Quantum AI

Searching for valuable applications of quantum computing at Google

Ryan Babbush

December 14, 2023

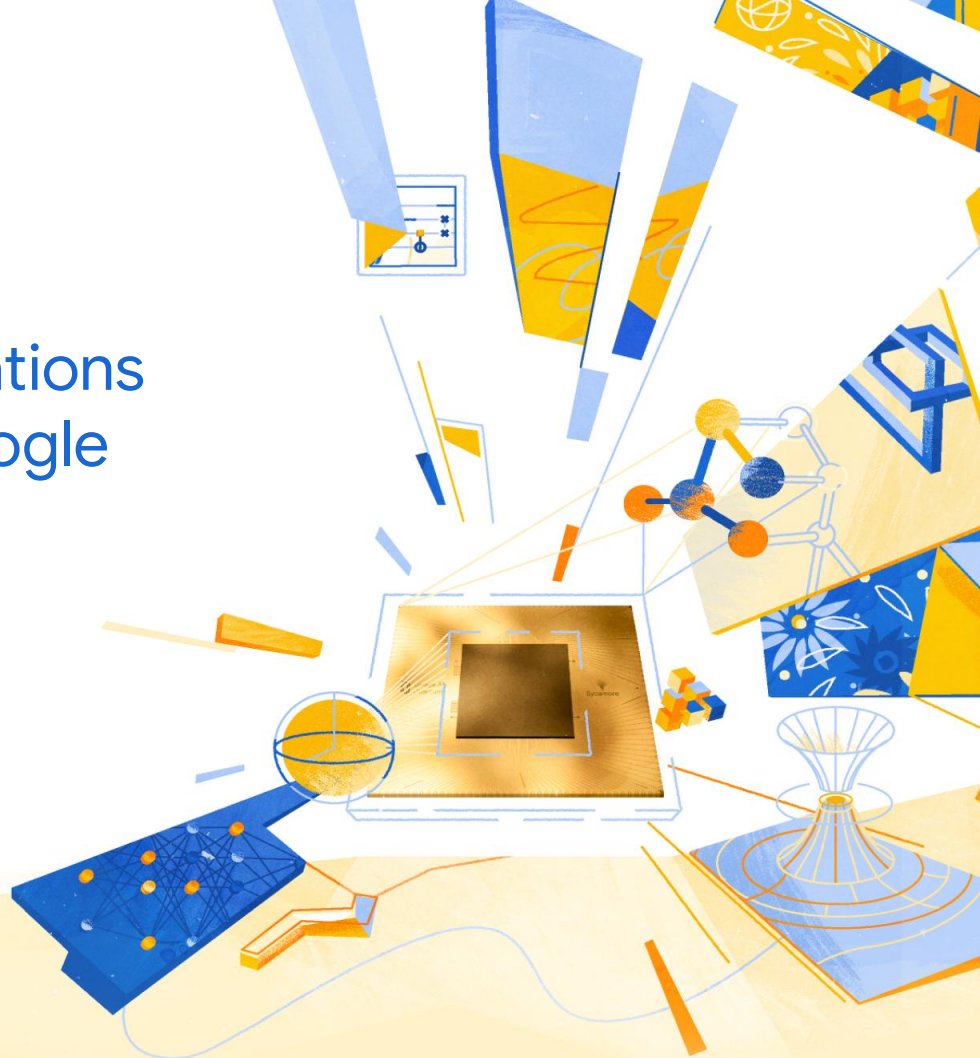

Google

Google is building an error-corrected quantum computer

Quantum AI

Targeting a device with ~1M physical qubits
that can execute billions, or trillions of gates

This is a challenging and expensive endeavor!

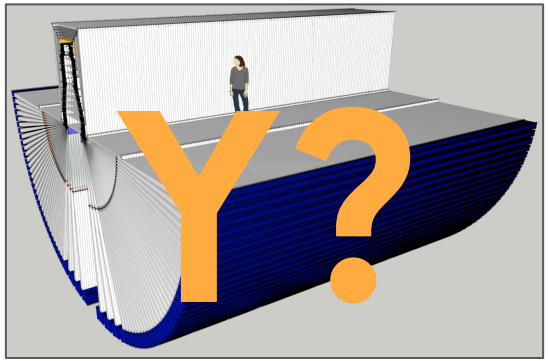

→ We hope these devices will solve important and otherwise intractable problems!

There are some use cases for 1M physical qubits, but fewer than we would hope

This talk is about research into the viability of promising quantum applications

Outline

1. Brief discussion of Google's roadmap towards quantum error-correction
2. Review of fault-tolerant quantum algorithms for chemistry
3. Results on identifying and assessing viability of valuable chemical applications
4. The viability of quantum advantage in topological data analysis
5. Brief overview of some other interesting results and directions in applications

Quantum AI

Part I: Google's roadmap towards quantum error-correction

Google

Google's roadmap towards practical quantum computing

We are here

Nature 574, 505-510 (2019)

Nature 614, 676–681 (2023)

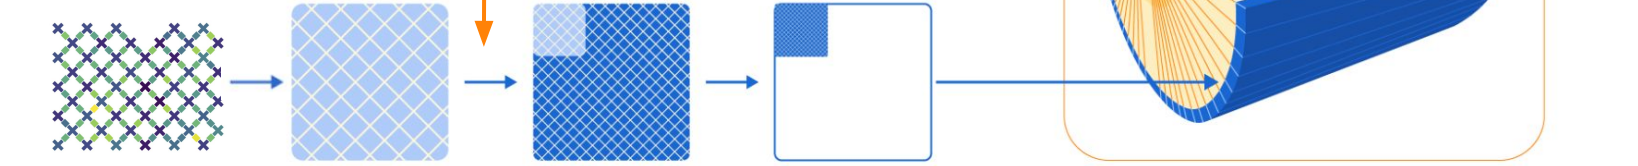

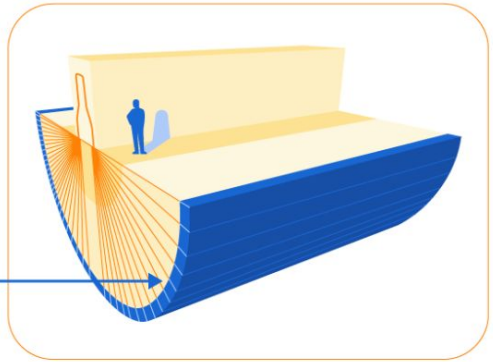

54	100	1,000	10,000	100,000	1 M qubits
Beyond classical ✔	**Logical qubit prototype** ✔	1 long-lived logical qubit	Tileable module (logical gate)	Engineering scale up	Error-corrected quantum computer
Milestone 1 (2019)	**Milestone 2** (2023)	**Milestone 3** (2025)	**Milestone 4**	**Milestone 5**	**Milestone 6** (in ~10 years)

Quantum AI

Milestone 2: Logical qubit prototype (plan)

Run experiments to implement surface codes

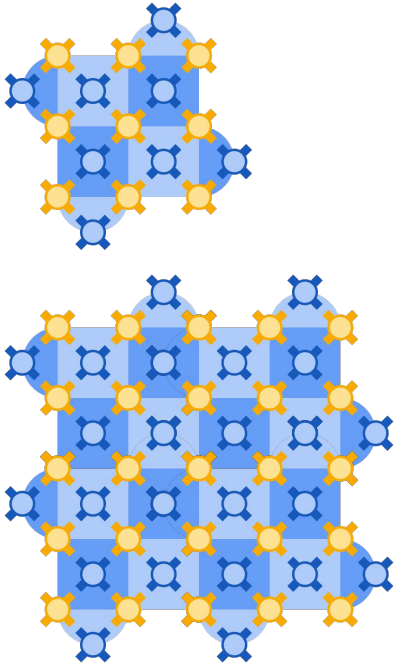

d = 3
surface code

d = 5
surface code

Analyze data and see if failure probability is lower with larger code

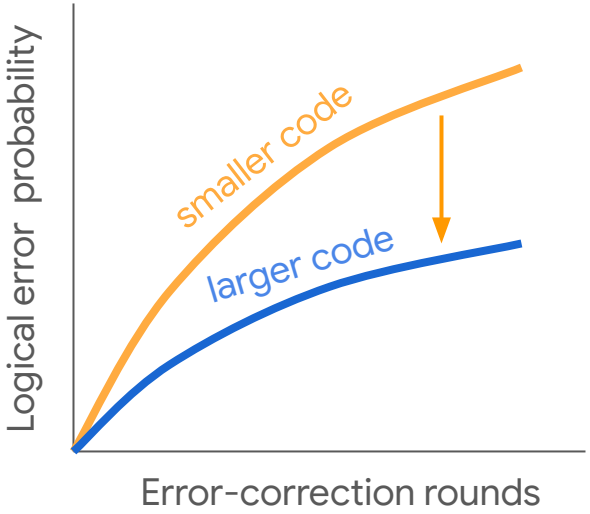

Milestone 2: Logical qubit prototype (experimental data)

Nature 614, 676–681 (2023)

