Phase transitions in the complexity of simulating random shallow quantum circuits

John Napp, Rolando La Placa, Alex Dalzell, Fernando Brandão, Aram Harrow

1 Apr 2021

arXiv:2001.00021

quantum computers

qubit = two-level system


spin-1/2

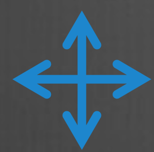
photon polarization

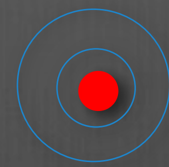
ion e⁻ states

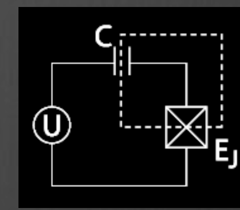

# of Cooper pairs in box

...

1 qubit = 2 dimensions $|\psi\rangle \in \mathbb{C}^2$

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \qquad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

n qubits = 2^n dimensions $|\psi\rangle \in \mathbb{C}^{2^n}$

$$|0,1\rangle = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}$$

Schrödinger's eq:

$$\frac{d|\psi\rangle}{dt} = -iH(t)|\psi\rangle$$

quantum gates

$$I \otimes \ldots \otimes I \otimes U \otimes I \otimes \ldots \otimes I$$

one- or two-qubit

How hard is it to simulate a quantum computer?

“Nature isn’t classical… and if you want to make a simulation of nature, you'd better make it quantum mechanical…” (1981)

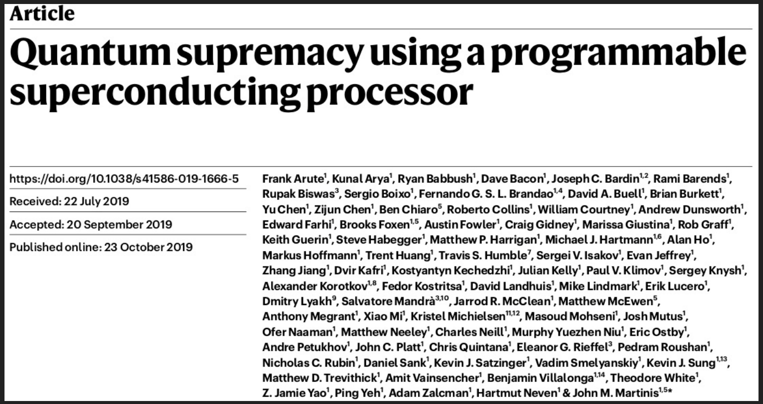

Article

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]*

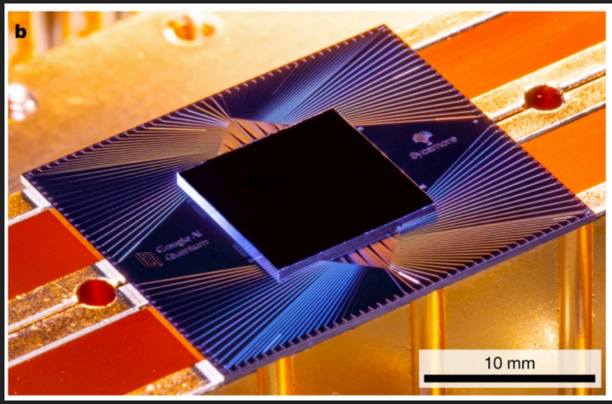

When does this sentiment *not* hold?

Is random circuit sampling truly classically intractable?

the simulation frontier

easy ? hard

- Clifford circuits
- Limited entanglement
- Noise rate $p \geq \approx 0.45$
- Random circuits with noise $p>0$
- 1D log depth

- 2D constant depth (MBQC)
- Random circuits (Google)
- Linear optics (boson sampling)
- Noise rate $p \leq \approx 0.01$ (FTQC)

theorists
look here

easier quantum simulation

Lightly entangling dynamics

product states + non-interacting gates are easy.
Cost grows exponentially with # of entangling gates.

Stabilizer circuits

Poly-time simulation of stabilizer circuits,
growing exponentially with # of non-stabilizer gates.

Likewise for matchgates / non-interacting fermions.

Ground states of 1-D systems

Effort grows exponentially with correlation length.

