Quantum Algorithmic Breakeven: on scaling up with noisy qubits

ISQST Quantum Symposium

Purdue

April 23, 2019

Daniel Lidar

USC

Algorithmic Scaling

Run an algorithm with exponential quantum speedup (e.g., quantum simulation) on quantum hardware

ideal case: no decoherence

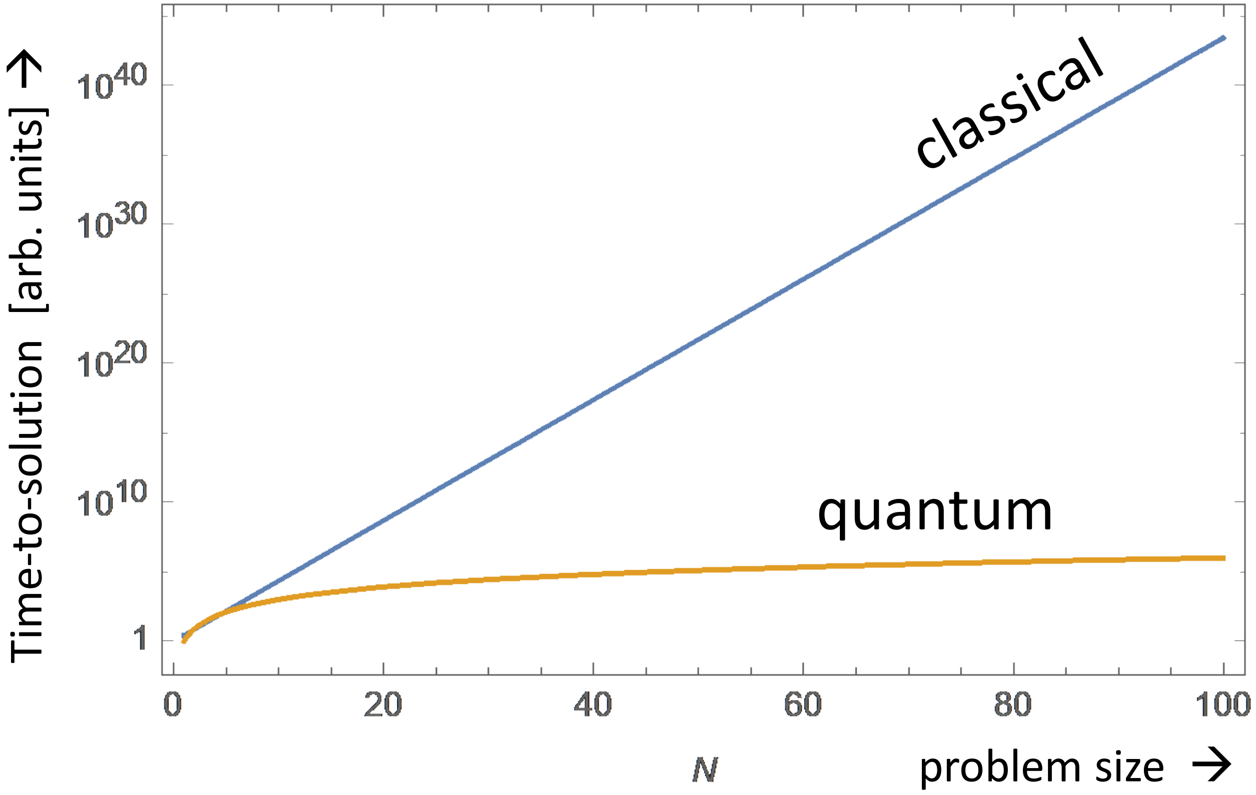

quantum scaling advantage clear for all *N*

So far never observed on quantum hardware

Algorithmic Scaling

Run an algorithm with exponential quantum speedup (e.g., quantum simulation) on quantum hardware

with decoherence

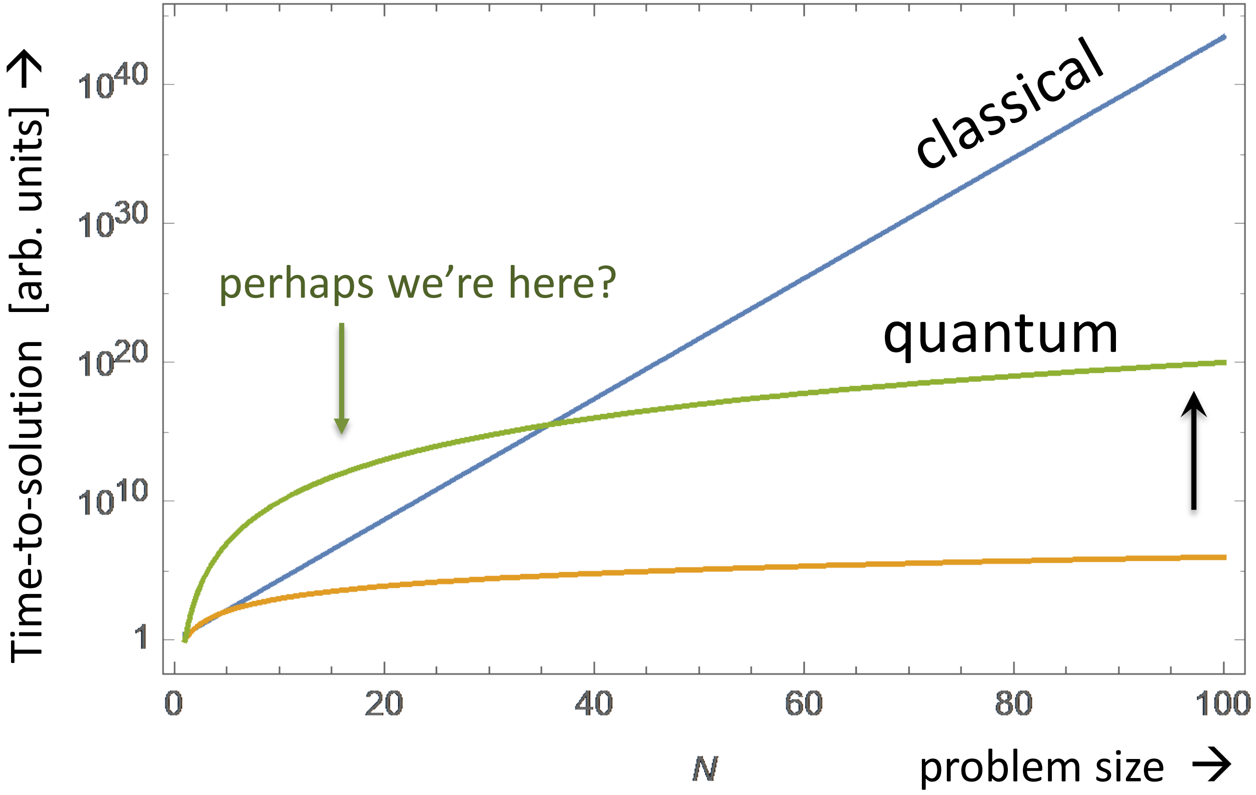

quantum scaling advantage only clear for large N

Algorithmic Scaling

Run an algorithm with exponential quantum speedup (e.g., quantum simulation) on quantum hardware

with more decoherence

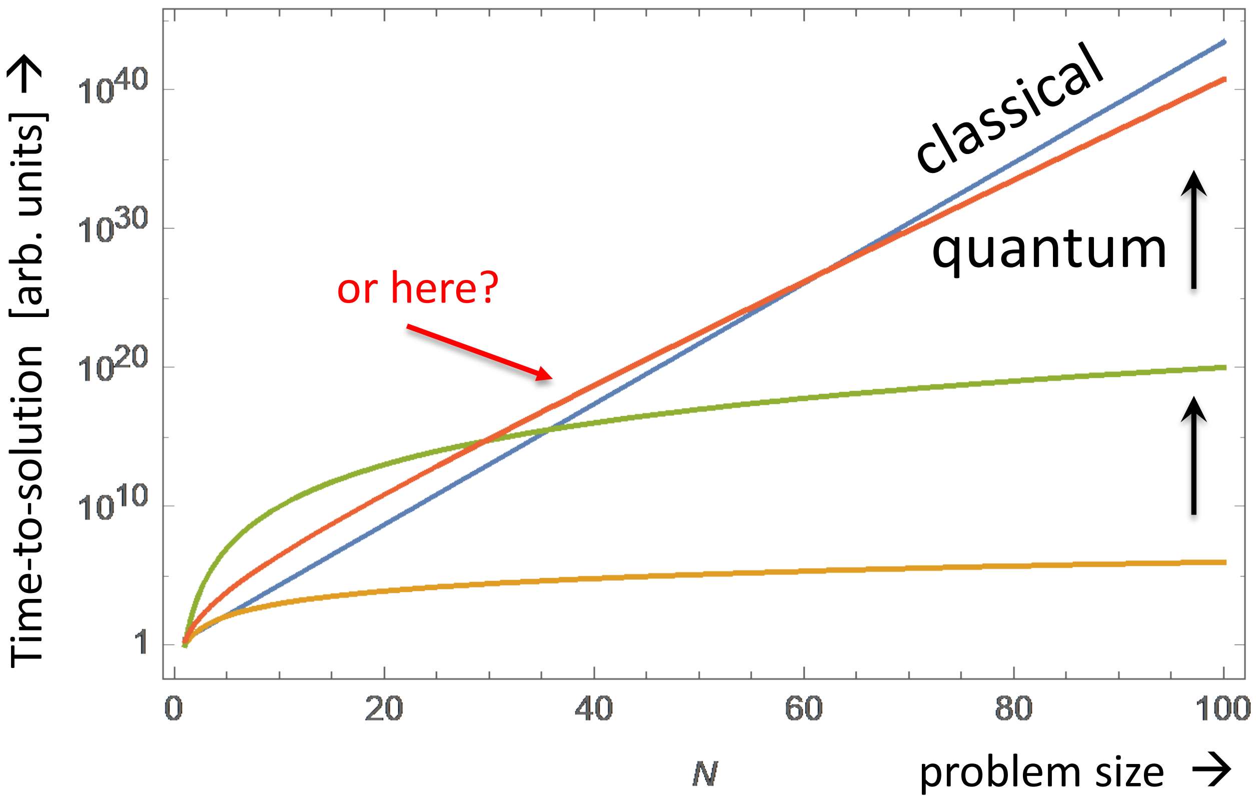

quantum scaling advantage only clear for even larger *N*

Algorithmic Scaling

Run an algorithm with exponential quantum speedup (e.g., quantum simulation) on quantum hardware