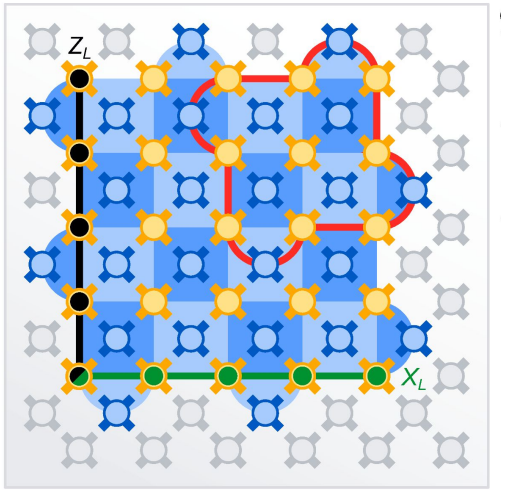

Data qubit (d^2)

Measure qubit (d^2 - 1)

Unused

Subset distance-3

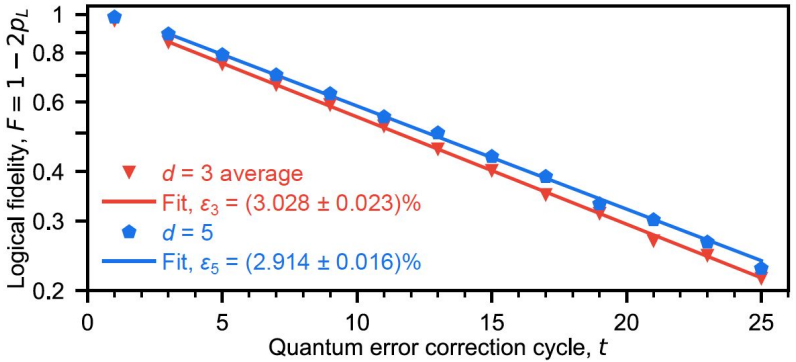

λ = error(d) / error(d + 2)

On the importance of super-quadratic speedups

PRX Quantum 2, 010103 (2021)

Quantum AI

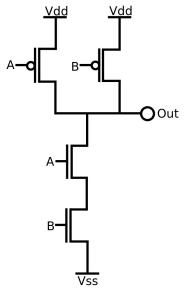

classical NAND gate (CMOS)
$<10^{-9}$ “transistorseconds”

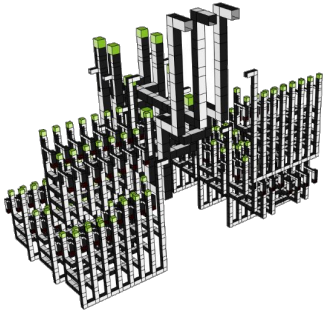

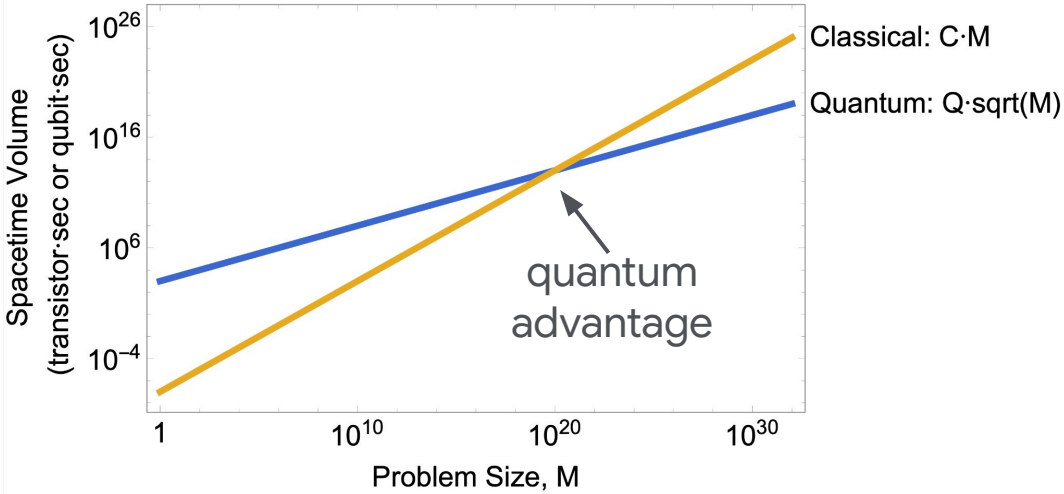

“quantum NAND” gate (distillation of Toffoli state) >10 “qubitseconds”

PRX Quantum 2, 010103 (2021) argues quadratic speedups will not enable error-corrected advantage until devices MUCH larger than 1MM physical qubits

Quantum AI

Part II: review of fault-tolerant quantum algorithms for chemistry

Google

Quantum AI

Spectrum of quantum simulation difficulty

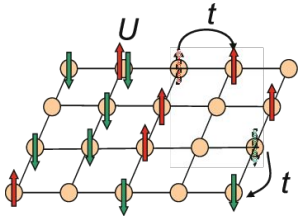

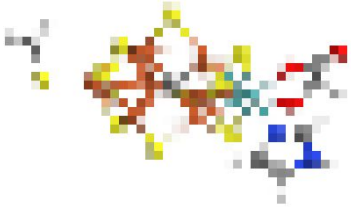

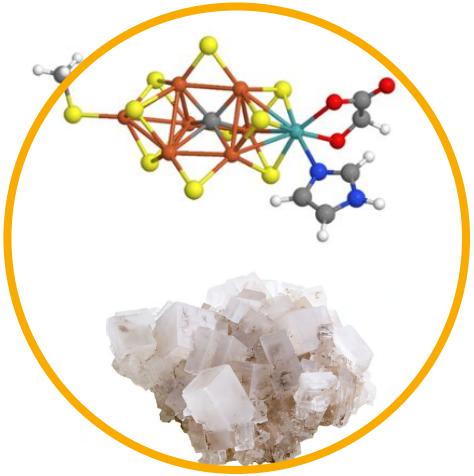

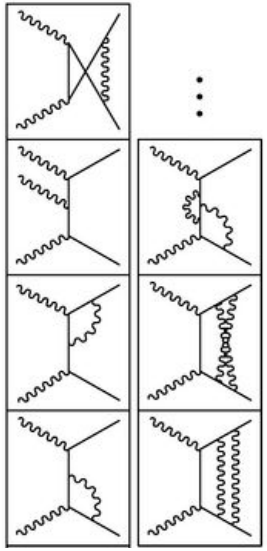

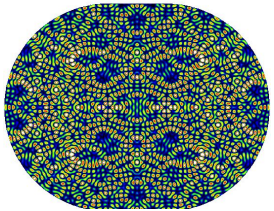

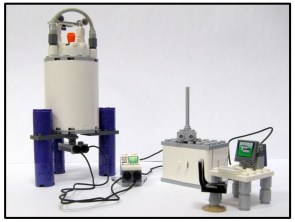

application difficulty

physical qubits required	25k-50k	50k - 250k	250k - 1MM	1MM - 5MM	???

The molecular electronic structure problem

Quantum AI

Solve for the energy of molecule under the Born-Oppenheimer approximation

$$H = \sum_{pq} (T_{pq} + U_{pq}) a_p^\dagger a_q + \sum_{pqrs} V_{pqrs} a_p^\dagger a_q^\dagger a_r a_s$$

$$T_{pq} = \int dr\, \varphi_p^*(r) \left(-\frac{\nabla^2}{2} \right) \varphi_q(r)$$

$$U_{pq} = \int dr\, \varphi_p^*(r)\, U(r)\, \varphi_q(r)$$

$$V_{pqrs} = \int dr_1\, dr_2\, \varphi_p^*(r_1)\, \varphi_q^*(r_2)\, \frac{1}{|r_1 - r_2|}\, \varphi_r(r_2)\, \varphi_s(r_1)$$

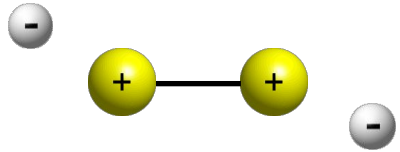

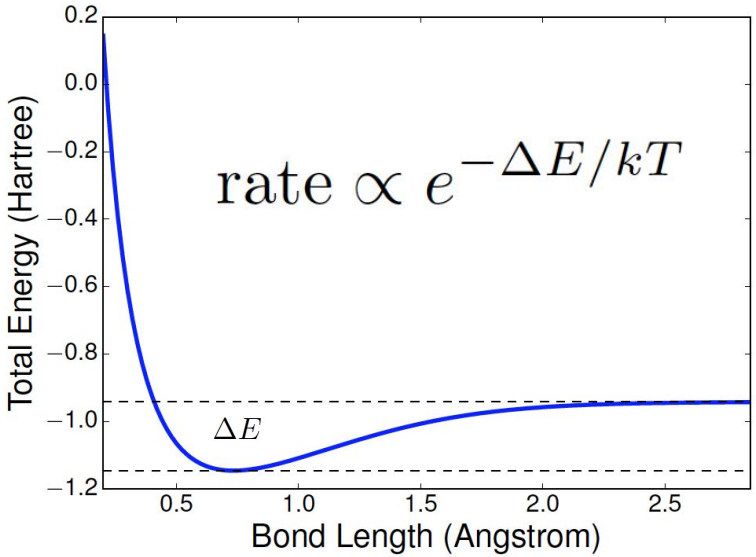

Energy surfaces allow us to understand reactions
Need chemical accuracy (1 kcal/mol) for rates