quantum circuits

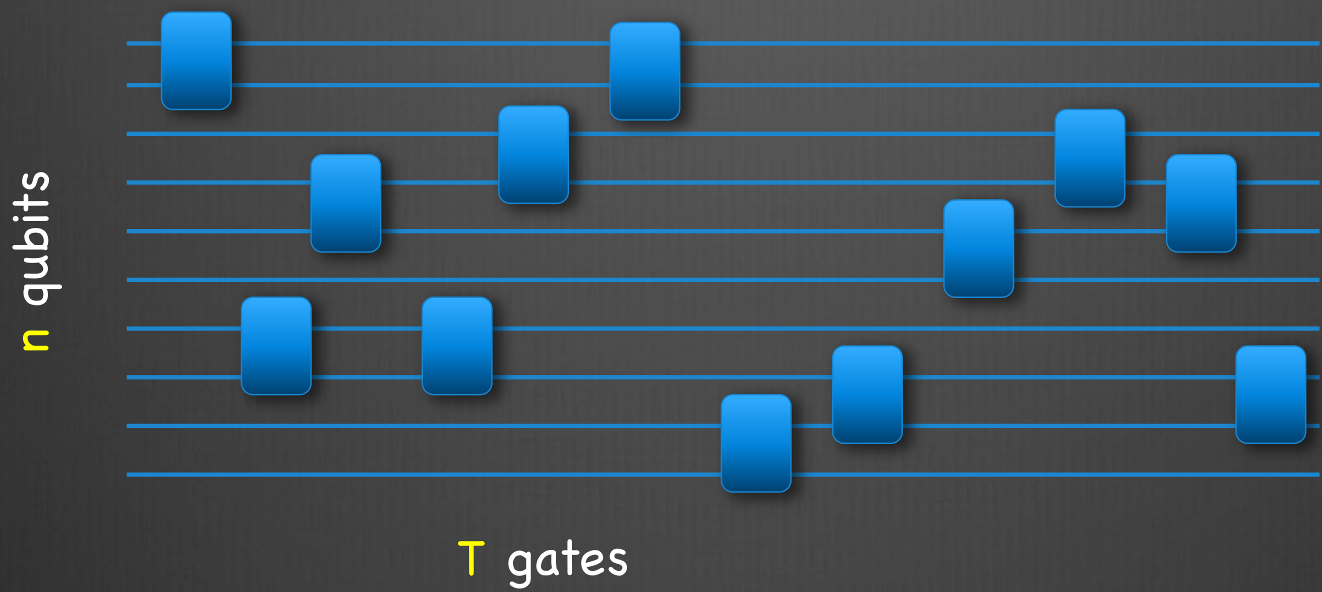

Classical simulation possible in time O(T)·exp(k), where

- k = treewidth [Markov-Shi '05]
- k = max # of gates crossing any single qubit [Yoran-Short '06, Jozsa '06]

+ Complexity interpolates between linear and exponential.

– Treating all gates as "potentially entangling" is too pessimistic.

noisy dynamics?

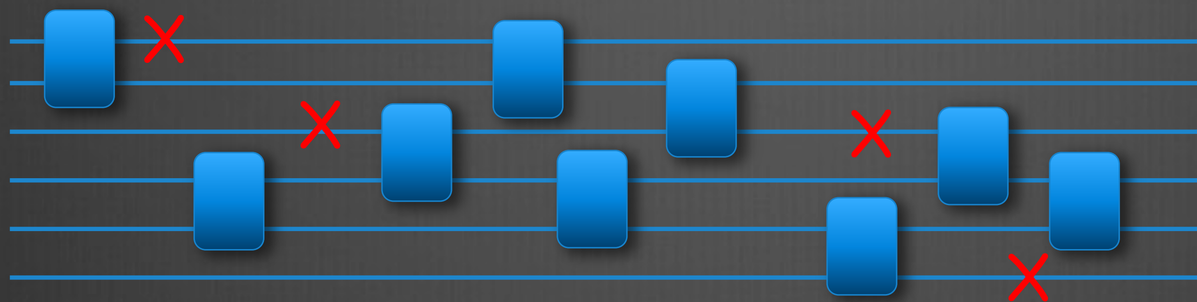

noise rate / gate rate

0

ideal QC

10^{-2}–ish

FTQC possible

?

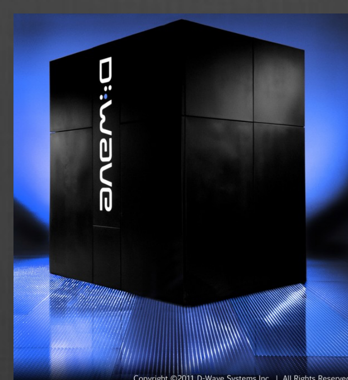

?

≈0.3

classical simulation possible

1

conjectured to exhibit phase transition
(possibly with intermediate phases)

phase transitions?

Complexity smoothly increases with
-entanglement
-correlation length
-# of non-stabilizer gates

Complexity jumps discontinuously with
-noise rate

Today: what about circuit depth?