with too much decoherence

quantum scaling
disadvantage

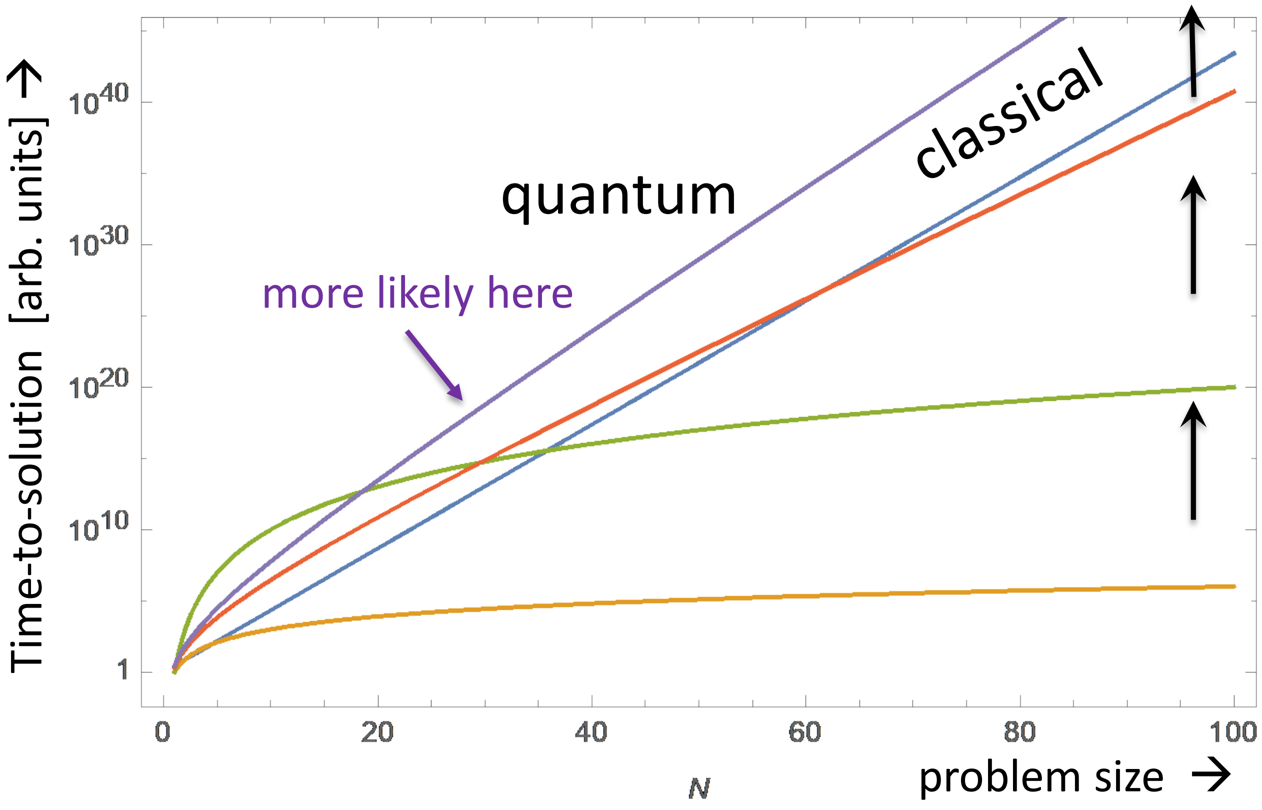

Algorithmic Scaling

Run an algorithm with exponential quantum speedup (e.g., quantum simulation) on quantum hardware

with too much decoherence

quantum blowup

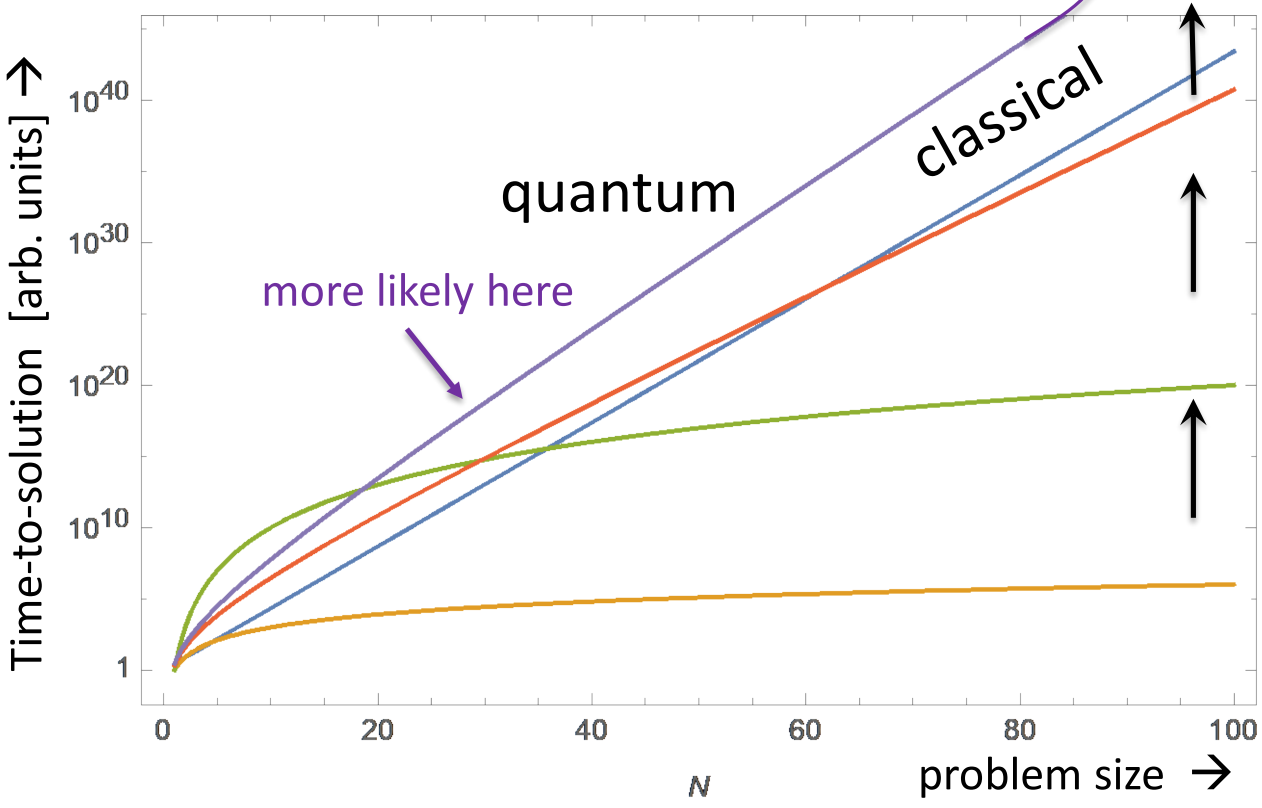

Trying for speedup: D-Wave 2000Q vs simulated annealing

Problem: find the ground state of a certain family of hard spin-glass instances of $N \leq 2048$ spins

Trying for speedup: D-Wave 2000Q scaling advantage against simulated annealing

Problem: find the ground state of a certain family of hard spin-glass instances of $N \leq 2048$ spins

Simulated Annealing (SA) with single-spin updates
D-Wave 2000Q unequivocally beats SA

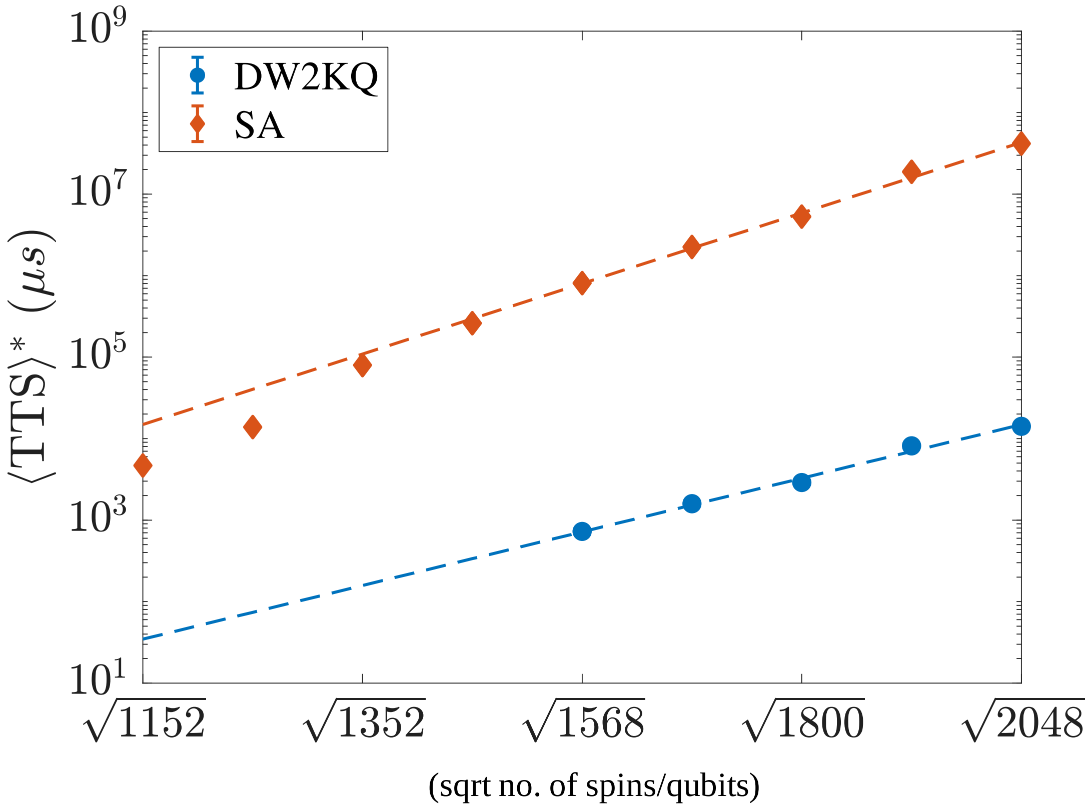

Fit curves to

$$a \exp(bL)$$

Solver	b [95% CI]
DW2KQ	0.760 ± 0.017
SA	1.002 ± 0.066

scaling of the median of the instance distribution

First example of an experimental quantum scaling advantage on a non-trivial problem relative to a generic classical algorithm