Such accuracy is often classically intractable
Especially for systems with strong correlation

Representing η electrons in N spin-orbitals

Quantum AI

Second quantization requires N qubits

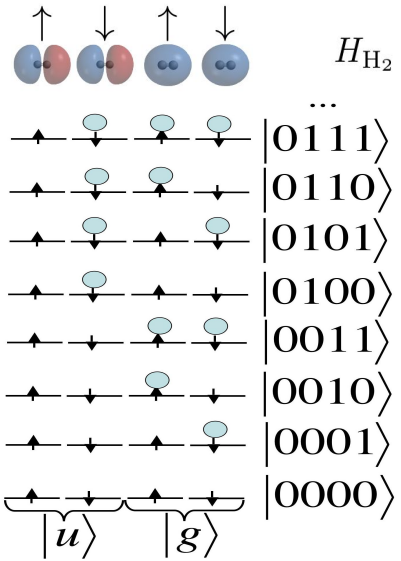

$$H_{\mathrm{H}_2} = a_0\,\mathbb{1} + a_1\,(Z_0 + Z_1) + a_2\,(Z_2 + Z_3)$$
$$+ a_6\,(Z_0Z_3 + Z_1Z_2) + a_7\,(X_0Y_1Y_2X_3$$
$$+ a_3\,Z_0Z_1 + a_4\,Z_2Z_3 + a_5\,(Z_0Z_2 + Z_1Z_3)$$
$$+ Y_0X_1X_2Y_3 - X_0X_1Y_2Y_3 - Y_0Y_1X_2X_3)$$

$$\{a_p, a_q^\dagger\} = a_p a_q^\dagger + a_q^\dagger a_p = \delta_{pq}\mathbb{1}$$
$$\{a_p, a_q\} = \{a_p^\dagger, a_q^\dagger\} = 0$$

- Anti-symmetry is “encoded in the operators”
- Good near half filling or with compact basis

First quantization requires η log N qubits

η registers of size log N index
which orbital the particle occupies

$$\sum_{\phi_p \in \{\phi\}} a_{\phi_1 \cdots \phi_\eta} |\phi_1 \cdots \phi_i \cdots \phi_j \cdots \phi_\eta\rangle$$
$$= (-1)^\pi \sum_{\phi_p \in \{\phi\}} a_{\phi_1 \cdots \phi_\eta} |\phi_1 \cdots \phi_j \cdots \phi_i \cdots \phi_\eta\rangle$$

- Anti-symmetry is “explicit in the state”
- Ideal for high precision calculations

Error-corrected quantum chemistry simulation

Science 309:5741 (2005), 1704-1707

1. Prepare an ansatz wavefunction ψ with "reasonable" support on the ground state

$$H\,|k\rangle = E_k\,|k\rangle \qquad\qquad |\langle\psi|0\rangle|^2 \text{ = not-too-small}$$

2. Form quantum circuit $U = e^{-i\,f(H)}$ that encodes Hamiltonian spectrum in its eigenvalues

e.g., for Trotter:

$$f(H) = H = \sum_\ell H_\ell \qquad\qquad U \approx \left(\prod_\ell e^{-iH_\ell/r}\right)^r$$

3. Application of U to ψ accumulates phases $f(E)$ encoding the spectrum

$$U\,|\psi\rangle = \sum_k \underbrace{\langle k|\psi\rangle}_{a_k} e^{-if(E_k)}\,|k\rangle$$

4. Phase estimation gives E_0 *with error* ϵ and probability $|a_0|^2$ using $\frac{1}{\epsilon}\left\|\frac{\partial f(E)}{\partial E}\right\|^{-1}$ queries to U

Linear combination of unitaries (LCU) simulation

H as an "LCU": $H = \sum_{\ell=0}^{L-1} w_\ell U_\ell$ with 1-norm $\lambda = \sum_{\ell=0}^{L-1} |w_\ell|$

"qubitization" (arXiv:1610.06546) synthesizes a quantum walk unitary: $\mathcal{W} = e^{-i \arccos(H/\lambda)}$

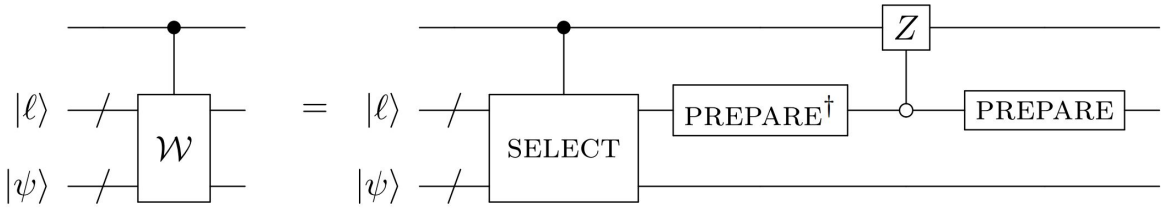

$$\text{PREPARE}\,|0\rangle^{\otimes \log L} \mapsto \sum_{\ell=0}^{L-1} \sqrt{\frac{w_\ell}{\lambda}}\,|\ell\rangle$$

$$\text{SELECT}\,|\ell\rangle\,|\psi\rangle \mapsto |\ell\rangle\,U_\ell\,|\psi\rangle$$

Repeat this circuit $\mathcal{O}(\lambda/\epsilon)$ times to estimate phase to within error ϵ

Quantum AI

But how does the quantum walk scale?

SELECT can be implemented at O(η) cost in first quantization, O(N) cost in second quantization

PREPARE is the hard part

- Cost proportional to computation required to compute (or “load”) Hamiltonian coefficients

Coulomb operator

$$\sum_{p,q,r,s=1} V_{pqrs} a_p^\dagger a_q^\dagger a_r a_s$$

total cost = $O(N^4 / \epsilon)$
(1902.02134)

Cholesky factorization

$$\sum_{\ell=1}^{L} \left(\sum_{\sigma \in \{\uparrow,\downarrow\}} \sum_{p,q=1}^{N/2} W_{pq}^{(\ell)} a_{p,\sigma}^\dagger a_{q,\sigma} \right)^2$$

total cost = $O(N^{7/2} / \epsilon)$
(1902.02134)

Tensor hypercontraction

$$V_{pqrs} = \sum_{\mu,\nu=1}^{\mathcal{O}(N)} \chi_p^{(\mu)} \chi_q^{(\mu)} \zeta_{\mu\nu} \chi_r^{(\nu)} \chi_s^{(\nu)}$$

total cost = $O(N^3 / \epsilon)$
(2011.03494)

First quantization and simple basis sets

Simple basis sets like grids, plane waves lead to analytic integrals!

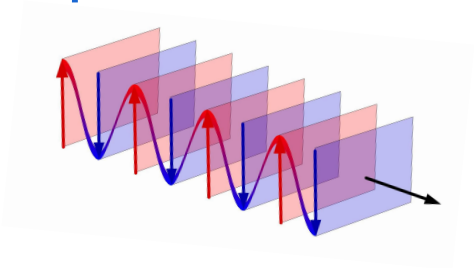

$$V = \frac{2\pi}{\Omega} \sum_{\substack{p \neq q \\ \nu \neq 0}} \frac{c^\dagger_{p+\nu} c^\dagger_{q-\nu} c_q c_p}{k_\nu^2}$$

But molecules need 100X - 1,000X more plane waves than MOs to reach chemical accuracy

- In second quantization, space complexity is O(N)
- Would need 100k logical qubits instead of 100!

In first quantization, space complexity is O(η log N)

- 60 electrons in 100k PWs needs ~1k logical qubits
- 1807.09802 + 2105.12767 scale as low as $O(\eta^{8/3} N^{1/3})$
- Particularly attractive for non-BO dynamics