https://www.pexels.com/photo/cold-glacier-iceland-melting-4797/

random circuit sampling

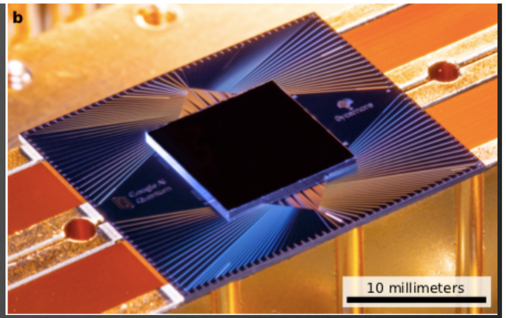

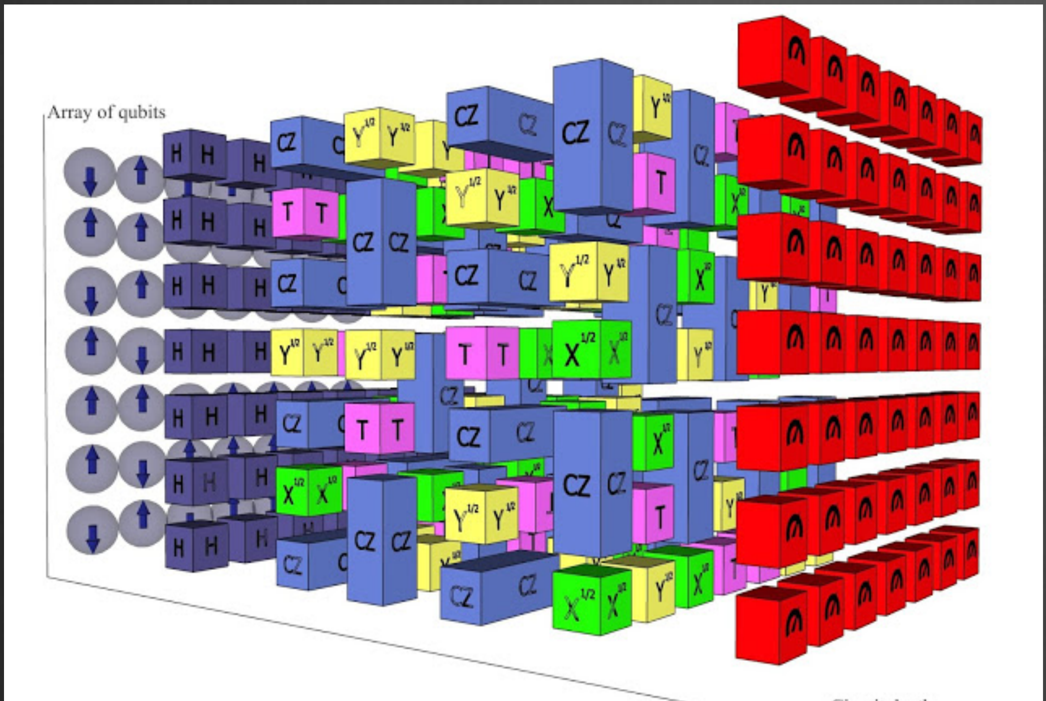

Photo credit: Google Quantum AI Lab

<u>Conjecture:</u>
Output distribution p(z) is hard to sample from on classical computer.

Google used N=53 qubits in 2D geometry with T=20.

Conjecture: $T \geq \sqrt{N}$ → classical simulation time exp(N).
[Aaronson, Bremner, Jozsa, Montanaro, Shepherd, ...]

low-depth circuits

Example: depth-3 “brickwork architecture”

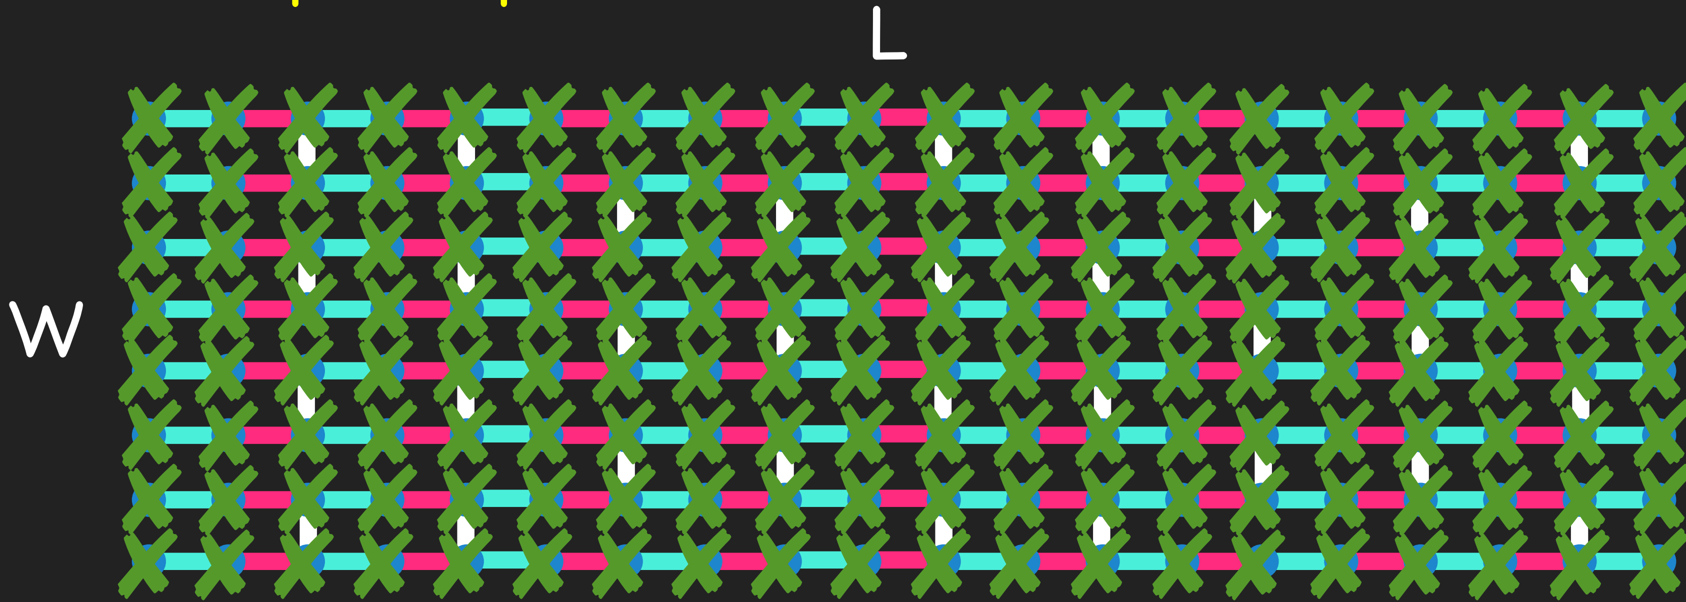

How hard is this sampling problem?