T. Albash, DL, Phys. Rev. X 8, 031016 (2018)

D-Wave 2000Q scaling disadvantage against better classical heuristic algorithms

Problem: find the ground state of a certain family of hard spin-glass instances of $N \leq 2048$ spins

Simulated Annealing (SA) with single-spin updates
D-Wave 2000Q unequivocally beats SA
Classical Spin Vector Monte Carlo (SVMC) beats both
Classical simulated quantum annealing (SQA) beats all

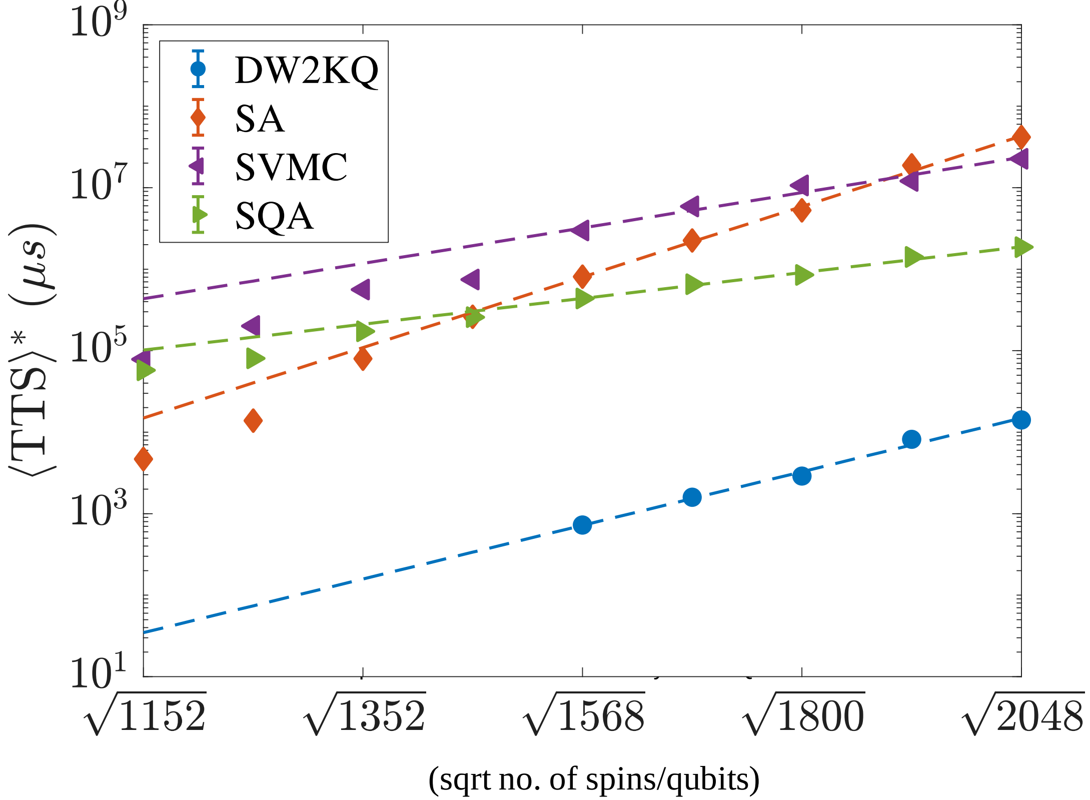

Fit curves to

$$a \exp(bL)$$

Solver	b [95% CI]
DW2KQ	0.760 ± 0.017
SA	1.002 ± 0.066
SVMC	0.500 ± 0.029
SQA	0.370 ± 0.052

scaling of the median of the instance distribution

T. Albash, DL, Phys. Rev. X 8, 031016 (2018)

Why no speedup?

D-Wave quantum annealers are NISQ-era devices
Current NISQ-era devices are indeed Noisy

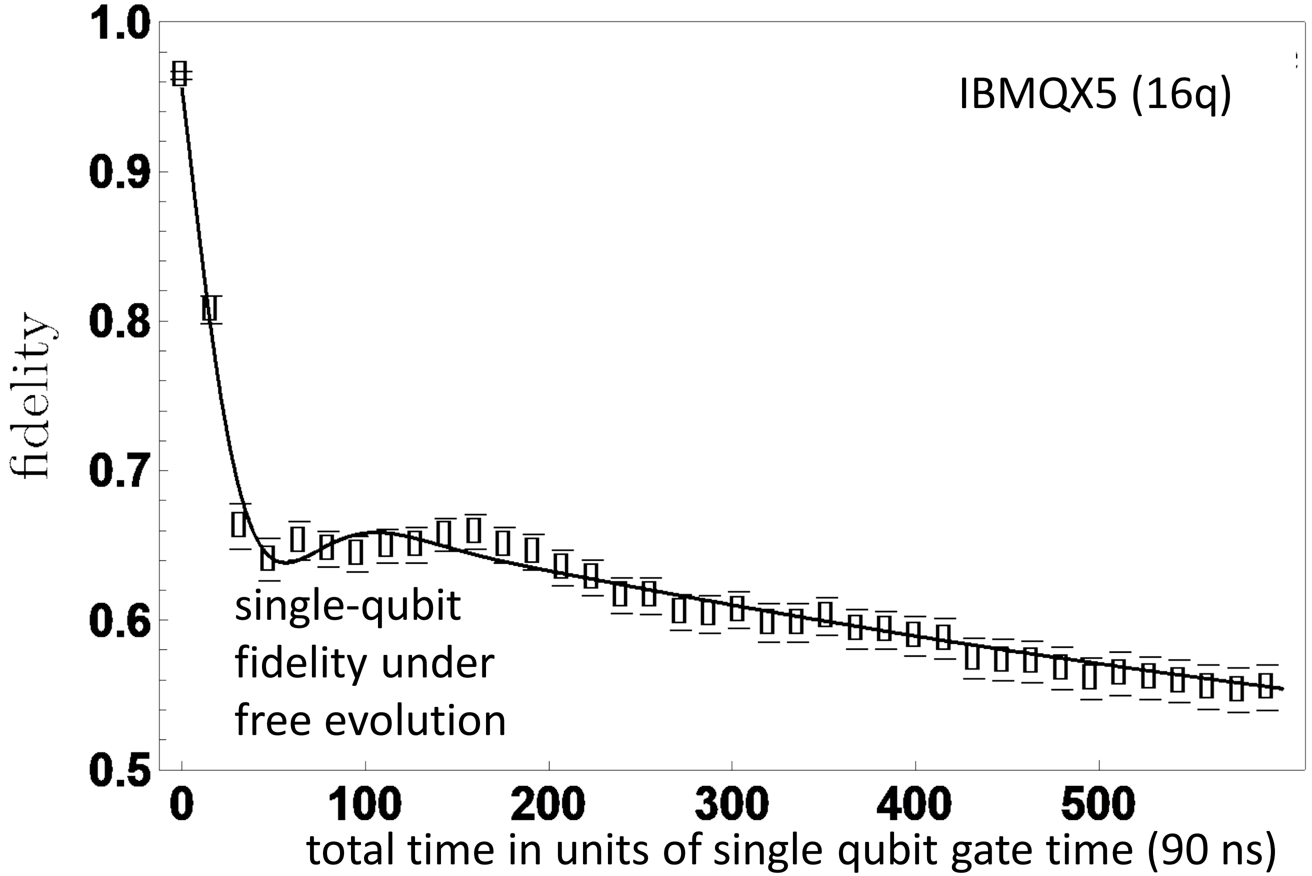

B. Pokharel, N. Anand, B. Fortman, DL, Phys. Rev. Lett. 121, 220502 (2018)

Why no speedup?

D-Wave quantum annealers are NISQ-era devices
Current NISQ-era devices are indeed noisy
But improvements are possible via error suppression methods