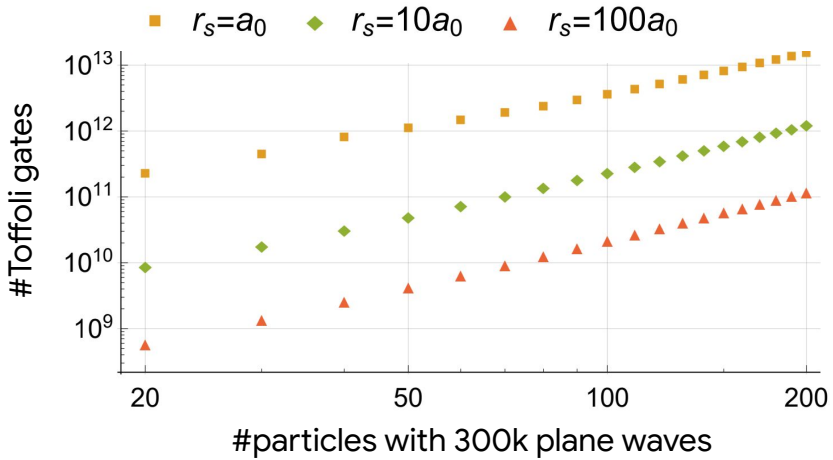

Compilation is tedious, often thankless work

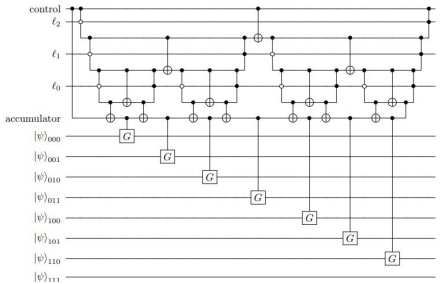

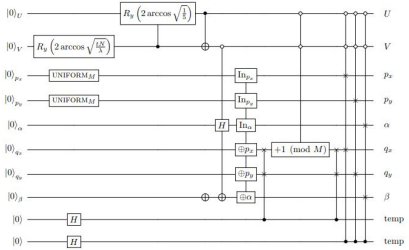

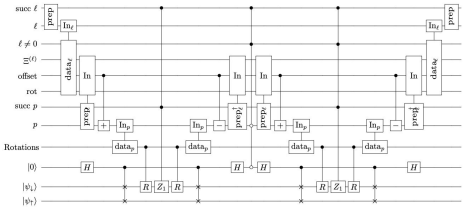

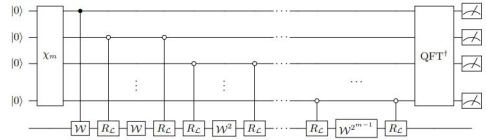

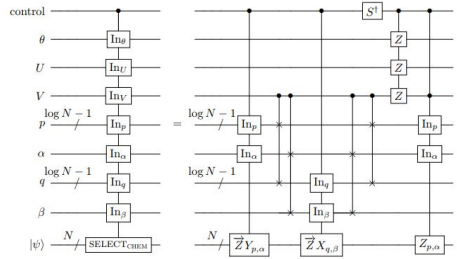

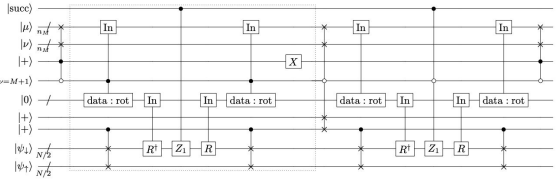

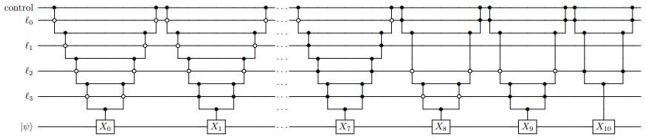

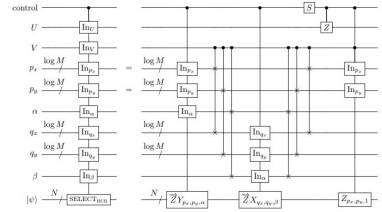

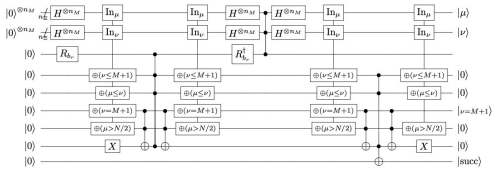

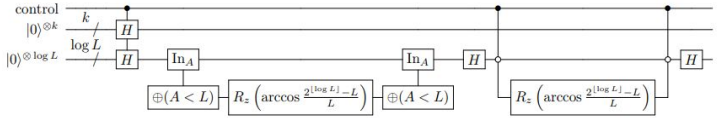

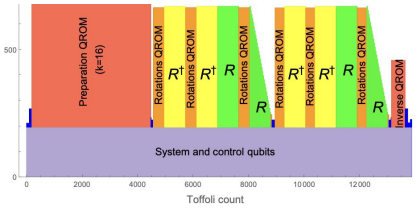

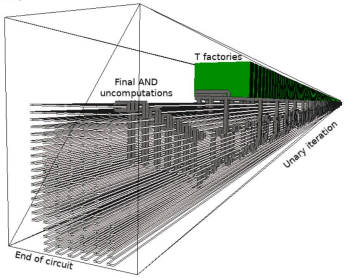

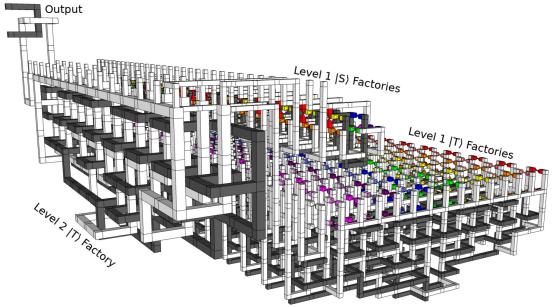

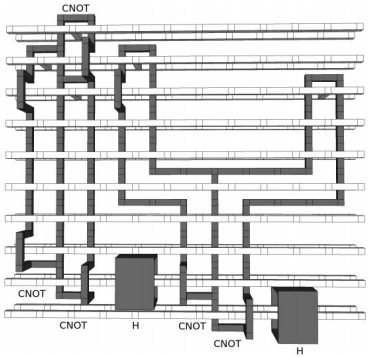

Quantum AI

Part III: searching for applications of these algorithms

Google

FeMoCo is great but we need more

Quantum AI

$N_2 + 3\,H_2 \longrightarrow 2\,NH_3$
2% of world energy

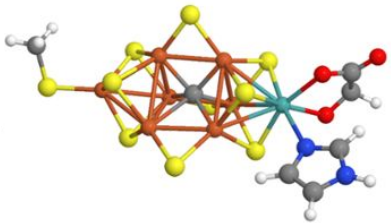

Year	arXiv	Method	Space	T Complexity	T Gates	Physical qubits
2005	0604193	Trotter	$\mathcal{O}(N)$	$\mathcal{O}(\text{poly}(N/\epsilon))$	Unknown	Unknown
2010	1001.3855	Trotter (first bounds on step complexity)	$\mathcal{O}(N)$	$\tilde{\mathcal{O}}(N^{11}/\epsilon^{3/2})$	Unknown	Unknown
2013	1312.1695	Trotter (first bounds on number of steps)	$\mathcal{O}(N)$	$\tilde{\mathcal{O}}(N^{9}/\epsilon^{3/2})$	Unknown	Unknown
2014	1406.4920	Trotter (tighter bounds)	$\mathcal{O}(N)$	$\tilde{\mathcal{O}}(N^{6}/\epsilon^{3/2})$	Unknown	Unknown
2015	1506.01020	Taylor series	$\mathcal{O}(N)$	$\tilde{\mathcal{O}}(N^{5}/\epsilon)$	Unknown	Unknown
2016	1605.03590	Trotter (first resource estimate)	$\mathcal{O}(N)$	$\tilde{\mathcal{O}}(N^{6}/\epsilon^{3/2})$	$\sim 10^{14}$	$\sim$ 20MM
2019	1902.02134	qubitization + single factorization	$\tilde{\mathcal{O}}(N^{3/2})$	$\tilde{\mathcal{O}}(N^{4}/\epsilon)$	$\sim 10^{11}$	$\sim$ 6MM
2020	2007.14460	qubitization + double factorization	$\tilde{\mathcal{O}}(N^{3/2})$	$\tilde{\mathcal{O}}(N^{7/2}/\epsilon)$	$\sim 10^{10}$	$\sim$ 4MM
2020	2011.03494	qubitization + tensor hypercontraction	$\tilde{\mathcal{O}}(N)$	$\tilde{\mathcal{O}}(N^{3}/\epsilon)$	$\sim 10^{9}$	$\sim$ 2MM

TABLE I. Best fault-tolerant algorithms for phase estimating chemistry in an arbitrary (e.g., molecular orbital) basis. N is number of basis functions and ϵ is target precision. Gate counts are for FeMoCo, physical qubit counts assume superconducting qubit surface code implementation (see most recent papers for further assumptions).

- Critical to flesh out more specifically what valuable technological problems might be practically solved with a few thousand logical qubits and < 10^{12} Toffoli gates

Assessing quantum/classical boundary for P450

PNAS 119, 2203533119 (2022)

Quantum AI

P450 is strongly correlated iron-porphyrin / drug anti-target (kind of like FeMoCo!)

