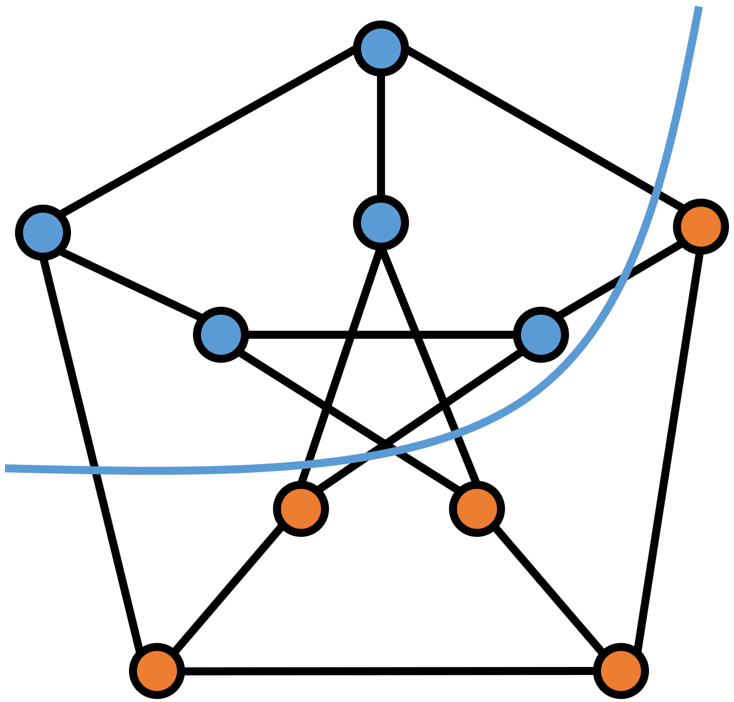

$G = (V, E)$

Max Cut

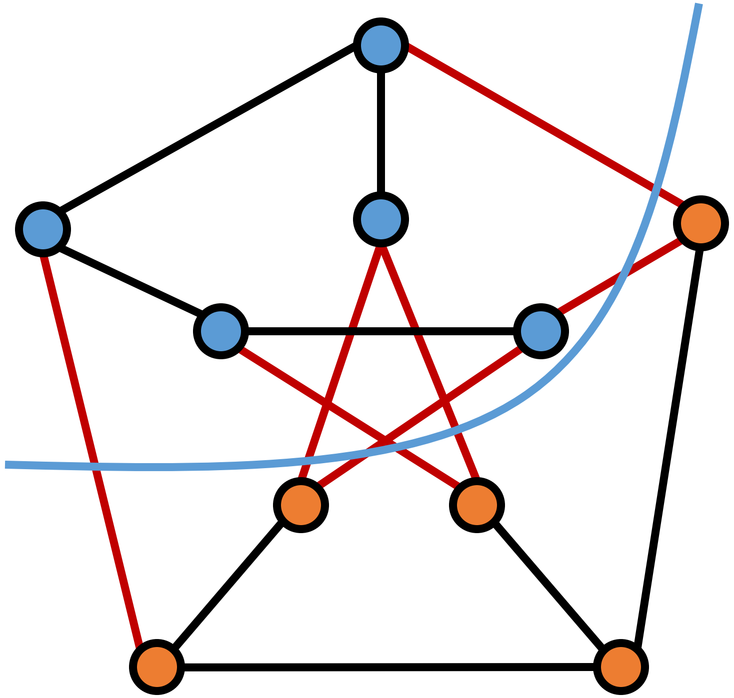

$G = (V, E)$

Max Cut

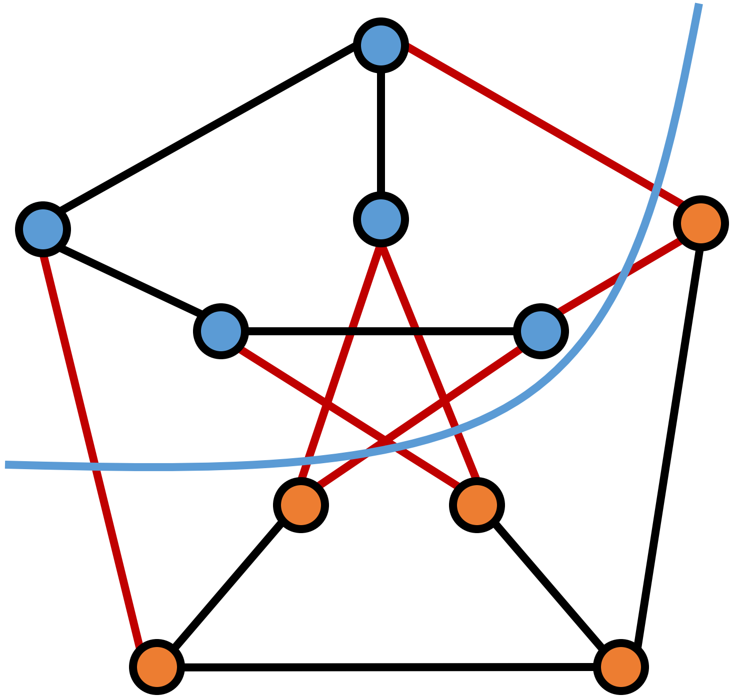