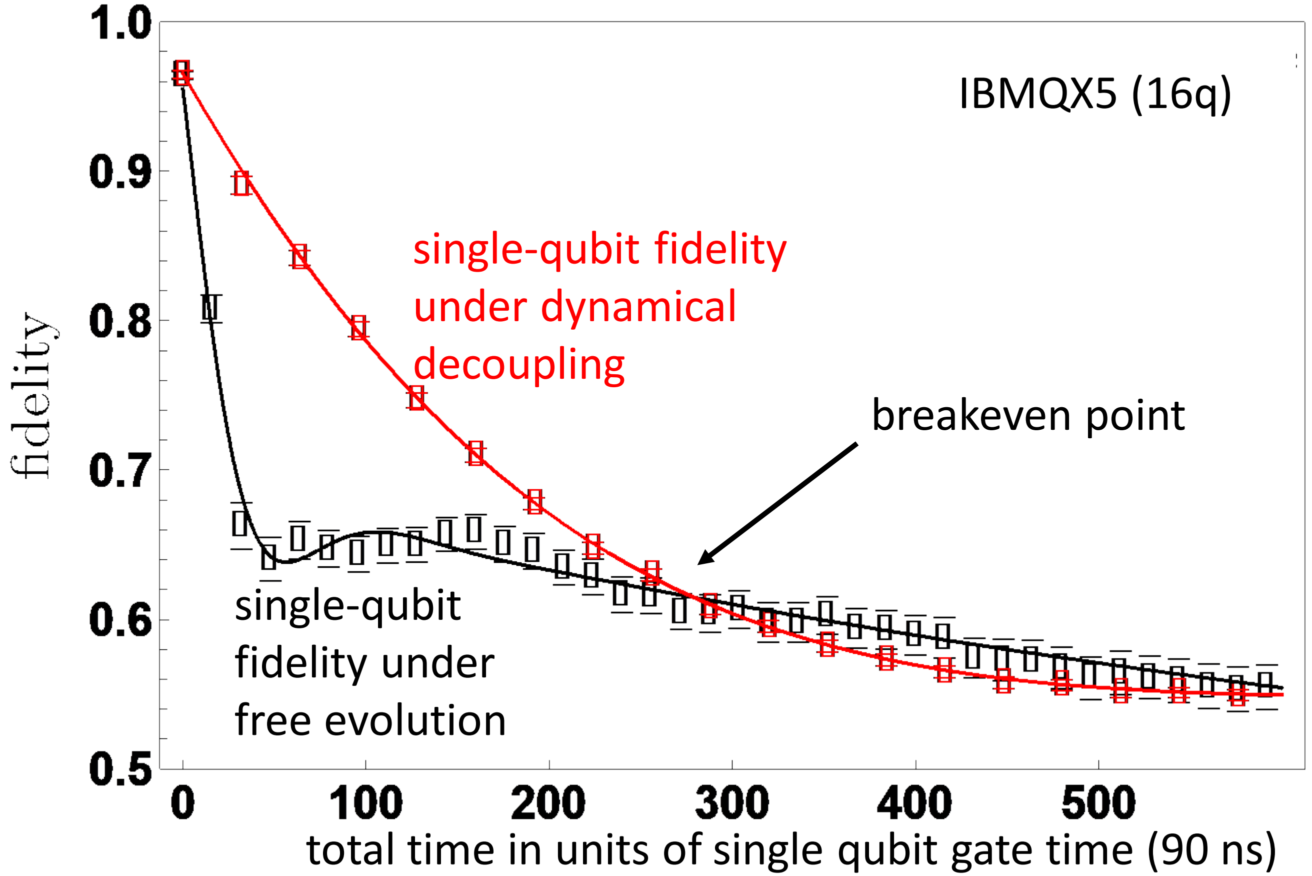

how well can we do with error correction?

B. Pokharel, N. Anand, B. Fortman, DL, Phys. Rev. Lett. 121, 220502 (2018)

Enter Quantum Error Correction (QEC)

Run an algorithm with exponential quantum speedup (e.g., quantum simulation) on quantum hardware

with decoherence + QEC

Algorithmic Success with QEC

Run an algorithm with exponential quantum speedup (e.g., quantum simulation) on quantum hardware

with decoherence + QEC

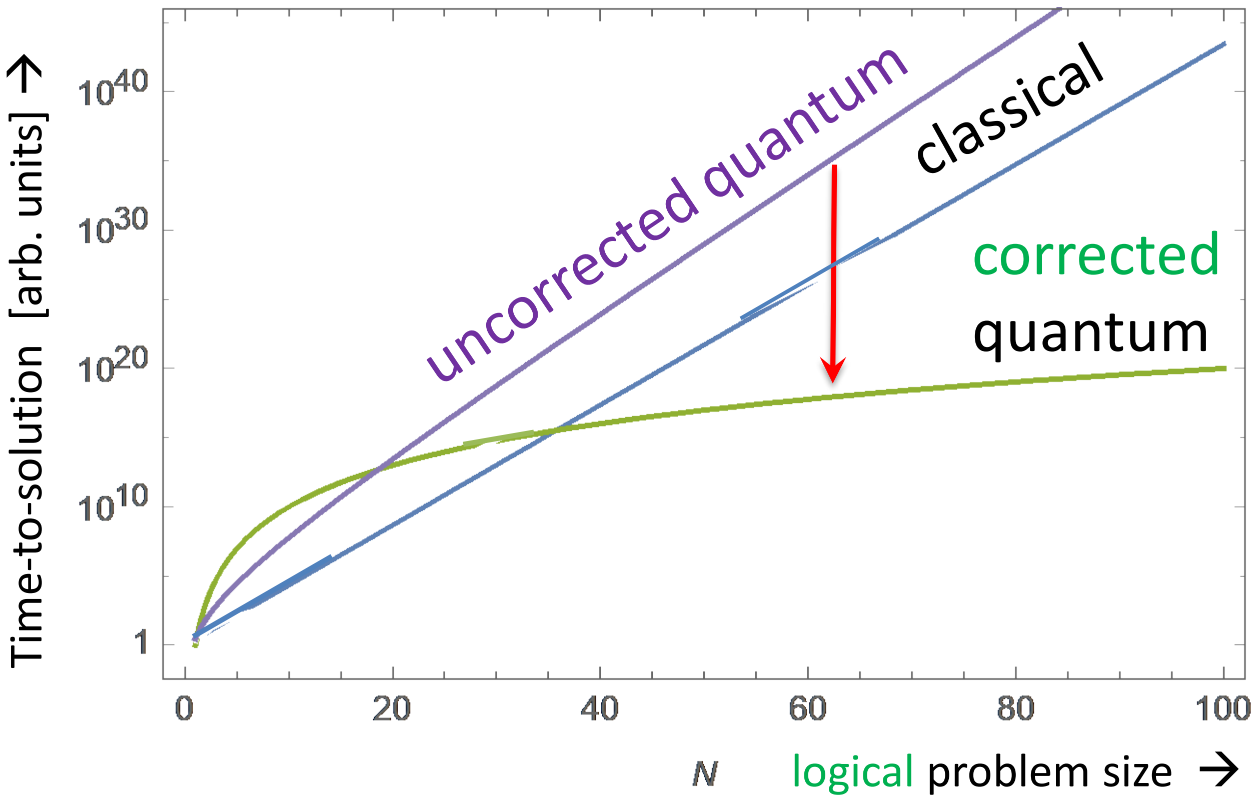

quantum scaling advantage returns

achieving this with quantum hardware =

Algorithmic success with QEC: corrected quantum scaling is better than both uncorrected quantum & classical

Algorithmic Breakeven with QEC

Run an algorithm with exponential quantum speedup (e.g., quantum simulation) on quantum hardware

with decoherence + QEC

More modest:

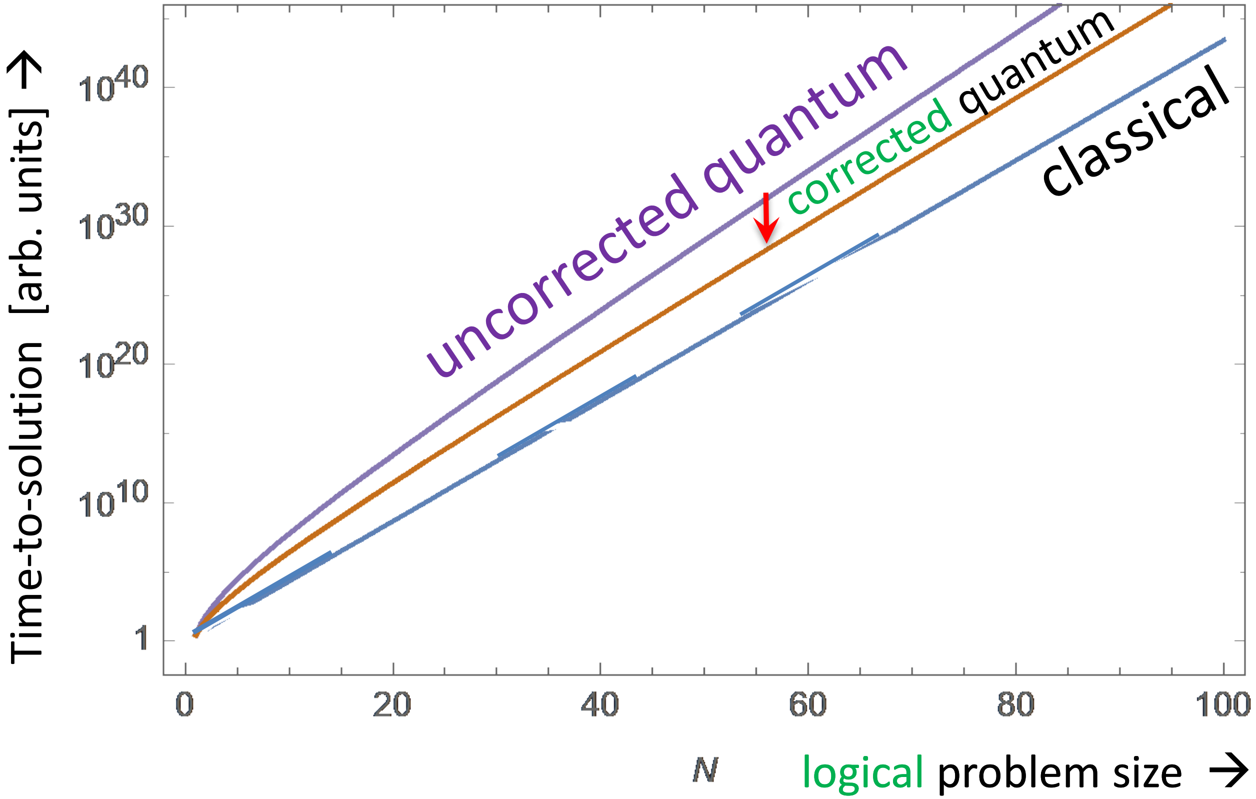

Can **this** be achieved with *existing* quantum hardware?