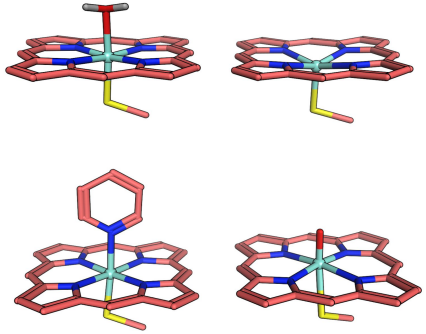

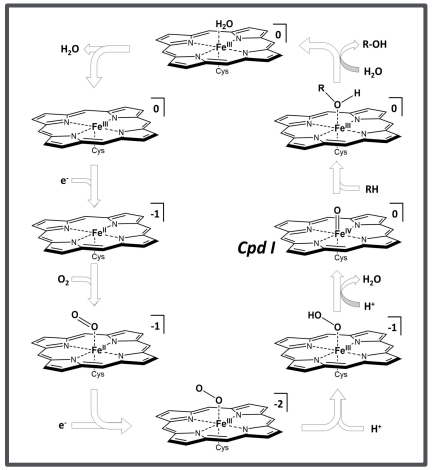

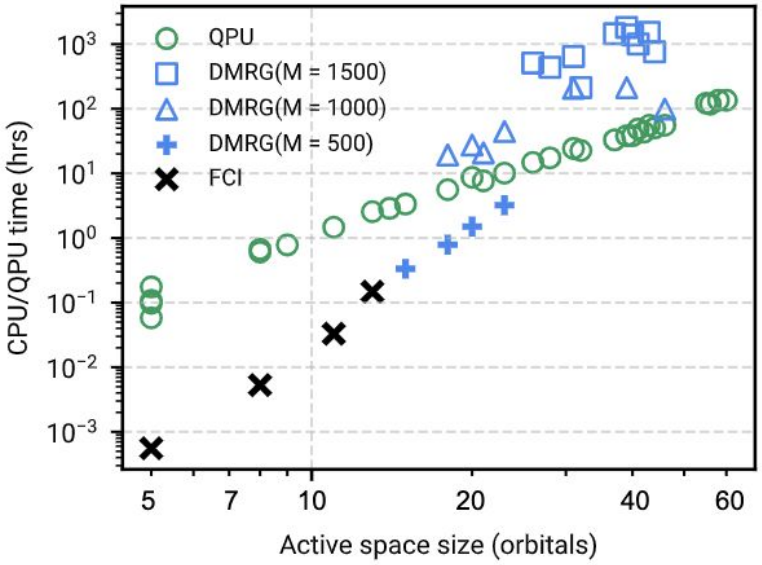

We observe onset of quantum advantage for active space sizes near 80 qubits

Chemically relevant (and classically intractable) calculations would require ~3k qubits, ~10^{10} Toffolis

Ab initio materials simulation is still very costly

PRX Quantum 4, 040303 (2023)

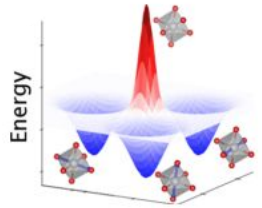

Want to get Co out of batteries, LNO is good candidate, why not Jahn-Teller distorted?

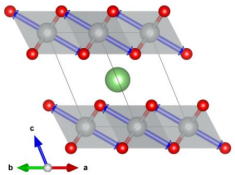

$C2/m$
(collinear Jahn-Teller)

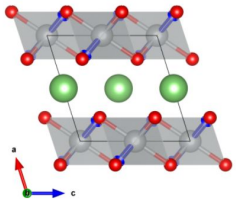

$P2_1/c$
(zig-zag Jahn-Teller)

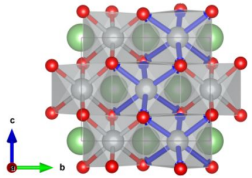

$P2/c$
(disproportionated)

- Classical many-body methods unreliable for metals
- Embedding difficult to converge finite size
- DFT disagrees between functionals

recent work symmetry adapts qubitization

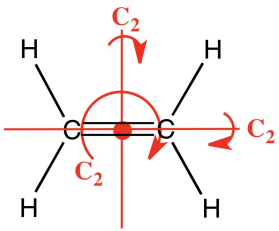

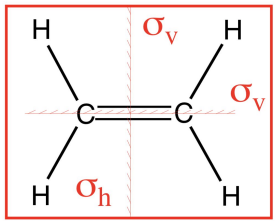

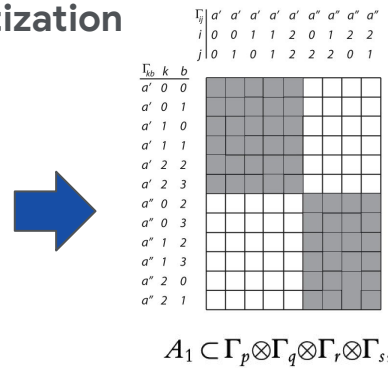

$A_1 \subset \Gamma_p \otimes \Gamma_q \otimes \Gamma_r \otimes \Gamma_s$

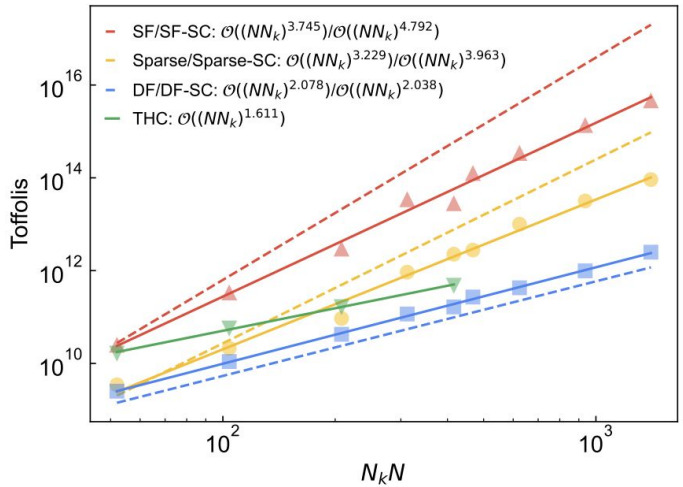

Gives linearly reduced scaling in number of k-points

Q algorithms must further improve for viability

Are only the most highly entangled systems promising applications of quantum simulation?

Quantum AI

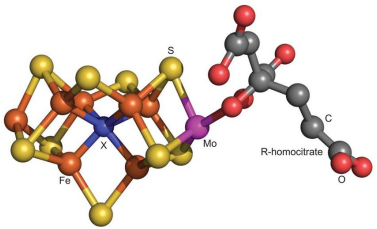

FeMoCo (fertilizer catalyst)
PRX Quantum 2, 030305 (2021)

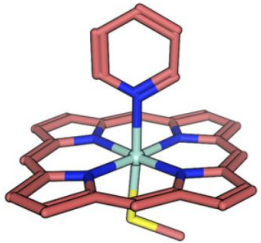

P450 (drug anti-target)
PNAS 119, 2203533119 (2022)

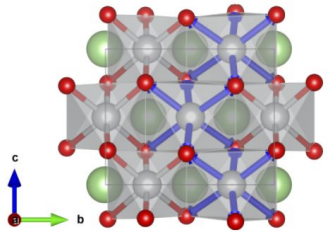

$LiNiO_2$ (battery cathode)
PRX Quantum 4, 040303 (2023)

- "classical competition" is only the most costly/accurate methods (e.g. AFQMC, DMRG)

most chemical computations do *not* require accurate treatment of strong correlation

- "classical competition" is highly efficient/approximate classical methods (e.g. mean-field, DFT)
- super-quadratic quantum advantage over mean-field would **dramatically** broaden applications

Super-quadratic quantum advantage over classical mean-field methods possible for electron dynamics

Quantum AI

Nat. Comm 14, 4058 (2023)

Usual quantum sim. advantage is resolution of entanglement

- storing wavefunction with η particles, N orbitals classically requires O(N choose η) bits
- mean-field has no particle correlation, only requires O(N η) bits and gate complexity:

$$N^{4/3}\eta^{7/3}t + N^{5/3}\eta^{4/3}t$$

or, for high temperature:

$$N^{10/3}\eta^{1/3}t + N^{11/3}t/\eta^{2/3}$$

quantum simulations need only O(η log N) qubits and gate complexity:

$$N^{1/3}\eta^{8/3}t$$

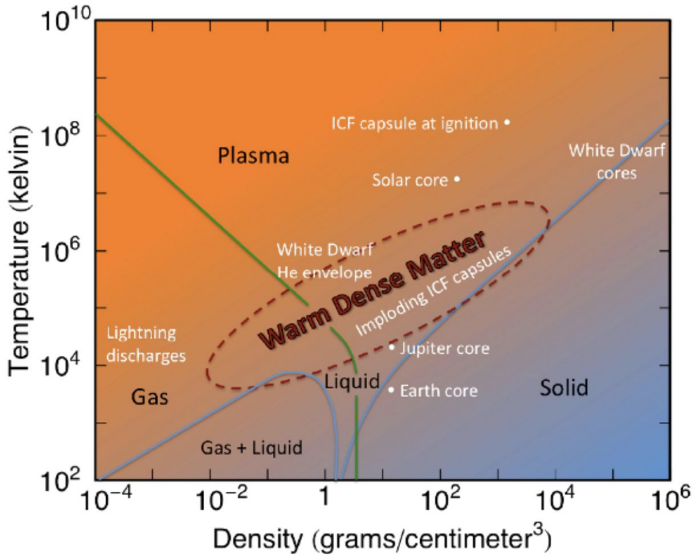

Quantum simulating heating of pre-ignition ICF

Quantum AI

“The essence of controlled laboratory thermonuclear fusion is to use the fusion product’s kinetic energy to self-heat the plasma, accelerating and perpetuating the burn”
Phys. Plasmas 26, 062701 (2019)

