How Much Structure Is Needed for Huge Quantum Speedups?

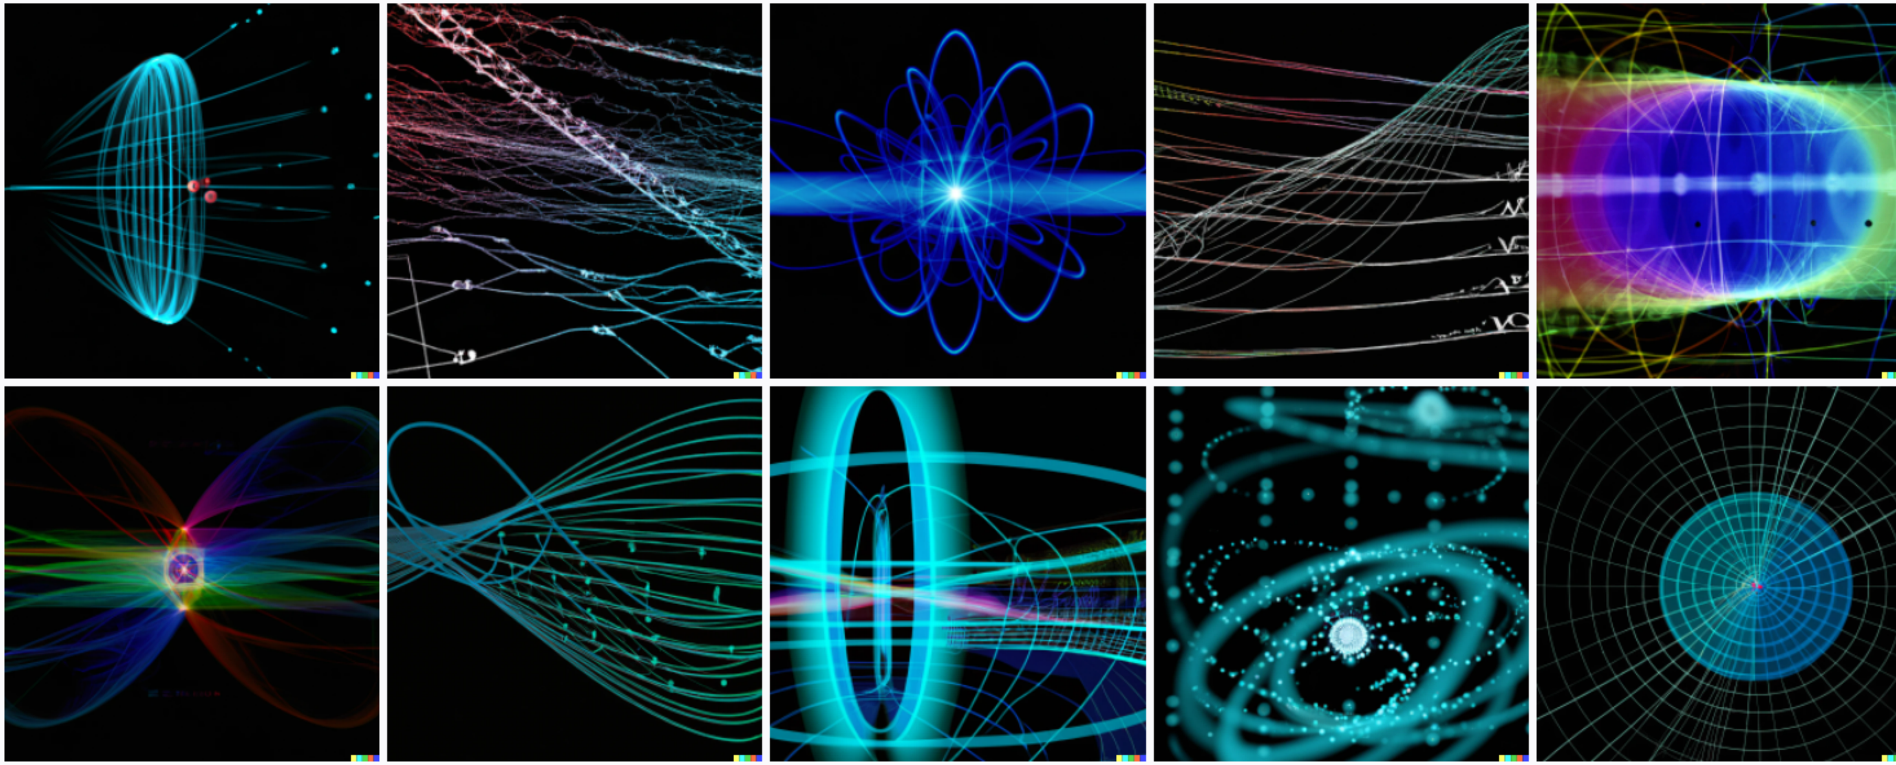

Scott Aaronson (UT Austin)

Solvay Conference, May 21, 2022

QC: "Weirder than any sci-fi writer would've had the imagination to invent"

