Quantum AI

Searching for valuable applications of fault-tolerant quantum computers in chemistry

Ryan Babbush

November 9, 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!

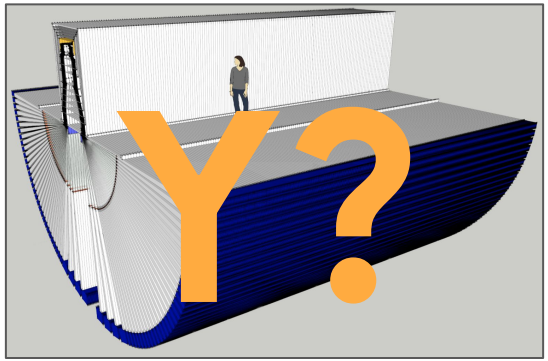

→ We hope these devices will solve important and otherwise intractable problems!

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

This talk is about research into the resources required for promising applications

- We need more researchers to study practical applications of small fault-tolerant devices!

Outline

1. Quick review of modern fault-tolerant quantum algorithms for chemistry
2. Results on identifying and assessing viability of valuable chemical applications
3. Quick survey of some recent non-electronic structure applications in chemistry

Quantum AI

Part I: a review of modern fault-tolerant quantum algorithms for chemistry

Google

Representing η electrons in N spin-orbitals

Quantum AI

Second quantization requires N qubits

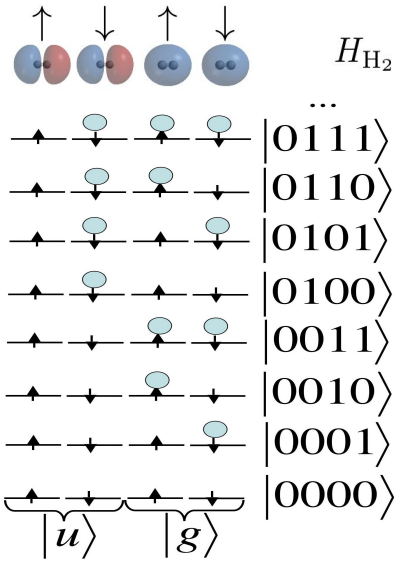

$$H_{\mathrm{H}_2} = a_0\,\mathbb{1} + a_1\,(Z_0 + Z_1) + a_2\,(Z_2 + Z_3)$$
$$+ a_6\,(Z_0Z_3 + Z_1Z_2) + a_7\,(X_0Y_1Y_2X_3$$
$$+ a_3\,Z_0Z_1 + a_4\,Z_2Z_3 + a_5\,(Z_0Z_2 + Z_1Z_3)$$
$$+ Y_0X_1X_2Y_3 - X_0X_1Y_2Y_3 - Y_0Y_1X_2X_3)$$

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

- Anti-symmetry is “encoded in the operators”
- Decomposition into local operators = NISQy
- 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”
- Operators sparse but not local ≠ NISQy
- Ideal for high precision calculations
- Low overhead for including nuclei,
  e.g. for performing dynamics

Methods for quantum simulation and their cost

Quantum AI

Each step of a product formula (Trotter) usually scales linearly in the Hamiltonian sparsity

Linear combination of unitaries (1202.5822) methods need not scale this way $H = \sum_{\ell=0}^{L-1} w_\ell U_\ell$

Qubitized quantum walks (1610.06546) act as $\mathcal{W}\,|n\rangle = e^{-i \arccos(E_n/\lambda)}\,|n\rangle$ where $H\,|n\rangle = E_n\,|n\rangle$