$G = (V, E)$

7 edges cut

Max Cut

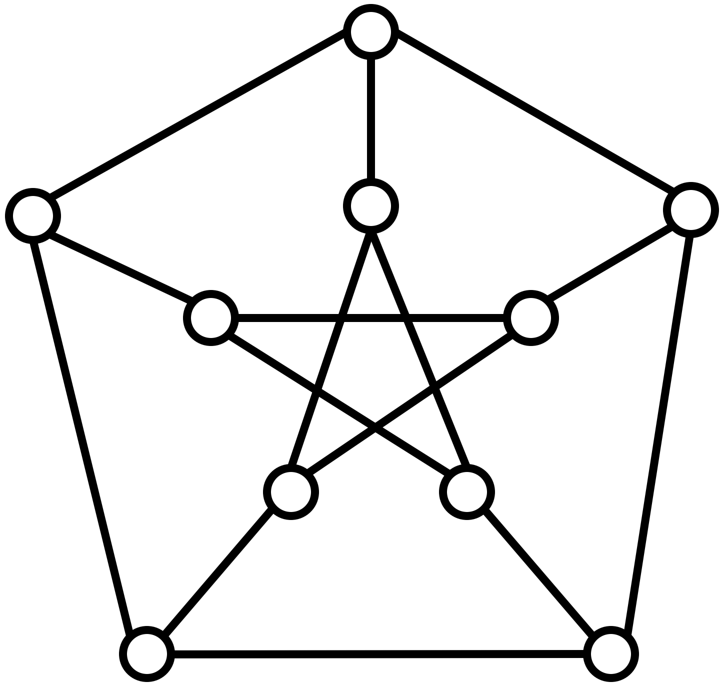

$G = (V, E)$

7 edges cut

Max Cut

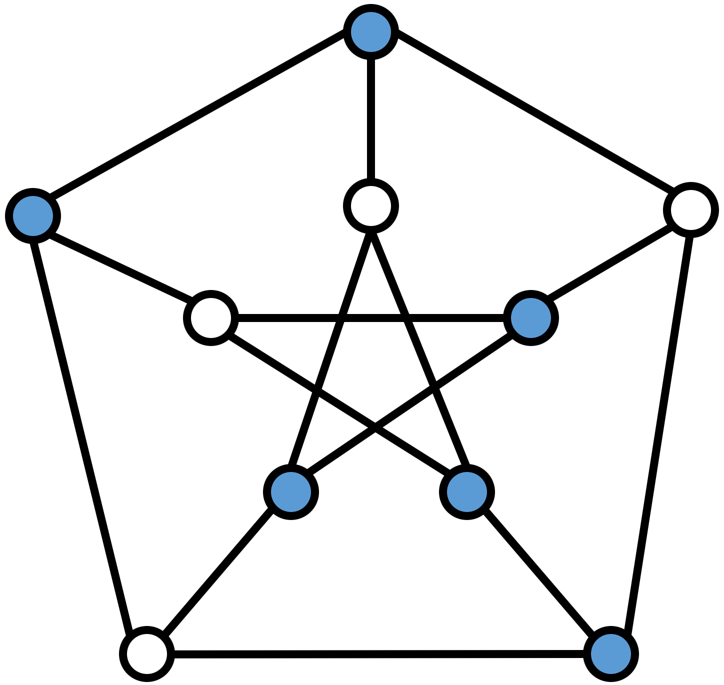

$G = (V, E)$