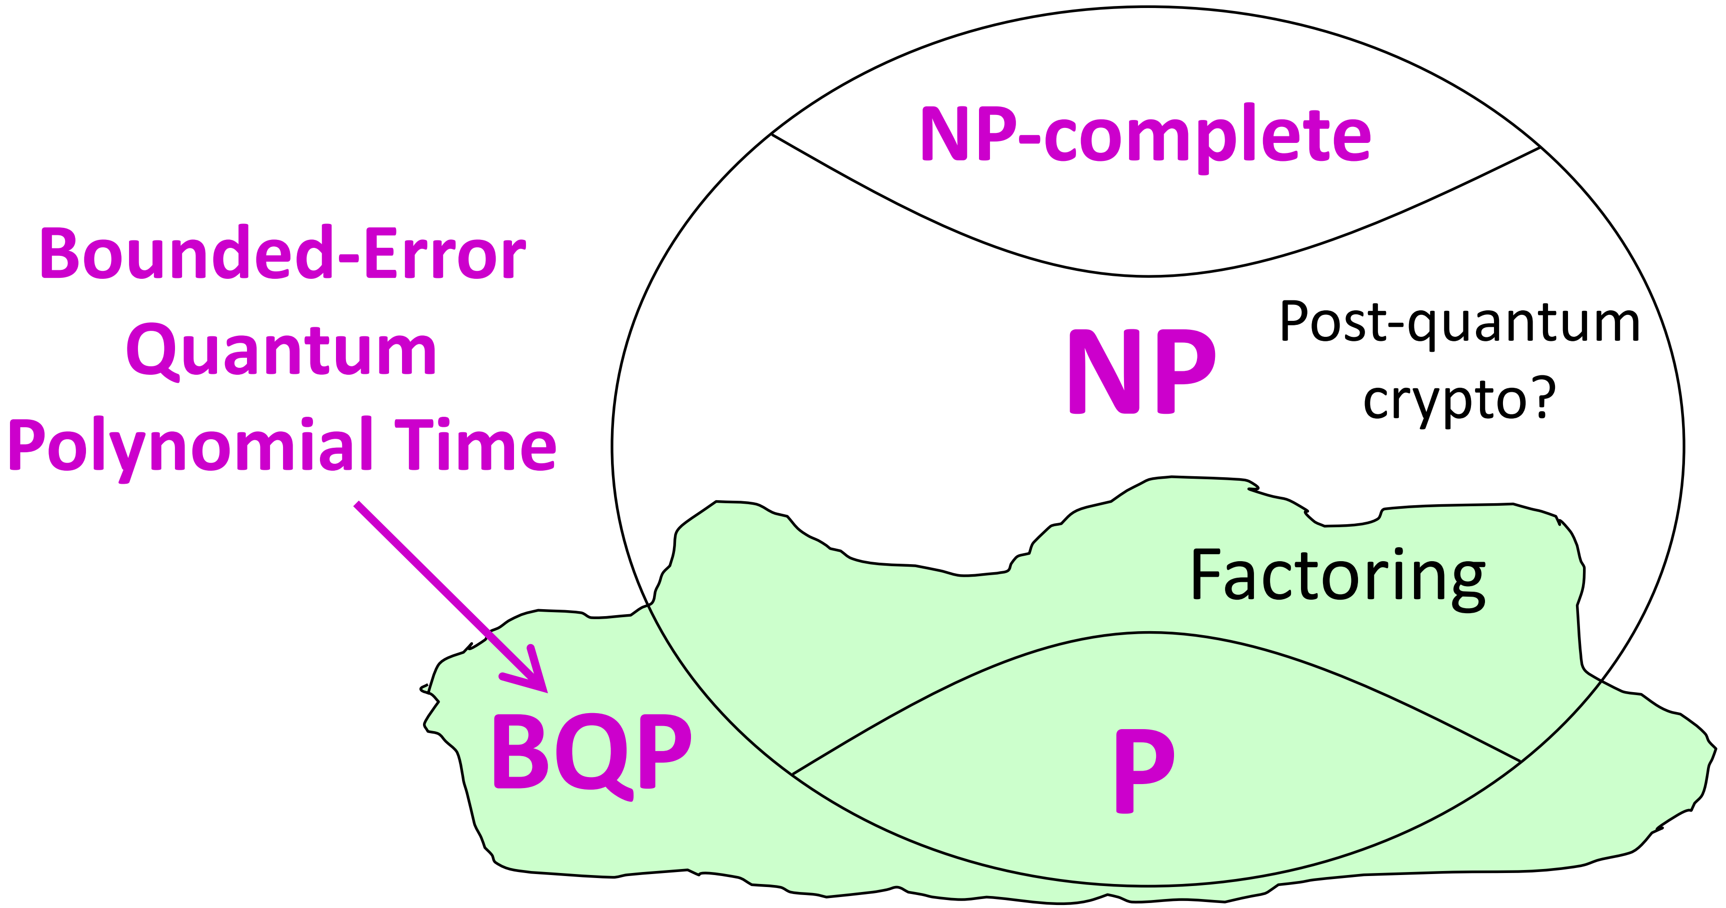

THE HAMMER OF HYPE

1. It’s not enough to create a superposition over all answers! We need **interference** to boost the amplitudes of right answers, and cancel the amplitudes of wrong ones

2. It’s not enough to do something quantumly—it has to **beat** the best that could’ve been done classically! And the classical side gets to **fight back**

