Title: Shallow circuits and the quantum-classical boundary

Speakers: David Gosset

Series: Colloquium

Date: September 22, 2021 - 2:00 PM

URL: https://pirsa.org/21090000

Abstract: In the last few years there have been demonstrations of quantum advantage using noisy quantum circuits that are believed to go beyond the limits of the classical computers that exist today. In this talk I will give an overview of a different type of quantum advantage that can be attained by shallow (short-depth) quantum circuits. I will discuss recent results which establish unconditionally that constant-depth quantum circuits can solve certain linear algebra problems faster than their classical counterparts. We will see that the reason quantum computers solve these problems provably faster (as measured by circuit depth) than classical computers is due to a strong form of quantum nonlocality that is present in their input/output statistics.

Zoom Link: https://pitp.zoom.us/j/96752851897?pwd=R29GWHovN0MwVXVYWklaNE1QZ1c5dz09

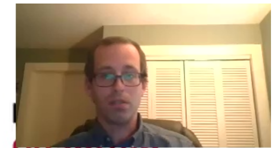

Shallow circuits and the quantum-classical boundary

David Gosset, University of Waterloo

Bravyi, DG, Koenig. *Science* 362 (6412), 2018.
Bravyi, DG, Koenig, Tomamichel. *Nature Physics 1-6*, 2020
DG, Grier, Kerzner, Schaeffer. *arXiv:2009.03218,* 2020

Perimeter Institute Colloquium September 22, 2021

Early quantum computing

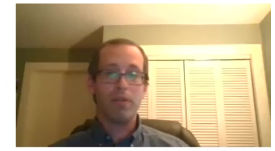

1930s

Classical computers.

1980s

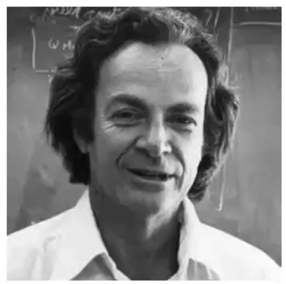

Quantum computers as quantum simulators.

1990s

Quantum algorithms for factoring and discrete log.

Quantum error correction.

Evidence for quantum advantage in computat

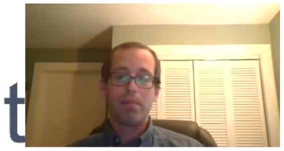

Quantum algorithms with speedups over classical

Shor's algorithm

Simulation of Hamiltonian dynamics

$$i\hbar \frac{d|\psi\rangle}{dt} = H|\psi\rangle$$

Sampling from classically hard distributions

Boson sampling

IQP circuits

Random quantum circuits

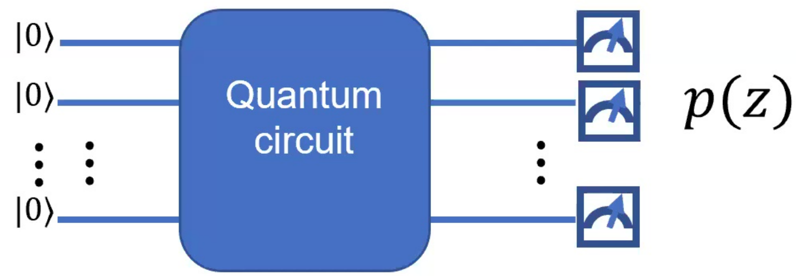

Provable speedups relative to an oracle

Bernstein-Vazirani

Simon's problem

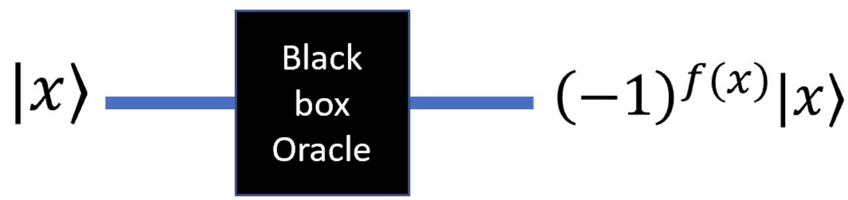

Noise can corrupt quantum information

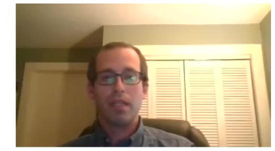

n

# of qubits

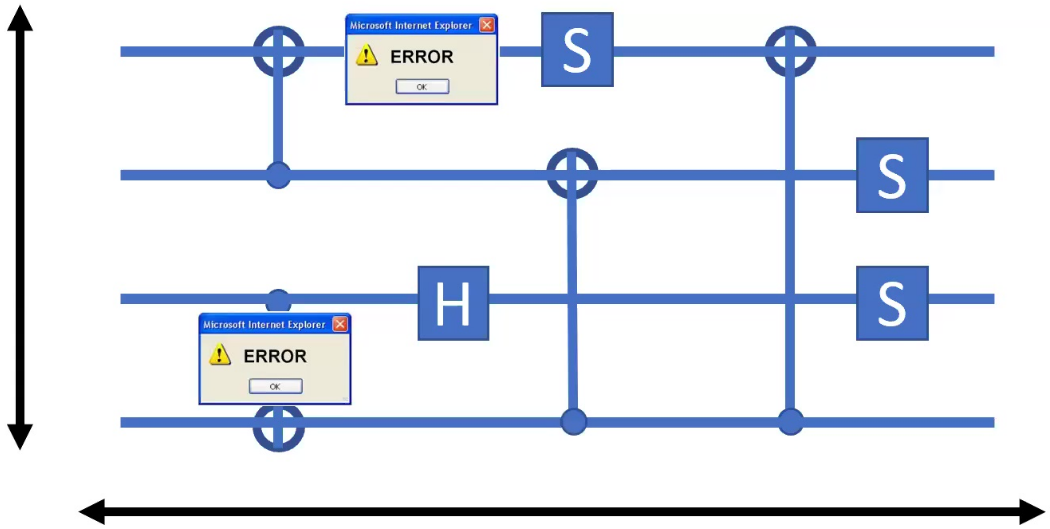

d : circuit depth

Expected number of errors scales with the total circuit size $\sim nd$

Error correction and fault-tolerance

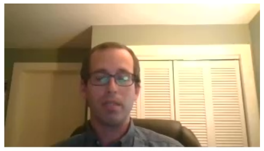

Quantum information can be protected using error correcting codes.

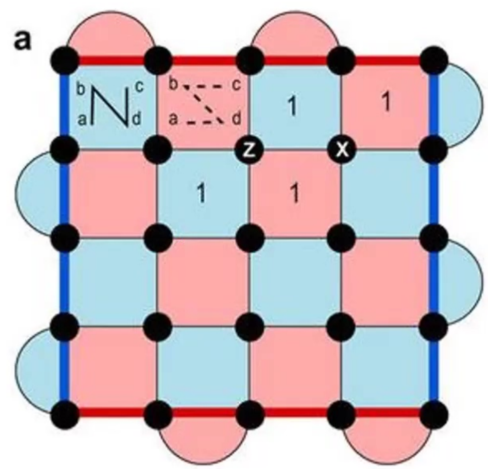

A logical qubit is composed of multiple physical qubits

Image source: [Gambetta, Chow, Steffen 2017]

Error correction and fault-tolerance

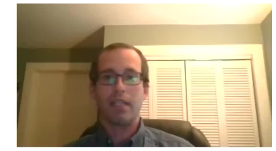

Quantum information can be protected using error correcting codes.

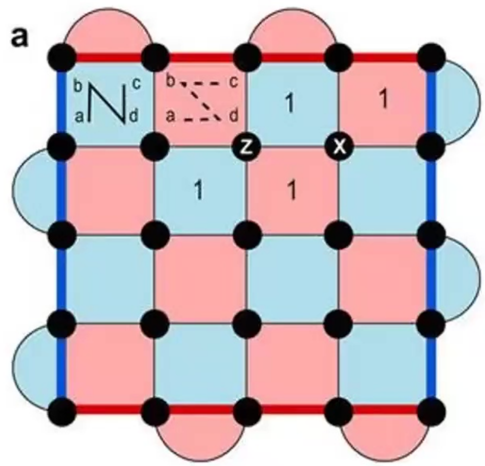

A logical qubit is composed of multiple physical qubits

Using quantum error correction it is possible to compute fault-tolerantly.
The overhead is impractical for now.

Image source: [Gambetta, Chow, Steffen 2017]

Quantum computers today

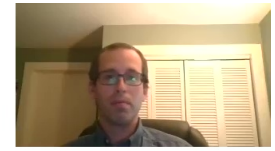

Current quantum computers do not incorporate error correction and are affected by noise.

IBM Quantum
https://quantum-computing.ibm.com/services?services=systems&system=ibmq_montreal

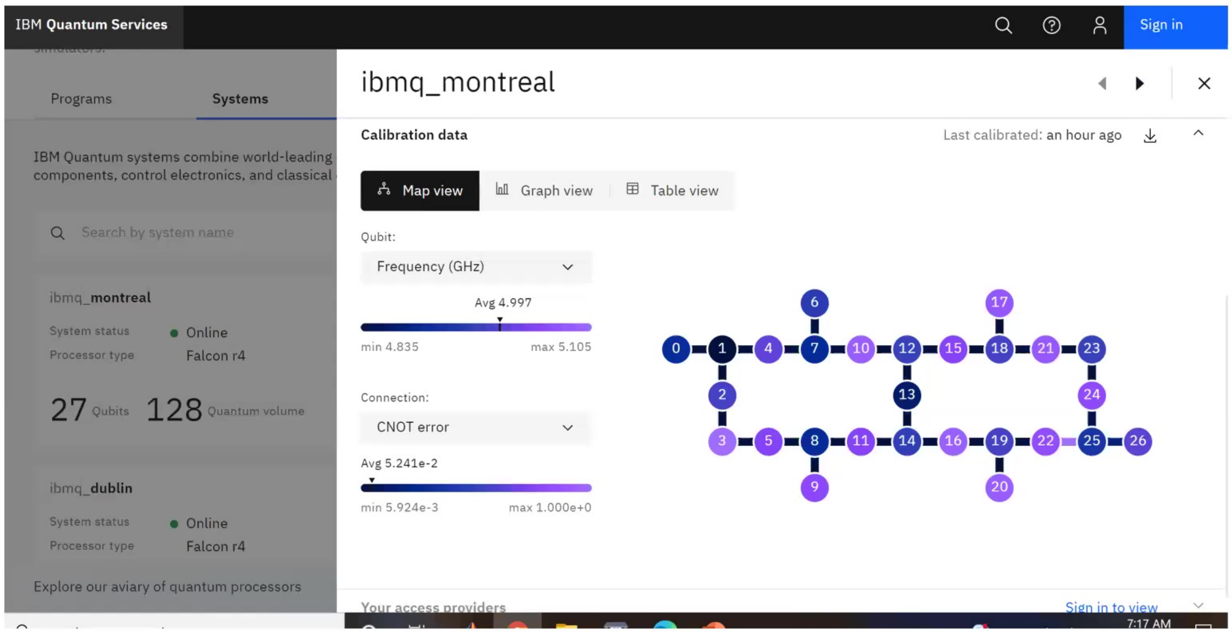

Quantum computers today

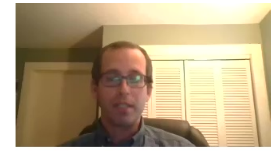

Current quantum computers do not incorporate error correction and are affected by noise.

IonQ
https://ionq.com/technology

Credit: IonQ

Performance Benchmarks†

Qubits
Single-qubit gates on
79 Qubits
Two-qubit gates on all pairs up to
11 Qubits

Average Fidelity
Single-qubit gates
>99%
Two-qubit gates
>98%*

Best Fidelity
Single-qubit gates
>99.97%
Two-qubit gates
>99.3%*

Minimum Fidelity
Single-qubit gates
>99%
Two-qubit gates
>96%*

Coming Soon: 32 Qubits
We are currently gathering detailed data on our latest system, which features a capacity of 32 fully-connected qubits and world-leading algorithmic performance.

Quantum computers today

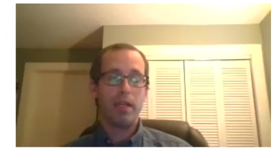

Current quantum computers do not incorporate error correction and are affected by noise.

Google

"Quantum supremacy using a programmable superconducting processor"

Arute et al, *Nature* 574.7779 (2019): 505-510

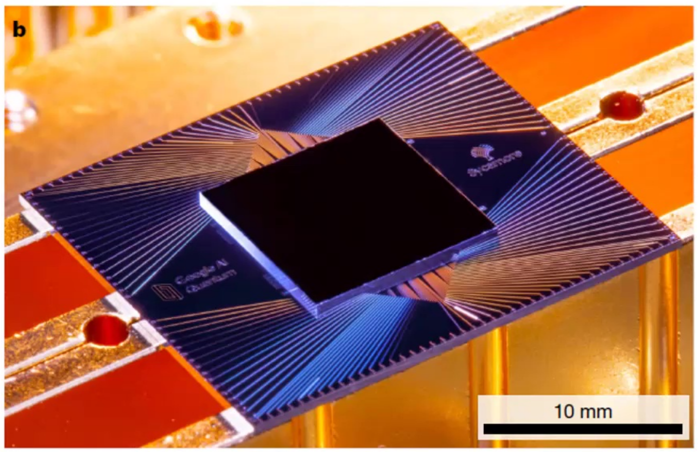

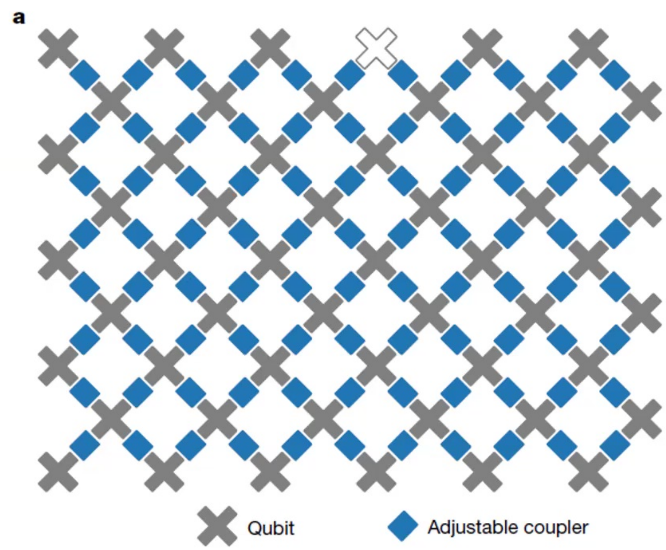

Fig. 1 | The Sycamore processor. a, Layout of processor, showing a rectangular array of 54 qubits (grey), each connected to its four nearest neighbours with couplers (blue). The inoperable qubit is outlined. **b**, Photograph of the Sycamore chip.

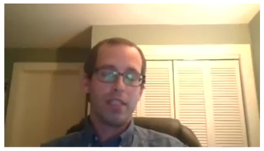

What can we do with **limited or no error correction**, using **short depth circuits** over a **gate set determined by architecture?**

Practical near-term quantum computing?

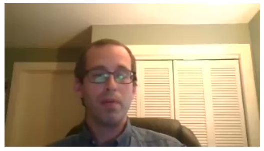

Most demonstrations that have been performed on real-world quantum devices, for "practical sounding problems" are based on heuristic or variational quantum algorithms.

Practical near-term quantum computing?

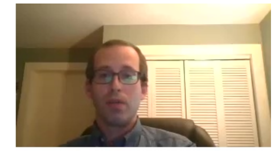

Most demonstrations that have been performed on real-world quantum devices, for “practical sounding problems” are based on heuristic or variational quantum algorithms.

Quantum chemistry

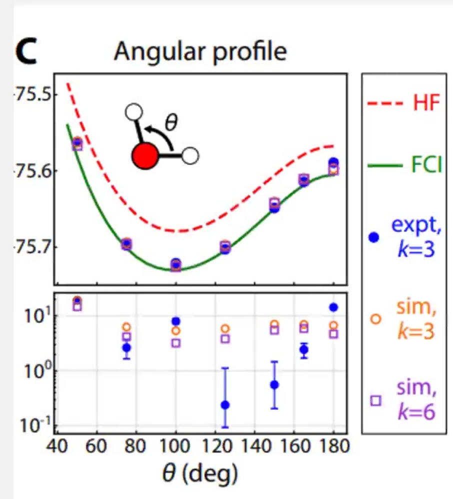

Ground state energy of water molecule
From Eddins et al. (IBM group) arXiv:2104.10220

Practical near-term quantum computing?

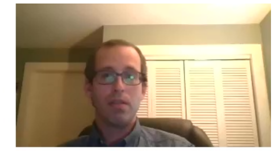

Most demonstrations that have been performed on real-world quantum devices, for "practical sounding problems" are based on heuristic or variational quantum algorithms.

Machine learning

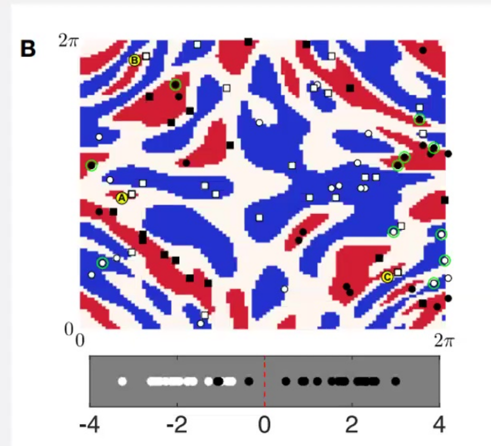

Binary classification using quantum Kernel methods
From Havlicek et al. (IBM group)
Nature **567**, 209-212 (2019).

Practical near-term quantum computing?

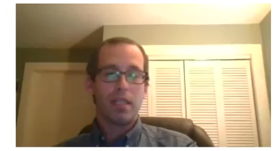

Most demonstrations that have been performed on real-world quantum devices, for "practical sounding problems" are based on heuristic or variational quantum algorithms.

For now these are proof of principle demonstrations. Even if the demonstrations can be scaled up, we don't know if there are quantum speedups for these problems.

Optimization

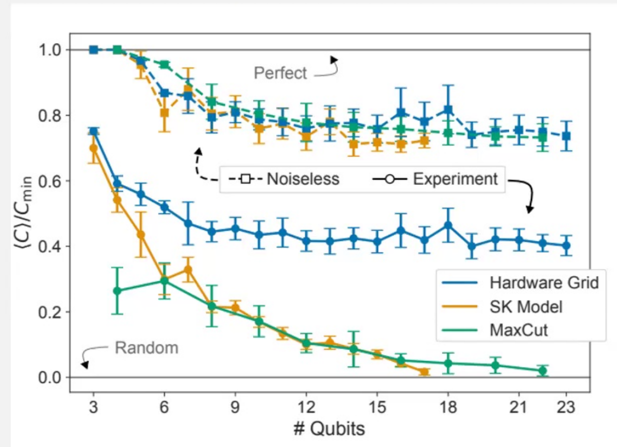

Quantum Approximate Optimization Algorithm
From Harrigan et al. (Google group)
Nature Physics **17**, 332-336 (2021).

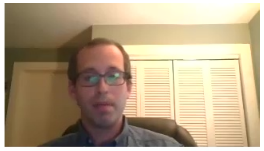

Other experiments have aimed only to convincingly **beat classical computers**, not to perform a useful task…

Beyond classical?

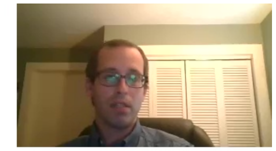

Random Circuit Sampling: simulating a random quantum circuit using a classical computer is believed to be computationally hard…
[Boixo et al. 2017]

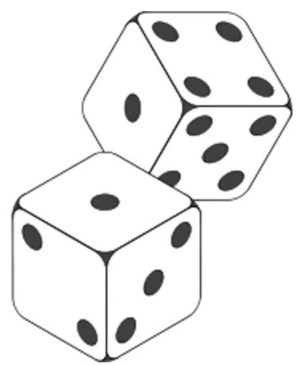

Beyond classical?

"Quantum supremacy using a programmable superconducting p
Arute et al, *Nature* 574.7779 (2019): 505-510

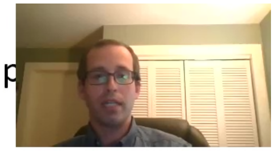

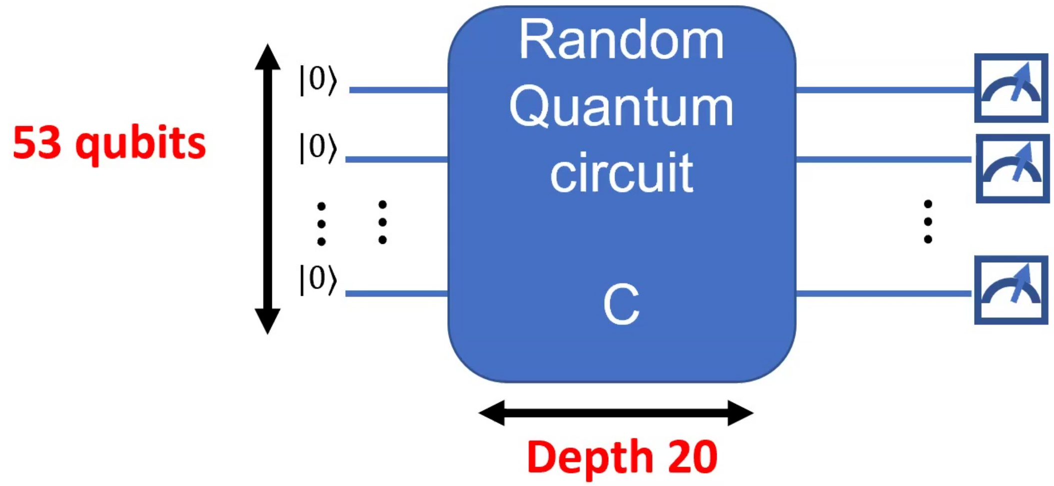

$z = z_1 z_2 \dots z_n$ **Output bits are correlated with *C*.**

Beyond classical?