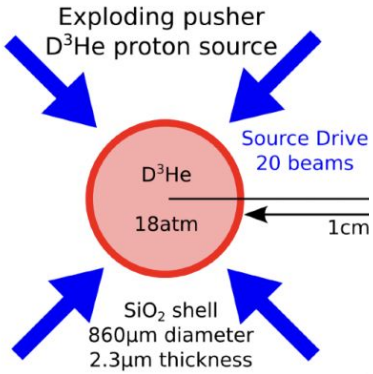

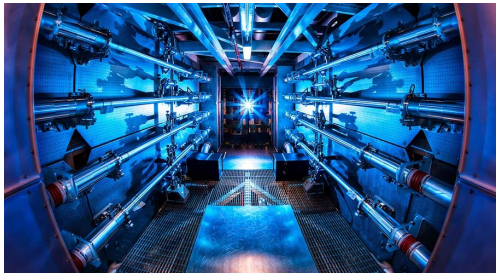

Stopping power = rate at which a material absorbs kinetic energy of charged particle passing through it

Multiscale ICF modeling depends sensitively on “stopping power” as a function of temperature, velocity

DOE spends billions of CPU hours per year running inaccurate TDDFT for stopping power

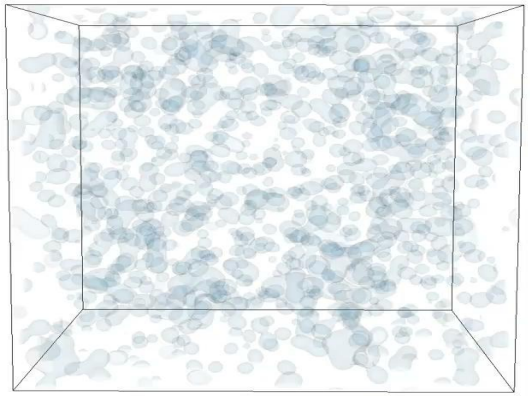

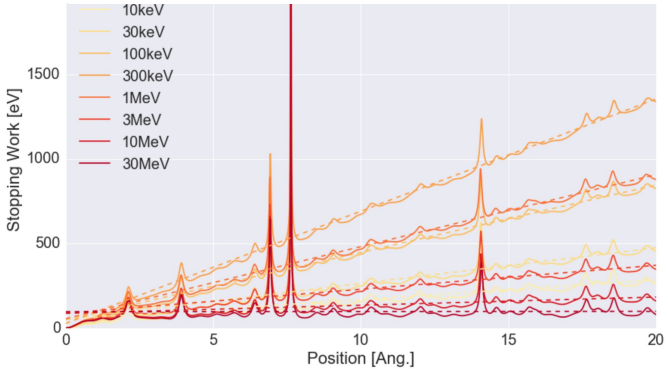

Resource estimates for stopping power

arXiv:2308.12352

Quantum AI

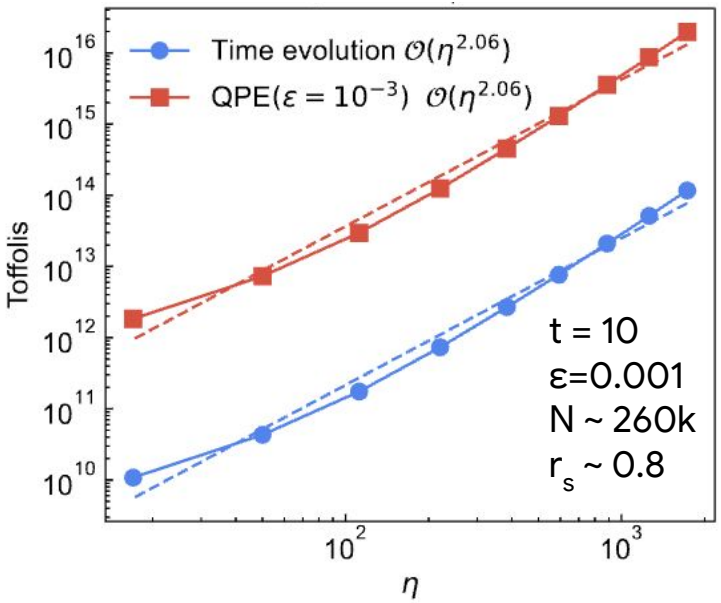

reference benchmarks

- **FeMoCO:** 2100 Qubits, 3.2 x 10^{10} Toffoli
- **P450:** 1500 Qubits, 7.0 x 10^{9} Toffoli

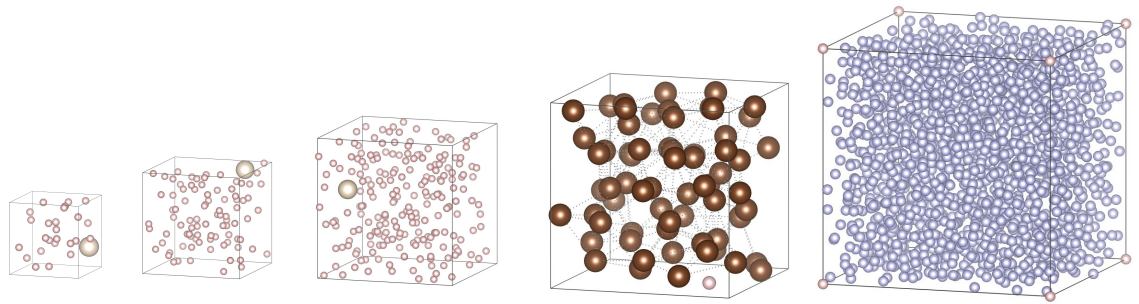

Projectile + Host	η	QSP Toffoli	Trotter Toffoli	QSP Qubits	Trotter Qubits
Helium + Hydrogen (50%)	28	1.166×10^{14}	2.143×10^{12}	1749	444
Helium + Hydrogen (75%)	92	4.239×10^{15}	5.852×10^{13}	3309	1680
Helium + Hydrogen	218	4.154×10^{16}	2.667×10^{14}	5650	3948
Hydrogen+Deuterium	1729	4.424×10^{19}	3.964×10^{16}	33038	31146
Hydrogen + Carbon	391	4.640×10^{17}	2.049×10^{15}	8841	7062

Quantum AI

Part IV: case study on the viability of quantum topological data analysis

Google

What is topological data analysis?

- Given n data points, make a graph G with data points as nodes, with edges between points wherever the points are within some distance (“filtration”) ϵ

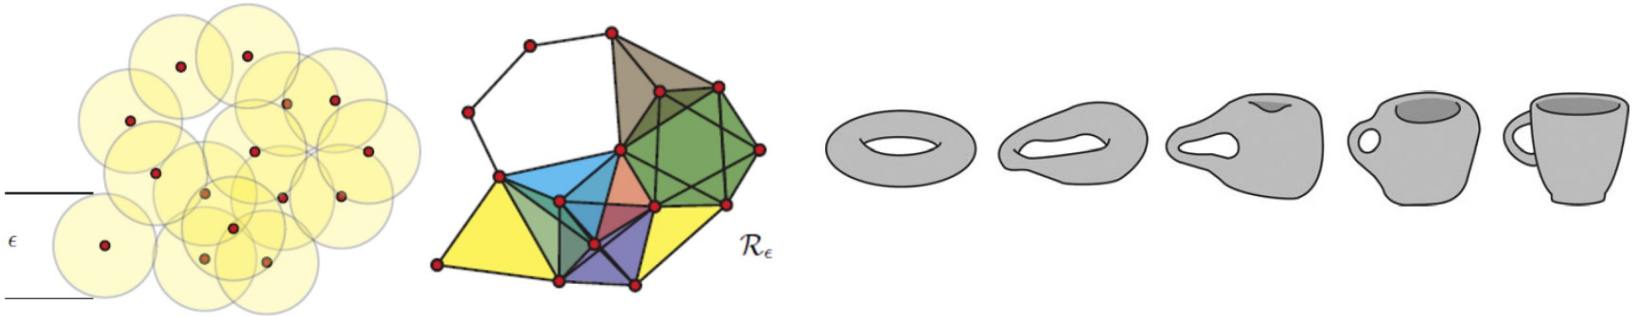

- The k^{th} order Betti number (β_k) is the number of k-dimensional holes in the simplicial complex
- Topological features are somewhat fundamental global properties, often robust to noise in data
- Real world applications of this in neuroscience, epidemiology, genetics, finance, even physics

ARTICLE

Received 17 Sep 2014 | Accepted 9 Nov 2015 | Published 25 Jan 2016

DOI: 10.1038/ncomms10138 OPEN

Quantum algorithms for topological and geometric analysis of data

Seth Lloyd[1], Silvano Garnerone[2] & Paolo Zanardi[3]

Extracting useful information from large data sets can be a daunting task. Topological methods for analysing data sets provide a powerful technique for extracting such information. Persistent homology is a sophisticated tool for identifying topological features and for determining how such features persist as the data is viewed at different scales. Here we present quantum machine learning algorithms for calculating Betti numbers—the numbers of connected components, holes and voids—in persistent homology, and for finding eigenvectors and eigenvalues of the combinatorial Laplacian. The algorithms provide an exponential speed-up over the best currently known classical algorithms for topological data analysis.