7 edges cut

Max Cut

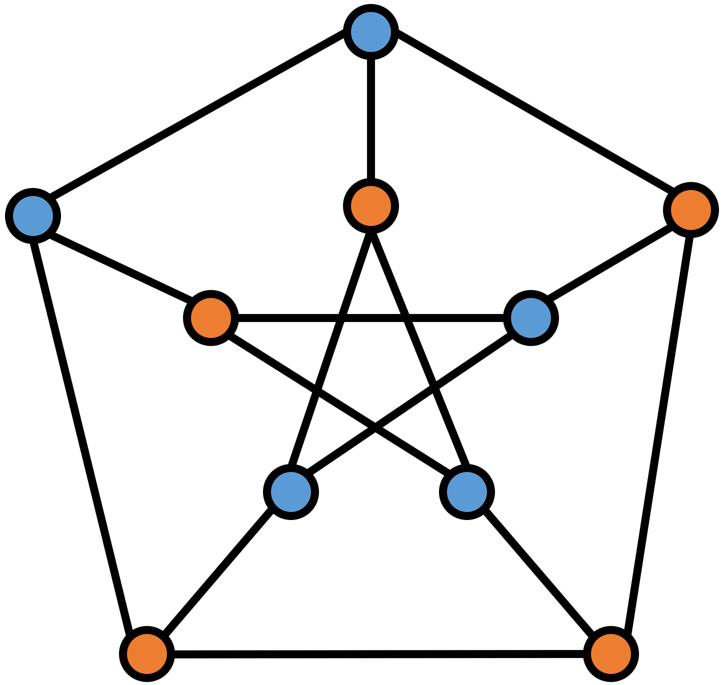

$G = (V, E)$

7 edges cut

Max Cut

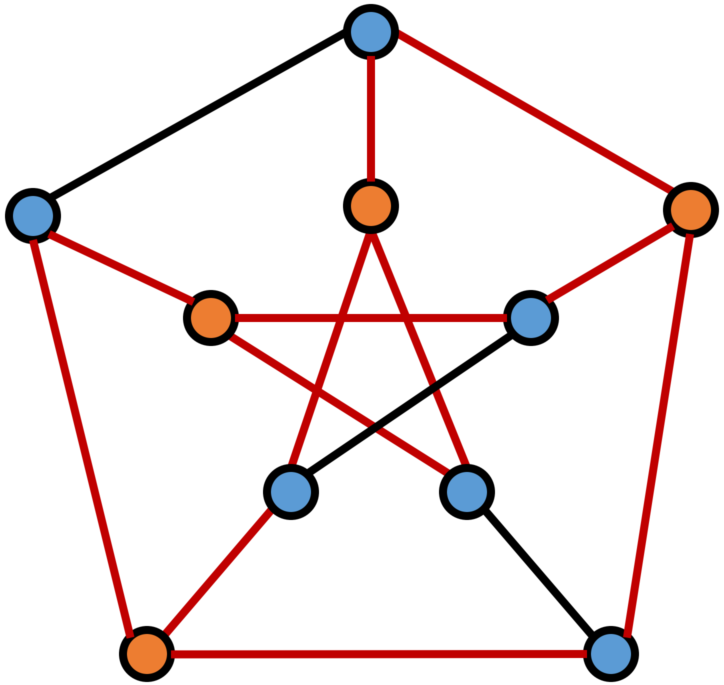

7 edges cut

$G = (V, E)$