"Quantum supremacy using a programmable superconducting p
Arute et al, *Nature* 574.7779 (2019): 505-510

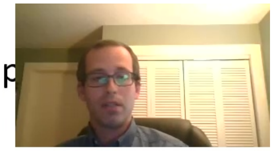

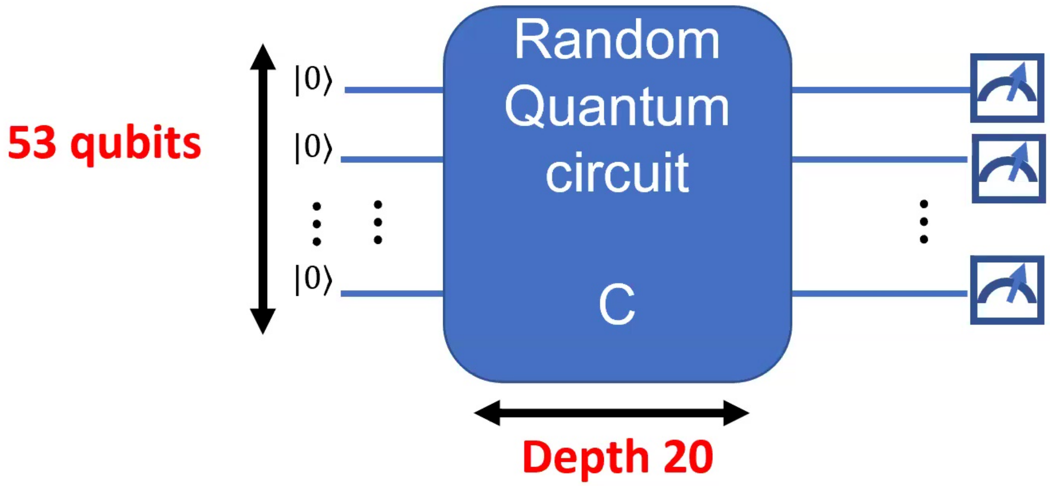

$$z = z_1 z_2 \ldots z_n$$

Output bits are correlated with C.
Correlation is still detectable even with noise in experiment

Theoretical evidence for classical hardness: rests on a **conjecture** that a certain family of complex temperature Ising model partition functions are hard to approximate in the average case.

Beyond classical?

"Quantum supremacy using a programmable superconducting p
Arute et al, *Nature* 574.7779 (2019): 505-510

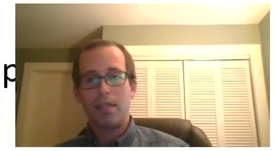

53 qubits

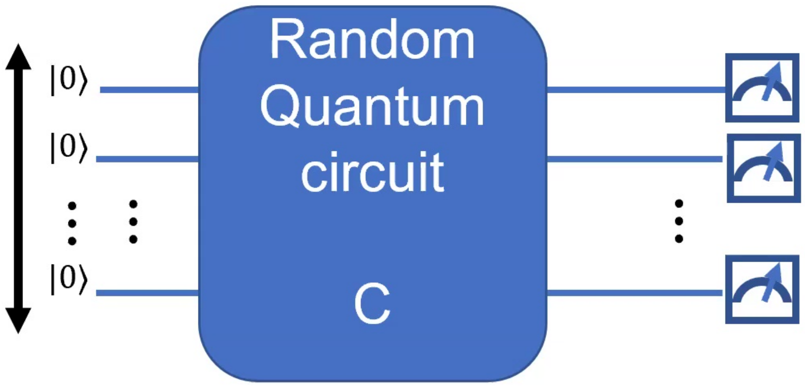

Depth 20

$z = z_1 z_2 \ldots z_n$

Output bits are correlated with C.

Correlation is still detectable even with noise in experiment

Empirical evidence for classical hardness: The best classical algorithm running on current classical computers takes much longer than the quantum experiment.

Beyond classical?

"Quantum supremacy using a programmable superconducting p
Arute et al, *Nature* 574.7779 (2019): 505-510

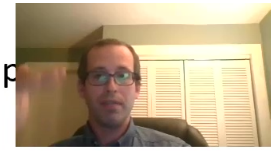

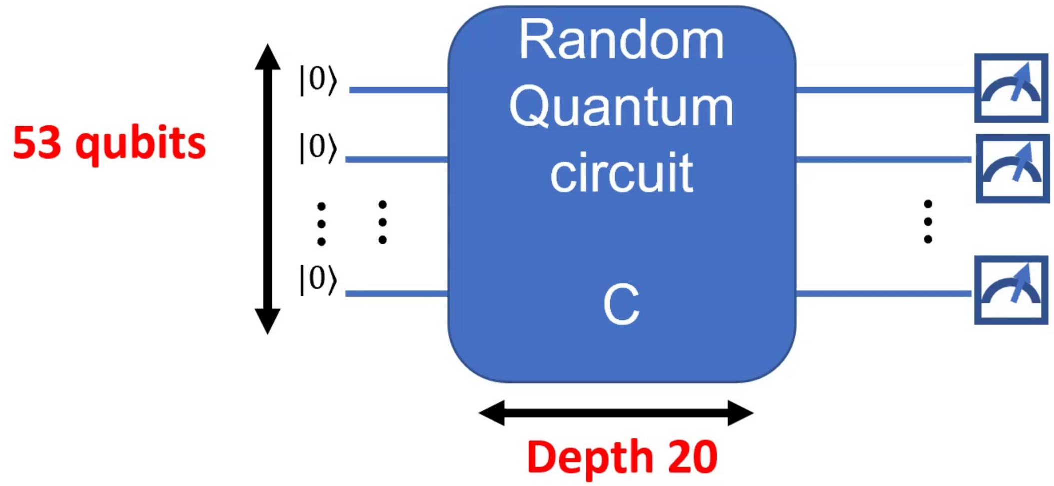

$$z = z_1 z_2 \dots z_n$$

Output bits are correlated with C. Correlation is still detectable even with noise in experiment

Empirical evidence for classical hardness: The best classical algorithm running on current classical computers takes much longer than the quantum experiment.

This is a moving target!

Recent classical progress: classical algorithm spoofs Google's correlation measure, using 5 days and 60 GPUs [Pan Zhang 2021]

Beyond classical?

"Quantum supremacy using a programmable superconducting p
Arute et al, *Nature* 574.7779 (2019): 505-510

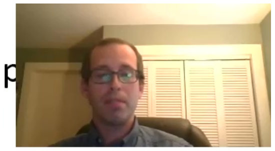

53 qubits

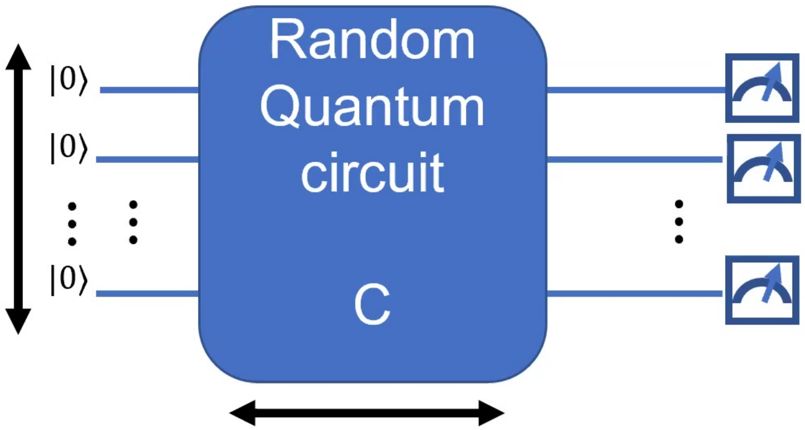

Depth 20

$z = z_1 z_2 \dots z_n$

Output bits are correlated with C.

Correlation is still detectable even with noise in experiment

Empirical evidence for classical hardness: The best classical algorithm running on current classical computers takes much longer than the quantum experiment.

This is a moving target!

Recent quantum progress: "Quantum computational advantage via 60-qubit 24-cycle Random Circuit Sampling" Zhu et al, *arXiv:2109.03494*

Beyond classical?

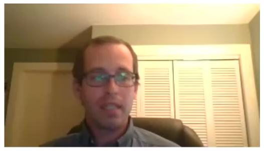

BosonSampling: systems of non-interacting bosons can also be used to sample from classically inaccessible distributions, assuming similar complexity theoretic conjectures. [Aaronson Arkhipov 2011]

Experiment at USTC implemented a variant called Gaussian Boson Sampling using a 100 mode linear optical network. Classical simulability is subject of debate.

"Quantum computational advantage using photons"
Zhong et al, *Science* 370.6523 (2020): 1460-1463

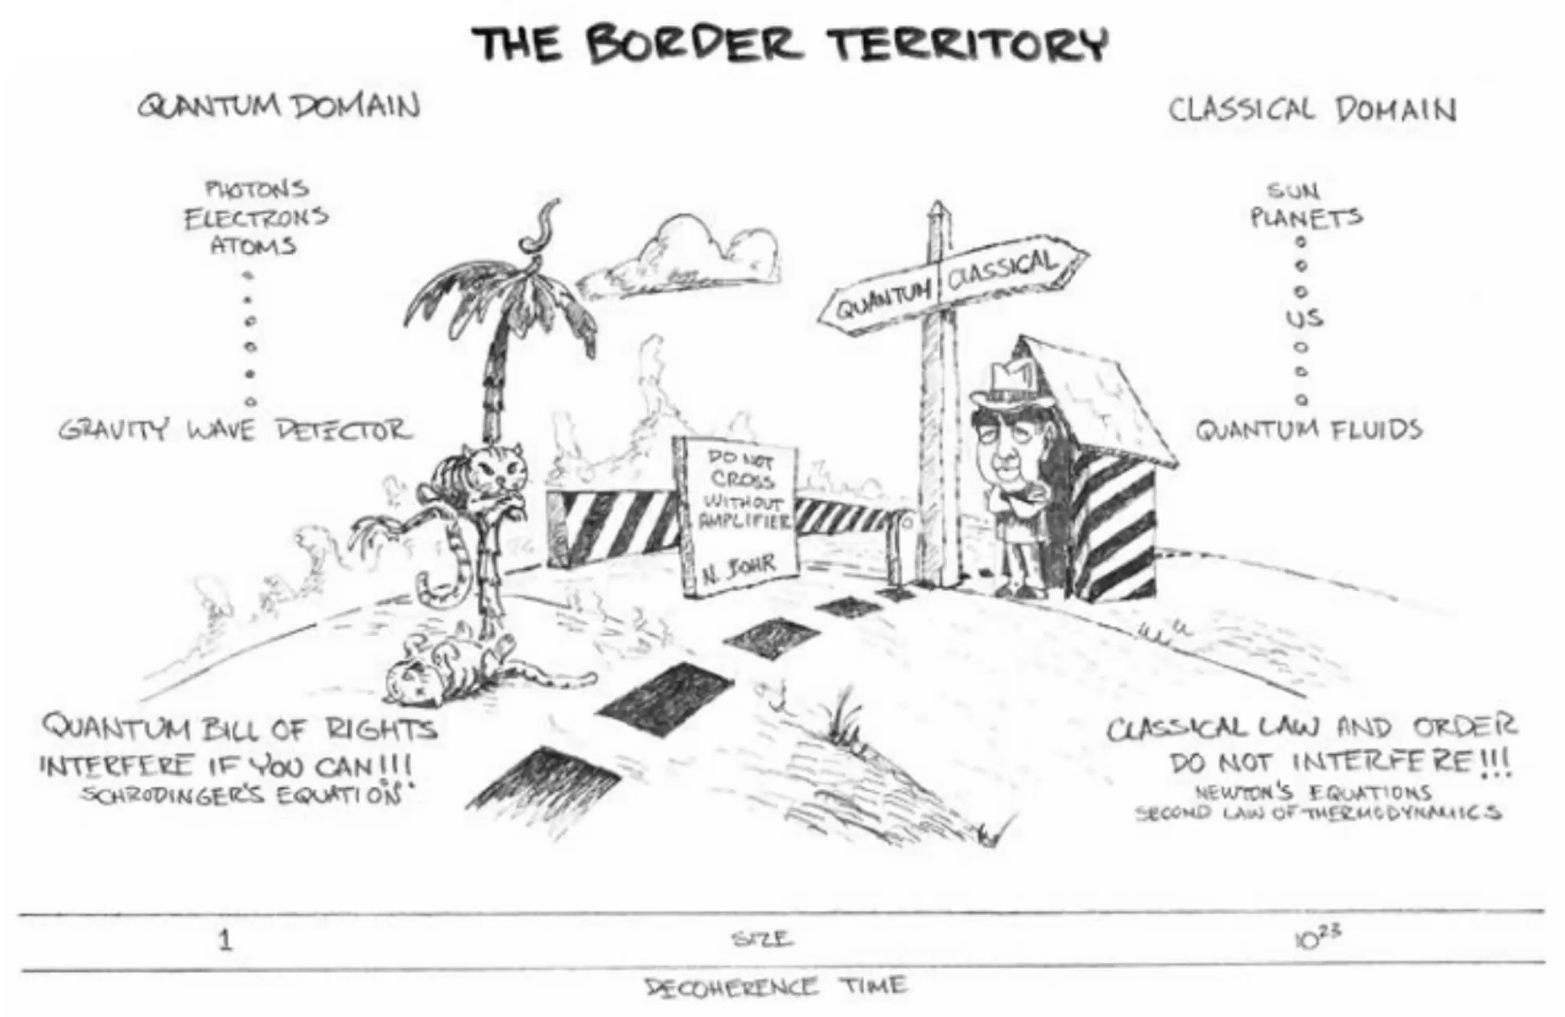

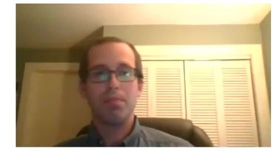

Image source: Zurek W.H. (2006) "Decoherence and the Transition from Quantum to Classical — Revisited"

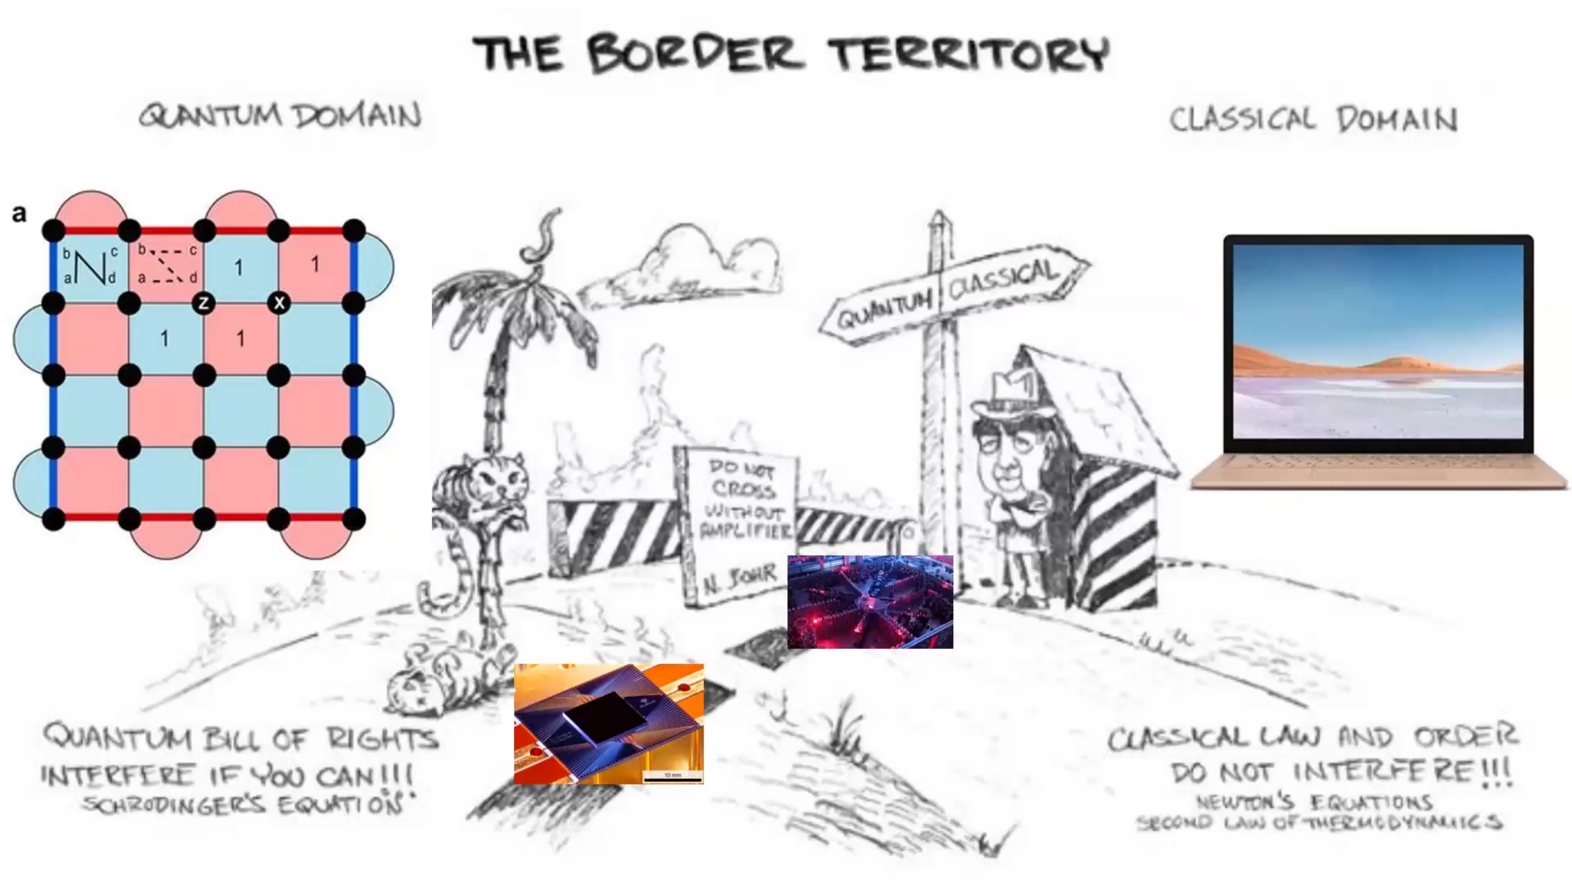

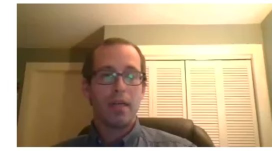

Image source: Zurek W.H. (2006) "Decoherence and the Transition from Quantum to Classical — Revisited"

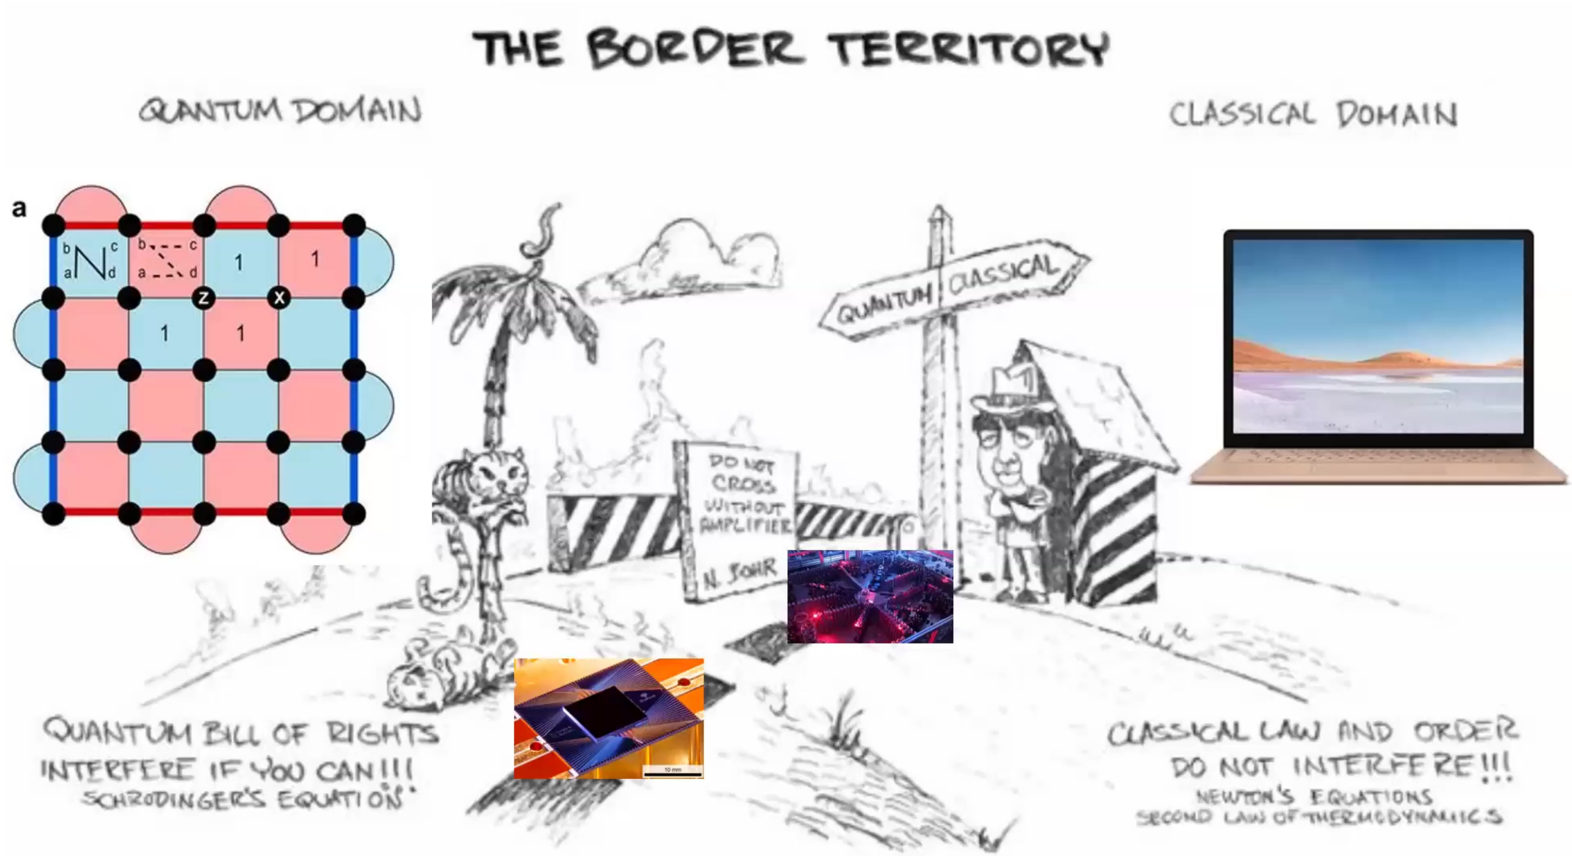

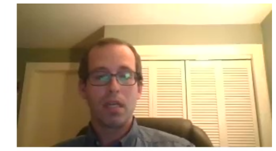

Which restricted forms of quantum computation can be more powerful than classical computers? Which are classically simulable?