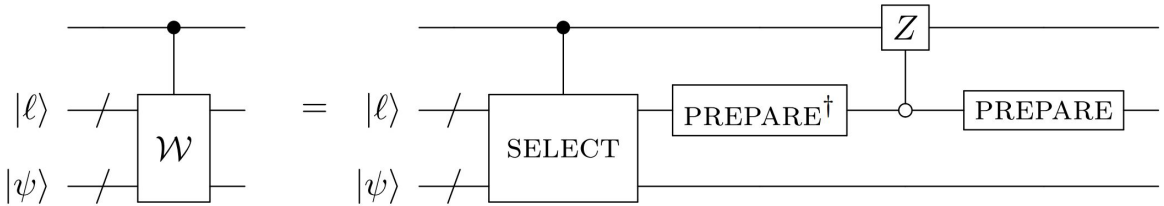

$\text{SELECT}\,|\ell\rangle\,|\psi\rangle \mapsto |\ell\rangle\,U_\ell\,|\psi\rangle$ $\qquad$ $\text{PREPARE}\,|0\rangle^{\otimes \log L} \mapsto \sum_{\ell=0}^{L-1} \sqrt{\frac{w_\ell}{\lambda}}\,|\ell\rangle$ $\qquad$ $\lambda = \sum_{\ell=0}^{L-1} |w_\ell|$

Can perform phase estimation directly on quantum walk (1711.11025) $\mathcal{O}\left(\frac{C_\mathcal{W}\,\lambda}{\epsilon}\right)$

Or, we can synthesize $e^{-i\,H\,t}$ via quantum signal processing (1606.02685) $\mathcal{O}\left(C_\mathcal{W}\lambda\,t + C_\mathcal{W}\log(1/\epsilon)\right)$

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 hard part, scaling as "cost of computing Hamiltonian coefficients from index"

- O(sqrt(Γ)) gates and ancilla where Γ is "amount of information to specify Hamiltonian"
- relies on "coherent alias sampling" (1805.03662) and "dirty QROM" (1812.00954)

no factorization

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

Γ = O(N^4)

PEA cost = O(N^4 / ϵ)

(1902.02134)

Cholesky factorization

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

Γ = O(L N^2) = O(N^3)

PEA cost = O($N^{7/2}$ / ϵ)

(1902.02134)

Tensor hypercontraction

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

Γ = O(N^2)

PEA cost = O(N^3 / ϵ)

(2011.03494)

Recall total PEA cost = O(λ sqrt(Γ) / ϵ) and for these systems λ = O(N^2)

First quantization and simple basis sets

Simple basis sets like grids, plane waves lead to analytic integrals!
Γ = O(1) - no QROM needed!

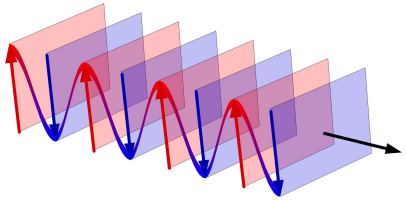

$$V = \frac{2\pi}{\Omega} \sum_{\substack{p \neq q \\ \nu \neq 0}} \frac{c^\dagger_{p+\nu} c^\dagger_{q-\nu} c_q c_p}{k_\nu^2}$$

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

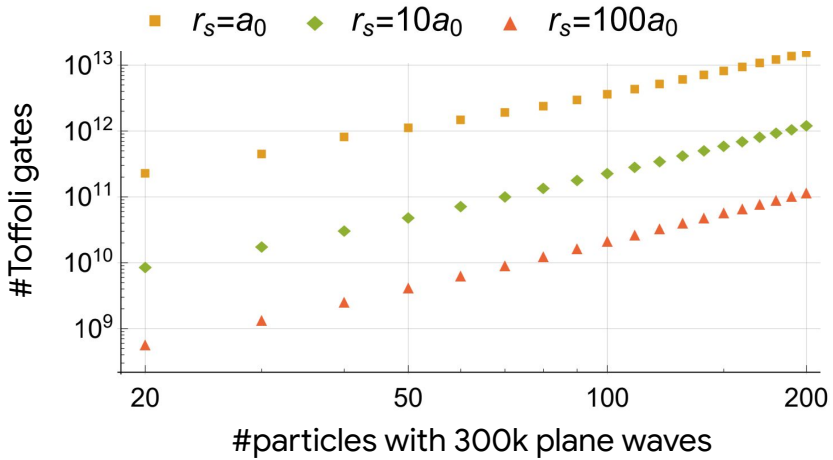

Quantum AI

Part II: the search for applications

Google

Anybody tired of FeMoCo yet?