So what superpolynomial quantum speedups do we know?

CIRCUIT MODEL

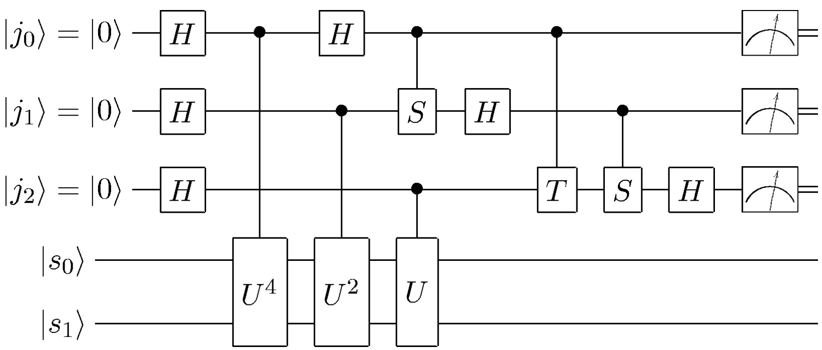

Resource: Gates

What we actually
about in real li

Almost no uncondit al
lower bounds!

BLACK-BOX MODEL

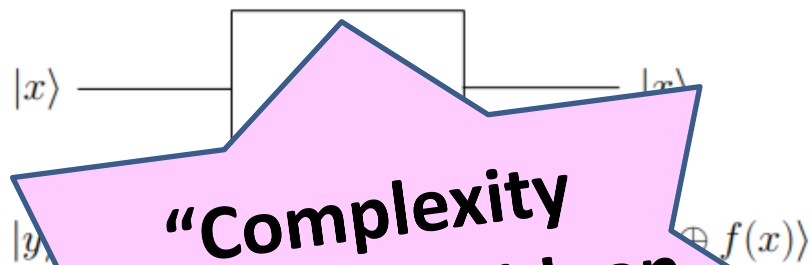

"Complexity theory's Euclidean path integral"

U_f

e know "oracle"
t its internals

Detailed understanding is achievable!