Image source: Zurek W.H. (2006) "Decoherence and the Transition from Quantum to Classical — Revisited"

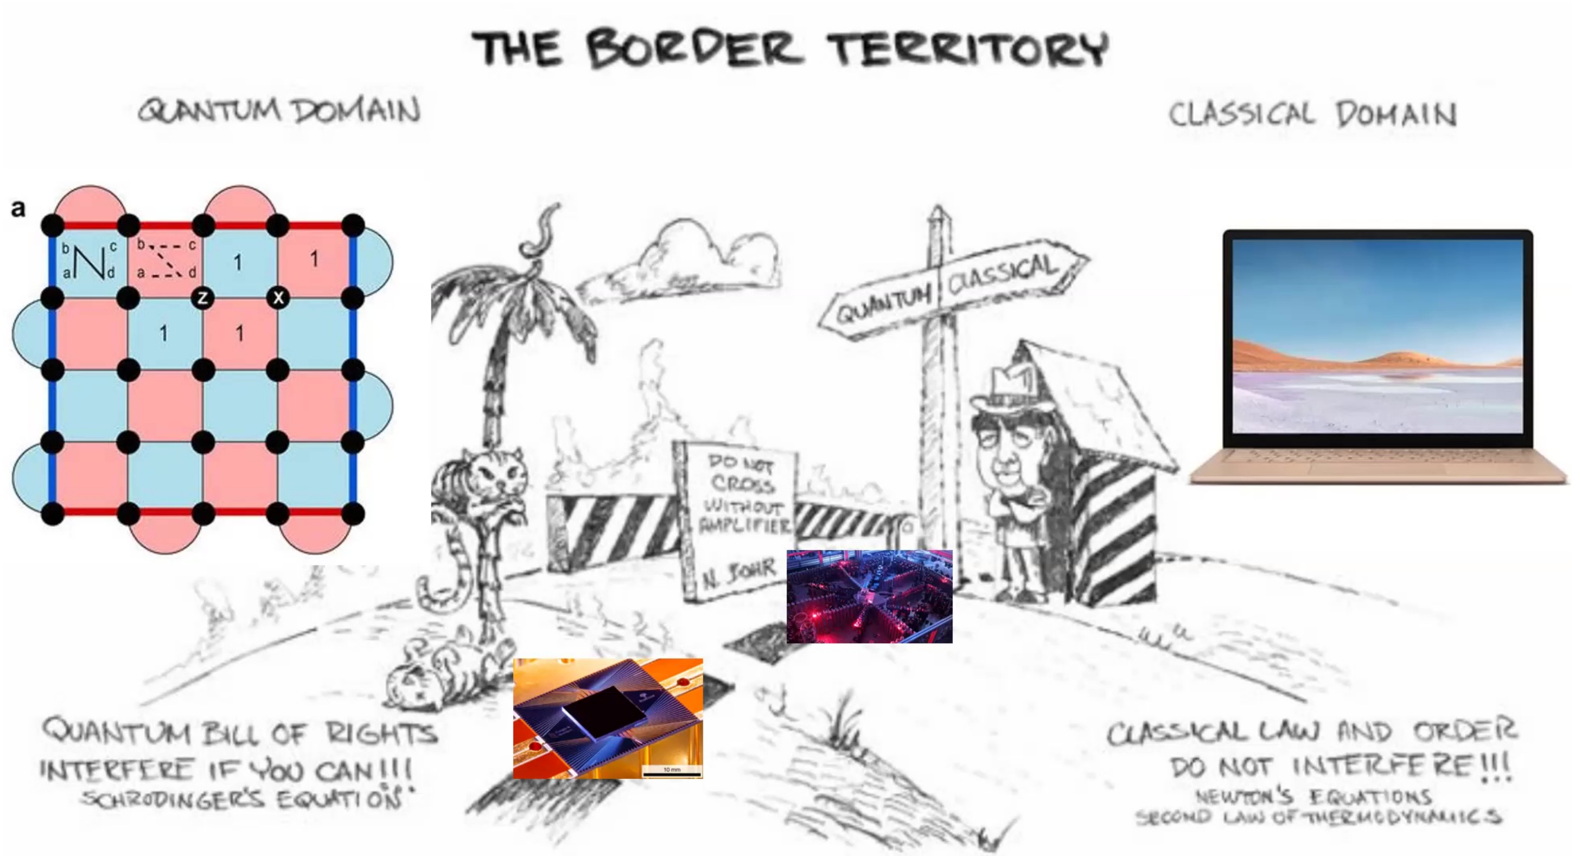

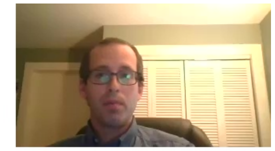

In this talk I will tell you about a new kind of quantum advantage with **shallow quantum circuits.**

Image source: Zurek W.H. (2006) “Decoherence and the Transition from Quantum to Classical — Revisited”

Shallow quantum circuits

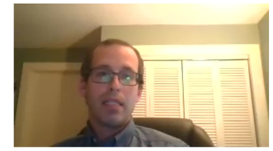

Circuit depth is the number of time steps allowing for parallel gates.

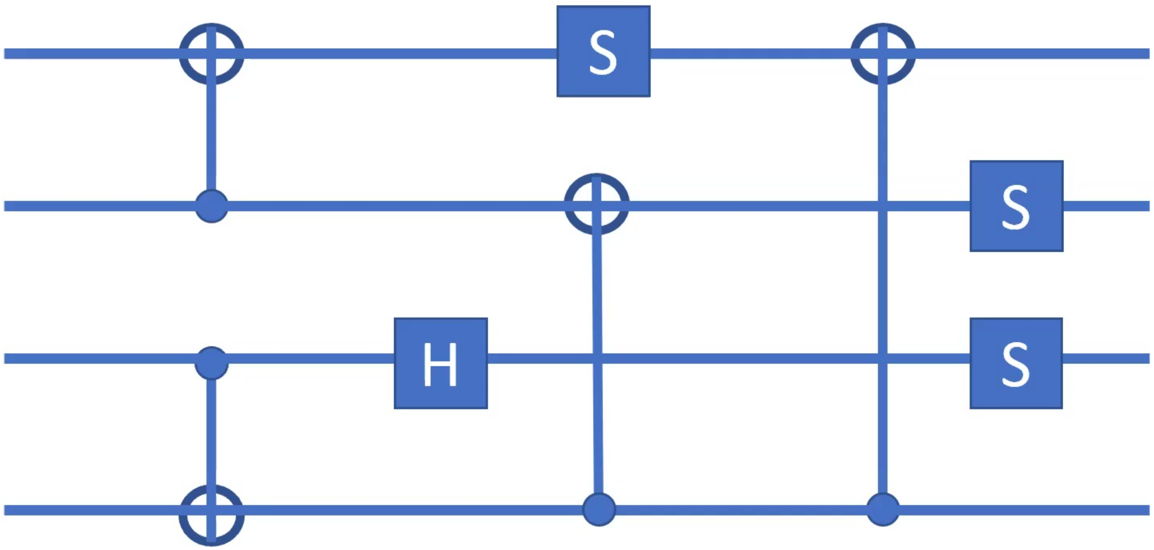

Shallow quantum circuits

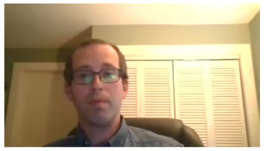

Circuit depth is the number of time steps allowing for parallel gates.

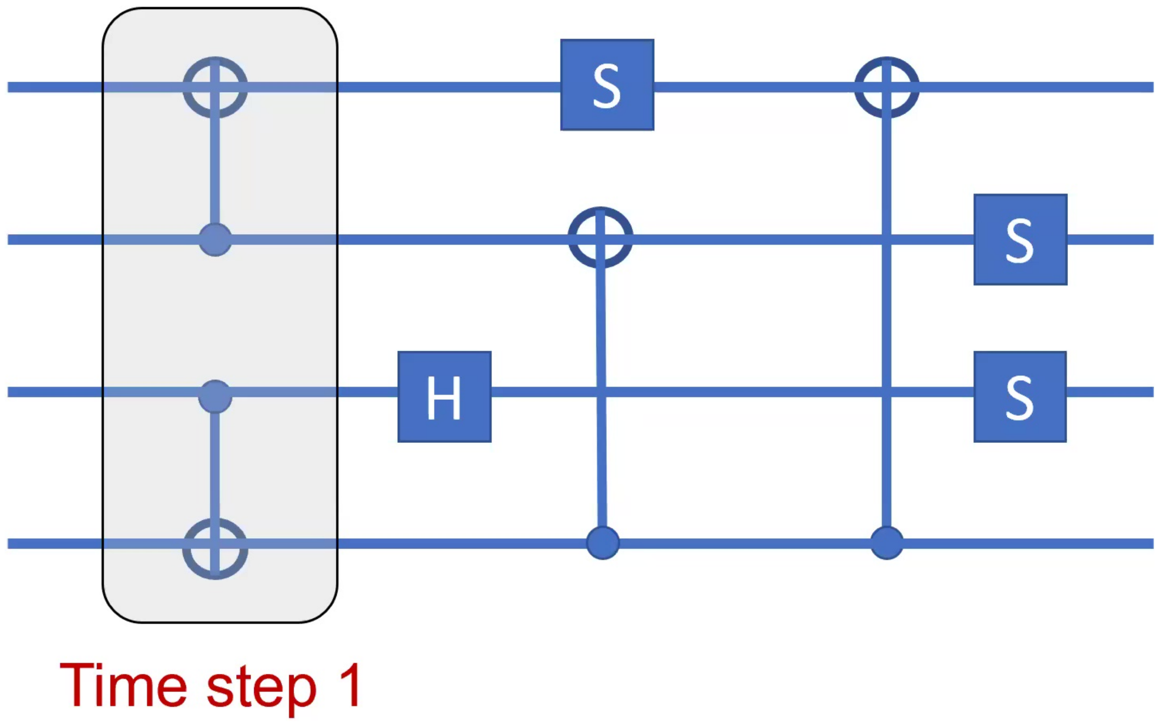

Shallow quantum circuits

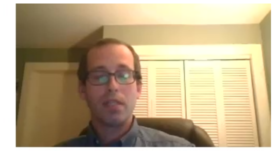

Circuit depth is the number of time steps allowing for parallel gates.

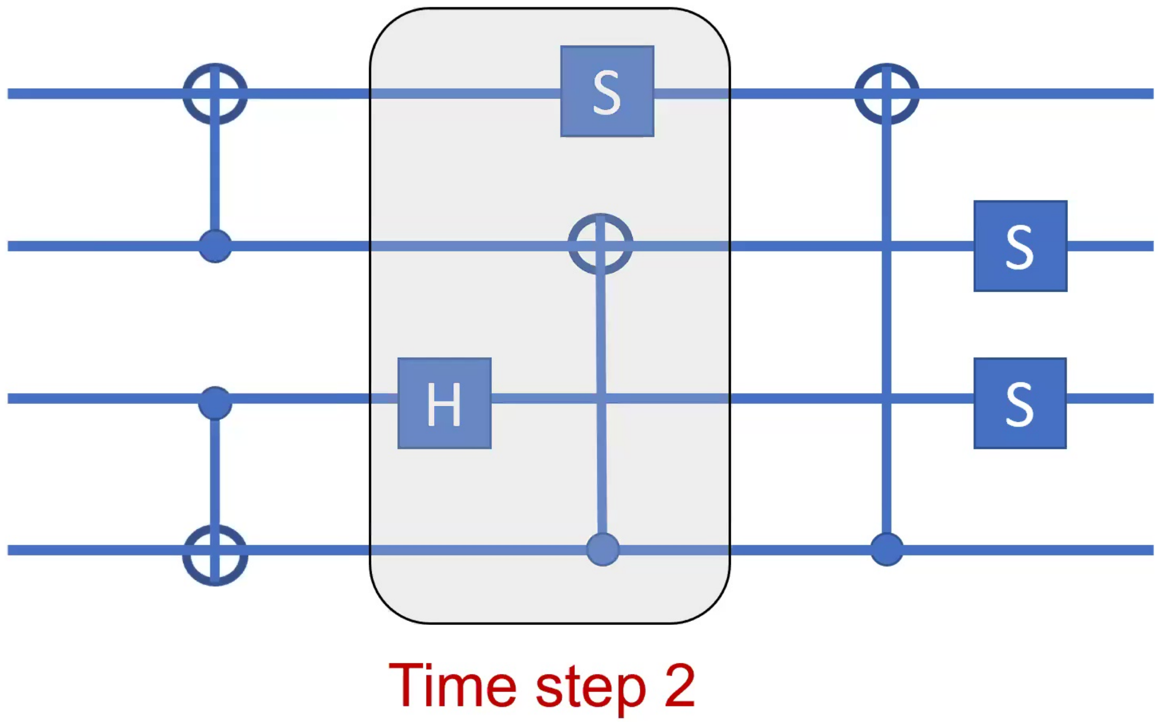

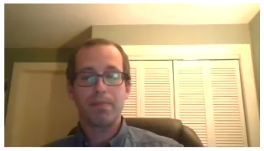

Shallow quantum circuits

Circuit depth is the number of time steps allowing for parallel gates.

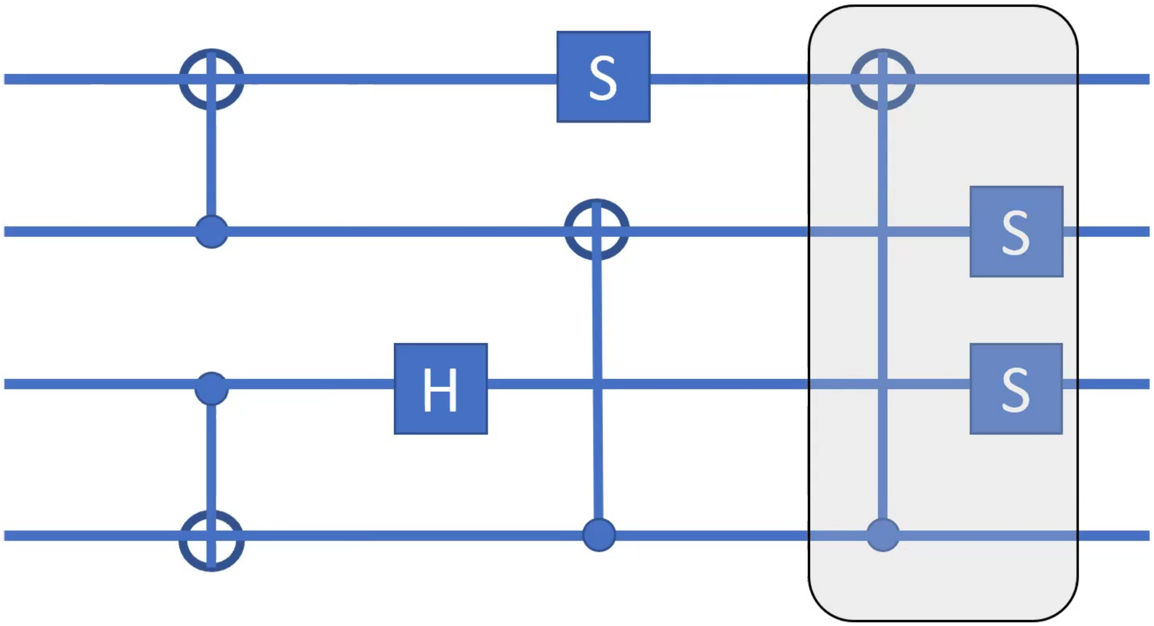

Time step 3

Shallow quantum circuits

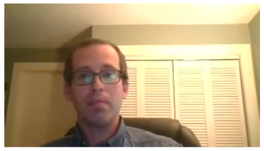

Circuit depth is the number of time steps allowing for parallel gates.

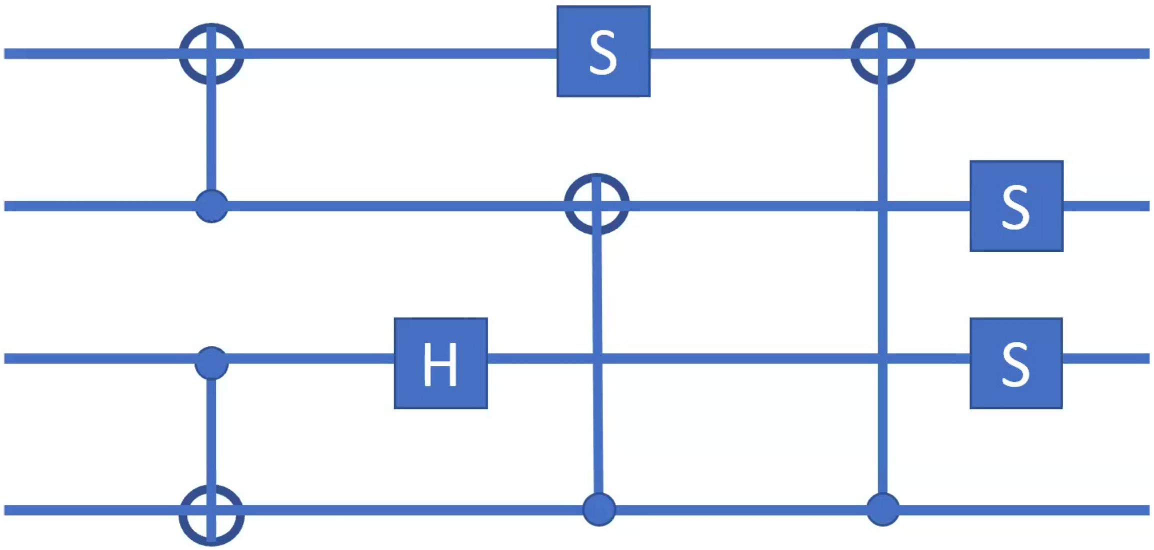

We are interested in quantum circuits with **constant depth**.

Can shallow quantum circuits beat classical computers?

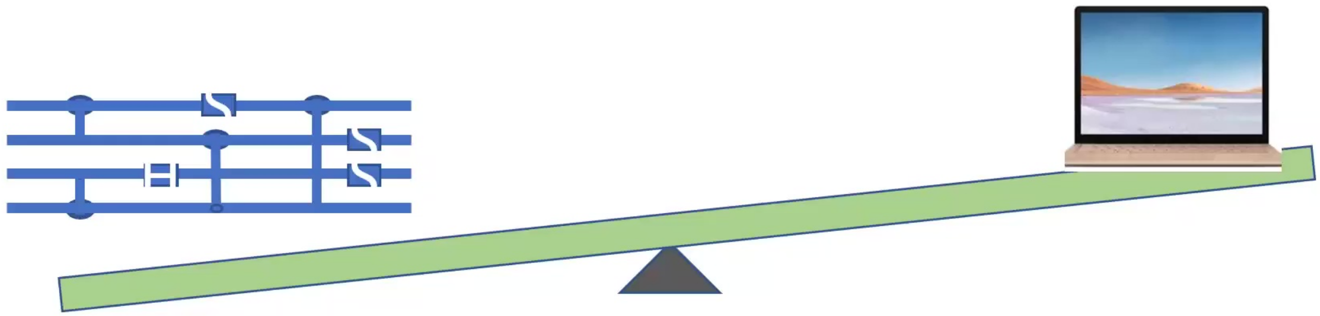

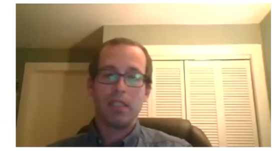

Can shallow quantum circuits beat classical computers?

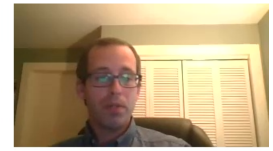

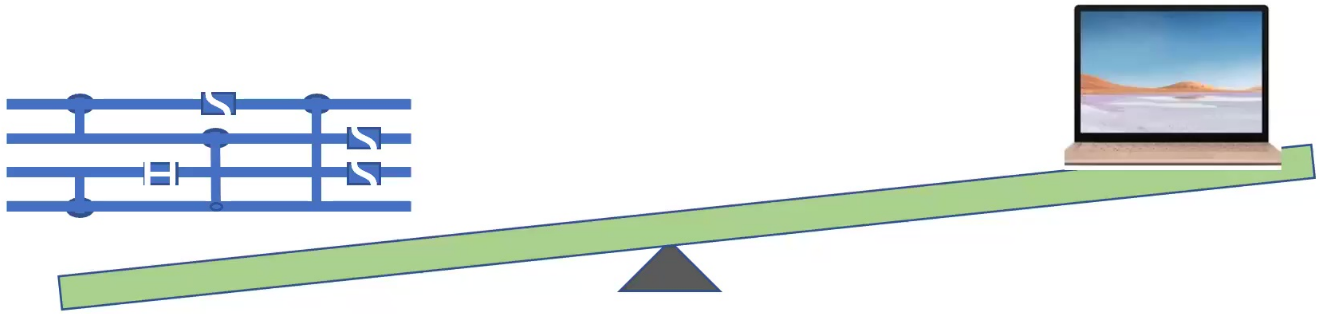

Big question: Can constant-depth quantum circuits perform a task that polynomial time classical computers can't? **We believe they can…**

[Terhal Divincenzo 2002][Gao et al 17]
[Bermejo-Vega et al. 17]

Smaller question: Can constant-depth quantum circuits solve a problem that constant-depth classical circuits can't?

Can shallow quantum circuits beat classical computers?