The quantum approach to TDA

Nat. Comms. 7: 10138 (2016)

- Every data point (node) is represented by a qubit, with every computational basis state representing a clique; e.g., |01101> is a clique between nodes 2,3,5

1. Prepare the equal mixed state over $Cl_k(G)$, all k-cliques in the simplicial complex
   → can be accomplished via amplitude amplification from Dicke state of order k
2. Apply phase estimation to the combinatorial Laplacian associated with the data graph G
3. Estimate β_k as the dimension of the kernel of that combinatorial Laplacian

- The most natural thing to estimate from this procedure is actually an approximate Betti number normalized by the number of k-cliques, i.e. $\beta_k / |Cl_k(G)|$

Algorithm improvements from our work

see arXiv:2209.13581

- New method based on threshold testing to prepare a mixture of fixed Hamming-weight states with garbage information that has significantly lower fault-tolerant cost
- Directly perform phase estimation on the quantum walk operator
- Use Kaiser windows functions to reduce the number of amplitude estimation steps
- Project onto the kernel of the boundary map by implementing a Chebyshev polynomial to optimally filter zero eigenvalues from phase estimation
- Use overlap estimation to estimate kernel dimension of block-encoded combinatorial Laplacian, leading to a quadratic improvement in precision over prior work

Quantum resource analysis

see arXiv:2209.13581

Classical cost: O((n choose k)

n = number of nodes/qubits
|E| = number of edges

β_k^G = kth order Betti number of graph G
δ = success probability

Cost in terms of α, the additive error in β_k^G:

$$6|E|\frac{\ln(1/\delta)}{\alpha}\sqrt{|\mathrm{Cl}_k(G)|\beta_{k-1}^G}\left[\frac{\pi}{2}\sqrt{\frac{\binom{n}{k}}{|\mathrm{Cl}_k(G)|}}+\frac{n}{\lambda_{\min}}\ln\left(\frac{4|\mathrm{Cl}_k(G)|}{\alpha}\right)\right]$$

Cost in terms of r, the multiplicative (e.g. percent) error in β_k^G:

$$6|E|\frac{\ln(1/\delta)}{r}\sqrt{\frac{|\mathrm{Cl}_k(G)|}{\beta_{k-1}^G}}\left[\frac{\pi}{2}\sqrt{\frac{\binom{n}{k}}{|\mathrm{Cl}_k(G)|}}+\frac{n}{\lambda_{\min}}\ln\left(\frac{4|\mathrm{Cl}_k(G)|}{r\beta_{k-1}^G}\right)\right]$$

Quantum resource analysis

see arXiv:2209.13581

Classical cost: O((n choose k)

n = number of nodes/qubits
|E| = number of edges

β_k^G = k^{th} order Betti number of graph G
δ = success probability

Cost in terms of α, the additive error in β_k^G:

$$3\pi|E|\frac{\ln(1/\delta)}{\alpha}\sqrt{\binom{n}{k}\beta_{k-1}^G}$$

- Never more than quadratic quantum speedup
- Implies practical advantage is unlikely

Cost in terms of r, the multiplicative (e.g. percent) error in β_k^G:

$$3\pi|E|\frac{\ln(1/\delta)}{r}\sqrt{\frac{\binom{n}{k}}{\beta_{k-1}^G}}$$

- Quantum speedup only when β_k^G is large
- Exponential iff $\beta_k^G = \Omega((n \text{ choose } k) / \text{poly}(n\ k))$

Problem instances with speedup

see arXiv:2209.13581

- We construct graphs with Betti number large enough for superpolynomial speedup
- Do real applications have speedup?
- Erdos-Renyi graphs with edge with probability p have mean Betti number (n choose k+1) / $p^{(k+1 \text{ choose } 2)}$
- Optimal p gives quartic speedup

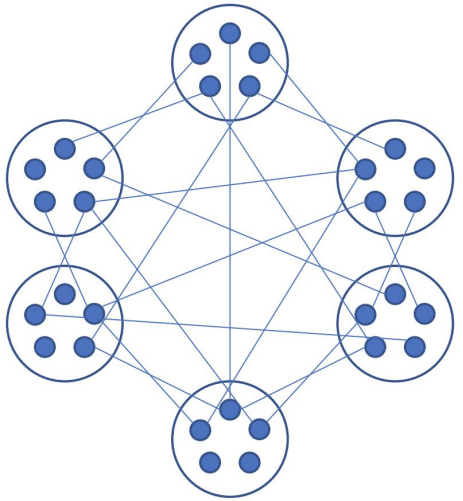

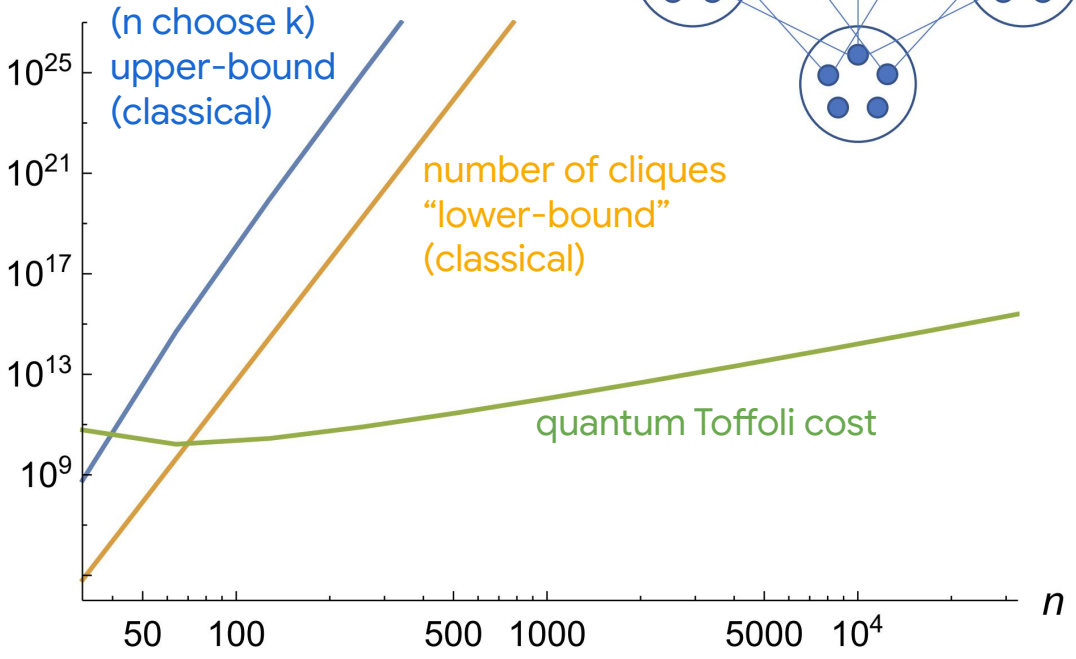

Quantum AI

Quantum AI

Part III: other interesting places to look for viable quantum applications

Google

Exponential speedup in simulating classical oscillators

Quantum AI

PRX 13, 041041 (2023) + *FOCS 2023*

Hooke's Law for coupled oscillators can be expressed as:

$$\ddot{\vec{y}}(t) = -A\vec{y}(t)$$

Adding $i\sqrt{A}\dot{\vec{y}}(t)$ to both sides gives Schrodinger equation:

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\dot{\vec{y}}(t) + i\sqrt{A}\vec{y}(t)\right) = i\sqrt{A}\left(\dot{\vec{y}}(t) + i\sqrt{A}\vec{y}(t)\right)$$

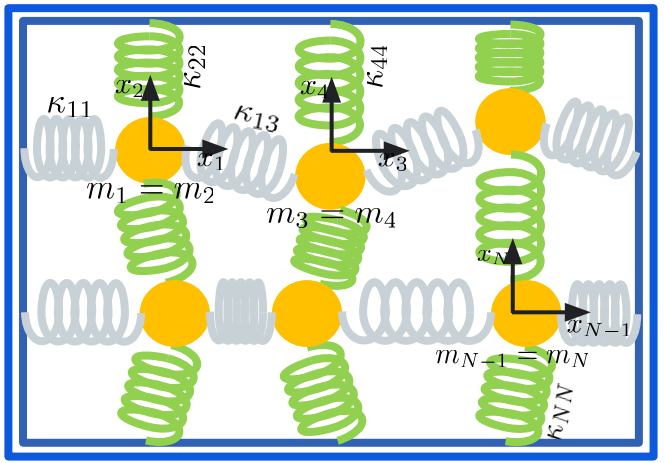

- If spring constants/masses efficiently computable, can simulate N oscillators in O(polylog(N))

- Many potential applications: modeling electrical grids, mechanical engineering, classical wave equation, molecular vibrations, statistical mechanics of fields, etc.