Algorithmic breakeven with QEC: corrected quantum scaling is better than uncorrected quantum, but not necessarily better than classical

Algorithmic breakeven with quantum annealing

Brief intro to D-Wave processors

They are programmable quantum annealers;
Designed to solve optimization problems
formulated as Ising spin-glass Hamiltonians:

Given the Hamiltonian

$$H_{\text{Ising}} = \sum_{i=1}^{N} h_i \sigma_i^z + \sum_{i<j} J_{ij} \sigma_i^z \sigma_j^z$$

Find the minimizing spin configuration $\{\sigma_i^z = \pm 1\}$

Solve by adiabatically evolving the transverse field Ising Hamiltonian

$$H(t) = A(t) \sum_{i=1}^{N} \sigma_i^x + B(t) H_{\text{Ising}}$$

from ground state of $\sum_{i=1}^{N} \sigma_i^x$ to ground state of H_{Ising}

4 generations so far:

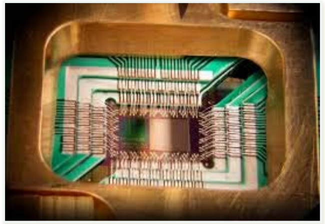

Oct 2011

D-Wave 1:
N=128 (USC)
K_{17} for ideal
K_{14} for actual

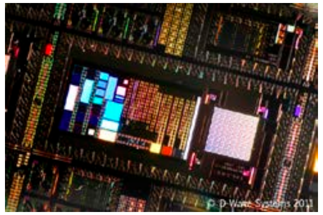

March 2013

D-Wave 2:
N=512 (USC, NASA)
K_{33} for ideal
K_{32} for actual

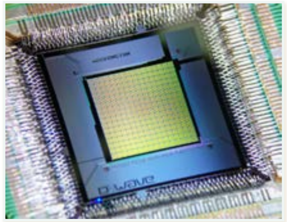

March 2016

D-Wave 2X:
N=1152 (USC, NASA, LANL)
K_{49} for ideal
K_{44} for actual

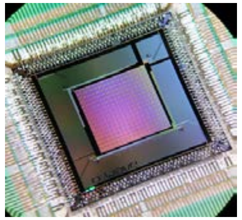

Sep 2017

D-Wave 2000Q
N=2048 (NASA, LANL)
K_{65} for ideal

Brief intro to D-Wave processors

- Use superconducting Nb flux qubits, each coupled to up to 6 other qubits ("Chimera graph")
- $T_1, T_2 \sim 10 - 100ns$, annealing time $t_f \geq 1\mu s$
- minimum gap(H) can be $\ll T \sim 10mK$
- "Stoquastic": efficient classical simulation possible in many cases using Quantum Monte Carlo

How quantum? "somewhat"

A testbed for algorithmic scaling with noisy qubits and error correction

Chimera graph:
square grid of $L \times L$ unit cells
number of qubits $N = 8L^2$

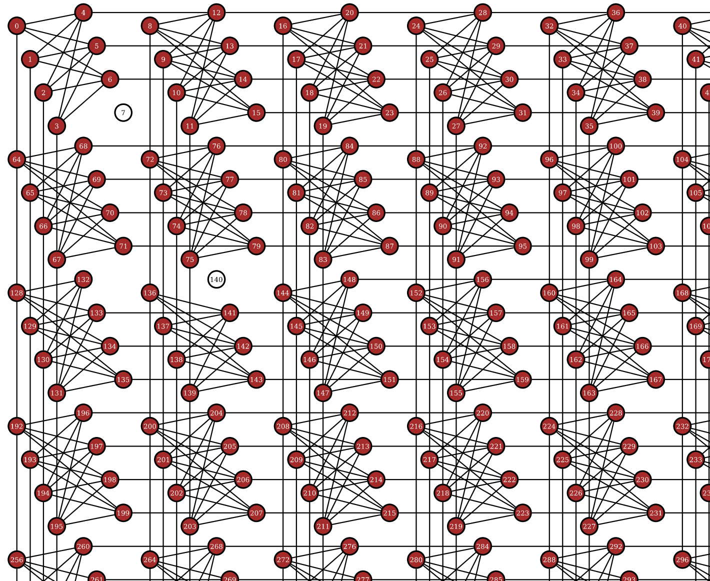

Algorithmic breakeven with quantum annealing

Problem:
find ground state energy of

$$H_{\text{Ising}} = \sum J_{ij}\sigma_i^z\sigma_j^z$$

with random

$$J_{ij} \in \left\{\pm\frac{1}{6},\cdots,\pm\frac{5}{6},\pm1\right\}$$

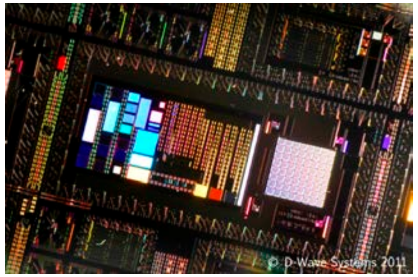

Run on a D-Wave 2
(503 qubits)

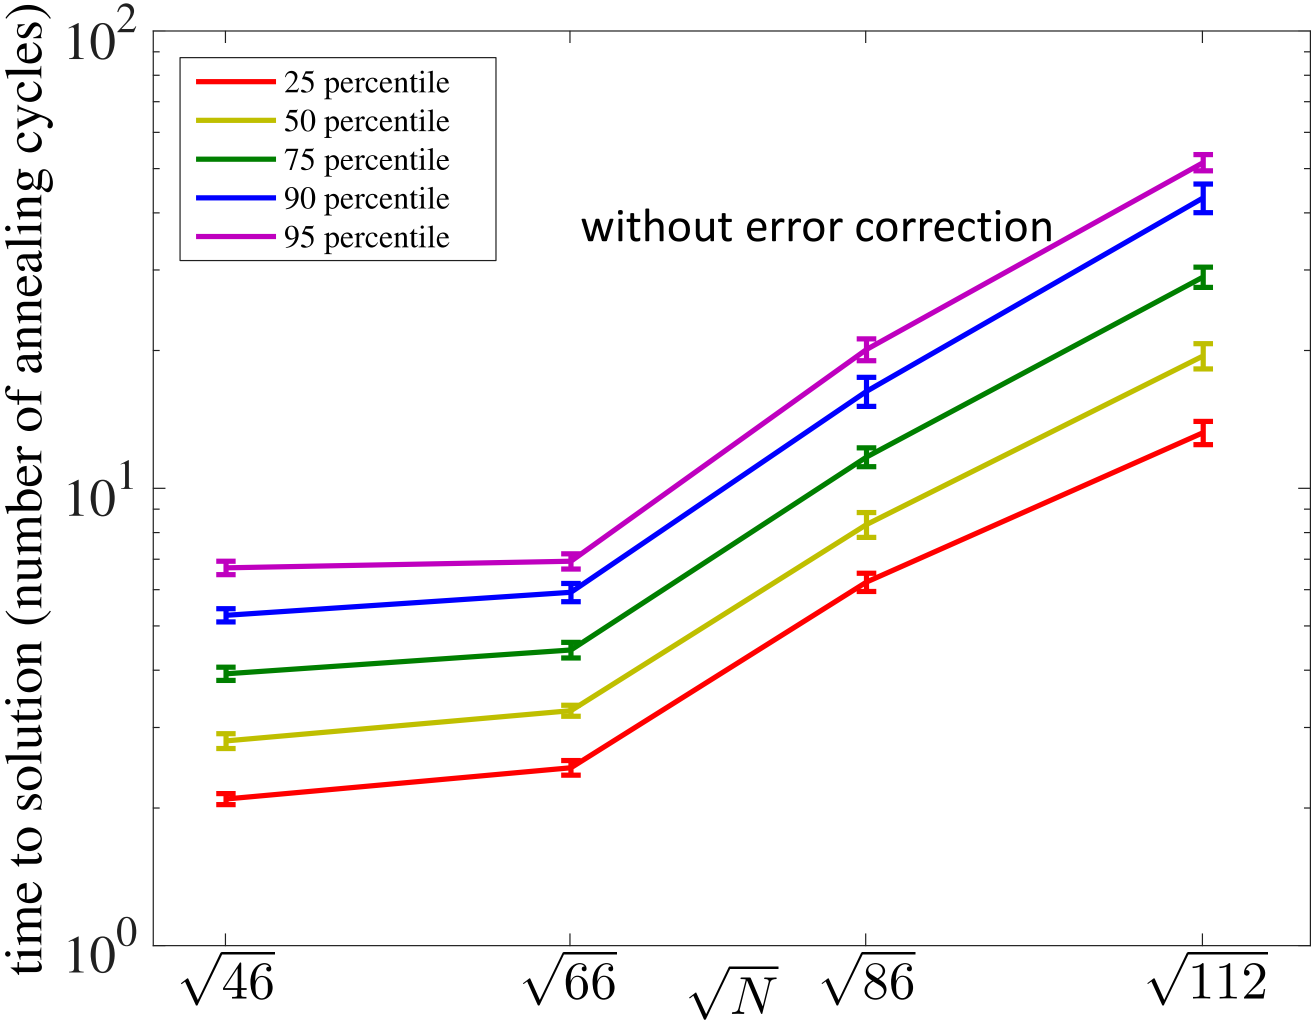

no. of physical qubits

K. Pudenz, T. Albash, DL, Nature Comm. 5, 3243 (2014); PRA 91, 042302 (2015)

Algorithmic breakeven with quantum annealing ✓

Problem:
find ground state energy of

$$H_{\text{Ising}} = \sum J_{ij} \sigma_i^z \sigma_j^z$$

with random

$$J_{ij} \in \left\{ \pm\frac{1}{6}, \cdots, \pm\frac{5}{6}, \pm 1 \right\}$$

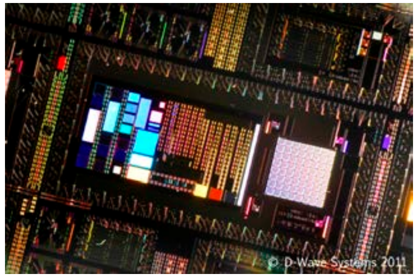

Run on a D-Wave 2
(503 qubits)