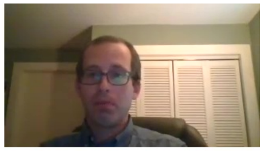

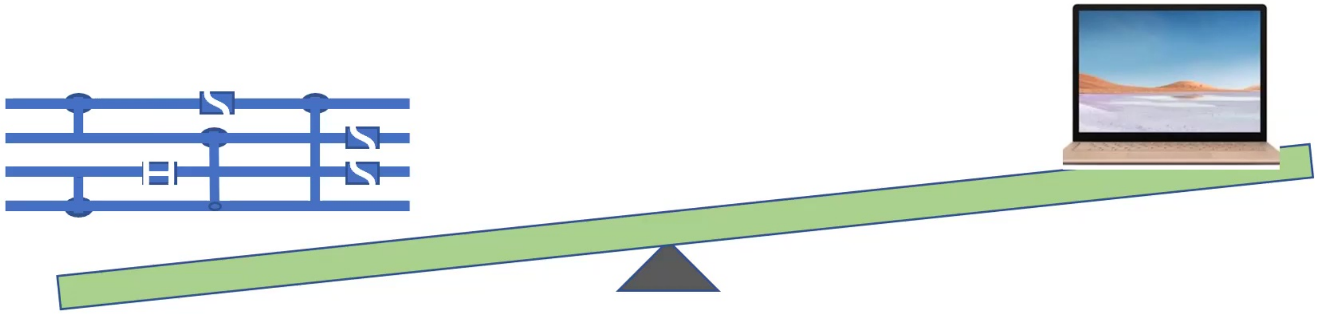

Big question: Can constant-depth quantum circuits perform a task that polynomial time classical computers can't? **We believe they can…**

[Terhal Divincenzo 2002][Gao et al 17]
[Bermejo-Vega et al. 17]

Smaller question: Can constant-depth quantum circuits solve a problem that constant-depth classical circuits can't? **YES…**

Shallow quantum beats shallow classical

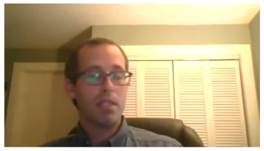

Plan for the rest of the talk:

Bravyi, DG, Koenig. *Science* 362 (6412), 2018.

Shallow quantum circuits can solve a linear algebra problem that provably can't be solved by shallow classical circuits.

Bravyi, DG, Koenig, Tomamichel. *Nature Physics 1-6*, 2020.

Noisy shallow quantum circuits can solve a linear algebra problem that provably can't be solved by shallow classical circuits.

Shallow quantum beats shallow classical

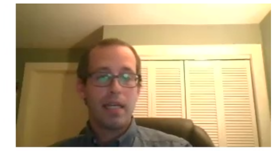

Plan for the rest of the talk:

Bravyi, DG, Koenig. *Science* 362 (6412), 2018.

Shallow quantum circuits can solve a linear algebra problem that provably can't be solved by shallow classical circuits.

Bravyi, DG, Koenig, Tomamichel. *Nature Physics 1-6*, 2020.

Noisy shallow quantum circuits can solve a linear algebra problem that provably can't be solved by shallow classical circuits.

The advantage is not very dramatic: log versus constant depth.

That said, it provides a new kind of unconditional evidence that quantum computers are more powerful than classical ones…

Evidence for quantum advantage in computat

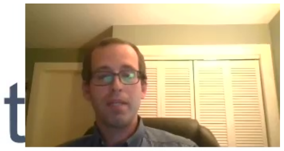

Quantum algorithms with speedups over classical

Shor's algorithm

Simulation of Hamiltonian dynamics

$$i\hbar \frac{d|\psi\rangle}{dt} = H|\psi\rangle$$

Sampling from classically hard distributions

Boson sampling

IQP circuits

Random quantum circuits

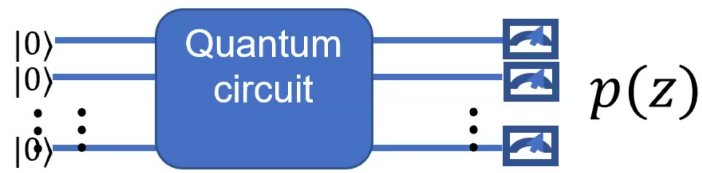

Provable speedups relative to an oracle

Bernstein-Vazirani

Simon's problem

Quantum advantage with shallow circuits [This talk]

A modest, but **provable and non-oracular** quantum speedup attained by constant-depth quantum circuits in a 2D architecture.

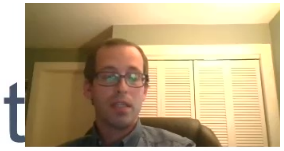

Evidence for quantum advantage in computat

Quantum algorithms w

Shor's algorithm

Simulation of Hamilton

> These speedups disappear if the classical algorithms can be improved

$H|\psi\rangle$

Sampling fro

Boson samp

IQP circuits

Random quantum circuits

> Assumes complexity-theoretic and other conjectures.

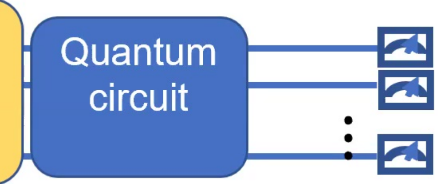

$p(z)$

Provable speedups

Bernstein-Vazirani

Simon's problem

> Oracles do not exist in the real world.

$(-1)^{f(x)}|x\rangle$

Quantum advantage with shallow circuits [This talk]

A modest, but **provable and non-oracular** quantum speedup attained by constant-depth quantum circuits in a 2D architecture.

Shallow quantum beats shallow classical

Bravyi, DG, Koenig. *Science* 362 (6412), 2018.

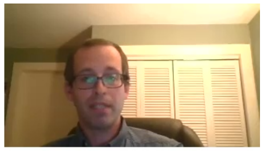

The Hidden Linear Function problem

We will define the problem in two equivalent ways

Simulate measurements
on a quantum **graph state**

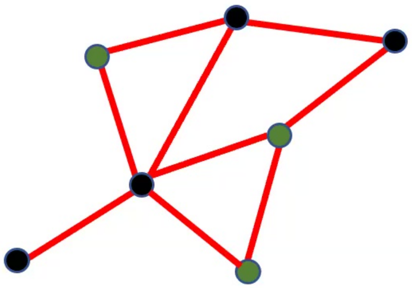

Compute a certain property
of a **quadratic form**
$q: \{0,1\}^N \to \mathbb{Z}_4$

$$q(x) = x^T A x \mod 4$$

The Hidden Linear Function problem

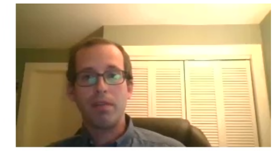

An instance of the problem is defined by a symmetric $n \times n$ binary matrix A

The off-diagonal part of A defines a graph

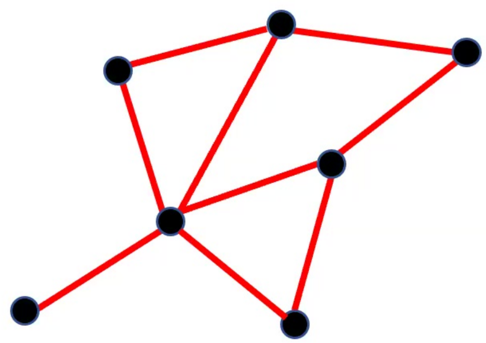

The Hidden Linear Function problem

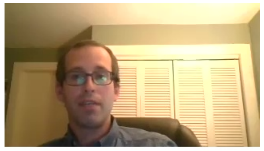

An instance of the problem is defined by a symmetric $n \times n$ binary matrix A

The off-diagonal part of A defines a graph

The diagonal part of A specifies a subset of marked vertices

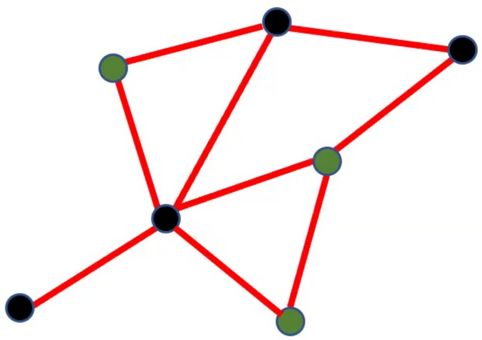

The Hidden Linear Function problem

Consider the corresponding quantum **graph state**

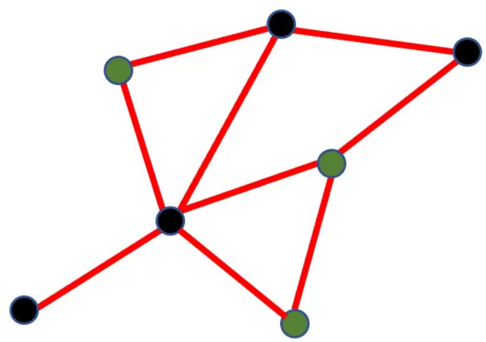

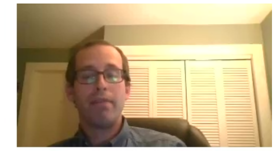

To prepare the quantum graph state:

Place a qubit at each vertex in $|0\rangle$ state

Apply single qubit H gates to all qubits

Apply two-qubit CZ gate on each edge

The Hidden Linear Function problem

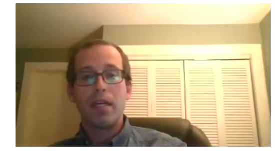

Consider the corresponding quantum **graph state**

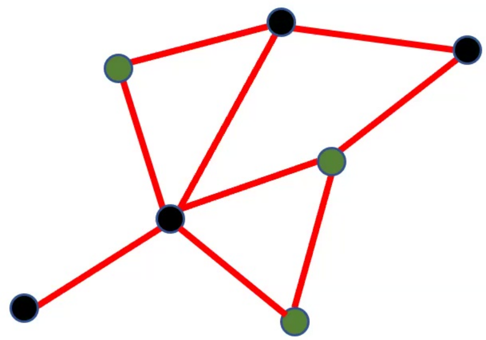

To prepare the quantum graph state:
Place a qubit at each vertex in $|0\rangle$ state
Apply single qubit H gates to all qubits
Apply two-qubit CZ gate on each edge

Imagine measuring each qubit in either the Pauli X basis (if vertex is unmarked)
or Pauli Y basis (if marked)

The Hidden Linear Function problem

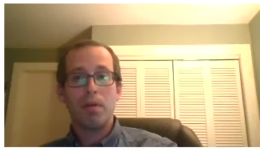

Consider the corresponding quantum **graph state**

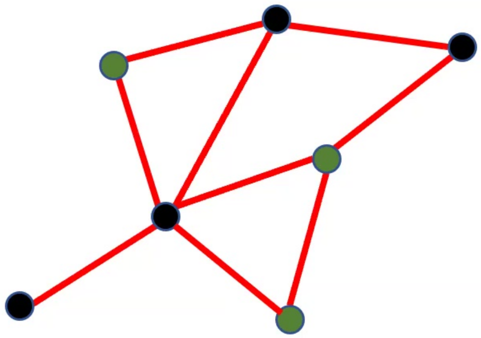

To prepare the quantum graph state:
Place a qubit at each vertex in $|0\rangle$ state
Apply single qubit H gates to all qubits
Apply two-qubit CZ gate on each edge

Imagine measuring each qubit in either the Pauli X basis (if vertex is unmarked)
or Pauli Y basis (if marked)

HLF problem: Given a graph and subset of marked vertices, output any measurement outcome $x \in \{0,1\}^n$ that occurs with nonzero probability in this experiment.

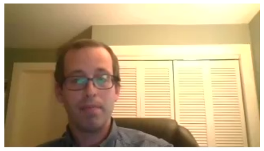

The Hidden Linear Function problem

Consider the corresponding quantum **graph state**

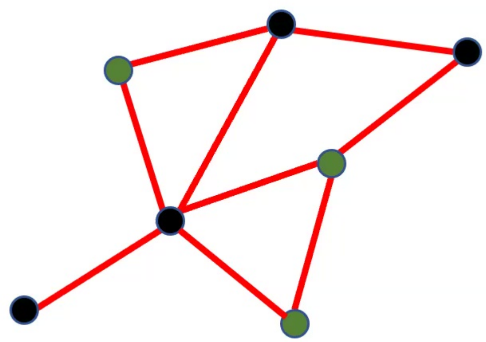

To prepare the quantum graph state:
Place a qubit at each vertex in $|0\rangle$ state
Apply single qubit H gates to all qubits
Apply two-qubit CZ gate on each edge

Imagine measuring each qubit in either the Pauli X basis (if vertex is unmarked)
or Pauli Y basis (if marked)

HLF problem: Given a graph and subset of marked vertices, output any measurement outcome $x \in \{0,1\}^n$ that occurs with nonzero probability in this experiment.

We have defined the problem in terms of a simple quantum circuit that solves it.

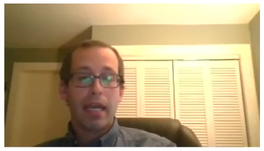

The Hidden Linear Function problem

Consider the corresponding quantum **graph state**

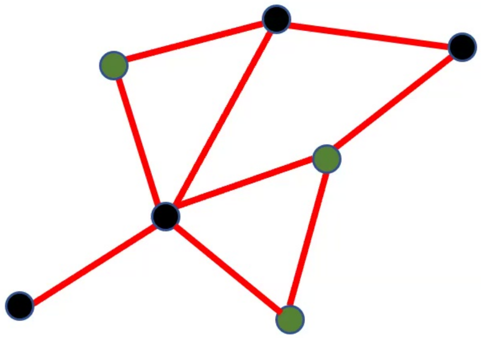

To prepare the quantum graph state:
Place a qubit at each vertex in $|0\rangle$ state
Apply single qubit H gates to all qubits
Apply two-qubit CZ gate on each edge

Imagine measuring each qubit in either the Pauli X basis (if vertex is unmarked)
or Pauli Y basis (if marked)

HLF problem: Given a graph and subset of marked vertices, output any measurement outcome $x \in \{0,1\}^n$ that occurs with nonzero probability in this experiment.

The quantum circuit has constant depth if the graph is a subgraph of a 2D grid

The Hidden Linear Function problem

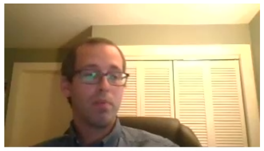

Consider the corresponding quantum **graph state**

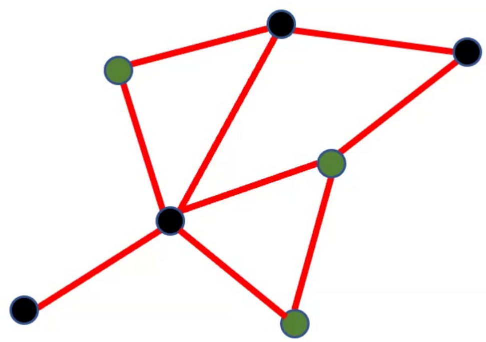

To prepare the quantum graph state:
Place a qubit at each vertex in $|0\rangle$ state
Apply single qubit H gates to all qubits
Apply two-qubit CZ gate on each edge

Imagine measuring each qubit in either the Pauli X basis (if vertex is unmarked)
or Pauli Y basis (if marked)

2D HLF problem: Given a subgraph of $\sqrt{n} \times \sqrt{n}$ grid and subset of marked vertices, output any measurement outcome $x \in \{0,1\}^n$ that occurs with nonzero probability in this experiment.

The quantum circuit has constant depth if the graph is a subgraph of a 2D grid

Constant depth quantum circuit that solves 2D HLF

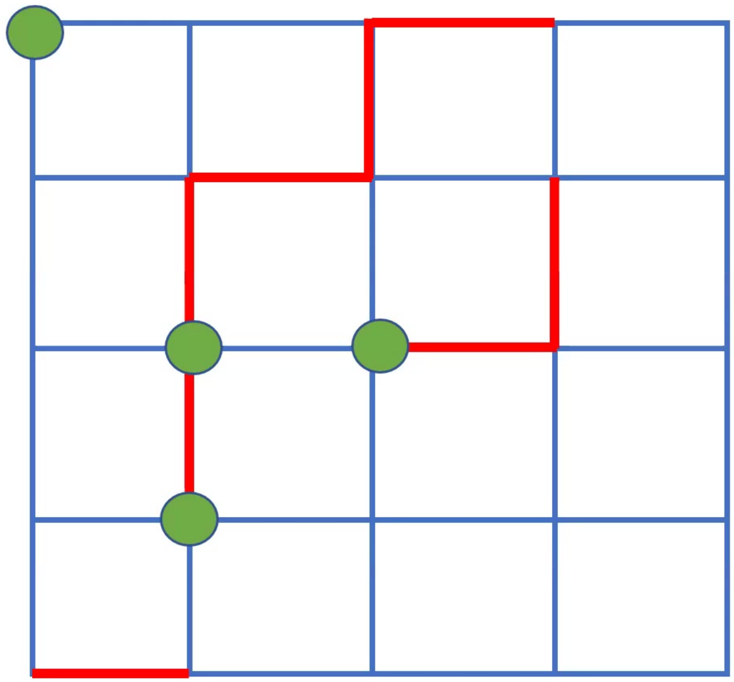

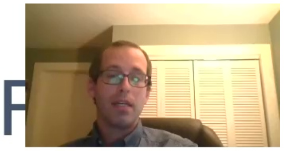

Place a qubit at each vertex
Place input bits on vertices and edges:

v —— w : edge of graph

● : marked vertex
v

Constant depth quantum circuit that solves 2D HLF

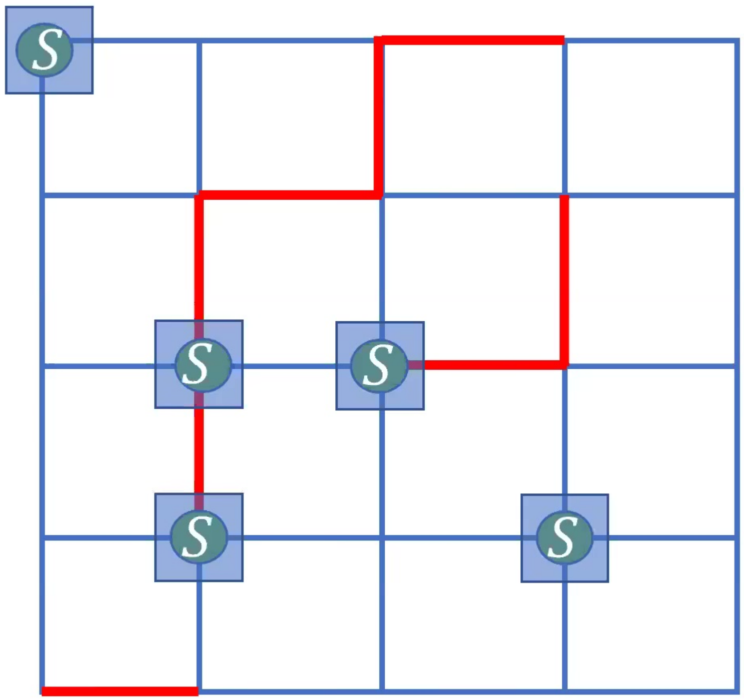

v —— w : edge of graph

● : marked vertex
v

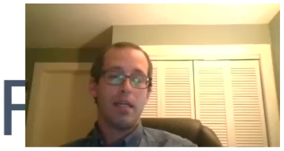

Constant depth quantum circuit that solves 2D HLF

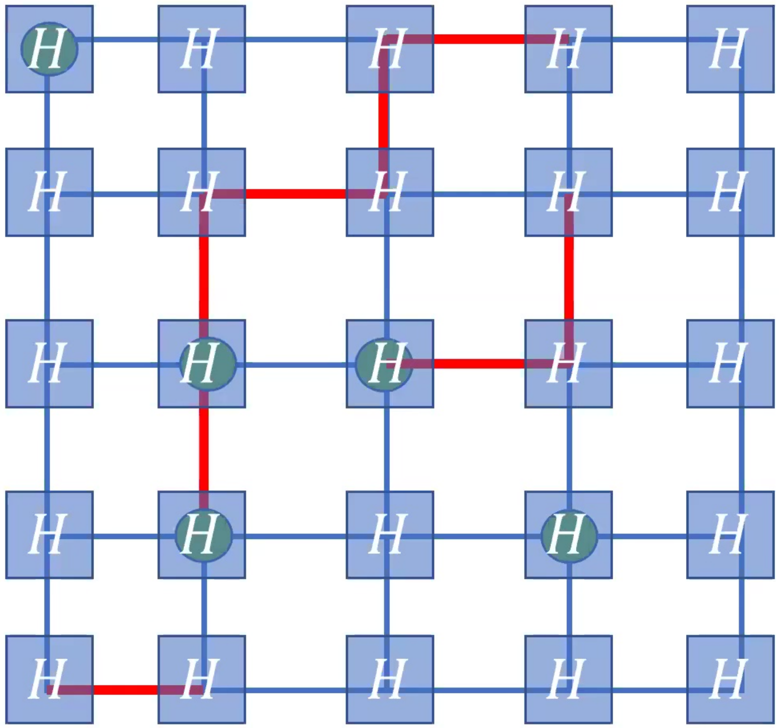

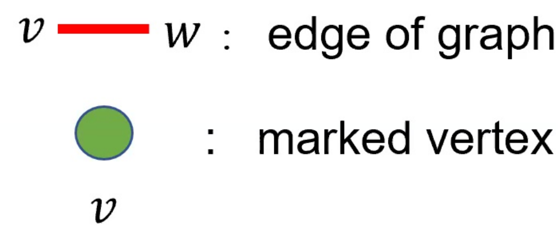

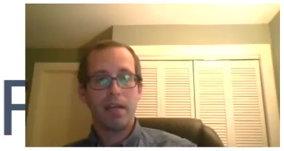

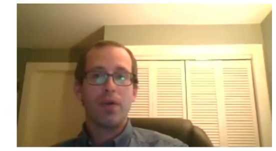

So the quantum circuit that solves the 2D HLF problem is **constant depth** and all gates act between **nearest neighbor qubits in a 2D geometry**.

It also has another very special feature….

The quantum circuit is Clifford

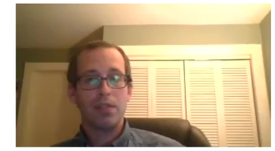

It's a **Clifford circuit:** built from 1- and 2-qubit gates

$$H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \qquad S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix} \qquad CZ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}$$

The quantum circuit is Clifford

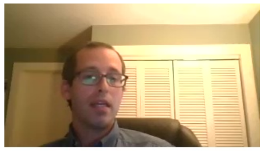

It's a **Clifford circuit:** built from 1- and 2-qubit gates

$$H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \qquad S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix} \qquad CZ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}$$

Clifford circuits are not powerful enough to implement most quantum algorithms.
They are special because…

Gottesman-Knill Theorem [Gottesman 1997]
Quantum circuits composed only of Clifford gates can be **efficiently** simulated on a classical computer using linear algebra.