Quantum AI

- Most studied quantum application is PEA applied to 2nd quantized MO-basis Hamiltonian

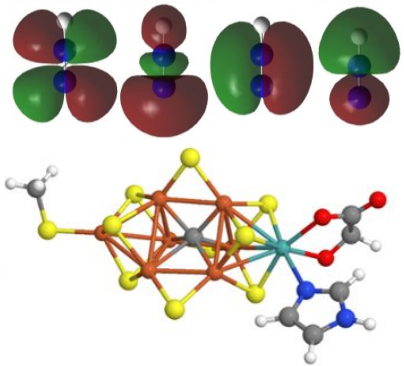

Year	arXiv	Method	Space	T Complexity	T Gates	Physical qubits
2005	0604193	Trotter	$\mathcal{O}(N)$	$\mathcal{O}(\text{poly}(N/\epsilon))$	Unknown	Unknown
2010	1001.3855	Trotter (first bounds on step complexity)	$\mathcal{O}(N)$	$\tilde{\mathcal{O}}(N^{11}/\epsilon^{3/2})$	Unknown	Unknown
2013	1312.1695	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

- Claims of quantum advantage are easiest to assess for highly specific proposals (e.g. compute this observable of this system to enough precision to resolve this)

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!)

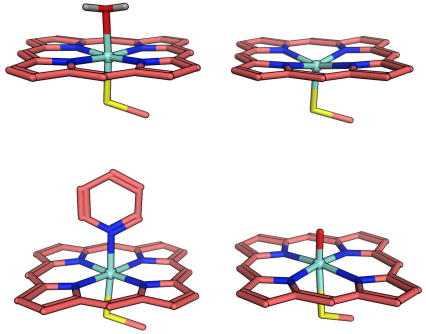

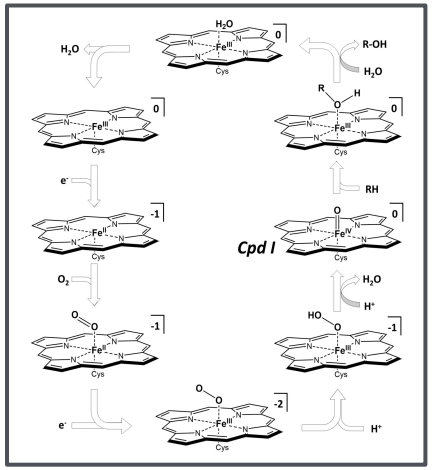

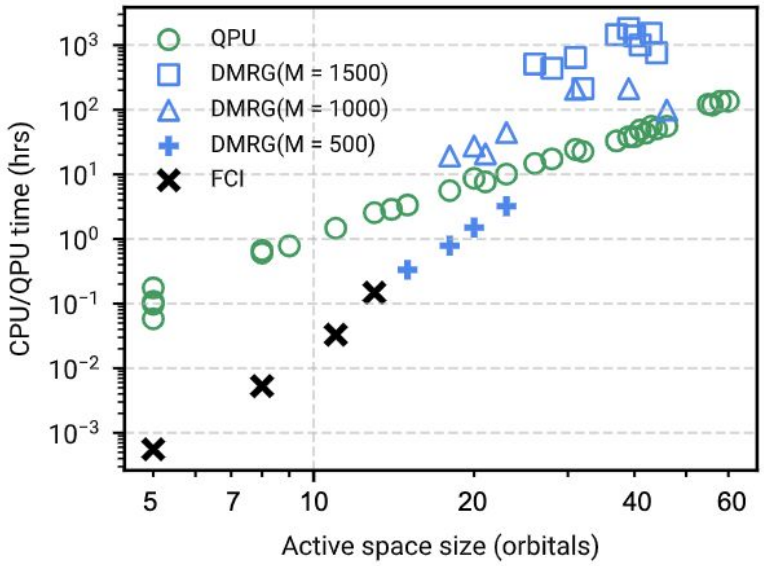

We observe onset of quantum advantage for active space sizes near 80 qubits

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)

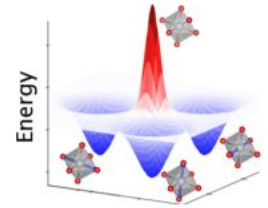

Want to get Co out of batteries, LNO is good candidate, why not Jahn-Teller distorted?