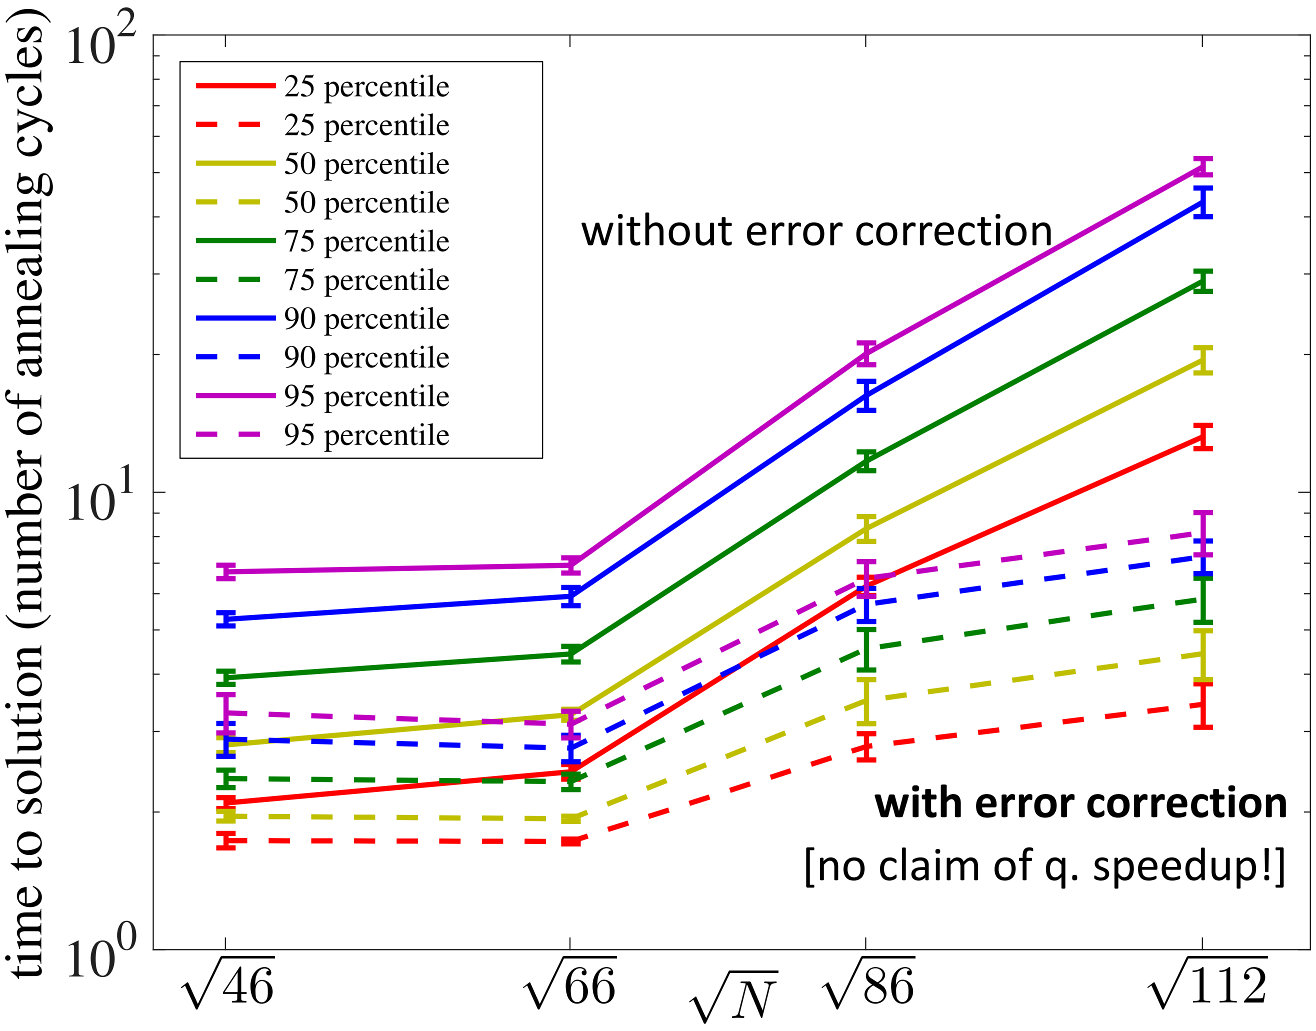

no. of physical (unc.) or logical (cor.) qubits

Quantum
Annealing
Correction

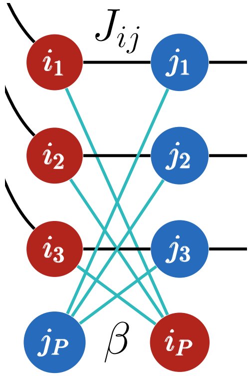

QAC works by:
- introducing energy penalty against excitations
- majority-vote decoding of logical qubits

K. Pudenz, T. Albash, DL, Nature Comm. 5, 3243 (2014); PRA 91, 042302 (2015)

Up the ante: analog control errors

Ideally, a quantum annealer evolves adiabatically according to $H(t) = A(t)H_X + B(t)H_P$

$$\downarrow H_X = H_{\text{transverse}} = \sum_{j \in V} \sigma_j^x$$

$$\uparrow H_P = H_{\text{Ising}} = \sum_{j \in V} h_j \sigma_j^z + \sum_{(i,j) \in E} J_{ij} \sigma_i^z \sigma_j^z$$

programmable

In reality: analog errors → **J-chaos**

H implemented ≠ H intended

$$+ \sum_j \delta h_j \sigma_i^z + \sum_{ij} \delta J_{ij} \sigma_i^z \sigma_j^z$$

$$\delta h_i, \delta J_{ij} \sim \mathcal{N}(0, \eta)$$

D-Wave 2X chip intrinsic analog errors: $\frac{\delta h}{h_{\max}}, \frac{\delta J}{J_{\max}} \sim N(0, 0.03)$

Let's add artificial Gaussian noise $\delta J_{ij} \sim N(0, \eta_{int} + \eta)$ on top of intrinsic analog errors: η_{int}=0.03