From Clifford circuits to quadratic forms

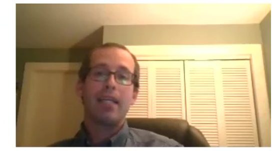

The quantum states prepared by such Clifford circuits ("stabilizer states") are associated with quadratic forms

$$|\psi\rangle \propto \sum_{x \in V} (-1)^{Q(x)} i^{\ell(x)} |x\rangle$$

$\boldsymbol{V}$: affine subspace of $\mathbb{F}_2^n$

$\boldsymbol{q}$: quadratic function $\mathrm{Q}(x) = x^T B x \mod 2$

$\boldsymbol{\ell}$: linear function $\ell(x) = d^T x \mod 2$

From Clifford circuits to quadratic forms

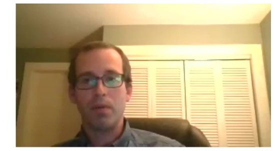

The quantum states prepared by such Clifford circuits ("stabilizer states") are associated with quadratic forms

$$|\psi\rangle \propto \sum_{x \in V} (-1)^{Q(x)} i^{\ell(x)} |x\rangle$$

V: affine subspace of $\mathbb{F}_2^n$

q: quadratic function $\mathrm{Q}(x) = x^T B x \mod 2$

ℓ: linear function $\ell(x) = d^T x \mod 2$

Using this connection we get an equivalent definition of the Hidden Linear Function problem based on quadratic forms…

HLF as a nonstandard linear algebra problem

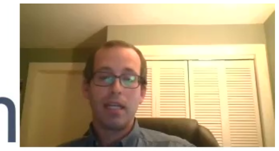

A Symmetric $n \times n$ binary matrix

$$\ker(A) = \{x : Ax = 0 \bmod 2\}$$

$$q(x) = x^T A x \mod 4$$

Fact: There is a secret bit string z such that

$$q(x) = 2z^T x \qquad x \in \ker(\mathrm{A})$$

The hidden linear function!

HLF as a nonstandard linear algebra problem

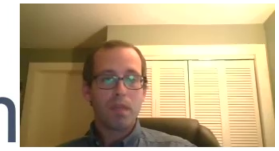

A Symmetric $n \times n$ binary matrix

$$\ker(A) = \{x : Ax = 0 \bmod 2\}$$

$$q(x) = x^T A x \mod 4$$

Hidden Linear Function problem: Given A, find a secret bit string z such that

$$q(x) = 2z^T x \qquad x \in \ker(\mathrm{A})$$

Now we can directly see how to solve the problem in polynomial time on a classical computer using linear algebra…

HLF as a nonstandard linear algebra problem

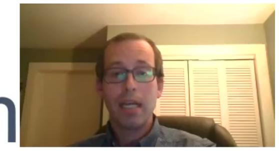

A Symmetric $n \times n$ binary matrix

$$\ker(A) = \{x : Ax = 0 \bmod 2\}$$

$$q(x) = x^T A x \mod 4$$

Hidden Linear Function problem: Given A, find a secret bit string z such that

$$q(x) = 2z^T x \qquad x \in \ker(\mathrm{A})$$

Efficient classical algorithm:

Solve for a basis $x_1, x_2, \ldots x_k$ of $\ker(A)$

Solve for z in system of linear equations $q(x_i) = 2z^T x_i$

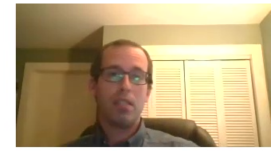

So far we have defined an unusual linear algebra problem, the 2D HLF problem, that can be solved in constant depth by a quantum computer.

On the other hand we show that it can’t be solved in constant depth by a classical computer…

Shallow classical circuits

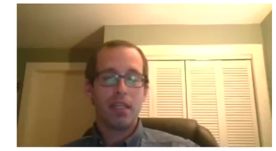

What family of shallow classical circuits is fair to compare against?

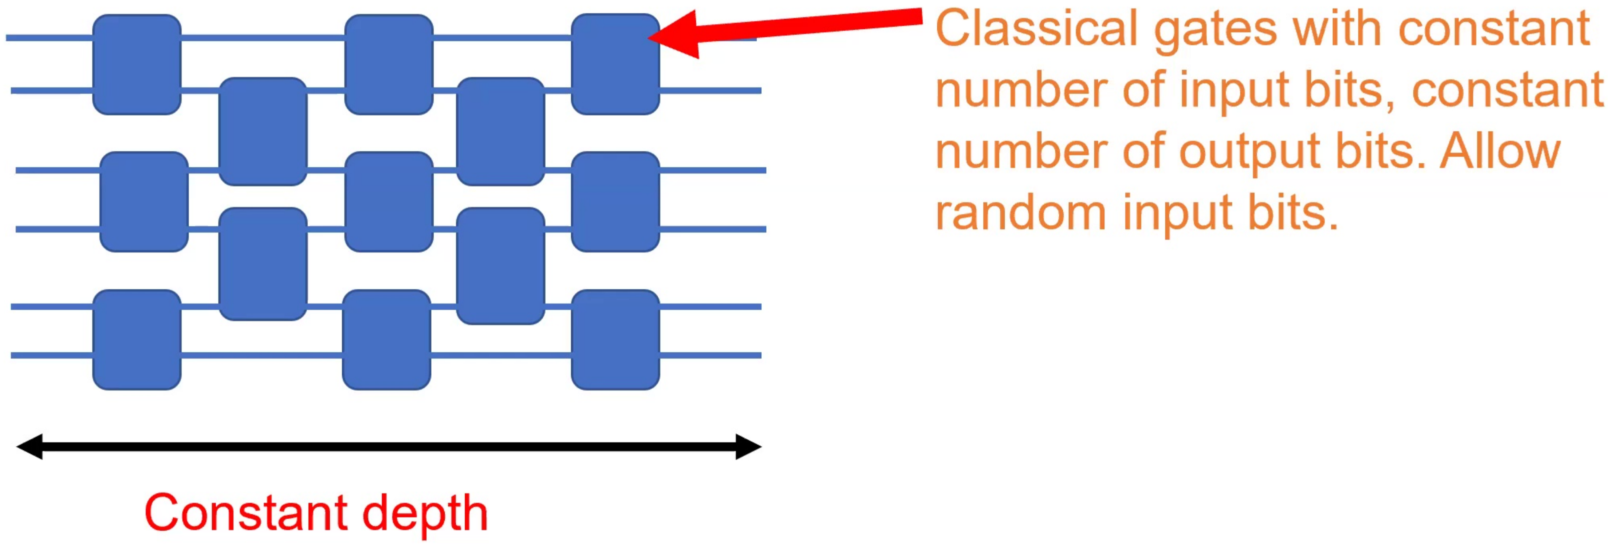

Bounded fan-in gates

We only require the gates to have **bounded fan-in** (number of inputs)

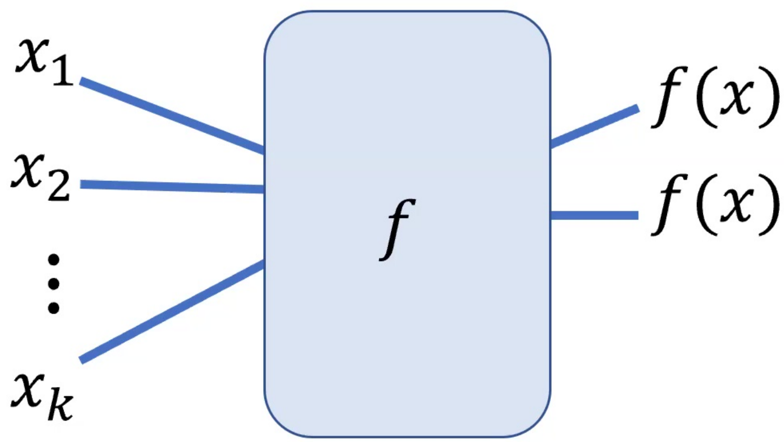

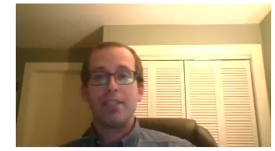

Classical circuits require increasing depth

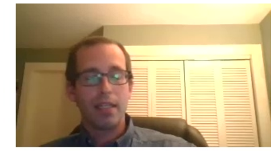

Any classical probabilistic circuit with bounded fan-in gates that solves the 2D HLF problem with high probability has a depth that increases at least logarithmically.

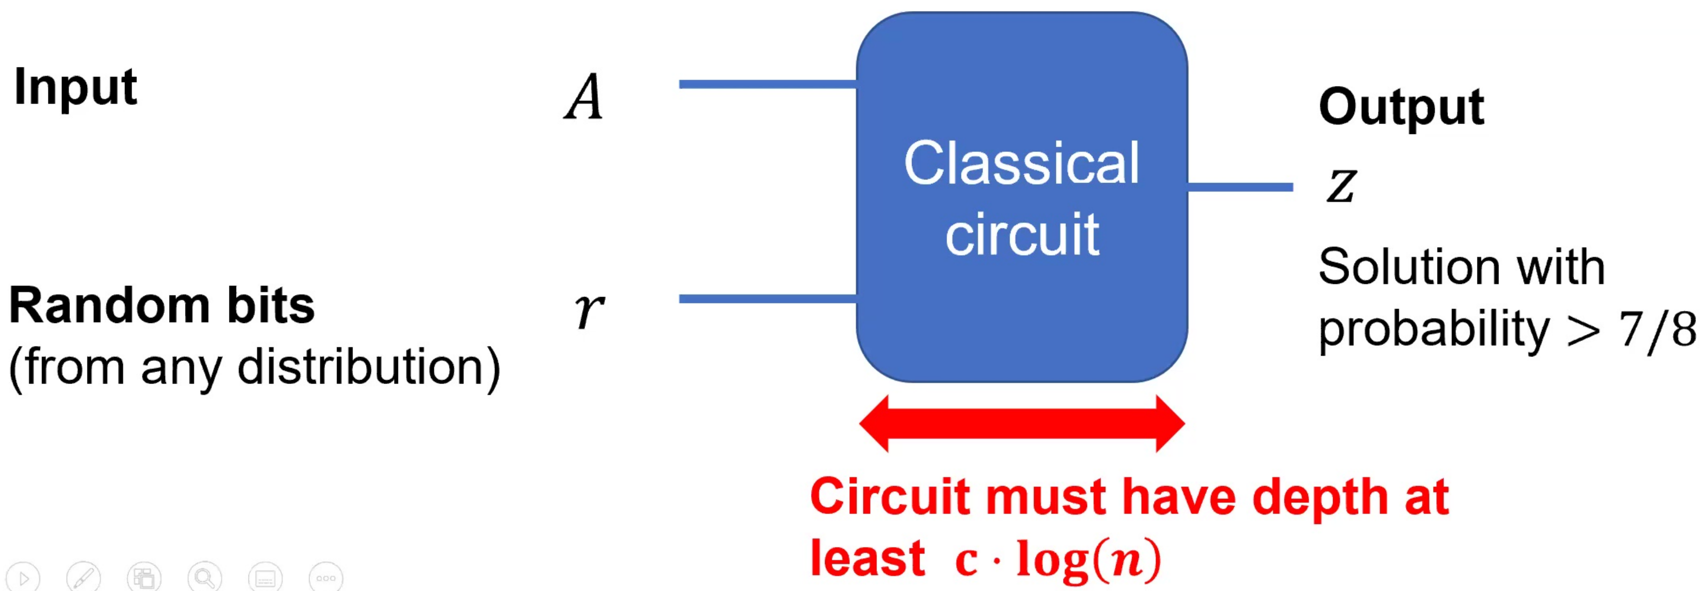

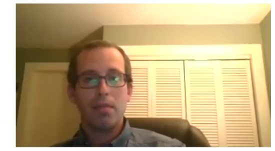

Proof ideas

Locality in shallow classical circuits

Each output bit can only depend on $O(1)$ input bits.

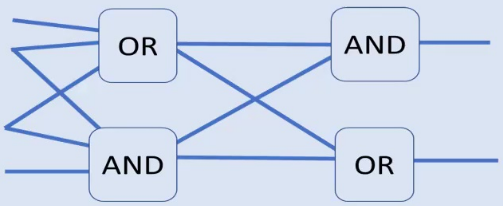

Shallow circuits generalize local hidden variable models

Vs.

Quantum nonlocality

Measurement statistics of entangled quantum states cannot be reproduced by local hidden variable models

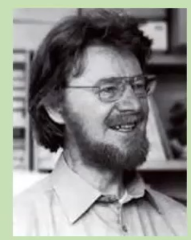

Outputs of constant depth quantum circuits have a strong form of quantum nonlocality

Locality in classical circuits

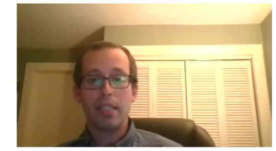

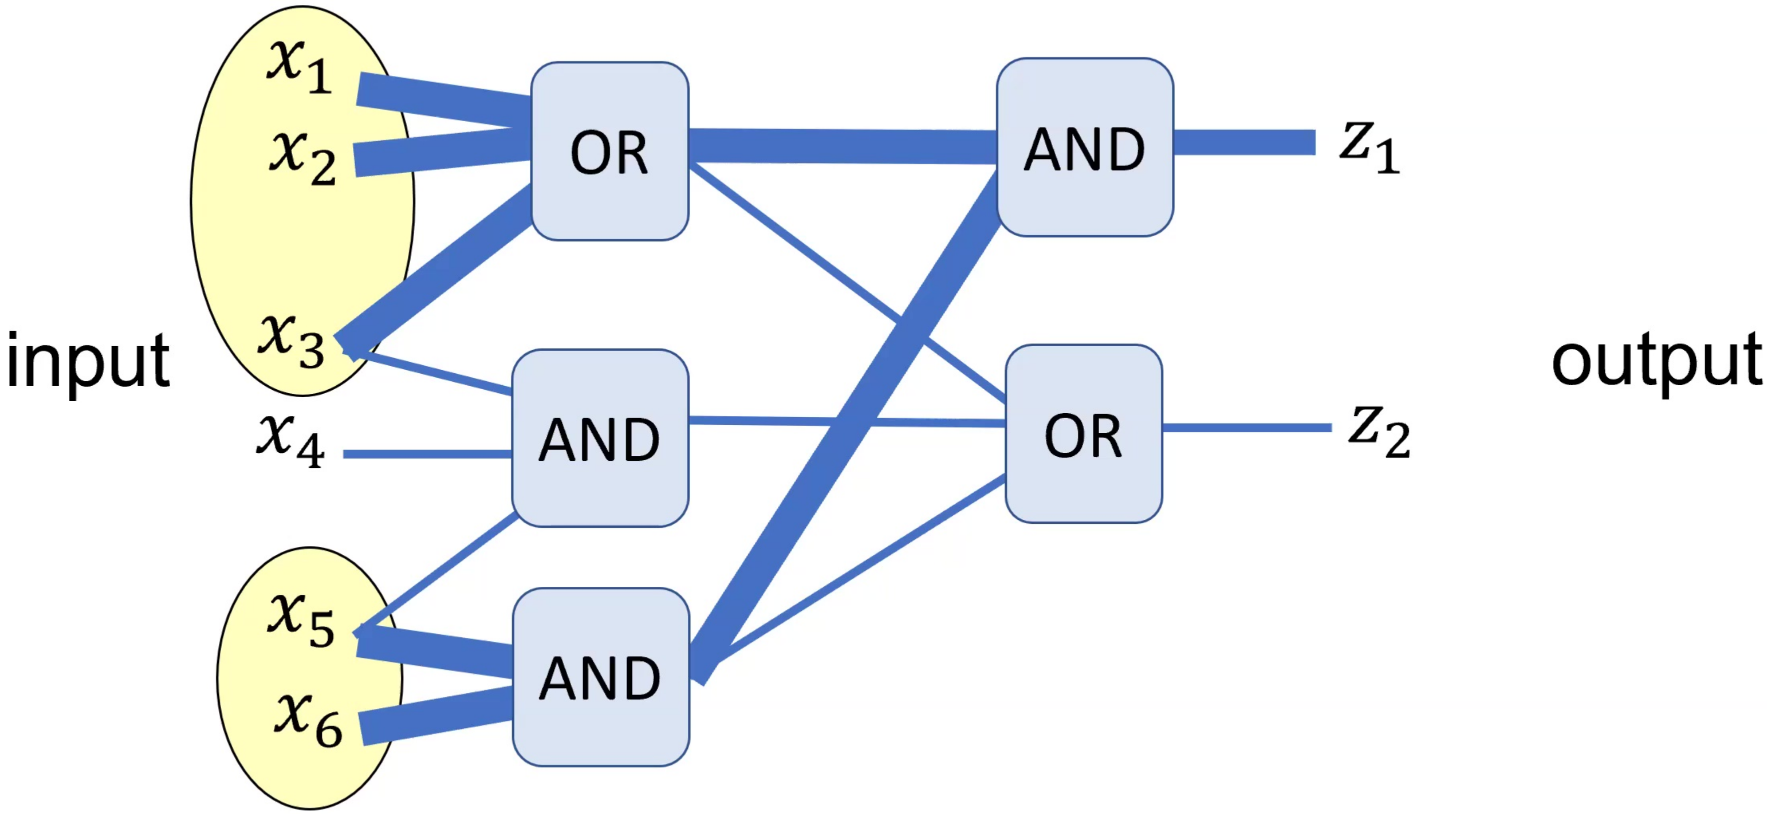

The **lightcone** $L(z_k)$ of an output bit z_k is the set of input bits x_i that are causally connected to z_k.

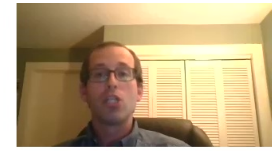

Locality in classical circuits

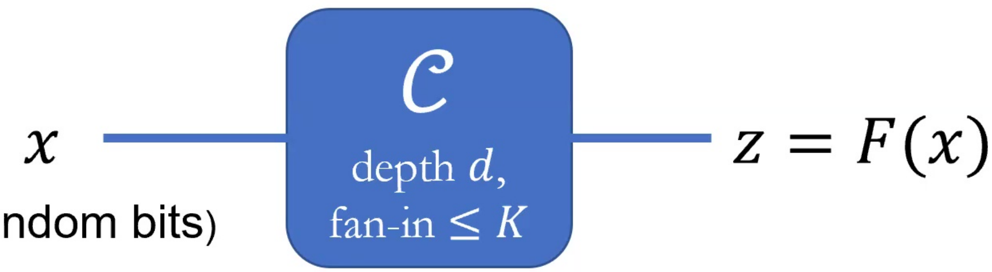

(Includes random bits)

"Constant-depth locality": Lightcones of output bits have constant size

$$|L(z_k)| \leq K^d$$

We'll see that the 2D Hidden Linear Function problem cannot be solved by "constant-depth local" circuits. First consider a simpler form of locality…

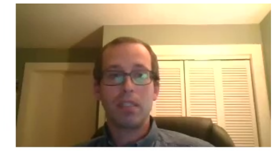

Quantum nonlocality beats completely local circuits

[Greenburger et al. 1990][Mermin 1990]

A **completely local** classical circuit.

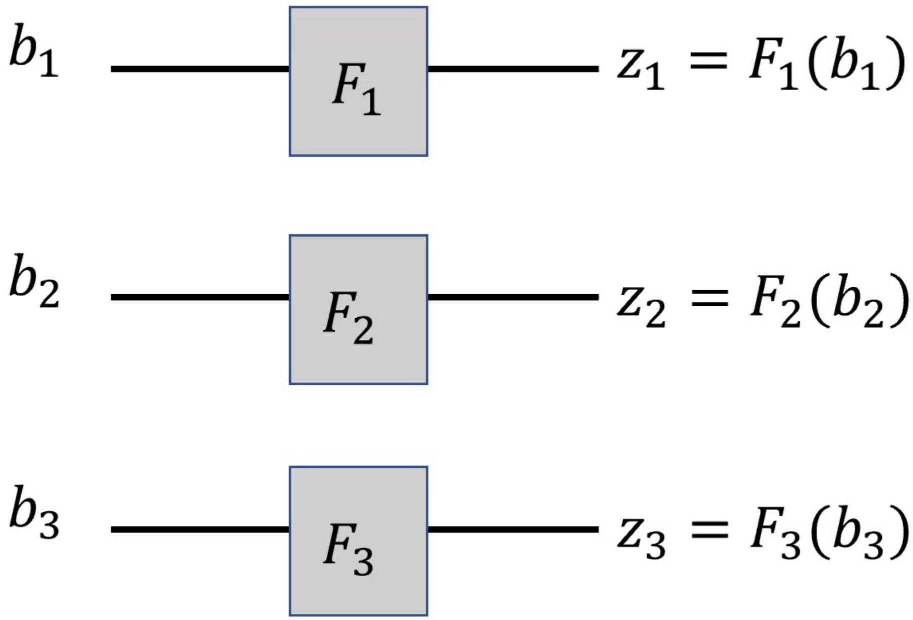

Inputs $\mathbf{b_1, b_2, b_3} \in \{\mathbf{0}, \mathbf{1}\}$

Outputs $\mathbf{z_1, z_2, z_3} \in \{-\mathbf{1}, \mathbf{1}\}$

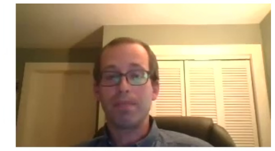

Quantum nonlocality beats **completely local** circuits

[Greenburger et al. 1990][Mermin 1990]

A **completely local** probabilistic classical circuit ⟷ Local hidden variable model