INSIGHTS

Exponential Speedups in Circuit Model?

Factoring, disc
group prob
nonab

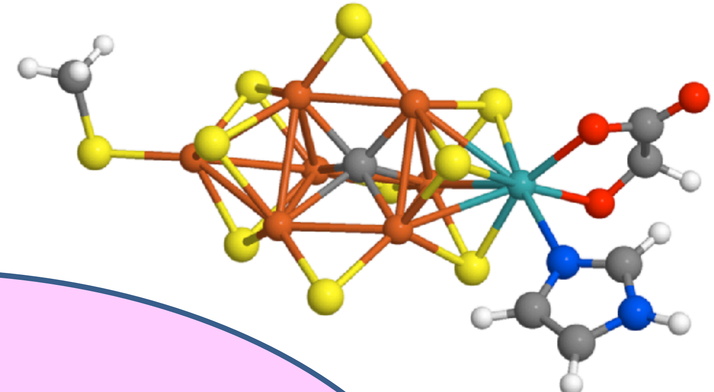

n of quantum
l chemistry

WAIT …
THAT'S IT?!

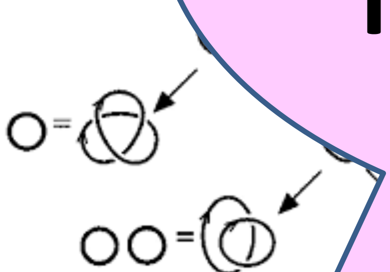

Approximating e Jones
polynomial at ots of unity

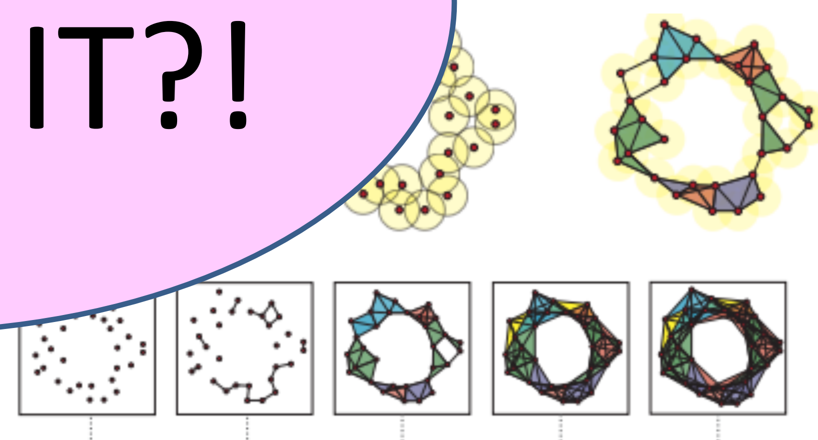

Some special machine learning
problems (e.g. Betti numbers)

Yamakawa and Zhandry's April 2022 breakthrough!

Task: Given a pseudorandom function $f:\{0,1\}^n \rightarrow \{0,1\}$, find a list of n-bit strings, $x_1,\ldots,x_n$, such that $f(x_1)=\ldots=f(x_n)=0$ **and** $x_1\ldots x_n$ is an n^2-bit codeword of a certain error-correcting code

They give a poly(n)-time quantum algorithm for this task, but show that any classical algorithm that treats f as a black box requires exponential time

First speedup for an NP search problem from an **unstructured** f. Alas, still won't work on a near-term QC!