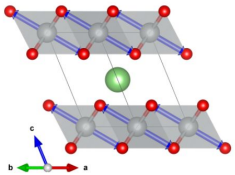

$C2/m$
(collinear Jahn-Teller)

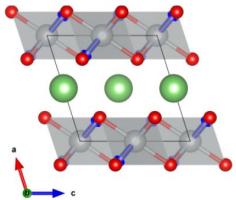

$P2_1/c$
(zig-zag Jahn-Teller)

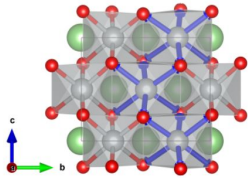

$P2/c$
(disproportionated)

- 340 basis functions (3x FeMoCo)
- 132 electrons in primitive cell (2x FeMoCo)
- Classical many-body methods unreliable for metals
- Embedding difficult to converge finite size
- DFT disagrees between functionals

recent work symmetry adapts qubitization

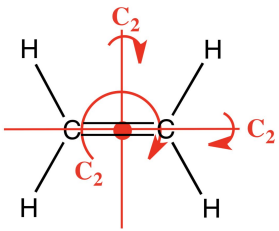

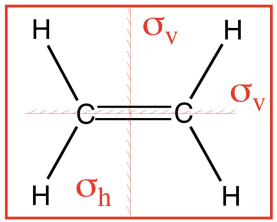

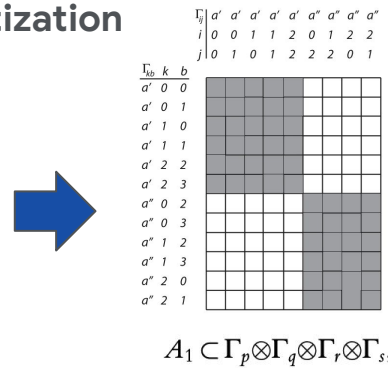

$A_1 \subset \Gamma_p \otimes \Gamma_q \otimes \Gamma_r \otimes \Gamma_s$

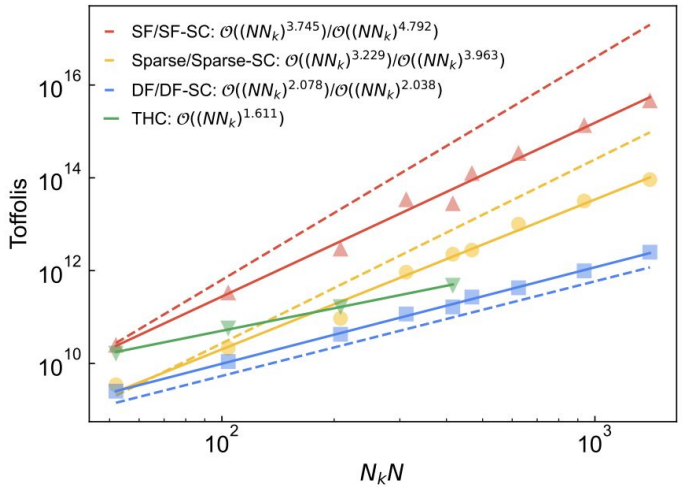

Gives linearly reduced scaling in number of k-points

Q algorithms must further improve for viability

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

Quantum AI

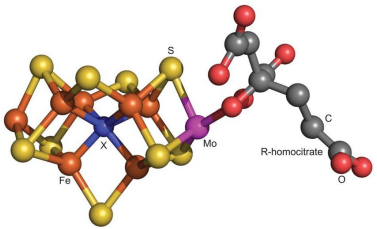

FeMoCo (fertilizer catalyst)
PRX Quantum 2, 030305 (2021)

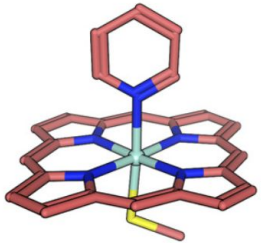

P450 (drug anti-target)
PNAS 119, 2203533119 (2022)