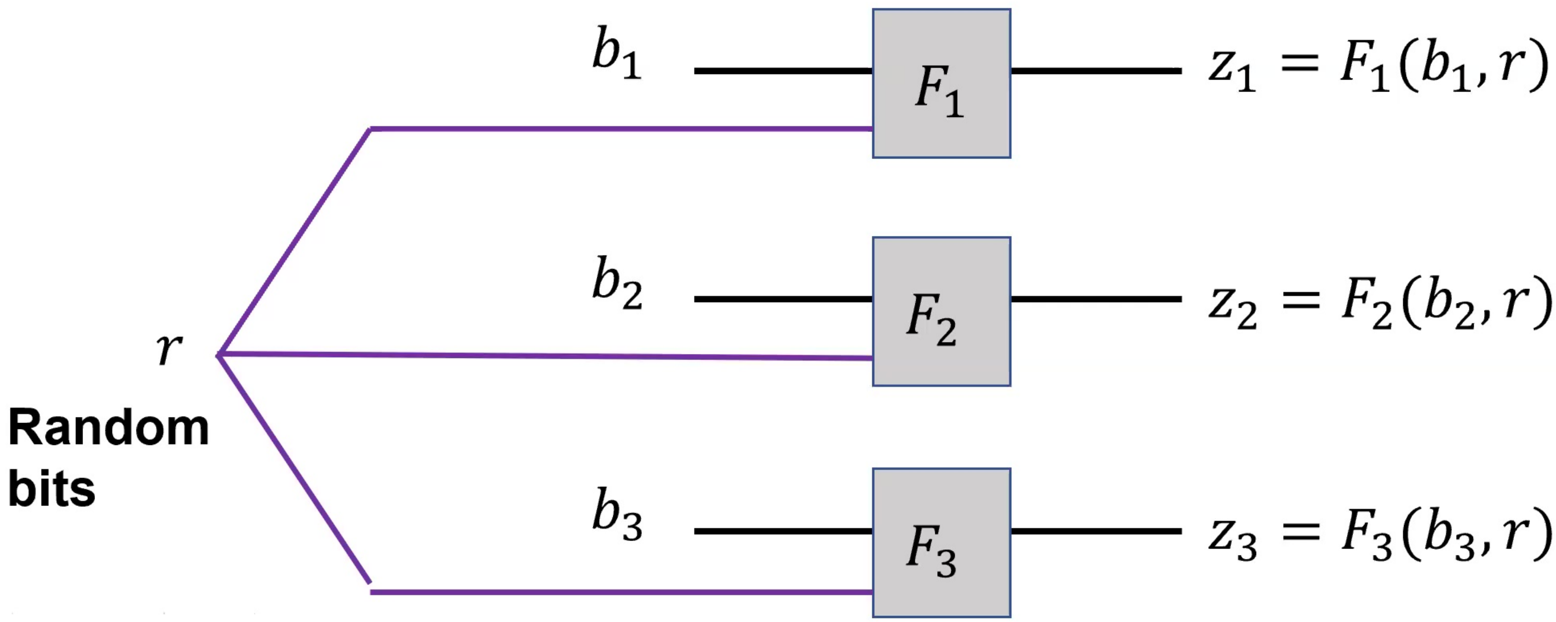

Quantum nonlocality beats completely local circuits

[Greenburger et al. 1990][Mermin 1990]

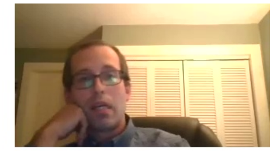

The following input/output relation cannot be realized by a **completely local** probabilistic classical circuit.

b_1	b_2	b_3	$z_1 z_2 z_3$
0	0	0	1
1	1	0	−1
0	1	1	−1
1	0	1	−1

“GHZ relation”

Quantum nonlocality beats completely local circuits

[Greenburger et al. 1990][Mermin 1990]

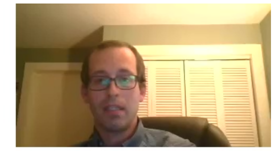

A **completely local** probabilistic classical circuit

Local hidden variable model

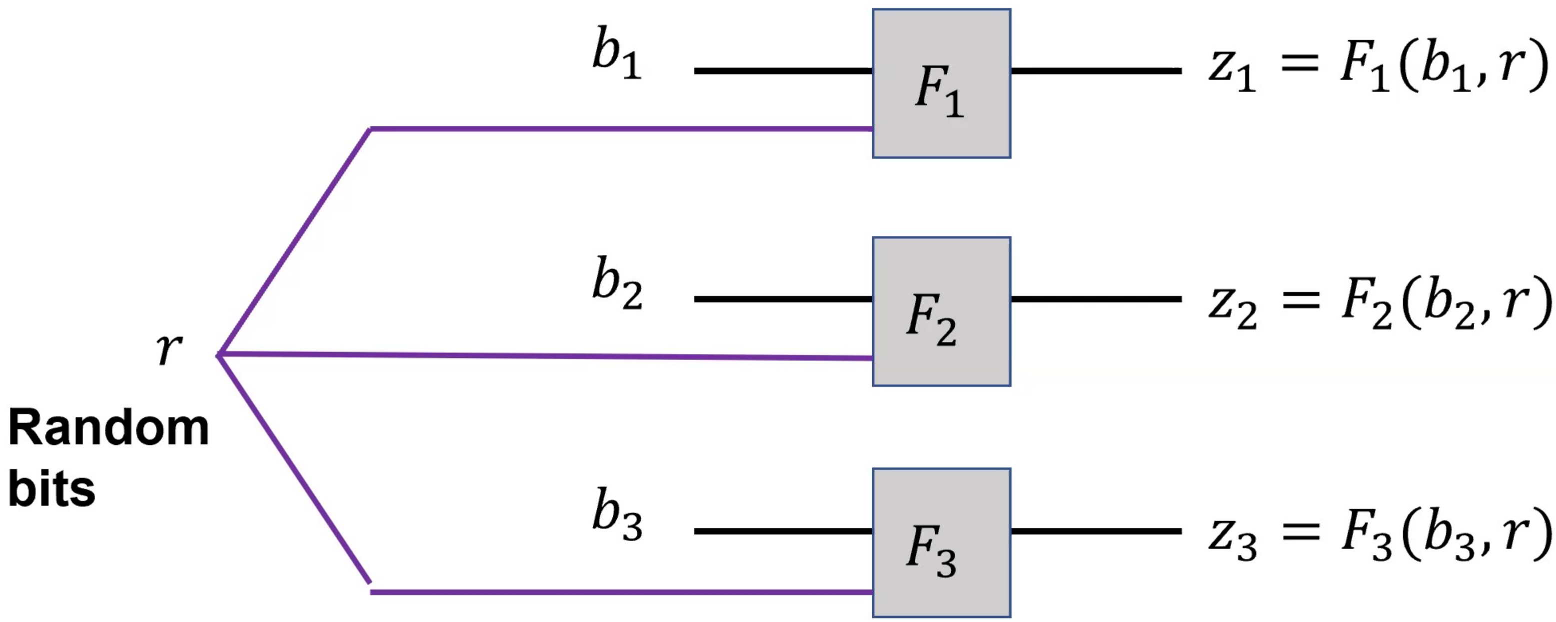

Quantum nonlocality beats completely local circuits

[Greenburger et al. 1990][Mermin 1990]

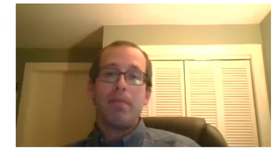

The GHZ relation can be realized by a **completely local quantum circuit** :

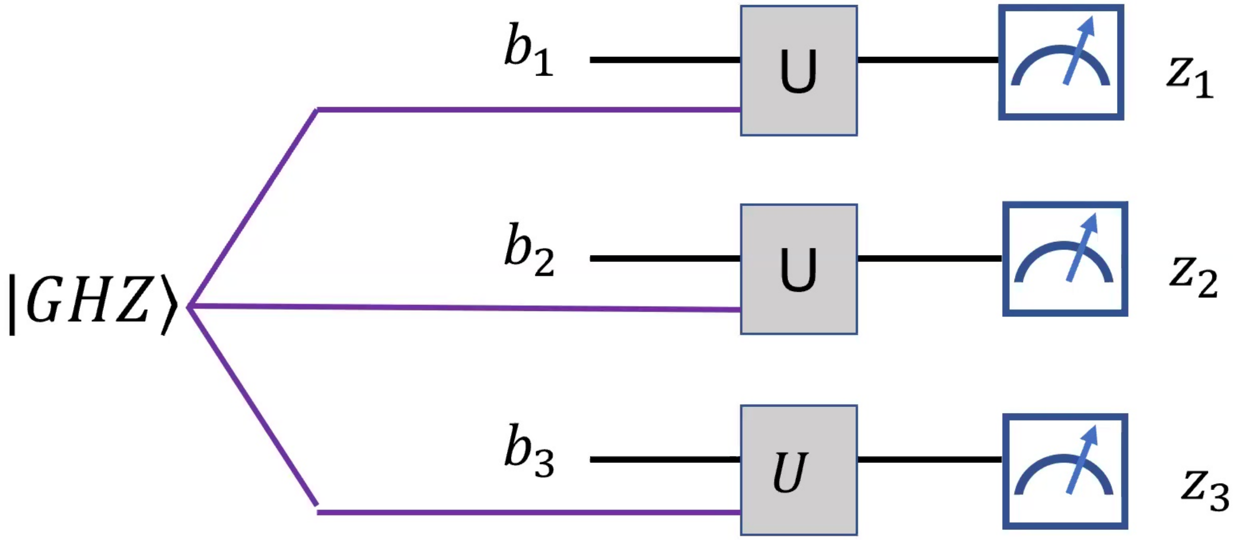

$$|GHZ\rangle = \frac{1}{\sqrt{2}}(|000\rangle + |111\rangle)$$

Quantum nonlocality beats completely local circuits

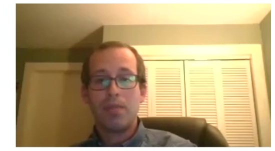

[Greenburger et al. 1990][Mermin 1990]

The GHZ relation can be realized by a **completely local quantum circuit** :

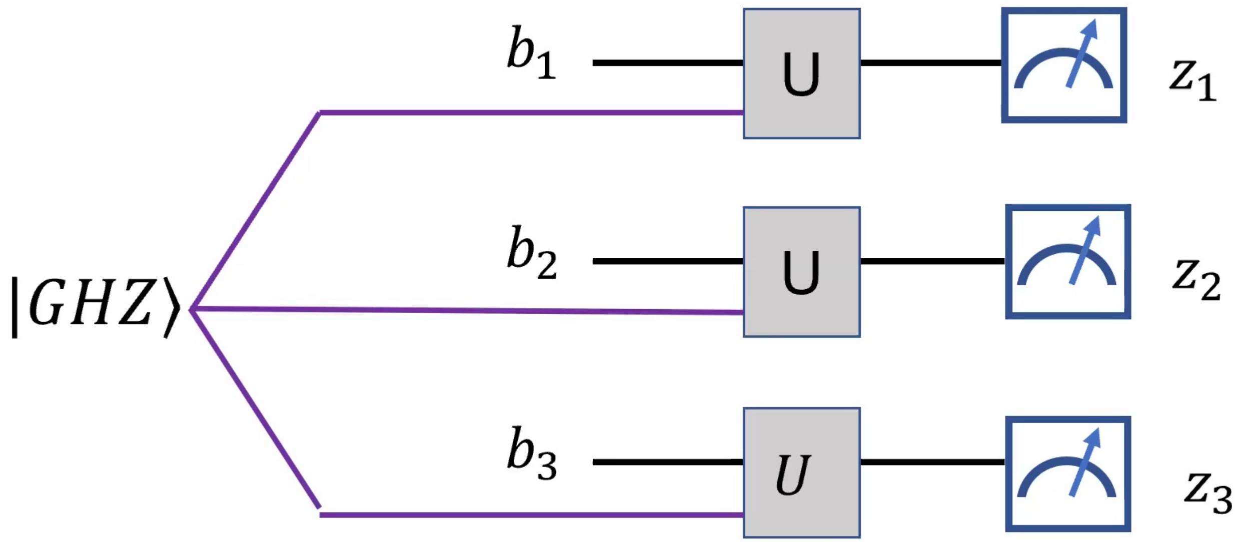

$$|GHZ\rangle = \frac{1}{\sqrt{2}}(|000\rangle + |111\rangle)$$

The circuit measures each qubit in the X or Y basis depending on the corresponding input bit

Quantum nonlocality beats completely local circuits

[Greenburger et al. 1990][Mermin 1990]

The GHZ relation can be realized by a **completely local quantum circuit** :

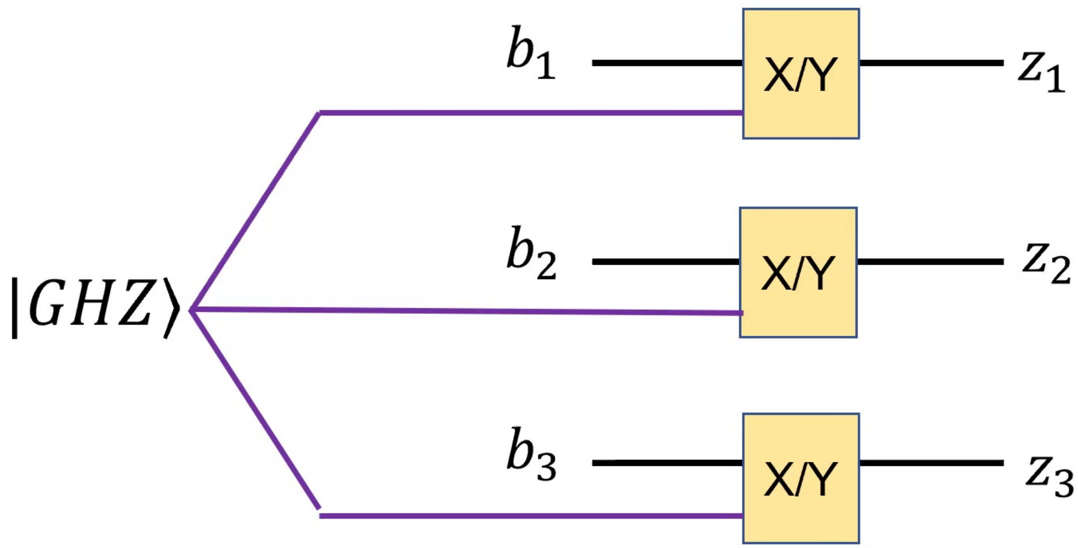

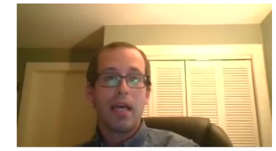

$$|GHZ\rangle = \frac{1}{\sqrt{2}}(|000\rangle + |111\rangle)$$

The circuit measures each qubit in the X or Y basis depending on the corresponding input bit

Quantum nonlocality beats geometrically local circuits

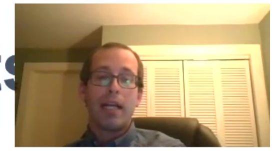

Barrett, Caves, Elliott, Pironio. Physical Review A 75(1):012103, 2007.

Barrett et al. (2007) describe a special family of HLF instances that can't be solved with geometrically local classical circuits:

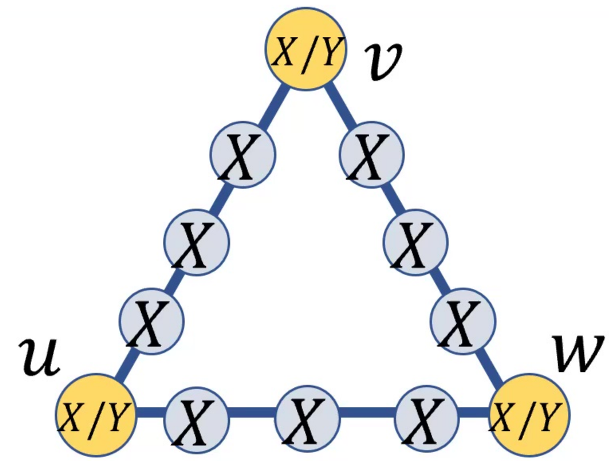

Graph state on an M-cycle (M even).

Choose 3 qubits u, v, w on the even sublattice. Measure u, v, w in X or Y basis and all other qubits in X basis.

Quantum nonlocality beats geometrically local circuits

Barrett, Caves, Elliott, Pironio. Physical Review A 75(1):012103, 2007.

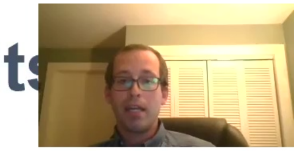

Barrett et al. (2007) describe a special family of HLF instances that can't be solved with geometrically local classical circuits:

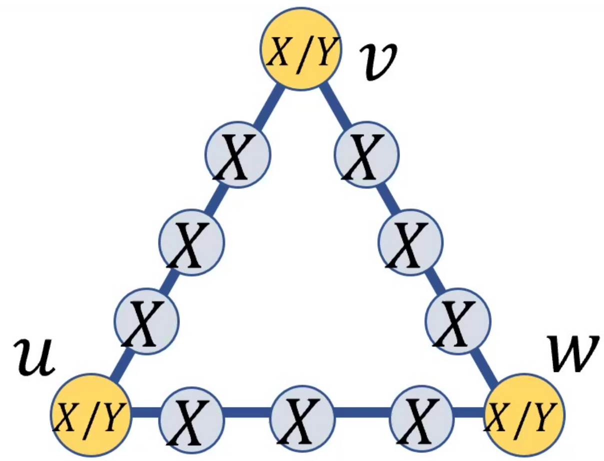

Graph state on an M-cycle (M even).

Choose 3 qubits u, v, w on the even sublattice. Measure u, v, w in X or Y basis and all other qubits in X basis.

Geometrically nonlocal correlations are necessary

To solve these instances of HLF, some output bit z_k must be correlated with a **distant** input bit b_u, b_v or b_w. (i.e., not the nearest vertex of the triangle)

Quantum nonlocality beats "constant-depth local" circuits

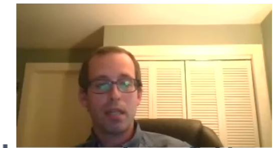

Getting back to the 2D HLF problem….there are instances corresponding to any subgraph of the grid. Let's focus on instances corresponding to cycles.

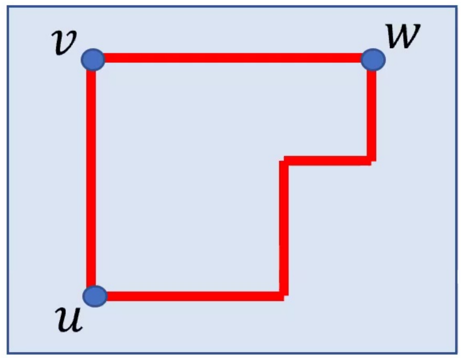

A classical circuit that solves 2D HLF must have geometrically nonlocal correlations with respect to every such cycle. (Barrett et al. example)

A probabilistic argument shows that this correlation structure is not possible unless there are output bits with lightcones of size at least $n^{1/8}$. This translates to a depth lower bound $\frac{\log(n)}{8\log(K)}$.

Extensions and recent results

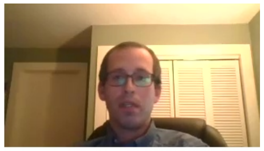

What if we make the classical circuit more powerful?

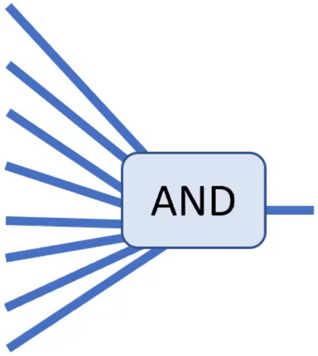

Bene Watts, Kothari, Schaeffer, Tal. In *Proceedings of STOC 2019.*

Classical shallow circuits *still* can't solve the problem even if we allow AND/OR gates with any number of inputs (unbounded fan-in)

Extensions and recent results

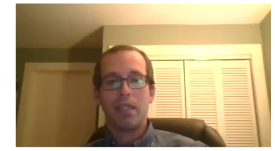

What if we make the classical circuit more powerful?

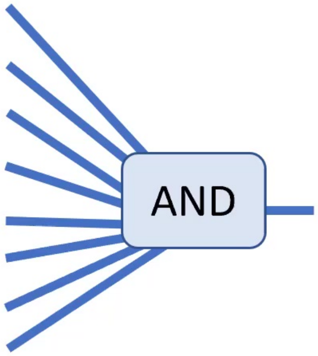

Bene Watts, Kothari, Schaeffer, Tal. In *Proceedings of STOC 2019.*

Classical shallow circuits *still* can't solve the problem even if we allow AND/OR gates with any number of inputs (unbounded fan-in)

Grier, Schaeffer. In *Proceedings of STOC 2020.*

Certain two-round interactive linear algebra tasks can be solved by shallow quantum circuits but not by even more powerful classical circuits.

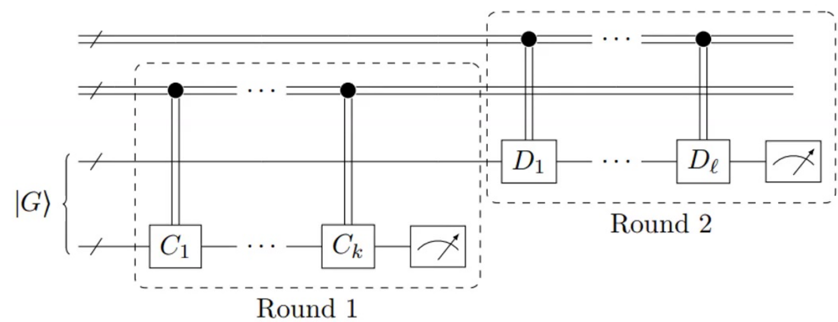

Extensions and recent results

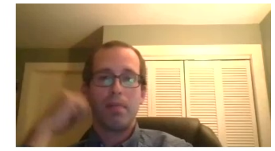

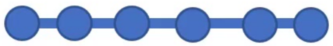

Is there a quantum advantage with shallow circuits in 1D?

Bravyi, DG, Koenig, Tomamichel. *Nature Physics 1-6*, 2020.

Yes: 1D quantum advantage via a multi-player variant of Mermin magic square game.

Extensions and recent results

**We know there is a speedup as measured by depth.
What about the total number of gates?**

DG, Grier, Kerzner, Schaeffer. *arXiv:2009.03218, 2020.*

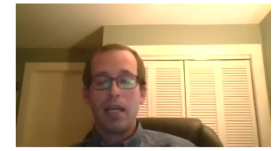

Extensions and recent results

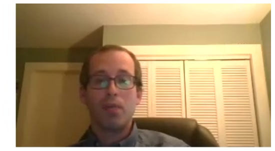

We know there is a speedup as measured by depth.
What about the total number of gates?

DG, Grier, Kerzner, Schaeffer. *arXiv:2009.03218, 2020.*

The shallow quantum circuit solves the 2D HLF problem using $O(n)$ gates.
A classical computer can solve it with linear algebra using $O(n^3)$ gates.
Can this be improved?

We first show that **any** shallow Clifford circuit can be simulated with runtime

$$O(n^{\omega})$$

$\mathbf{2 \leq \omega \leq 2.3729}$

[Le Gall 2014]
[Strassen 1969]

Then we give a recursive divide-and-conquer algorithm with improved runtime

$$O(n^{\omega/2})$$

for shallow Clifford circuits in planar geometries.