So then what **does** work on a near-device? So far, sampling-based quantum supremacy experiments…

BosonSampling

(A.-Arkhipov 2011, ~100-photon experiments by USTC team 2020)

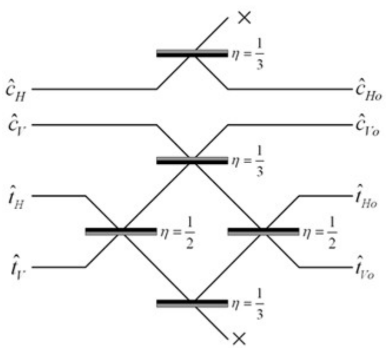

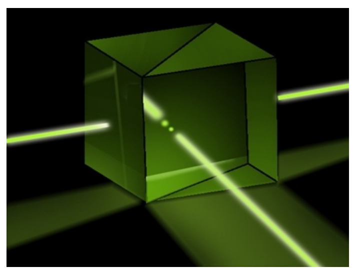

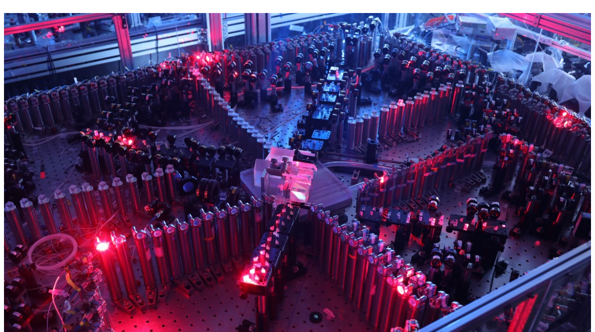

Random Circuit Sampling

(53-qubit experiment by Google team 2019; another by USTC team 2020)

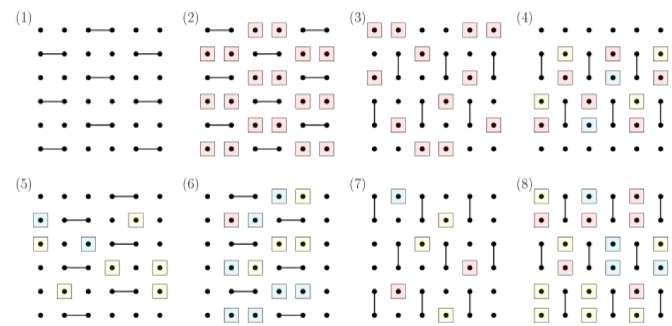

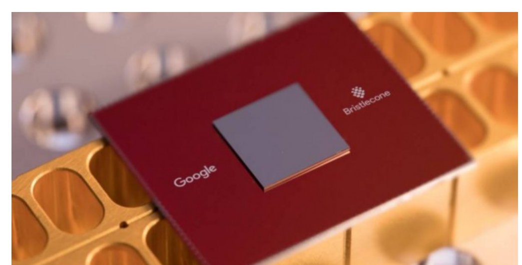

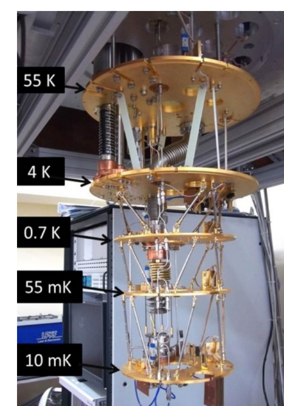

Advantages of the sampling tasks:

Can be implemented today

Assuming an ideal QC, exponential speedup over classical computing based on **very** secure complexity assumptions

Disadvantages:

Takes exponential time to verify the results with a classical computer!

Amazing recent verification protocols (e.g. Mahadev 2018), but not yet near-term implementable

Unclear whether there are any applications!

A. 2018: Cryptographically certified random numbers?

What Other Exponential Speedups Are Known In The Black-Box Model?

$f:\{0,1\}^n \rightarrow \{0,1\}^n$

$f(x) = f(x \oplus s)$

Find s

Simon's Problem (1994)

$f,g:\{0,1\}^n \rightarrow \{1,-1\}$

Decide if $f \approx \hat{g}$

Forrelation

A. 2009

A.-Ambainis 2015: Maximal quantum/classical separation

Raz-Tal 2018: Separation between BQP and PH

What Other Exponential Speedups Are Known In The Black-Box Model?