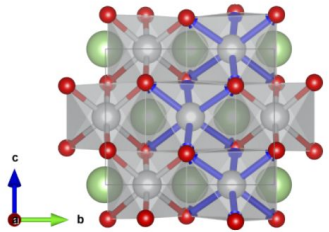

$LiNiO_2$ (battery cathode)
PRX Quantum 4, 040303 (2023)

- "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

On the importance of super-quadratic speedups

PRX Quantum 2, 010103 (2021)

Quantum AI

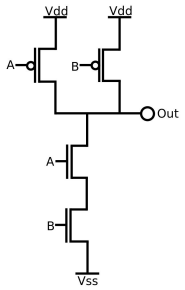

classical NAND gate (CMOS)
$<10^{-9}$ “transistorseconds”

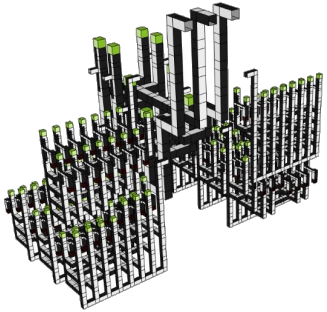

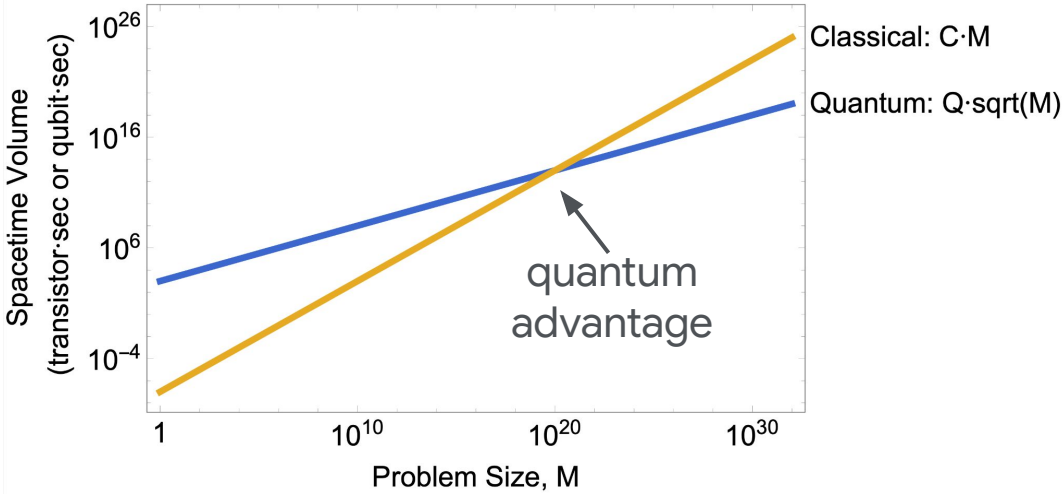

“quantum NAND” gate (distillation of Toffoli state) >10 “qubitseconds”

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

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

Quantum AI

Nat. Comm 14, 4058 (2023)

N = number of basis functions (e.g. plane waves, grid points) η = number of electrons << N

Usual quantum simulation advantage is resolution of entanglement between particles

- storing arbitrary wavefunction 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 temp $$N^{10/3}\eta^{1/3}t + N^{11/3}t/\eta^{2/3}$$

first quantized quantum simulations need only O(η log N) qubits rather than O(N)

$$N^{1/3}\eta^{7/3}t + N^{2/3}\eta^{4/3}t$$ or, for very large N $$N^{1/3}\eta^{8/3}t$$

Prospects for super-quadratic speedup over mean-field

Nat. Comm 14, 4058 (2023)

Estimates below for sampling from dynamics; additional costs from state prep, measurement