$\eta \in$ {0.03, 0.05, 0.07, 0.10, 0.15}

Time-to-solution as a function of problem size and noise

Time-to-solution (# of repetitions R at $t_f = 5\mu s$) for median instances

without error correction

number of runs required to find the ground state

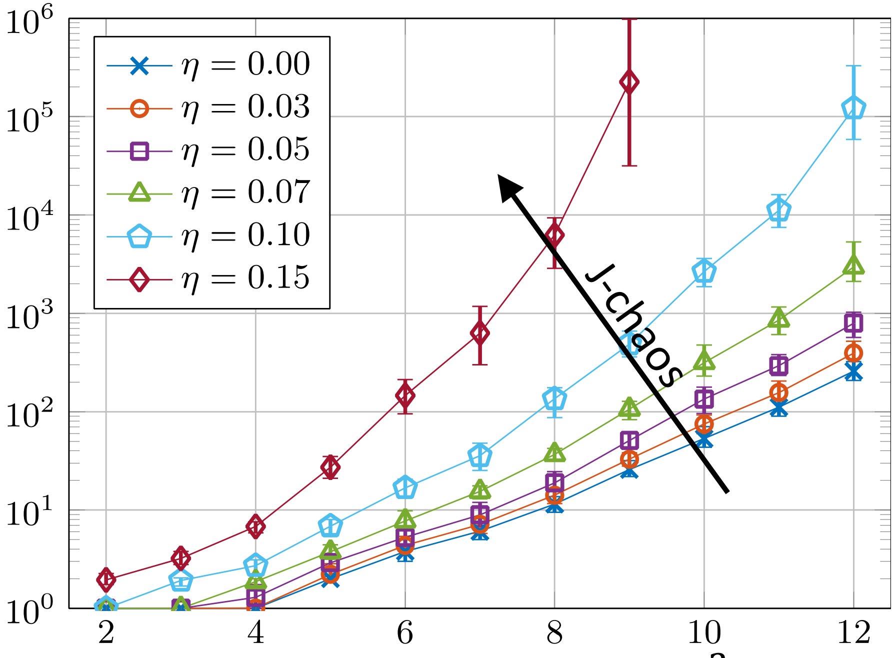

L $[N = 8L^2]$ $N = 1152$

run on USC's DW2X

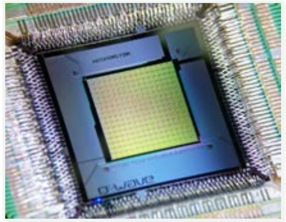

A. Pearson, A. Mishra, DL, I. Hen, arXiv:1907.12678

Time-to-solution as a function of problem size and noise

Time-to-solution (# of repetitions R at $t_f = 5\mu s$) for median instances

number of runs required to find the ground state

without error correction

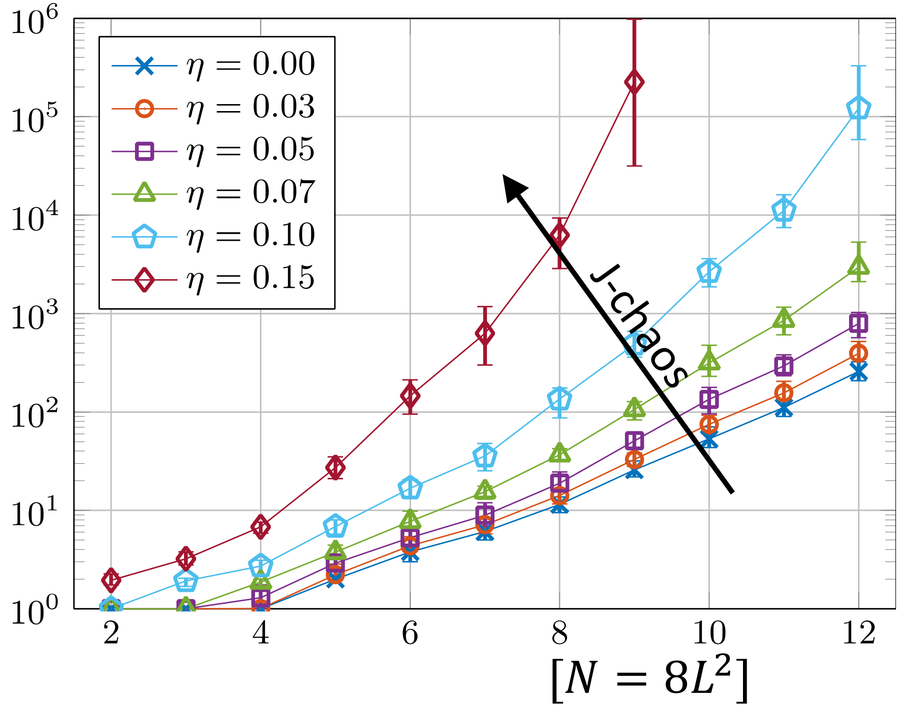

with error correction

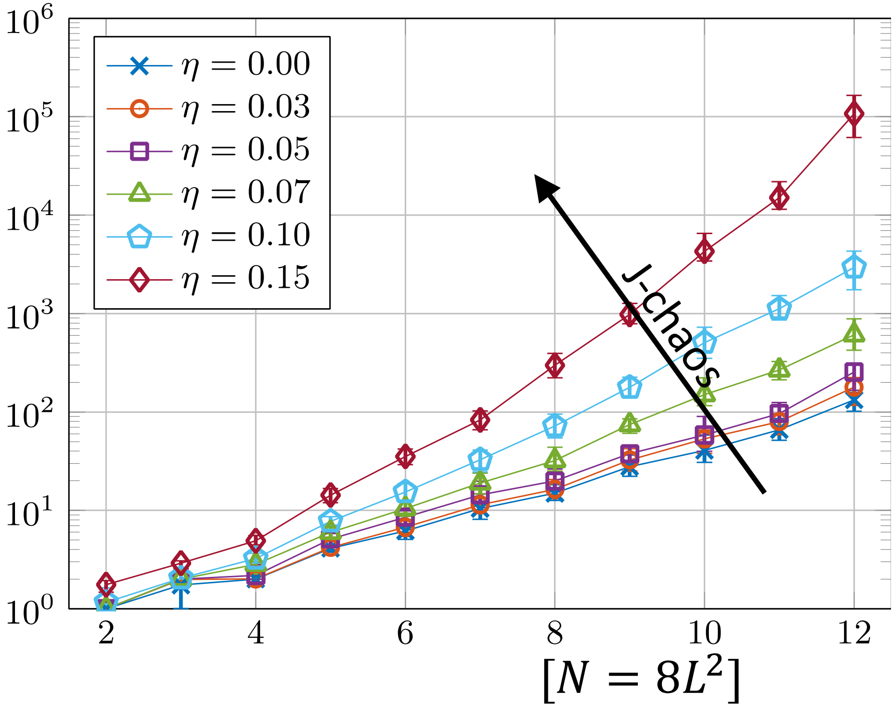

Quantum Annealing Correction

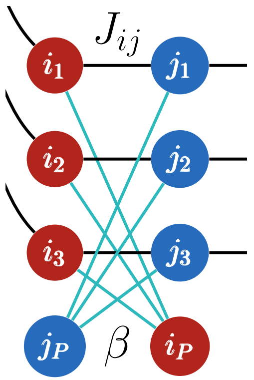

run on USC's DW2X

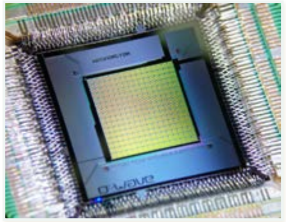

scaling with Quantum Annealing Correction is better than without
but how good is this scaling?

perform data collapse,
extract finite-size scaling

algorithmic breakeven
error-corrected scaling that is better than uncorrected on a non-trivial computational problem

A. Pearson, A. Mishra, DL, I. Hen, arXiv:1907.12678

Time-to-solution data collapse & finite-size scaling:

fit to $R = e^{a(\eta^2+b^2)^c} L^d$

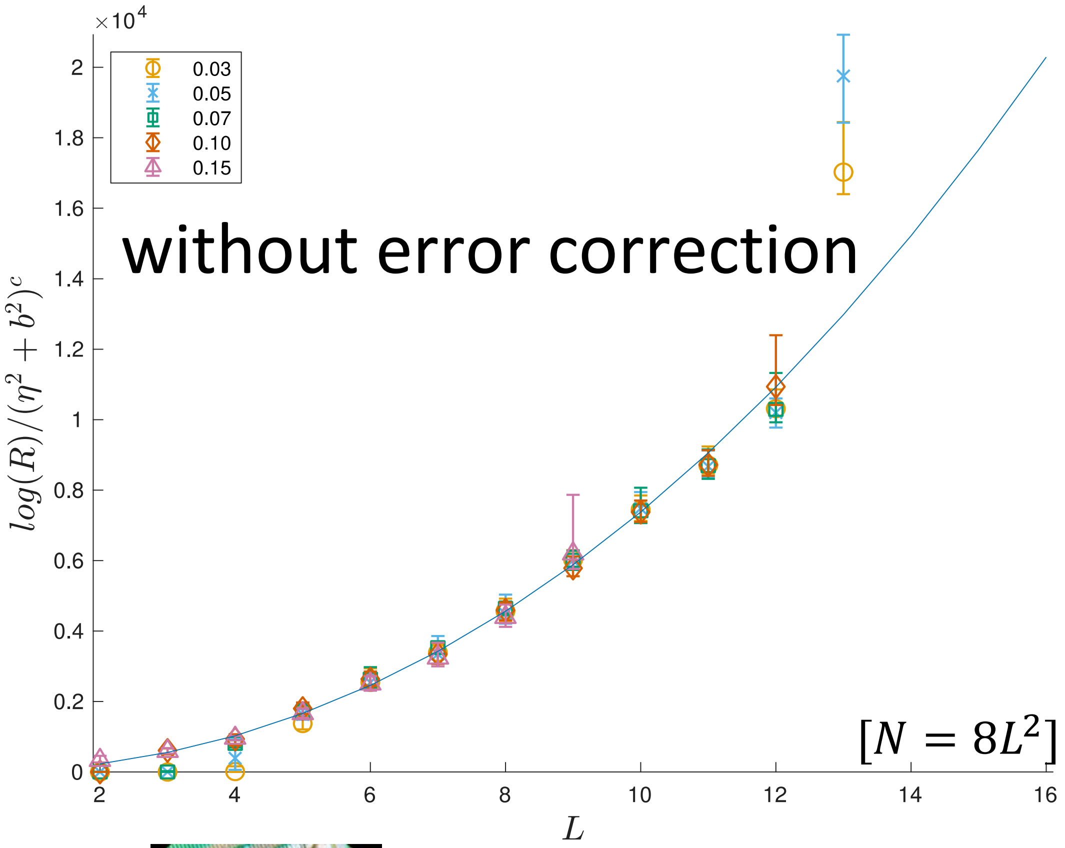

run on
NASA's
DW2000Q

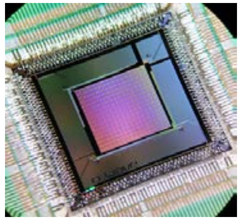

$a = 51.97$
$b = 0.16$
$c = 2.04$
$\boldsymbol{d = 2.15}$

worse than classical *deterministic* upper bound
$R \sim \exp(\xi L^2)$!

A. Pearson, A. Mishra, DL, I. Hen, arXiv:1907.12678

Time-to-solution data collapse & finite-size scaling:

fit to $R = e^{a(\eta^2+b^2)^c} L^d$

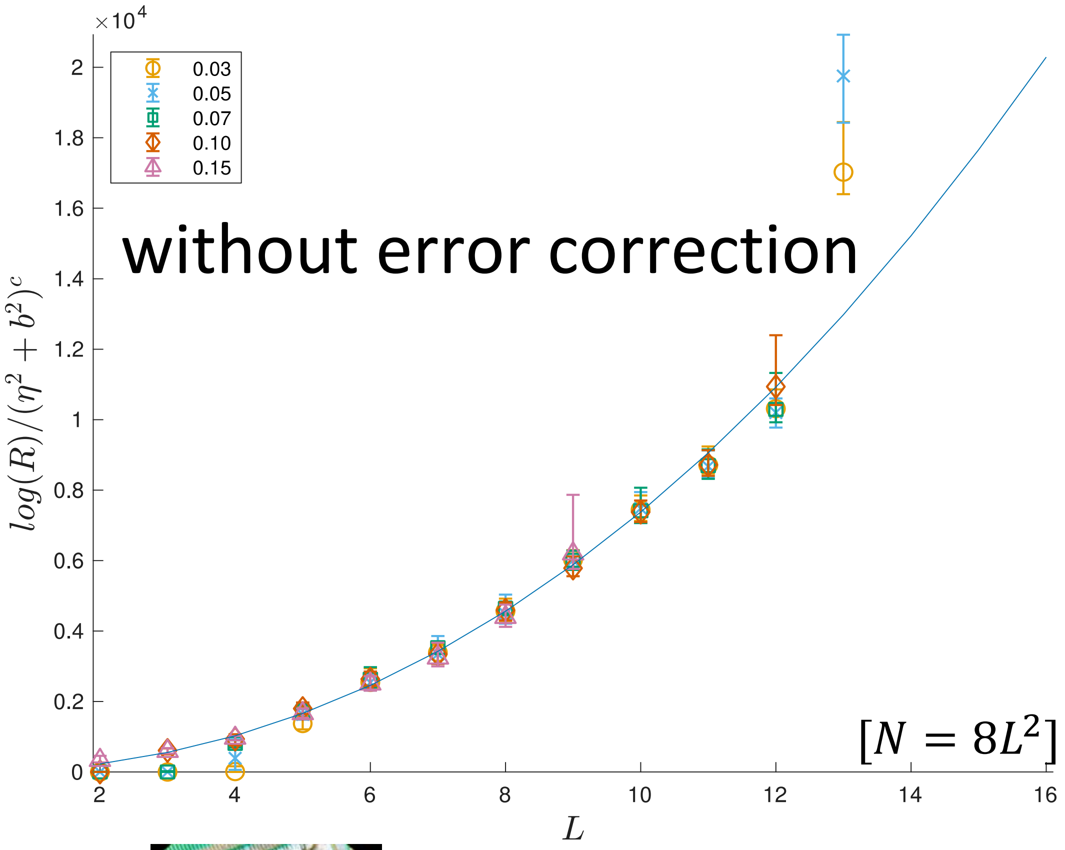

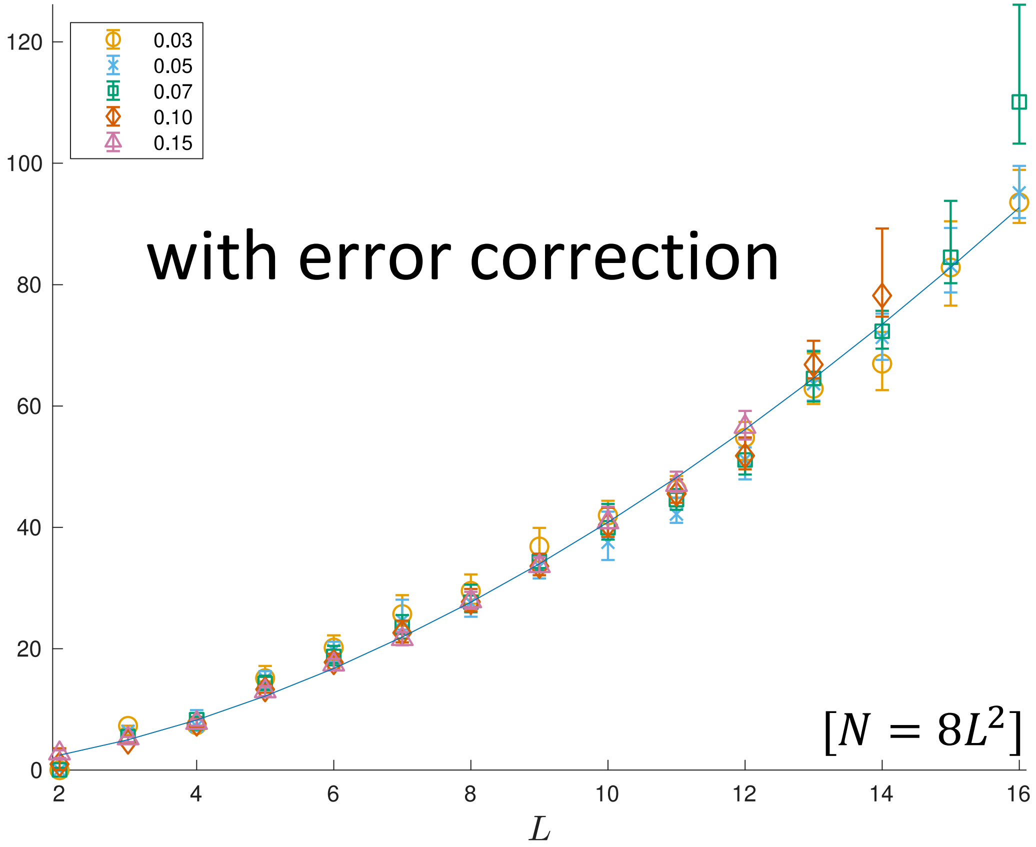

run on NASA's DW2000Q

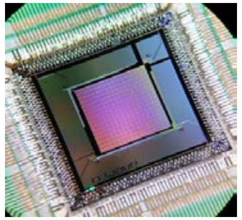

$a = 51.97$
$b = 0.16$
$c = 2.04$
$\boldsymbol{d = 2.15}$

better than classical upper bound
(lower bound unknown)

$a = 0.74$
$b = 0.06$
$c = 0.44$
} all smaller
$\boldsymbol{d = 1.74}$

error correction restores hope :-)

A. Pearson, A. Mishra, DL, I. Hen, arXiv:1907.12678

Beyond D-Wave: The IARPA Quantum Annealing Consortium

IARPA Quantum Enhanced Optimization (QEO) Program

Goal: **find out the ultimate capabilities of quantum annealing. Is there a quantum speedup?**

Building a 100-qubit quantum annealer using *high-coherence* (Al) superconducting flux qubits, for quantum optimization and sampling applications. *Built-in error suppression and correction.*

IARPA

USC University of Southern California

UCSC

MIT

Berkeley UNIVERSITY OF CALIFORNIA

NASA Ames Research Center

MIT Lincoln Laboratory

UNIVERSITÄT DES SAARLANDES

UCL

東京工業大学 Tokyo Institute of Technology

UNIVERSITY OF WATERLOO

LOCKHEED MARTIN

NORTHROP GRUMMAN

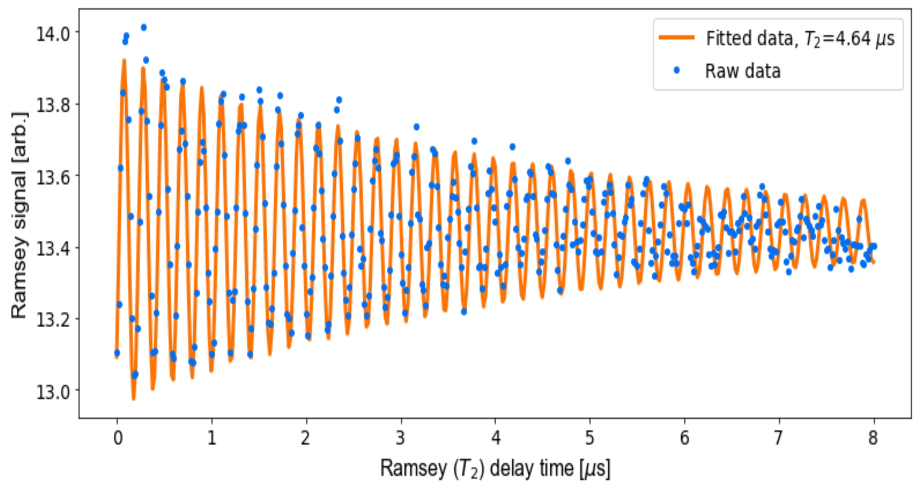

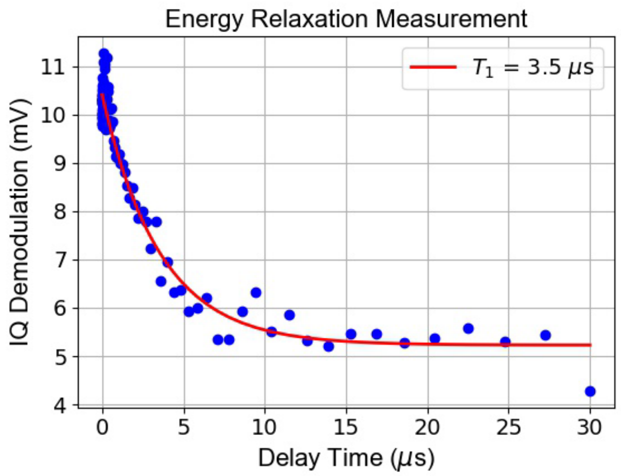

$T_1, T_2 \sim 4\ \mu s$ for annealing-capable qubits

coupling $|J_{ij}| \sim$ 2 GHz

S. Novikov et al., arXiv:1809.04485

Quantum Algorithmic Breakeven

Intermediate goal for QC, similar in spirit to "quantum supremacy":

Demonstrate error-corrected scaling that is better than uncorrected on a non-trivial computational problem

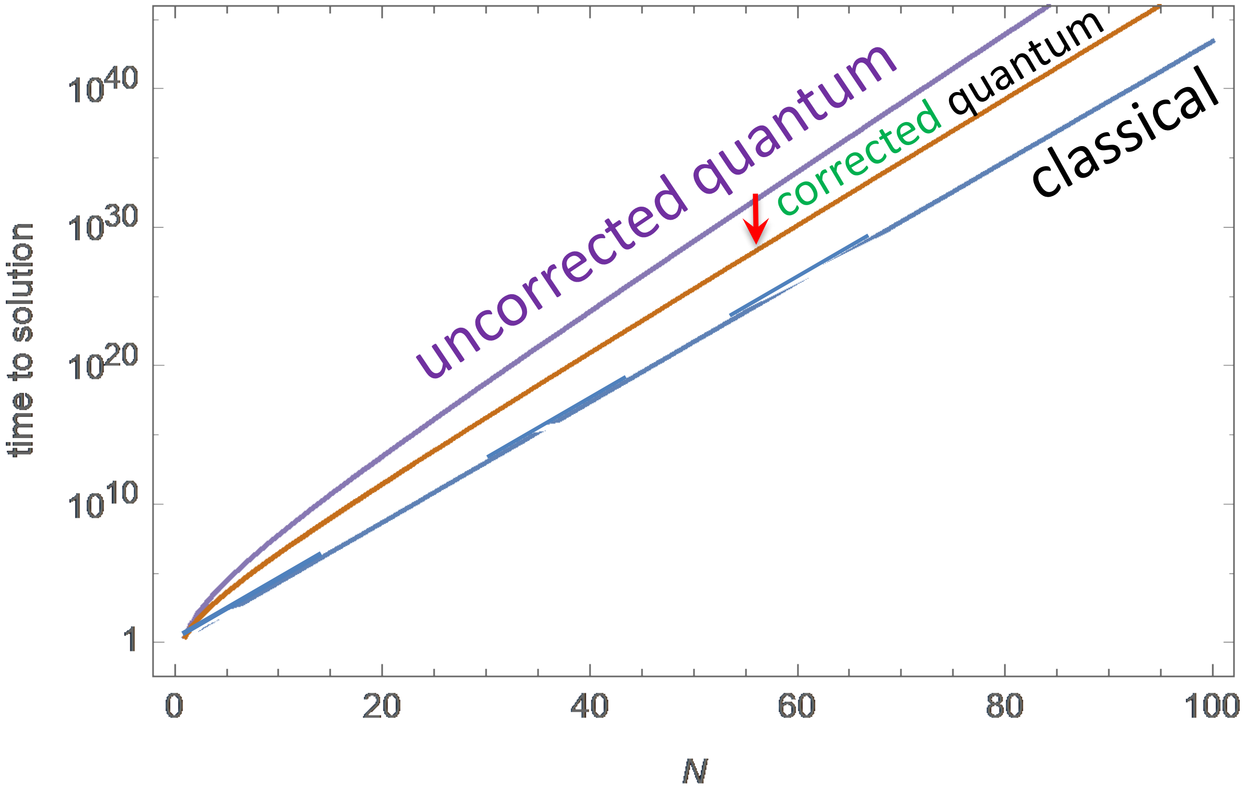

Already achievable in quantum annealing

An invitation for gate-model quantum computation

Thanks!