- Assuming non-collapse of PH, for algorithms that make exponentially small error…
  - Sampling is hard in the worst case [Terhal-DiVincenzo `02]
  - Computing *output probabilities* is hard in the average case [Bouland-Fefferman-Nirkhe-Vazirani `18]
- Is this evidence that approximate, average-case sampling (i.e. RCS) is hard?
  - Implication of this work: not necessarily!

tensor contraction

$\langle z_1 z_2 z_3 z_4 z_5 |\ U_9 U_8 U_7 U_6 U_5 U_4 U_3 U_2 U_1\ |00000\rangle =$

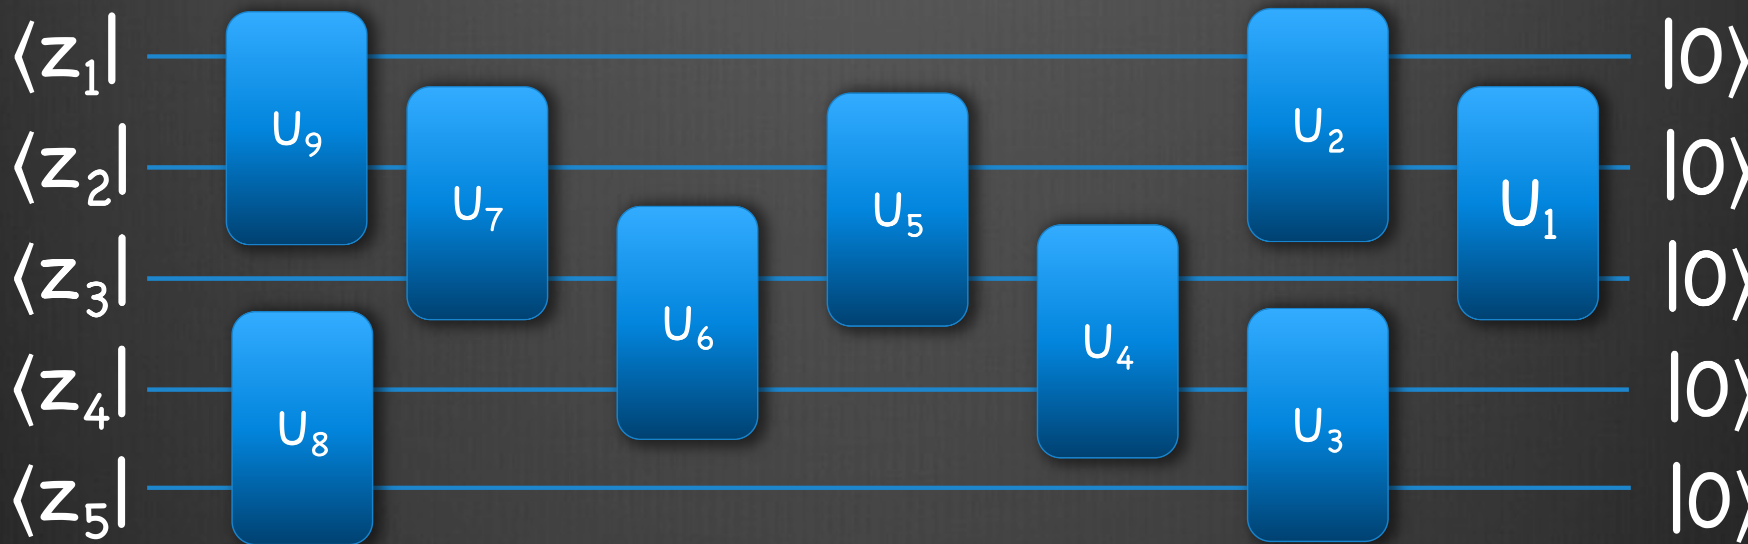

= tensor with 4 indices, each dim 2

tensor contraction:
sum over

tensor contraction in 1-D

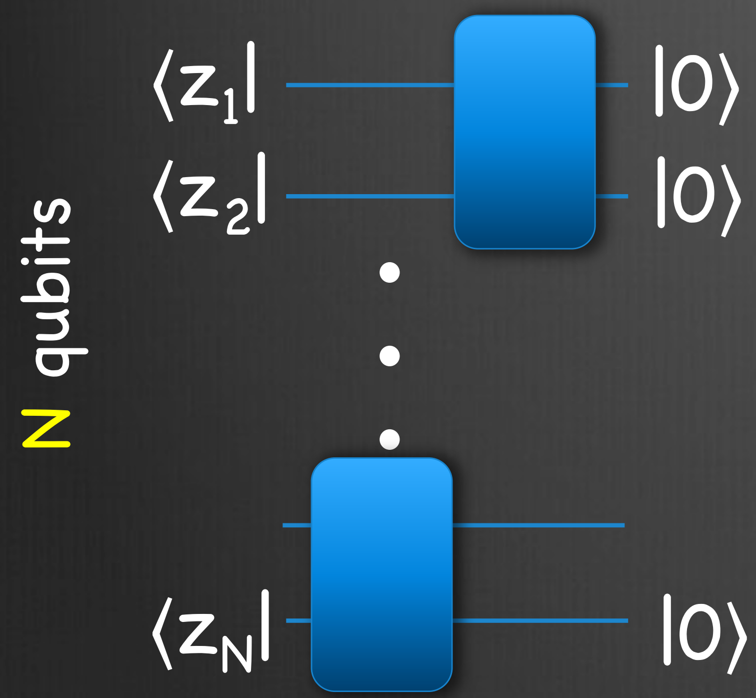

depth T

intermediate tensors can be

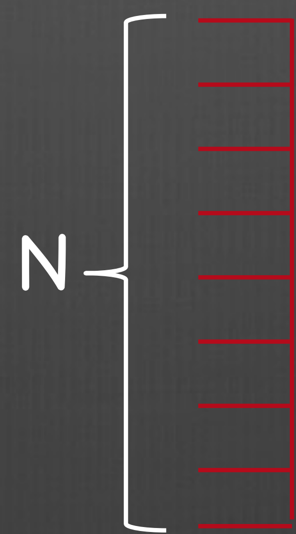

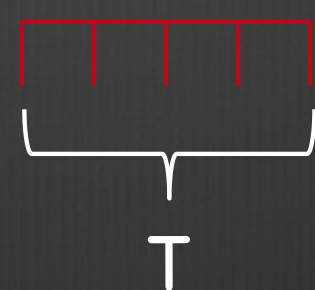

run time:
T exp(N) or
N exp(T)