Max Cut

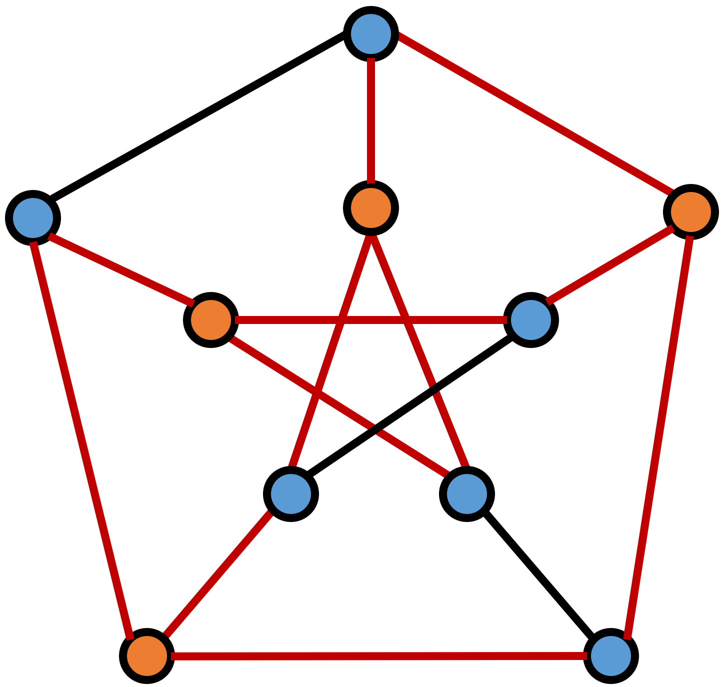

12 edges cut

$G = (V, E)$

Max Cut

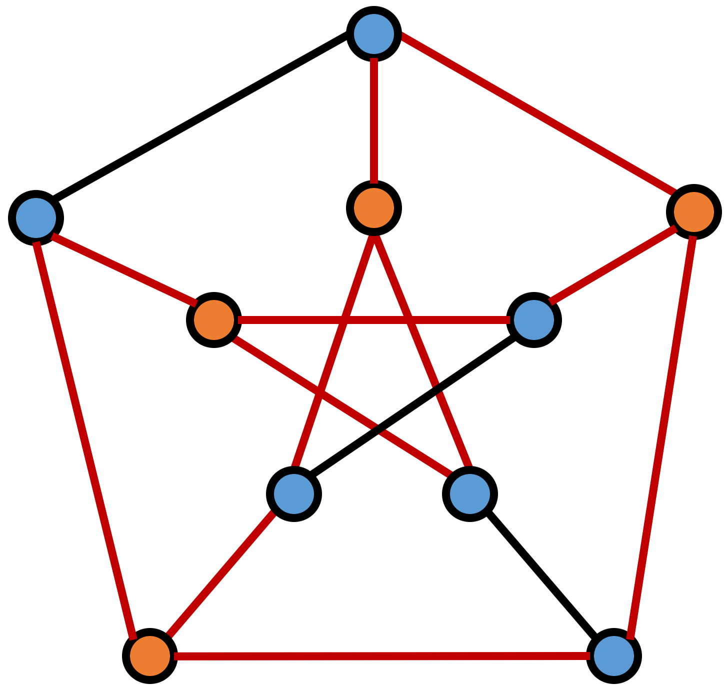

$G = (V, E)$

12 edges cut
(the **max cut**)