Exponential speedup in simulating classical oscillators

Quantum AI

PRX 13, 041041 (2023) + *FOCS 2023*

- We prove that measuring kinetic energy of an oscillator is BQP-Complete

  Even when all masses are the same and there are two spring constants

- We also show relativized exponential speedup in an oracle model, leveraging glued trees

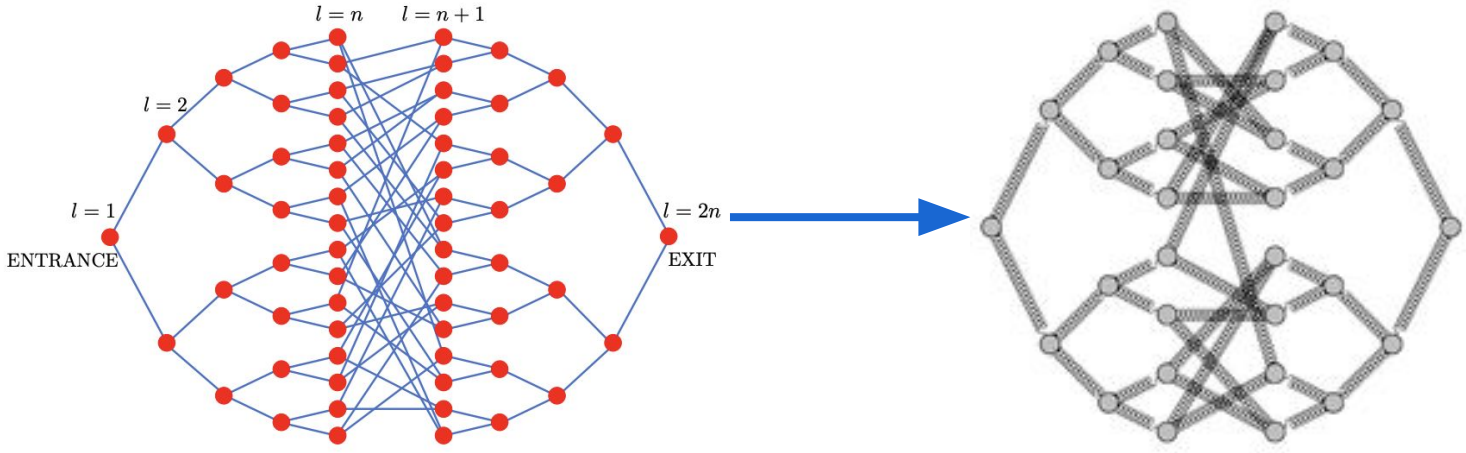

Searching for applications of quantum linear algebra

- Quantum computers can solve systems of N linear equations (A **x** = **b**) “querying” A only O(polylog(N)) times 😀
- But sometimes each queries to A or b take time O(poly(N)) 😒
- Focus on problems where linear system has underlying structure, e.g., physically inspired linear differential equations
- Should avoid situations where sampling the output is efficient via Monte Carlo (e.g. options pricing, Fokker-Planck, heat equation, etc.)
- Promising directions include acoustic or electromagnetic scattering, fluid dynamics with certain constraints, plasma physics, etc.

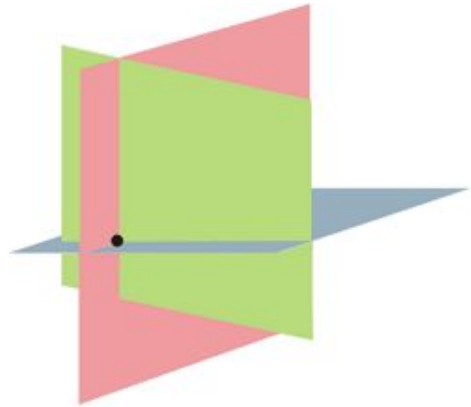

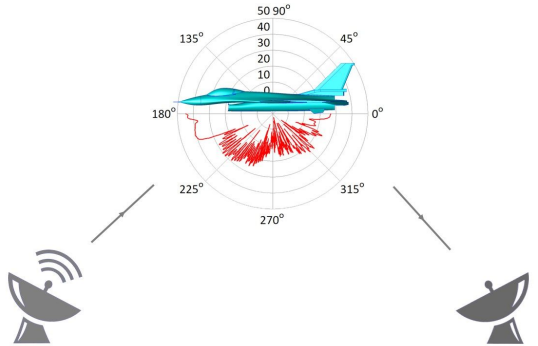

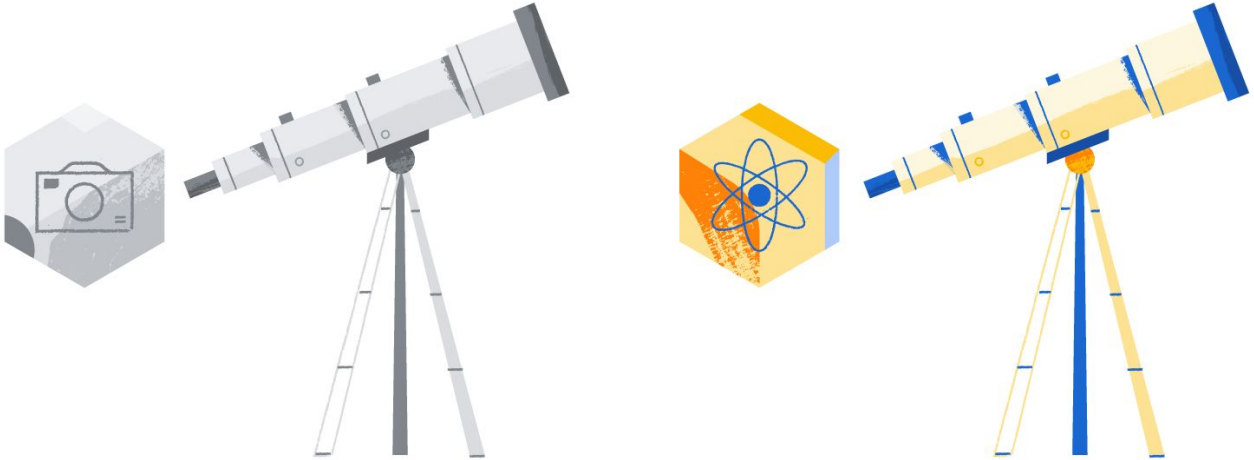

Classical Data

Quantum Data

Whether starlight or molecular spectra, thus far science has only used “classical data”

If data is collected by quantum sensor (rapidly progressing technology) and transduced to an error-corrected quantum memory, we can do amazing things with that data

Quantum states output from a quantum simulation (e.g. of chemistry) also constitute quantum data

Quantum enhanced experiments

Quantum AI

Science 376, 1182-1186 (2022)

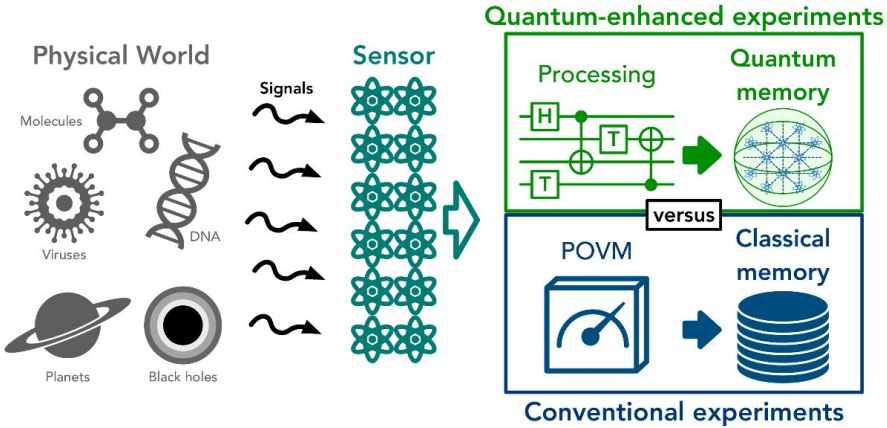

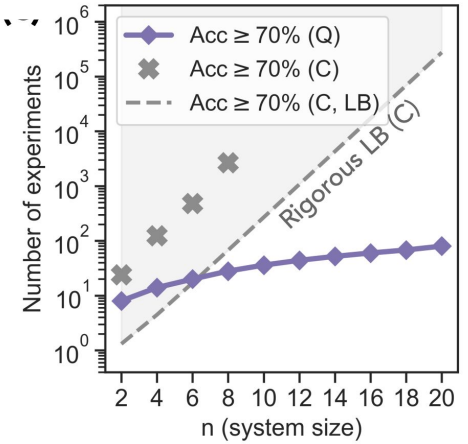

By entangling two copies of N qubit state (e.g. from a sensor), we can learn properties with 2^N fewer queries vs single copy

With limited data, one can achieve quantum advantage with very few qubits

How will a modest error-corrected quantum computer impact the world?

- We have a responsibility to answer this question with more clarity
- Simulations of molecules and materials remain promising
  - Molecular ground states in scope if strongly correlated
  - Many initio solid-state applications but resource requirements very high
  - Can outperform classical mean-field when goal is dynamics
- Computing stopping power of pre-ignition fusion fuels has decisive quantum advantage, corresponds to relevant real-world experiments for 10^{12} - 10^{13} Toffoli gates
- Interesting other applications include topological data analysis, simulating certain classical differential equations, analysis of data from quantum sensors or simulations

Quantum AI

Thank you!

Bill Huggins
Craig Gidney
David Gosset
Dar Gilboa
Dominic Berry
Edward Farhi
Fionn Malone
Joonho Lee
Kianna Wan
Marika Kieferova
Matthew Harrigan
Mike Freedman
Nicholas Rubin
Ramis Movassagh
Robert Huang
Robin Kothari
Rolando Somma
Stephen Jordan
Tanuj Khattar
Thomas O’Brien

Google