simulating 2-D circuits

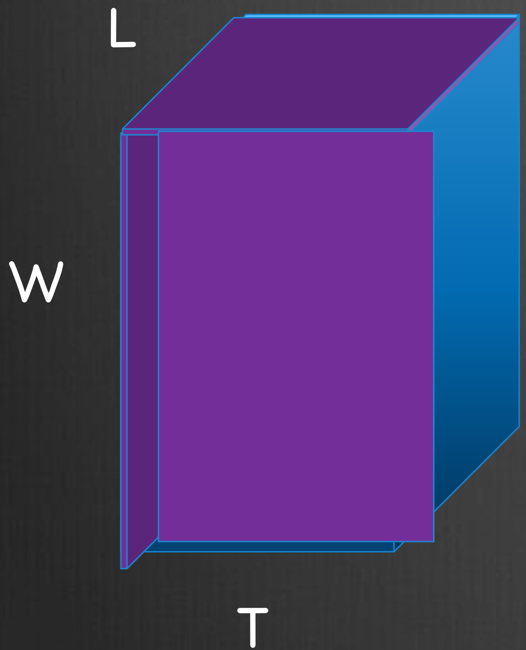

can be simulated in time 2^{LW} or 2^{LT} or 2^{WT}

Depth T=O(1) circuit on $\sqrt{N} \times \sqrt{N}$ grid

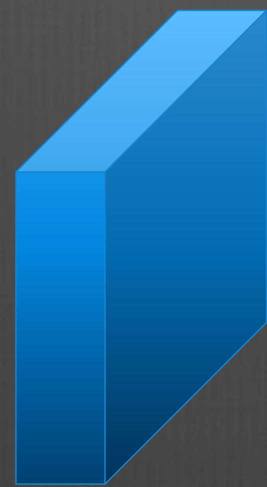

Naively takes time $2^{O(\sqrt{N})}$
≈ $\sqrt{N}$ qubits on line for time $\sqrt{N}$.

But 1-D effective evolution is not unitary.

Entanglement has phase transition from area law -> volume law.

T=3 in area law phase →
$N^{O(1)}$-time classical simulation for approximate sampling of random circuits.
Exact or worst-case is #P-hard.

effective 1-D dynamics

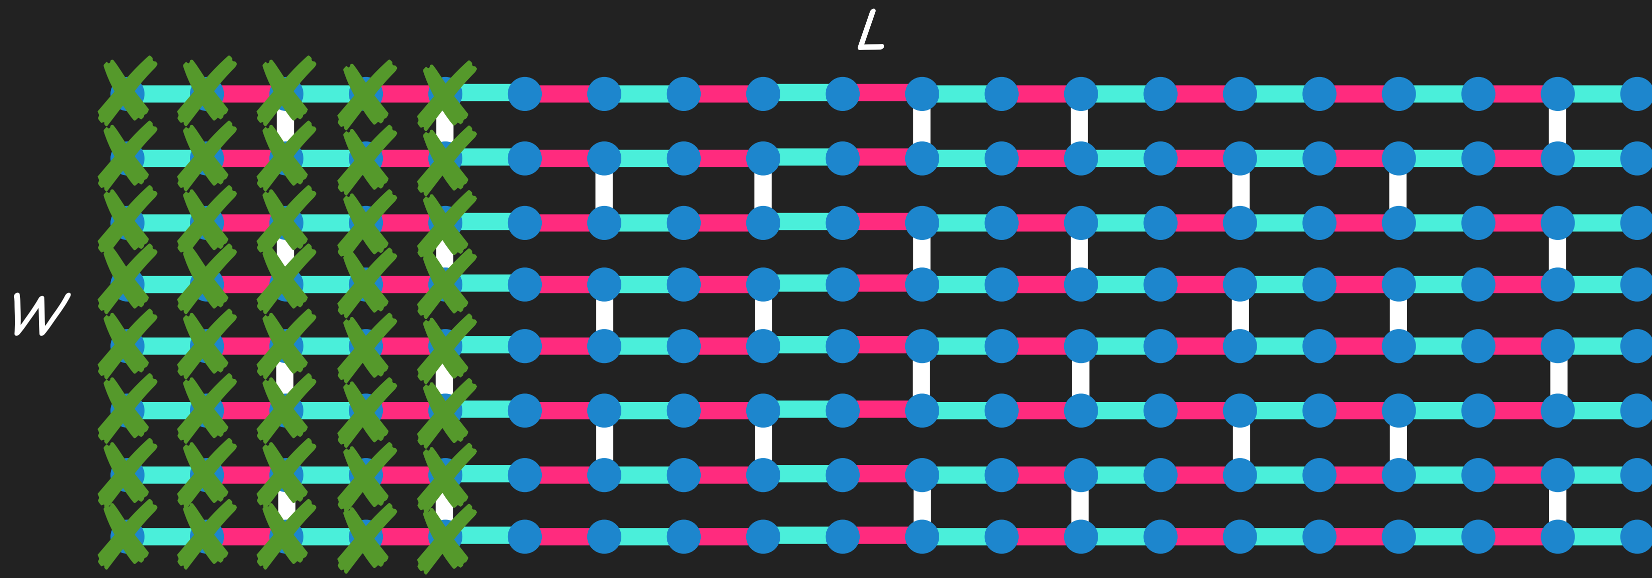

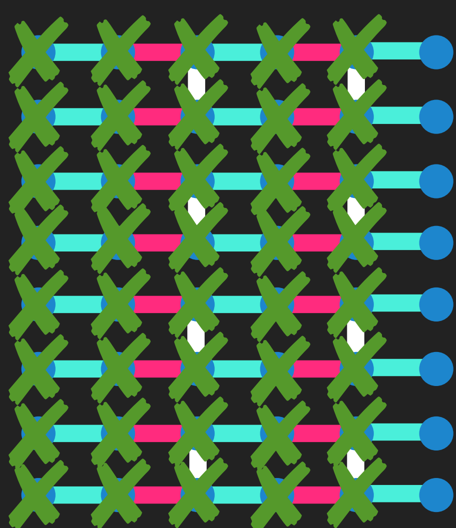

- Need only simulate an “effective 1D dynamics”:

$$|\psi_1\rangle \to |\psi_2\rangle \to \cdots \to |\psi_L\rangle$$

cheaper tensor contraction

T=O(1)

$\sqrt{N} \times \sqrt{N}$ grid

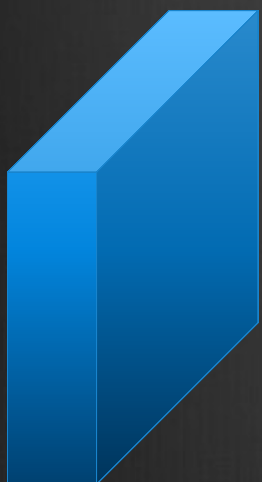

effective
time
evolution

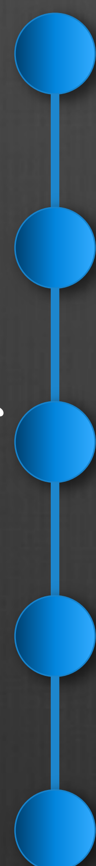

$\sqrt{N}$ qubits
evolving for
$\sqrt{N}$ time

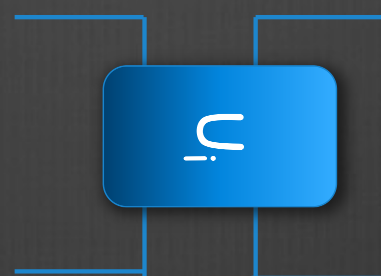

≈

weak
measurement

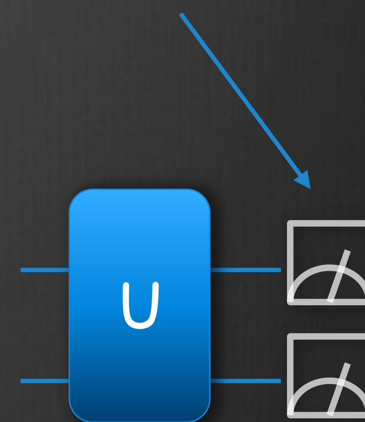

sideways gates
generically not
unitary

Approximate simulation

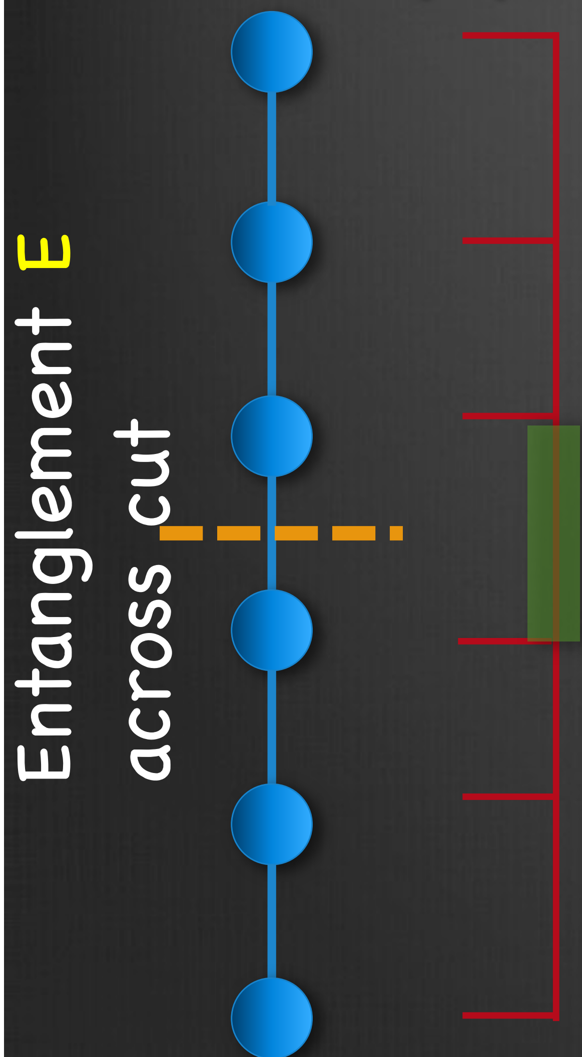

Matrix product state

bond dimension exp(E)

Example:

$$\sum_k |k\rangle|\alpha_k\rangle$$

$$\sum_k \lambda_k \langle k|\langle k|$$

$$\sum_k |k\rangle|\beta_k\rangle$$

$$|\psi\rangle = \sum_{i,j} c_{i,j}|i\rangle \otimes |j\rangle$$

$$= \sum_k \lambda_k |\alpha_k\rangle|\beta_k\rangle$$

Simulation algorithm:

- Do tensor contraction
- Truncate bonds to dim exp(O(E)).

Run-time is $N2^{O(E)}$.

Does the algorithm work?

"Beware of bugs in the above code; I have only proved it correct, not tried it."
--Donald Knuth

1. Yes.
   We tested it and simulated 400x400 grids on a laptop.

2. Probably.
   We proved a phase transition in something like the effective entanglement.

3. Sometimes.
   The extended brickwork architecture is #P-hard to simulate exactly but our algorithm is proven to work on it.

stat mech model

Qubits → dim-q particles.