Directed Max Cut

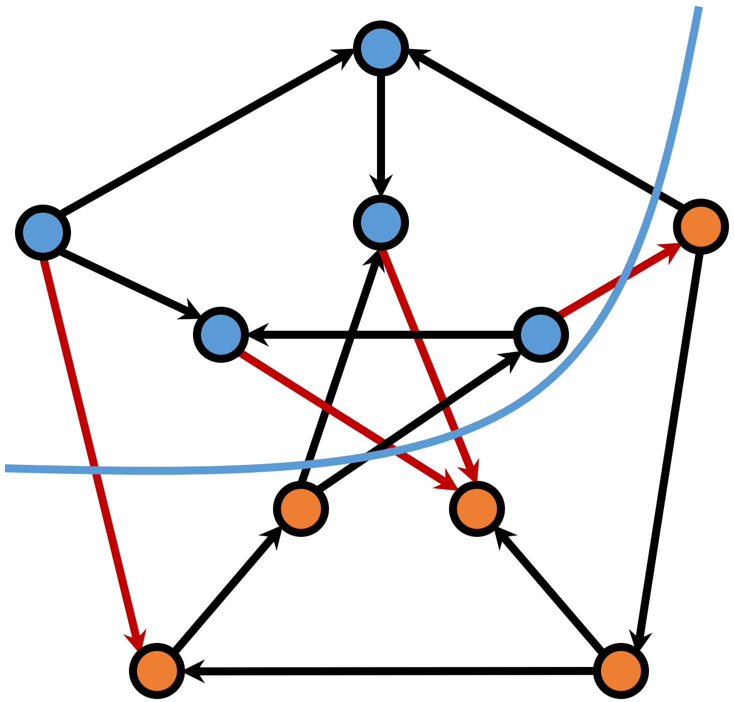

$G = (V, E)$

4 directed edges cut (among 7 edges crossing partition)

STREAMING APPROXIMATIONS FOR MAX CUT

Max Cut is hard to solve exactly, so we turn to approximations

0.878 is best possible for standard non-streaming version
[Goemans, Williamson 1995]

approx factor = ½ (random)

1

Quantum space: $\Omega(n)$
[Kallaugher, P 2022]

Classical space: $\Omega(n)$
[Kapralov, Krachun 2018]

No nontrivial approximations possible!

STREAMING APPROXIMATIONS FOR DIRECTED MAX CUT

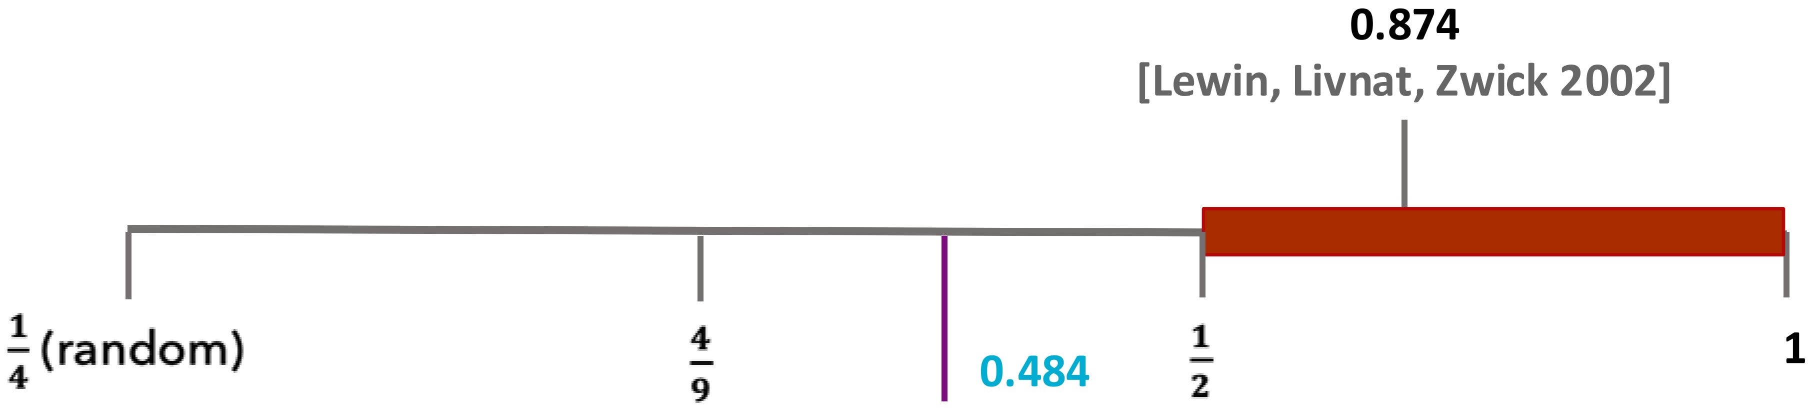

Quantum space:	$\mathbf{polylog}(n)$ [Chou, Golovnev, Velusami '20]	$\mathbf{polylog}(n)$ [This work]	$\Omega(n)$ [Kallaugher, P '22]
Classical space:	$\mathbf{polylog}(n)$ [Chou, Golovnev, Velusami '20]	$\Omega(\sqrt{n})$ [Saxena, Singer, Sudan, Velusami '23]	$\Omega(n)$ [Kapralov, Krachun '18]

Directed Max Cut leaves room for an exponential quantum advantage!