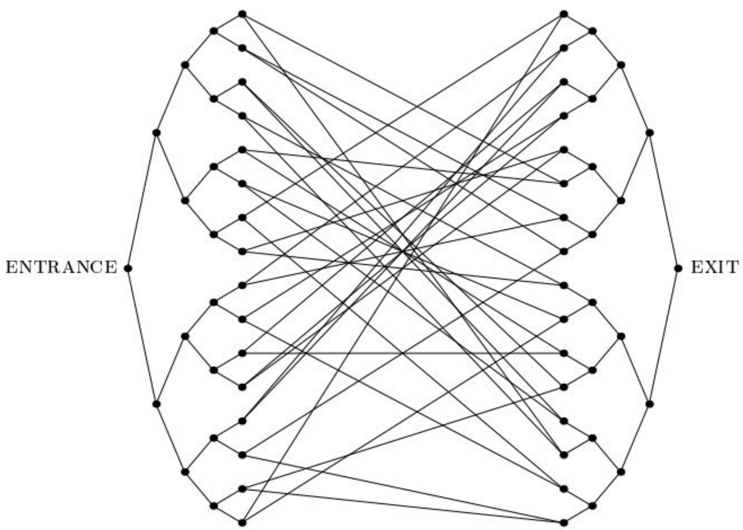

Quantum Walks

(Childs et al. 2002)

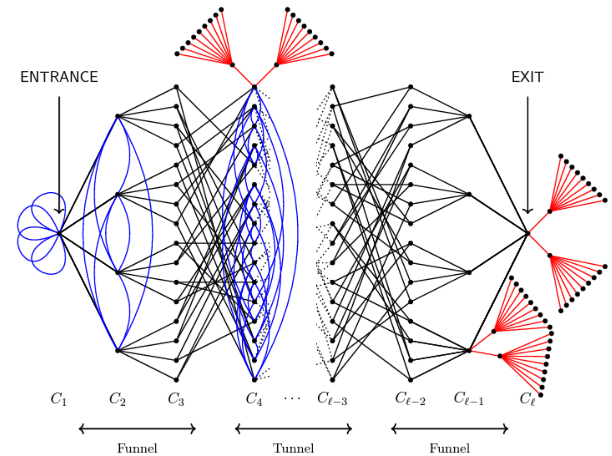

Adiabatic Optimization

(Separations by Hastings 2020, Gilyén-Vazirani 2020)

Limits to Black-Box Quantum Speedup!

Bennett, Bernstein, Brassard, Vazirani 1994: Searching an unstructured list of size N requires at least $\sim\sqrt{N}$ quantum queries—i.e., "Grover's algorithm is optimal"

- **For many other "unstructured" problems, like PARITY and MAJORITY, no asymptotic quantum speedup at all**
- **Proofs use linearity of QM + inability to notice a small change**

Generalization (Beals et al. 1998): For every **total** Boolean function $F:\{0,1\}^n \rightarrow \{0,1\}$, $\mathbf{D(F)=O(Q(F)^6)}$, where D, Q are deterministic and quantum query complexities

- **Recently improved to $D(F)=O(Q(F)^4)$, which was also proven tight (for R(F) vs. Q(F), best exponent is between 3 and 4)**
- **Explains why Simon's and Shor's problems needed "promises"**

Collision Problem

Given 2-to-1 function f:[n]→[n], find x,y with f(x)=f(y)

10 4 1 8 7 9 11 5 6 4 2 10 3 2 7 9 11 5 1 6 3 8

Birthday Paradox: Classically, $\Theta(\sqrt{n})$ queries to f are necessary and sufficient

Brassard-Høyer-Tapp 1997: Can cleverly combine birthday & Grover to get $O(n^{1/3})$ quantumly. **Better??**

$$\frac{1}{\sqrt{n}}\sum_{x=1}^{n}|x\rangle|f(x)\rangle \xrightarrow{\text{Measure 2nd register}} \frac{|x\rangle+|y\rangle}{\sqrt{2}}|f(x)\rangle$$

A.-Shi 2002: Alas, Brassard-Høyer-Tapp is optimal

Generalizing the Collision Lower Bound: Permutation Symmetry Precludes Exponential Quantum Speedups, Even For Promise Problems

A.-Ambainis 2011: For all partial functions F *that are symmetric under permuting inputs and outputs*, $R(F)=O(Q(F)^7)$, where R(F) and Q(F) are randomized and quantum query complexities respectively

Chailloux 2019: Improved to $R(F)=O(Q(F)^3)$ and generalized to input permutations only

Ben-David et al. 2020: At most polynomial quantum speedups for, e.g., partial functions of graph adjacency matrices

The Hierarchy of Structure?