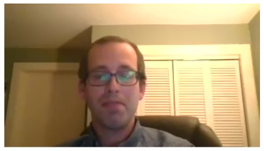

Quantum advantage with **noisy** shallow circuits

Bravyi, DG, Koenig, Tomamichel. *Nature Physics 1-6*, 2020.

Noise model

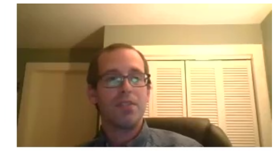

Each layer of gates is followed by a random (Pauli) error. Think of a simple independent noise model where each qubit is corrupted with probability p.

.

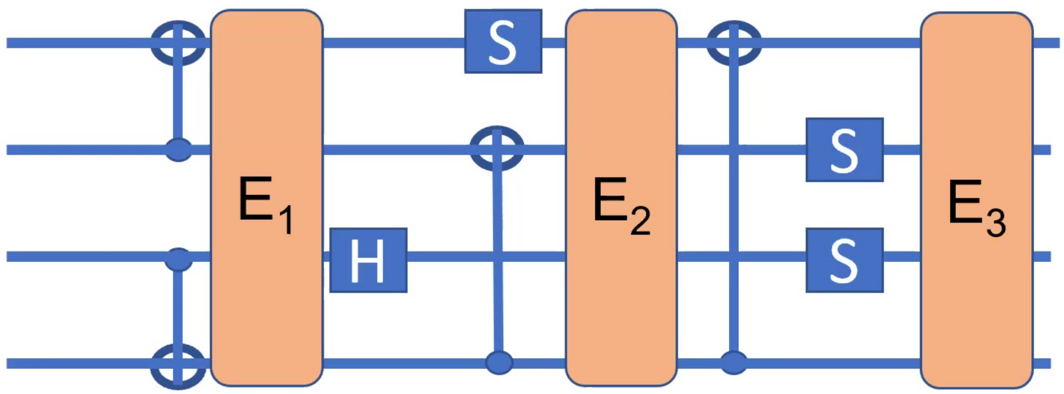

Why does noise cause a problem?

Circuit which solves the 2D HLF problem:

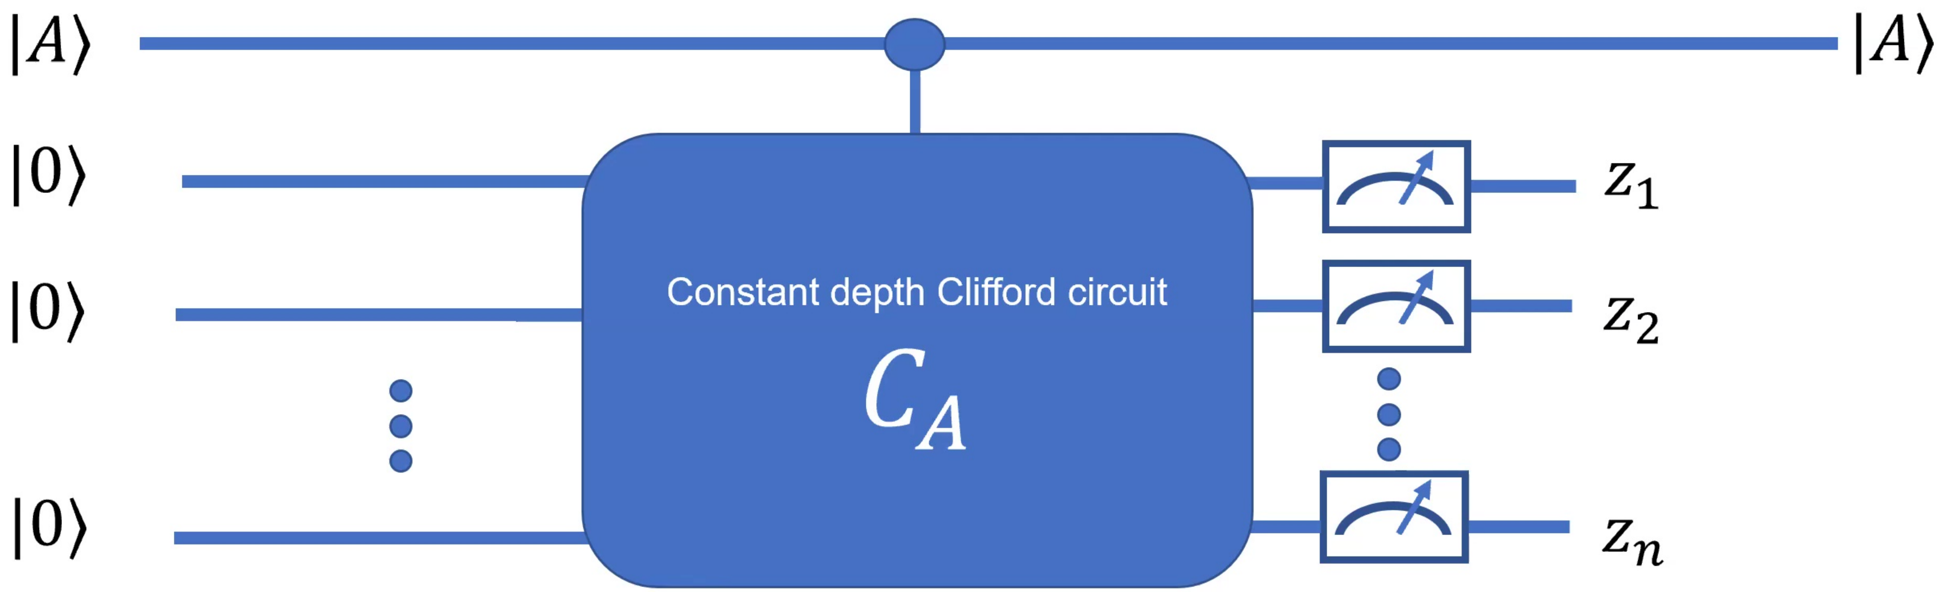

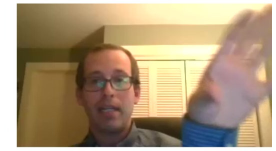

Why does noise cause a problem?

Circuit which solves the 2D HLF problem:

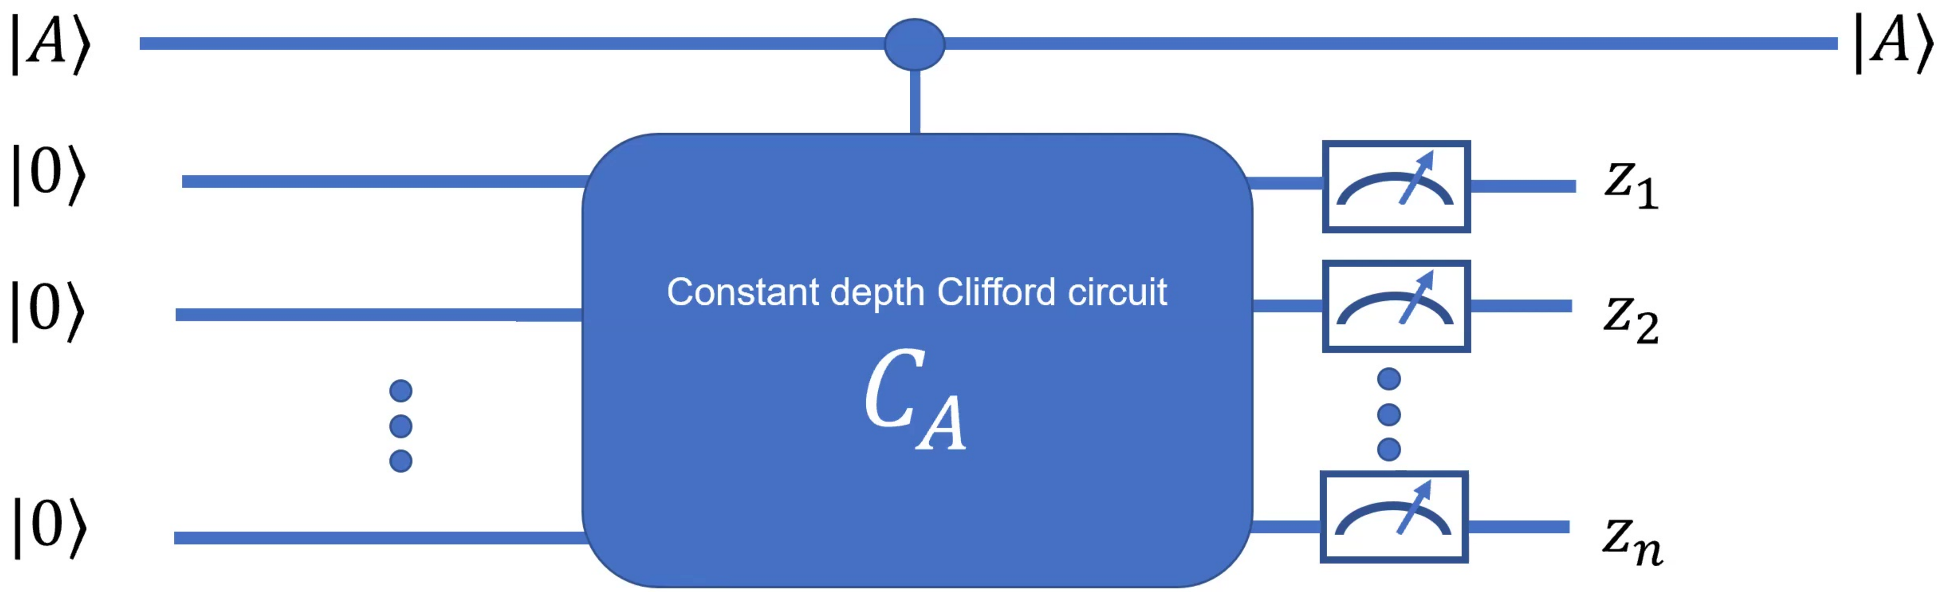

Unfortunately, noise with rate p corrupts a constant fraction $\sim pn$ of the output bits…

Challenges with naively using quantum error correction

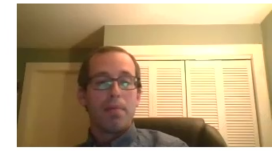

Imagine input A held in a noise-free classical memory.

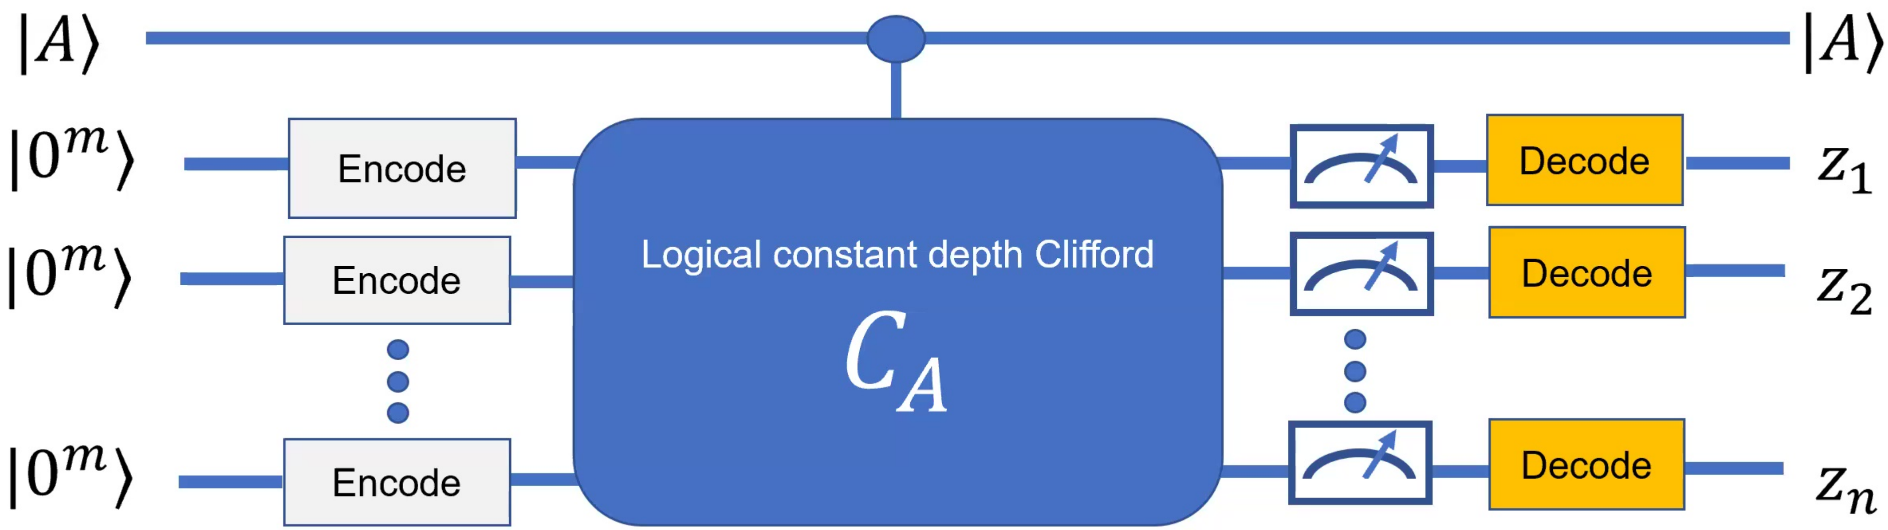

Choose an error correcting code where one- and two-qubit logical Clifford gates can be performed by constant-depth circuits.

Good news: this makes C_A constant depth.

Bad news: neither the **encoding** nor the **decoding** is constant depth.

Getting around the decoding problem

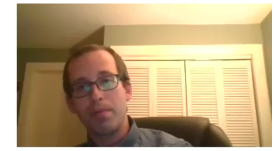

The decoding is a purely classical postprocessing step:

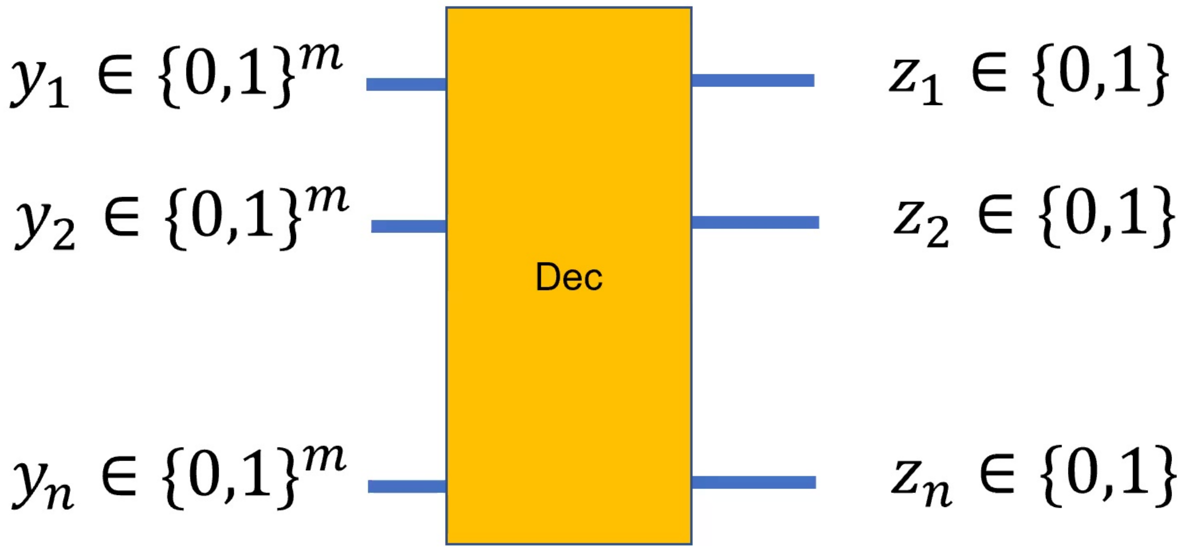

We can fold the decoding step into the problem definition:

Define a new computational problem so that y is a solution to the new problem if and only if $z = Dec(y)$ is a solution to the 2D HLF problem.

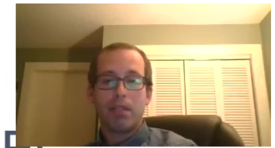

Getting around the encoding problem

The ideal encoding operation prepares a logical basis state encoded in the QEC.

Can't do this in constant depth

A related task which can be done in constant depth for LDPC stabilizer codes:

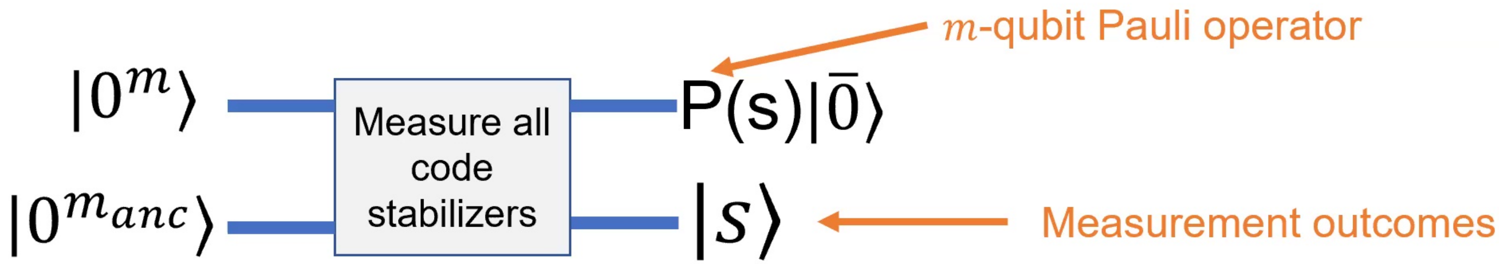

This is good enough! We can (again) modify the problem definition so that it incorporates information about the measurement outcomes s. This part uses the fact that the circuit is Clifford and plays nicely with the Pauli correction $P(s)$.

Errors in state preparation

Unfortunately a few errors in the measurement outcome s can spread to many errors in $\mathrm{P}(s)$ …

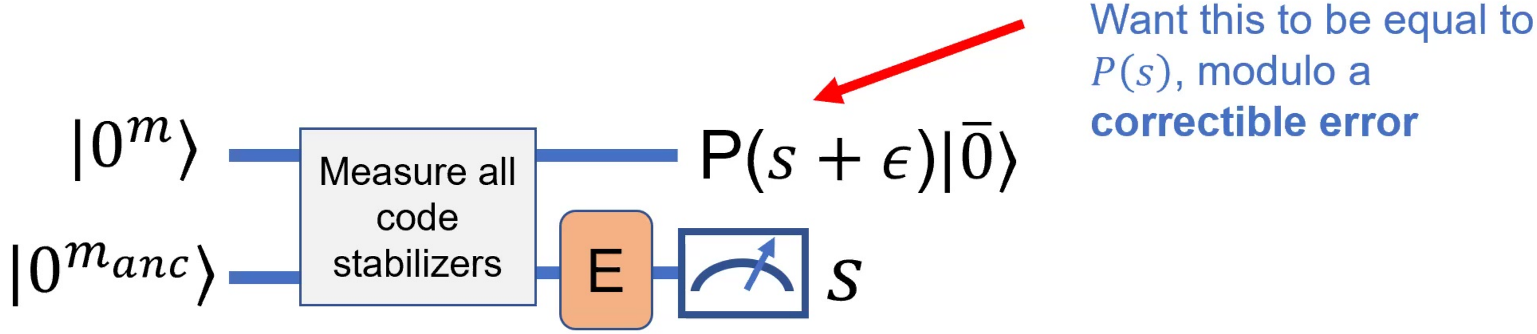

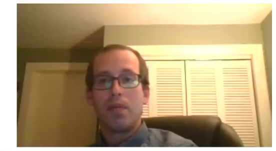

The property that we need is called **single-shot state preparation**.