THE QUEST FOR QUANTUM APPROXIMATION ADVANTAGES

Despite considerable work in quantum approximation for discrete optimization problems, no provable advantages have been discovered

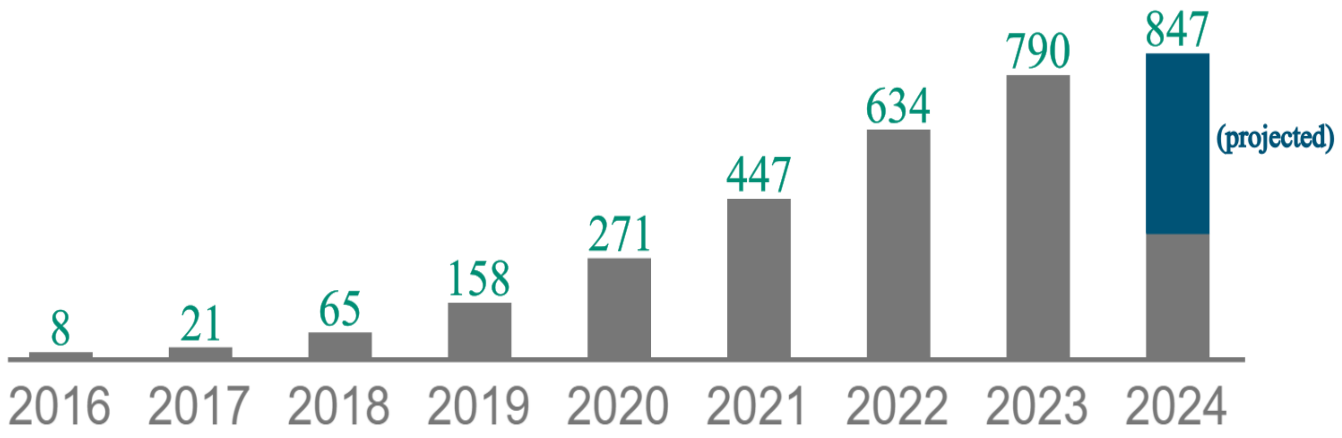

Citations of "A Quantum Approximate Optimization Algorithm" by Farhi, Goldstone and Gutmann

Directed Max Cut work is first such result, albeit it is a space advantage

TAKEAWAYS

- **Considering resources beyond runtime may open door to new quantum advantages**
- **Quantum advantages are highly sensitive to problem details:**
  e.g., input/output models and problem formulation (e.g. directed vs undirected)
- **Are we overlooking important quantum advantages because of focus or bias toward certain kinds of problems?**
- **Quantum advantages for the problems we *are currently* solving may not exist.**
  What overlooked problems *should* we be solving to impact science mission?

PRIORITY RESEARCH DIRECTIONS

Priority Research Direction: Software Toolchains

1. End-to-end software toolchains to program and control quantum systems and networks at scale

Key Questions: How can we design expressive programming models and languages to attract a broad user base and facilitate quantum algorithm design and implementation? How can we incorporate these into end-to-end toolchains to produce resource-efficient quantum programs?

Quantum computing and networking systems continue to grow in scale and complexity and will place an increased burden on the software stack to program, control, and manage these systems effectively. Software toolchains will be needed that integrate programming models with hardware-level control systems to maximize the performance and fidelity of quantum systems and to facilitate codesign of hardware, control systems, and algorithms across different technology platforms. Integrating quantum networking systems with quantum computing systems will be critical in advancing the delivery of distributed distributed quantum computing services. This integration will require a quantum networking stack that is compatible with the quantum computing software stack, ensuring that the combined system can be efficiently managed, controlled, and programmed.

Priority Research Direction: Quantum Advantages

2. Efficient algorithms delivering quantum advantages

Key Questions: What classes of existing and understudied scientific applications admit substantial quantum advantages over conventional classical computing paradigms? How can we design novel algorithms and supporting mathematical models to realize such advantages? Are there any provable or empirical barriers for quantum advantages? What are the physical resource requirements of practical implementations of these algorithms, including numbers of physical qubits and quantum circuit depth?