$E[p^{k-1}\mathrm{tr}[\rho_A{}^k]]$ = partition function

$Wg^{-1}(\sigma,\tau)$
$= G(\sigma,\tau)$
$= q^{-\mathrm{dist}(\sigma,\tau)}$

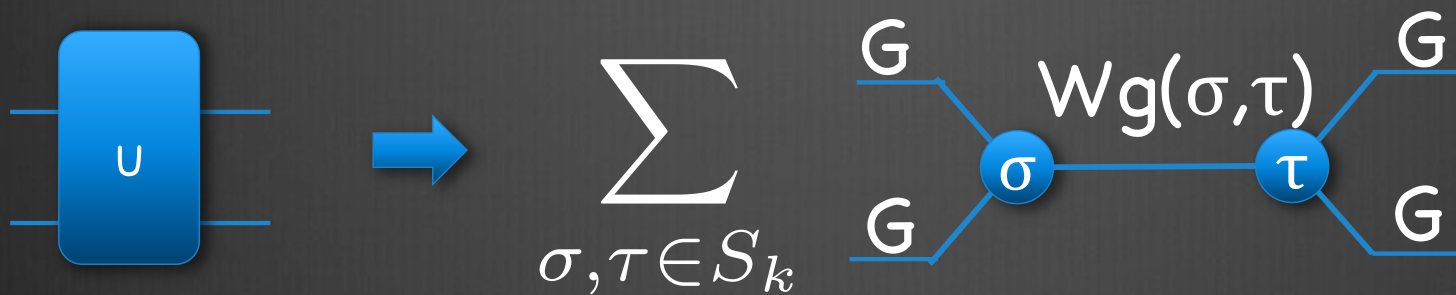

k=2 ↔ anisotropic Ising

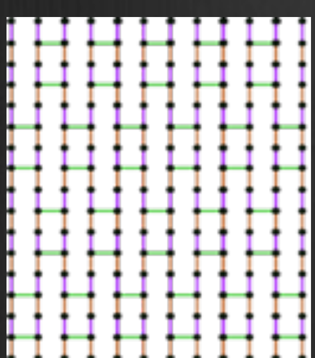

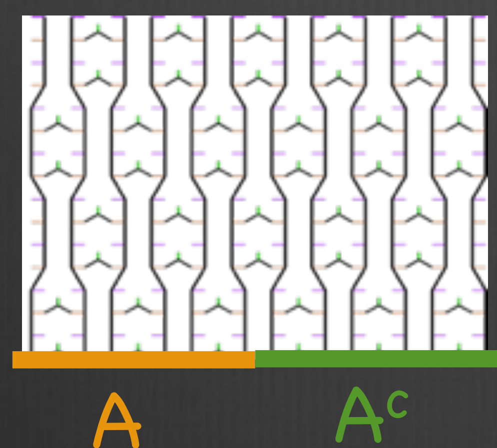

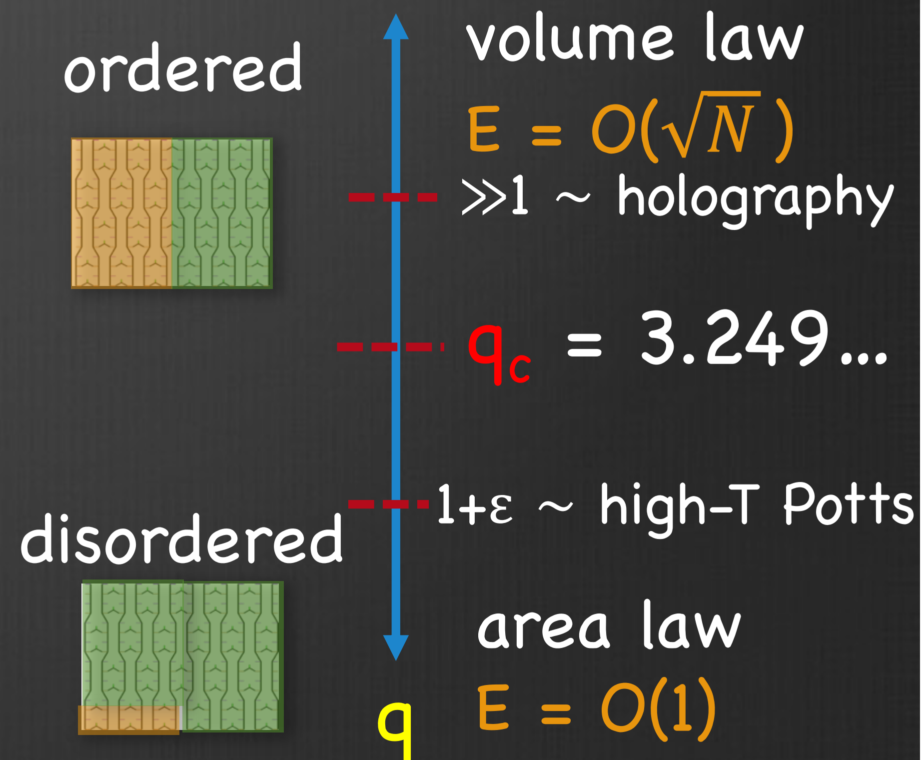

conjecture: decimating yields nonnegative weights for all k, q

random tensor networks

q = local dim, $E[\mathrm{tr}[\rho^k]]$

$\sum_{\sigma,\tau\in S_k}$

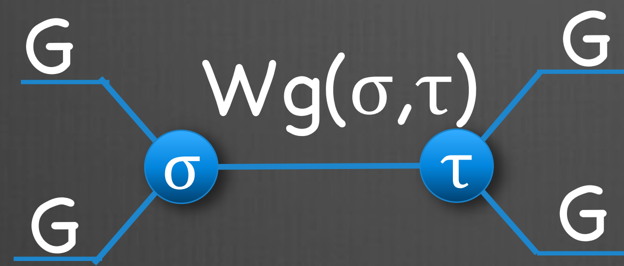

$Wg^{-1}(\sigma,\tau)$
$= G(\sigma,\tau)$
$= q^{-\mathrm{dist}(\sigma,\ \tau)}$

- $E_{U\sim Haar}[U^{\otimes k} \otimes (U^*)^{\otimes k}] = \Sigma_{\sigma,\tau}\, Wg(\sigma,\tau)\, |\sigma\rangle\langle\tau|$
- $E_{G\sim Gaussian}[G^{\otimes k} \otimes (G^*)^{\otimes k}] = \Sigma_{\sigma}\, |\sigma\rangle\langle\sigma|$
- $\langle\sigma|\tau\rangle = G(\sigma,\tau)$