Putting it together

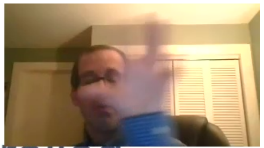

Quantum circuit for 2D HLF problem. Or any other problem defined by a controlled Clifford circuit:

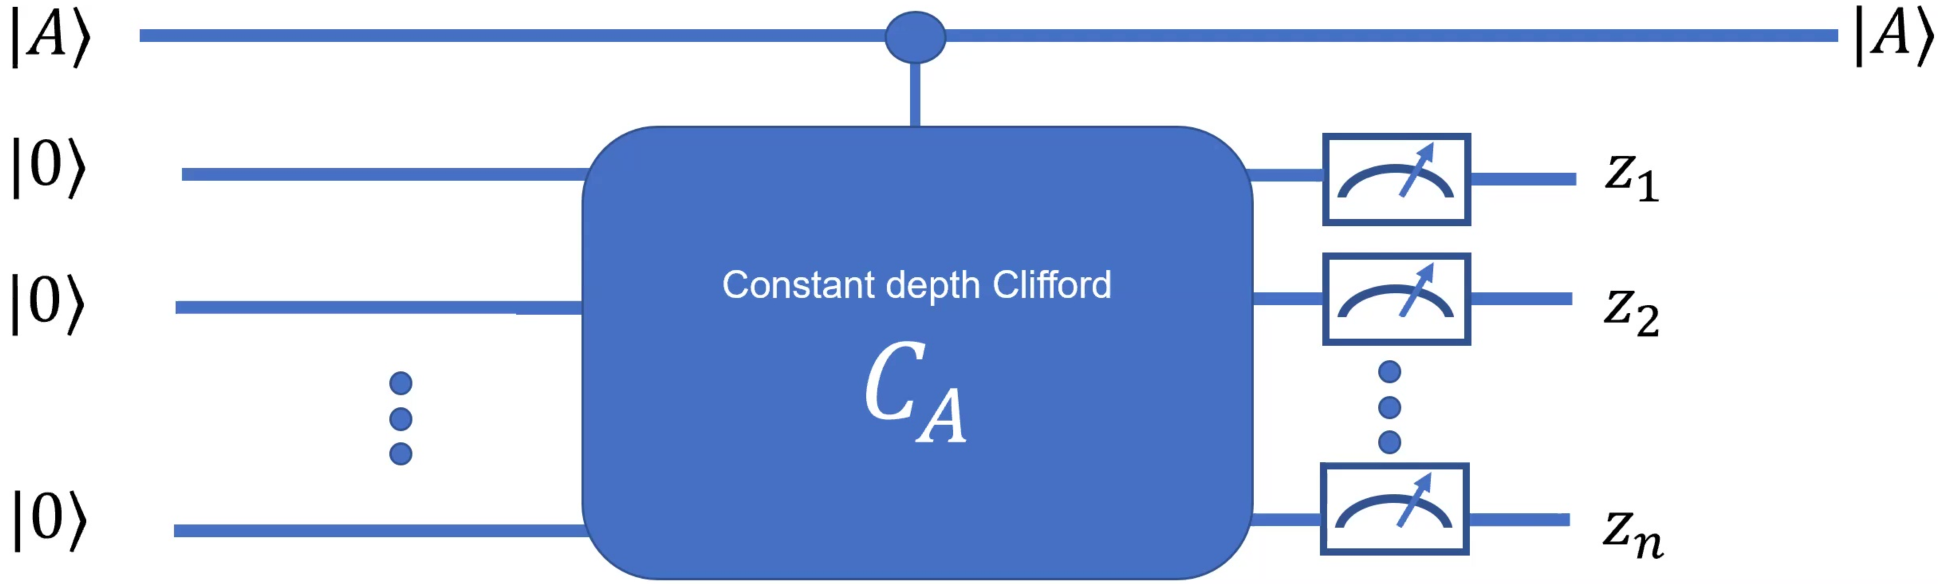

Input/output pairs produced by a noise-free implementation of the circuit satisfy a relation $R(A, z) = 1$

Putting it together: the noise-tolerant problem

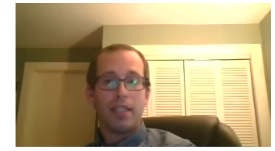

We transform the circuit as follows

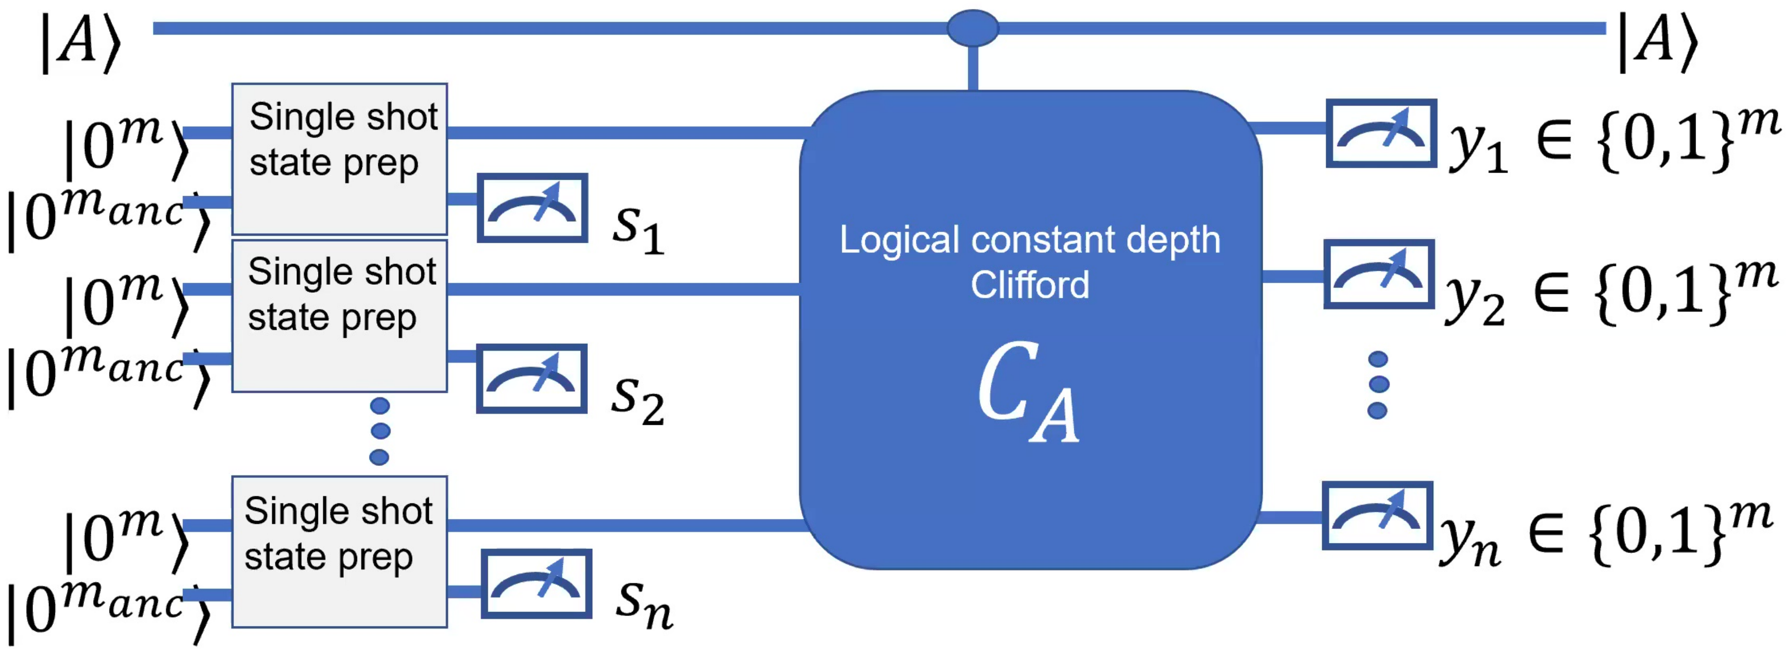

Putting it together

Quantum circuit for 2D HLF problem. Or any other problem defined by a controlled Clifford circuit:

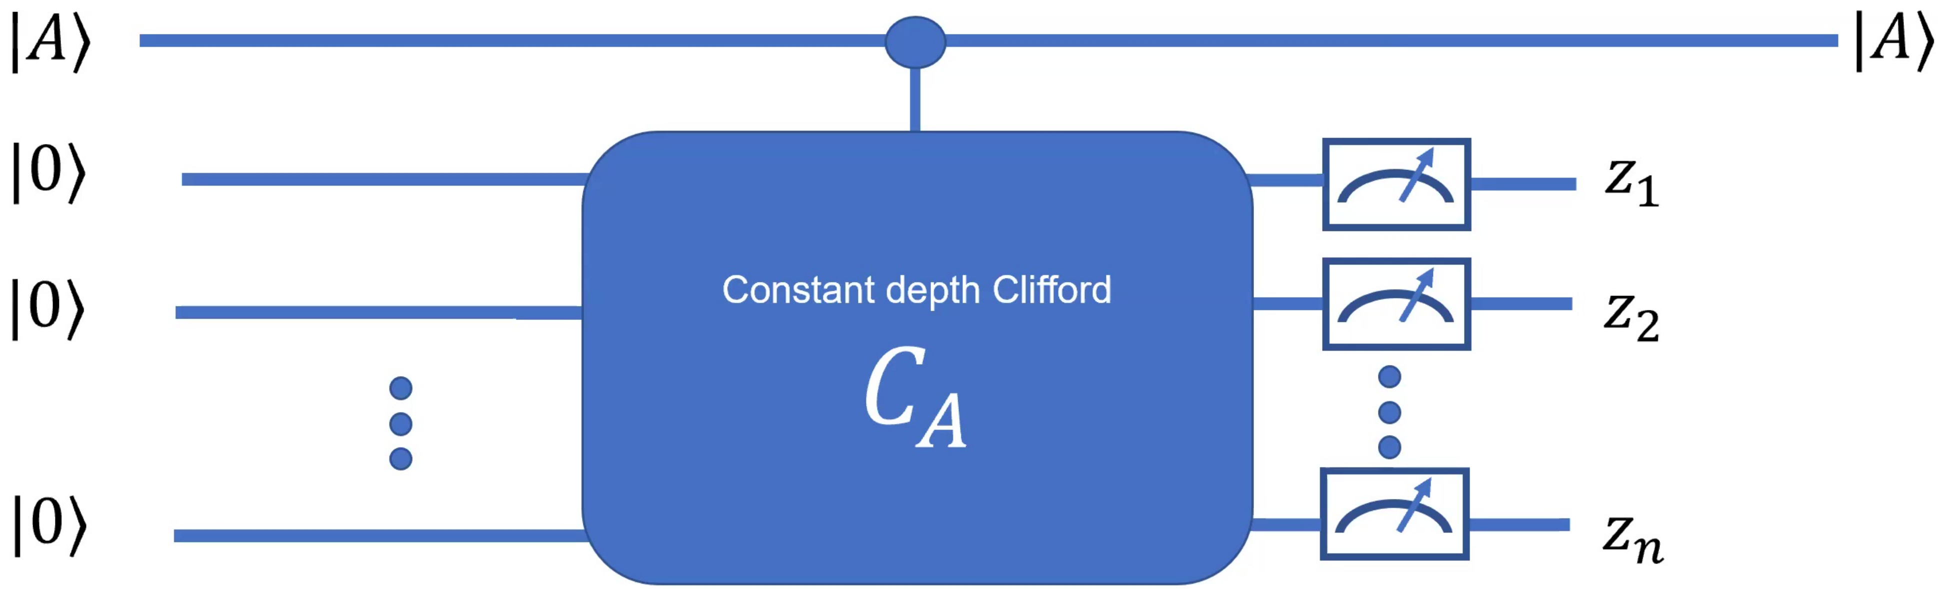

Input/output pairs produced by a noise-free implementation of the circuit satisfy a relation $R(A,z) = 1$

Putting it together: the noise-tolerant problem

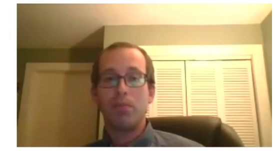

We transform the circuit as follows

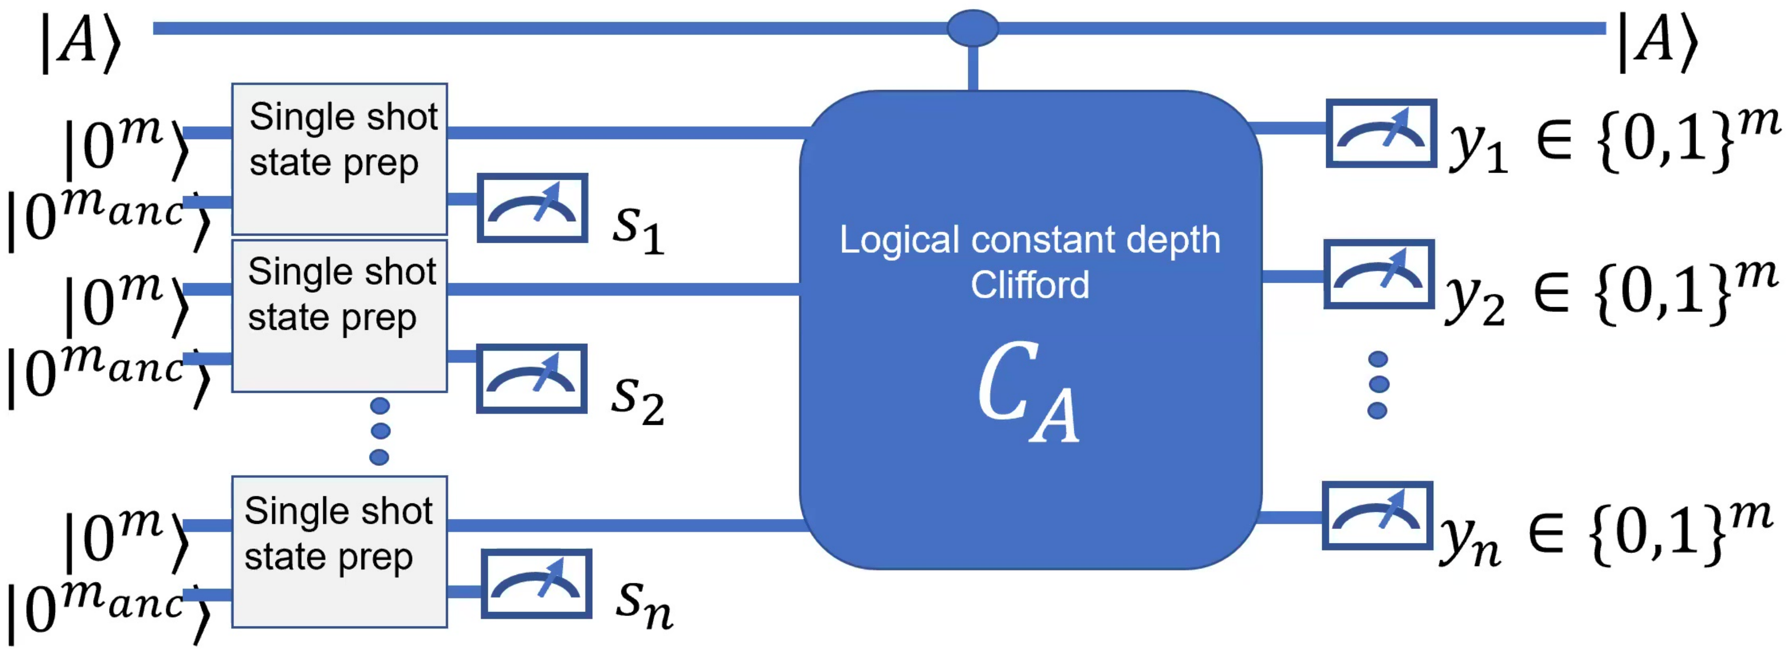

Now input/output pairs produced by a noisy implementation of the circuit satisfy a ``noise-tolerant" relation $\tilde{R}(A, y, s) = 1$.

The noise-tolerant problem is not much easier

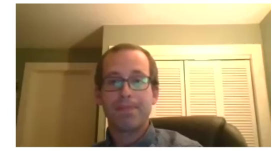

We prove that the new, noise-tolerant problem $\tilde{R}$ is not much easier than the original problem R.

Proof idea: Can compute a pair (A, z) satisfying R starting from a triple (A, y, s) satisfying $\tilde{R}$ using a low depth circuit with fan-in $K = O(poly(\log(n)))$

For the noise tolerant version of 2D HLF, we infer a classical depth lower bound

$$D \geq constant \cdot \frac{\log(n)}{\log(\log(n))}$$

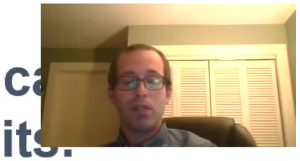

So…a problem defined by a constant-depth controlled Clifford circuit ca made noise-tolerant, without making it much easier for classical circuits.

It relies on the existence of a certain family of quantum error correcting codes…

We show that all of our requirements are met by the well-known **surface code.**
The new piece is the single-shot logical state preparation.

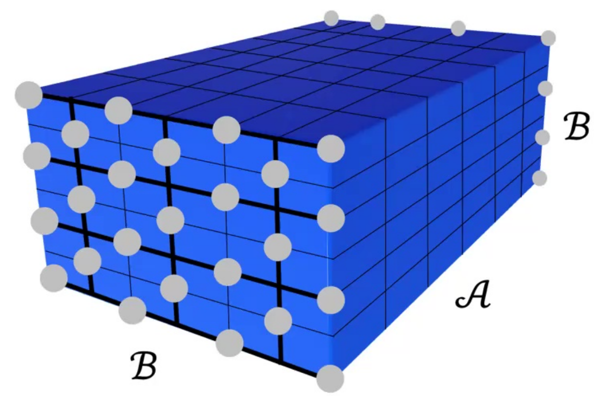

The single-shot state preparation protocol we give is an extension of **[Raussendorf, Bravyi, Harrington 2004].**

Concluding remarks and open problems

Studying shallow quantum circuits—motivated by near-term QCs—has taught us about a new kind of quantum advantage and drawn a connection between circuit complexity and quantum nonlocality.

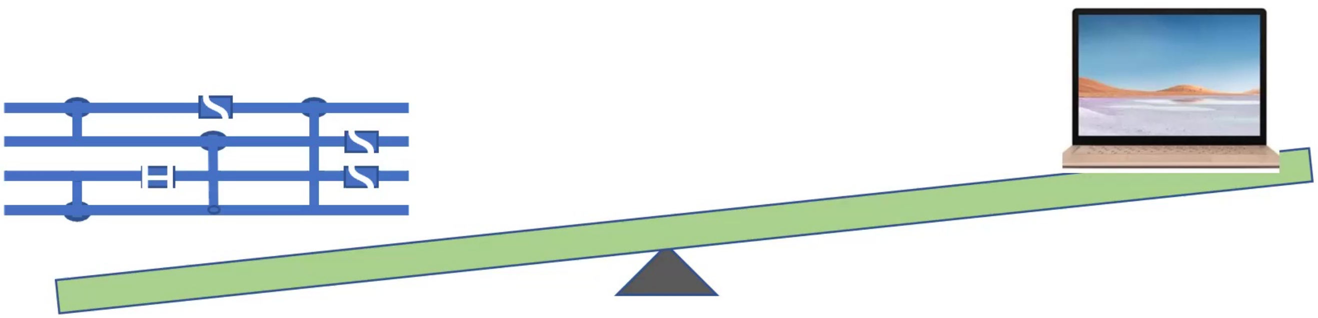

Experimental demonstrations?

Does quantum nonlocality play a role in quantum optimization algorithms such as QAOA?

Can quantum nonlocality give larger speedups?

So…a problem defined by a constant-depth controlled Clifford circuit can be made noise-tolerant, without making it much easier for classical circuits.

It relies on the existence of a certain family of quantum error correcting codes…

We show that all of our requirements are met by the well-known **surface code.**
The new piece is the single-shot logical state preparation.

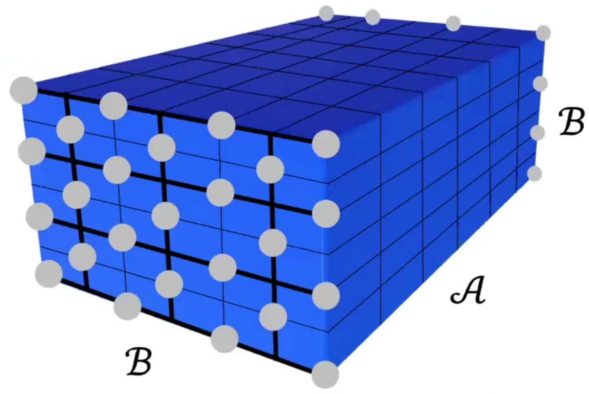

The single-shot state preparation protocol we give is an extension of **[Raussendorf, Bravyi, Harrington 2004].**

Putting it together: the noise-tolerant problem

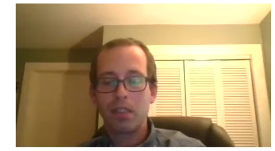

We transform the circuit as follows

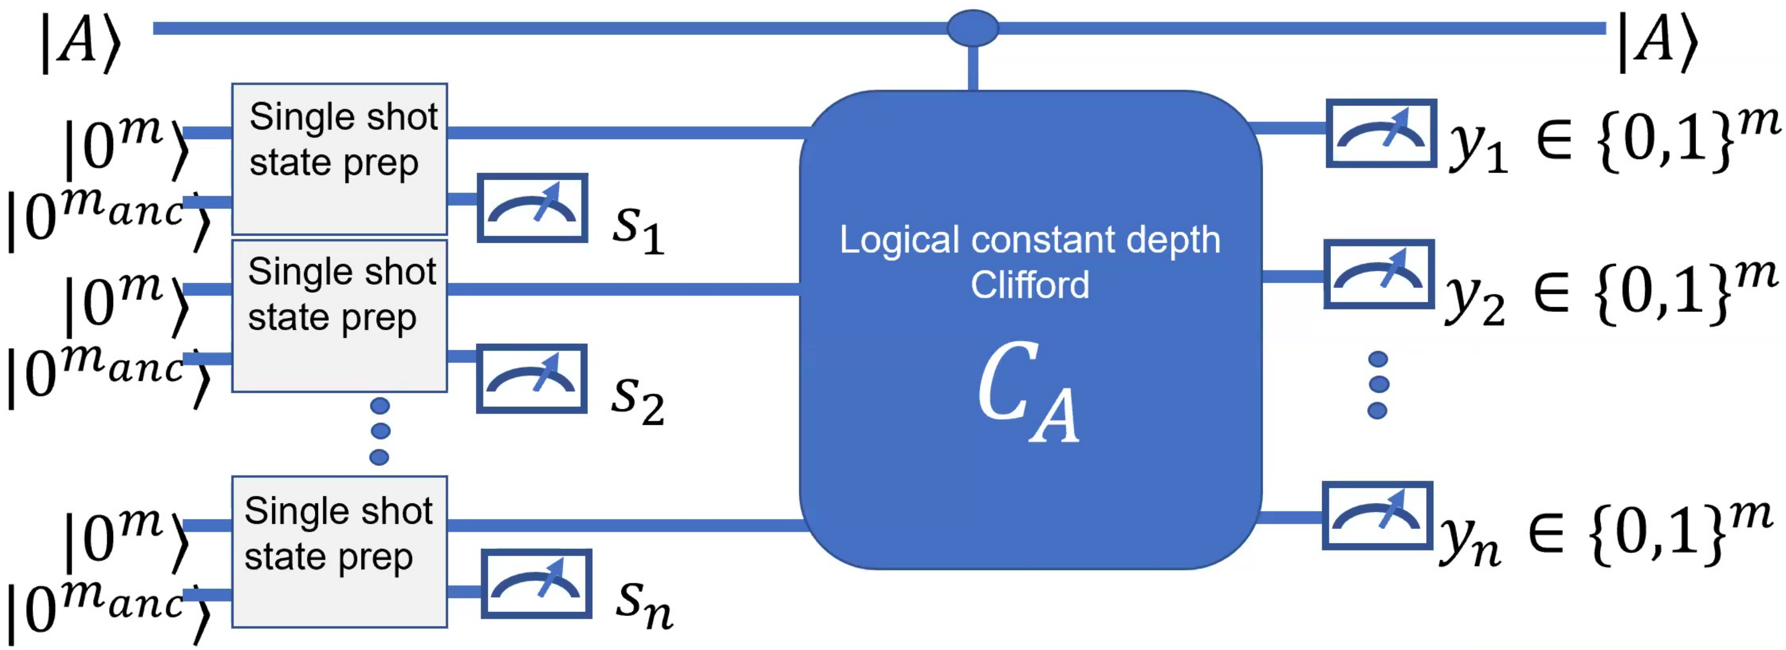

Now input/output pairs produced by a noisy implementation of the circuit satisfy a ``noise-tolerant" relation $\tilde{R}(A, y, s) = 1$.

Thanks!