Quantum computing is not expected to universally accelerate current computing tasks, and so identifying problems with special structure amenable to quantum advantages is a paramount goal. Taking a complementary perspective, broadening our understanding of foundational computational kernels admitting quantum advantages is equally important. While a variety of quantum advantages are currently known, they are subject to shortcomings such as a lack of known practical applications, near-term realization, rigorous provability of advantage, or efficient verifiability of advantage. In addition, quantum advantages have largely focused on improving execution time. Advantages with respect to other critical resources, such as quality/accuracy of solution, energy consumption, space/memory, or communication, are understudied, especially in the context of quantum networking.

Priority Research Direction: Benchmarking

3. Benchmarking, verification, and simulation methods to assess quantum advantages

Key Questions: How can we rigorously assess quantum advantage relative to classical capabilities as quantum computing and networking technologies evolve and scale? What metrics and evaluation methodologies faithfully reflect or enable the assessment of quantum advantage across the computing and networking stacks?

Assessing progress toward quantum advantage is a challenging and multi-faceted endeavor. Empirical evidence of advantages are expected to continue to rely on large-scale classical simulations of quantum systems as quantum technologies mature. A considerable hurdle is forecasting scalable quantum advantage based on limited results obtained from relatively small near-term quantum systems and classical simulations. On the one hand, while rigorous proofs of asymptotic quantum advantage are ultimately desirable and may be used to direct empirical studies, the former often rely on abstract or specialized models of quantum computing or otherwise impose additional restrictions. On the other hand, quantum advantage suggested by empirical assessments may not be sustainable as problems grow in scale or complexity, or as better classical algorithms are developed. Bridging this gap between theory and practice is essential for establishing sound and practical quantum advantages. Rigorous, informative, and efficiently verifiable performance metrics, at all levels of quantum computing and networking stacks need to be defined and developed. Ideally, such metrics should be integrated across the stack so that improvements can be quantified and predicted performance may be realized in practice.

Priority Research Direction: Error Resilience

4. Resilience through error detection, prevention, protection, mitigation, and correction

Key Questions: How can we enhance the resilience of quantum systems to noise and errors to relieve scalability and quantum advantage bottlenecks? What kinds of quantum algorithm codesign techniques can aid in yielding resilient quantum systems?

Scientists and engineers in national laboratories, academia, and industry continue to improve quantum computing and networking hardware, but despite these steady advances, these systems will be noisy and imperfect. In recent years, significant efforts characterizing errors and inserting error mitigation at various layers of the software stack have allowed the research community to cut through some of the noise and achieve reliable results in small-scale quantum experiments. To achieve reliable results with quantum systems at larger scale and complexity, more efficient and better methods characterizing, mitigating, preventing, or protecting against dynamical errors need to be integrated in the critical layers of the software stack. Steps are needed toward fault tolerance, codesign, and early demonstrations of quantum error correction that outperform the physical counterpart. Another approach would be to identify the error resistance mechanisms for quantum algorithms and codesign new hardware-aware algorithms and hardware controls that lead to error resistance.

Priority Research Direction: Quantum Networks

5. Hardware and protocols for next-generation quantum networks

Key Questions: Can quantum repeater hardware be built to achieve entanglement distribution rates higher than those of repeat-until-success direct transmission experiments? What software and hardware, besides the repeaters, is needed to build scalable quantum networks? What applications and advantages will those networks enable? What kinds of distributed quantum computing models will result in novel quantum applications and advantages?

To date, non-error-corrected quantum memories and entanglement distribution between them have been demonstrated with multiple qubit technologies. Moving forward and enabling scalable entanglement distribution networks will require progress in multiple directions. To create fault-tolerant quantum repeaters, it will be necessary to enhance quantum memories by integrating error detection and correction functionality during the design and implementation process. Photon sources, detectors, and time tagging hardware will need to improve to increase the fidelity of entanglement swapping operations. The quantum networking software stack will need to implement distributed error correction protocols and enable optimization across the stack. High-fidelity quantum information transduction methods and hardware need to be developed to enable the use of entanglement distribution networks in distributed quantum computing applications.

THANKS FOR KINDLY RESTRAINING YOURSELF FROM THROWING TOMATOES, ETC.!