- $\|\, G - I \,\|_{op}\, , \|\, Wg - I \,\|_{op} \le k^2 / q$

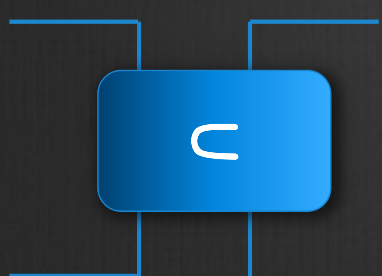

$\approx$

$k^2 \ll q$

Patching Algorithm

- Regions of size k can be computed in time exp(O(k))
- Zero correlation beyond distance O(d)
- Can fill in holes with error ~I(existing : new | border)
(following Brandao-Kastoryano, Swingle, Kim, Kitaev)

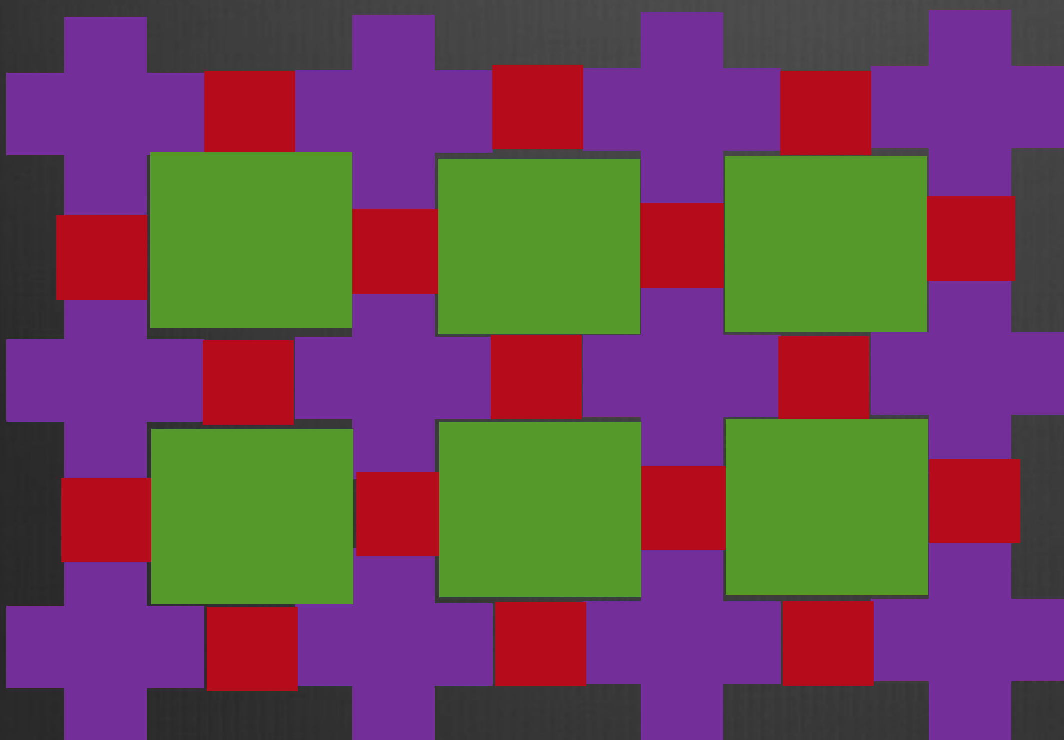

Open questions

- Rigorously prove the location of the phase transition and the correctness of the algorithm.
- Random tensor networks with low bond dimension
- Universality classes in random circuits?
- (time-independent) Hamiltonian versions?
- Where exactly is the boundary between easy and hard?

Thanks!

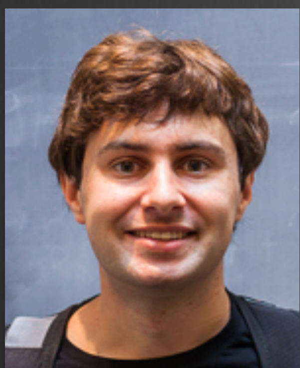

John Napp

Alex
Dalzell

Rolando
La Placa

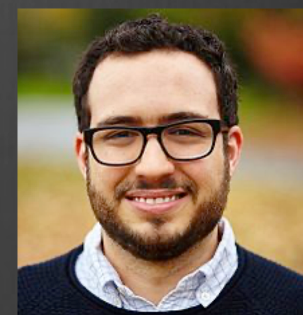

Fernando
Brandão

arXiv:2001.00021