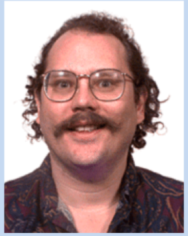

ABELIAN GROUP STRUCTURE

Exponential speedups!

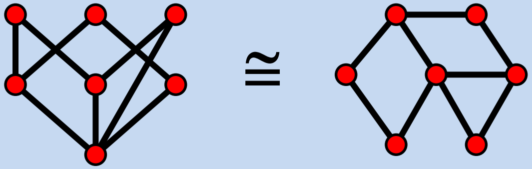

NONABELIAN GROUP STRUCTURE

Exponential speedups??

MINIMAL STRUCTURE (e.g. being 2-to-1)

Polynomial speedups only

NO GLOBAL STRUCTURE

Polynomial speedups only

Aaronson-Ambainis Conjecture

Conjecture: Let Q be a quantum algorithm that makes T queries to $X \in \{0,1\}^N$. Then for all $\varepsilon,\delta>0$, there's a classical algorithm that makes poly(T,1/ε,1/δ) queries to X, and approximates T's acc… …o $\pm\varepsilon$ on a 1-δ …s.

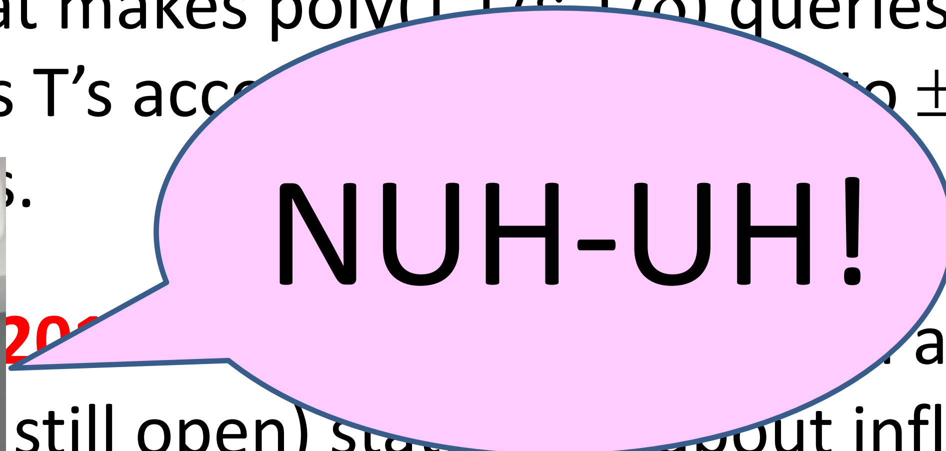

20… … an extremely … still open) stat… …bout influences in …-degree multivariate polynomials

Interpretation: "Relative to *random* oracles, only Grover-type speedups are possible. Exponential speedups require *structured* oracles, like periodic functions"

Urgent problem: Make near-term quantum supremacy verifiable!

Interactive protocol? Challenger creates a pseudorandom quantum circuit C that conceals a secret s, then sends C to QC, which has to find s by running C

- But how to implement this idea?? Bremner-Shepherd proposal killed by Kahanamoku-Meyer 2019; A.-Nguyen 2014 partial no-go theorem for BosonSampling

Question: Let C be a quantum circuit on n qubits with (say) n^2 gates, chosen uniformly at random *from among all such circuits such that* $|\langle 0^n|C|x\rangle|^2 \geq 0.1$ *for some basis state* $|x\rangle$. What does C look like? Is it hard to determine $|x\rangle$ given C?

Concluding Thought: The Law of Conservation of Weirdness?

“For every problem that admits an exponential quantum speedup, there must be some weirdness in its detailed statement, which the quantum algorithm exploits to focus amplitude on the rare right answers”