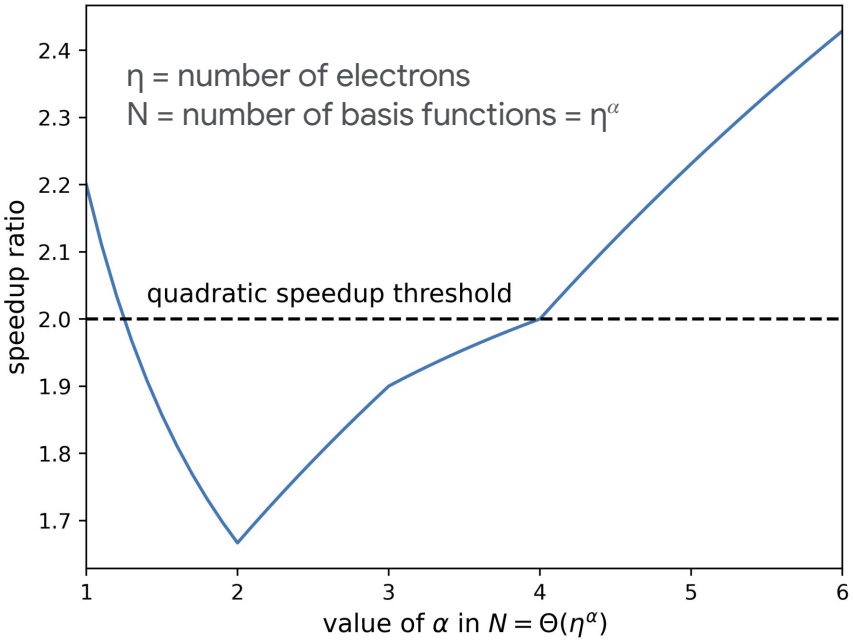

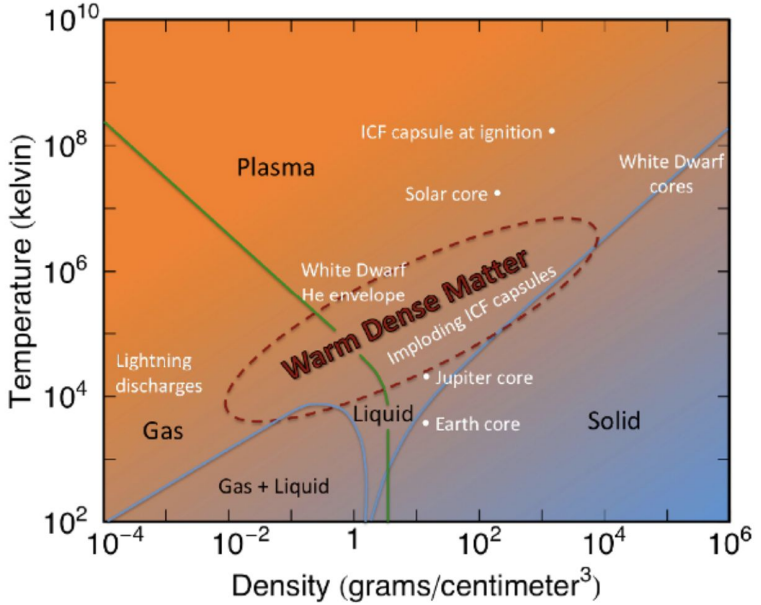

prospects much better when systems require at least *some* correlation, or for systems at finite temperature

Quantum simulating heating processes of pre-ignition ICF

“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)

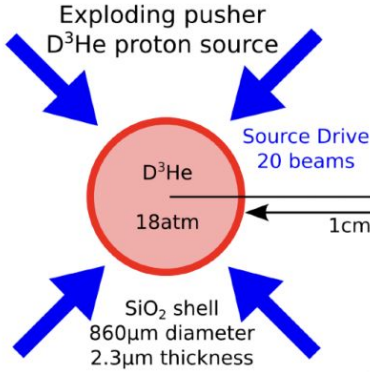

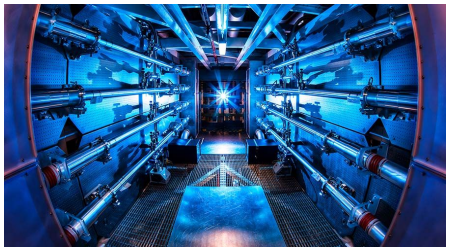

Drag of on-transport through in the pre-ignition plasma leads to heating

Difficult regime because electron correlation and temperature are about equal contributions

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

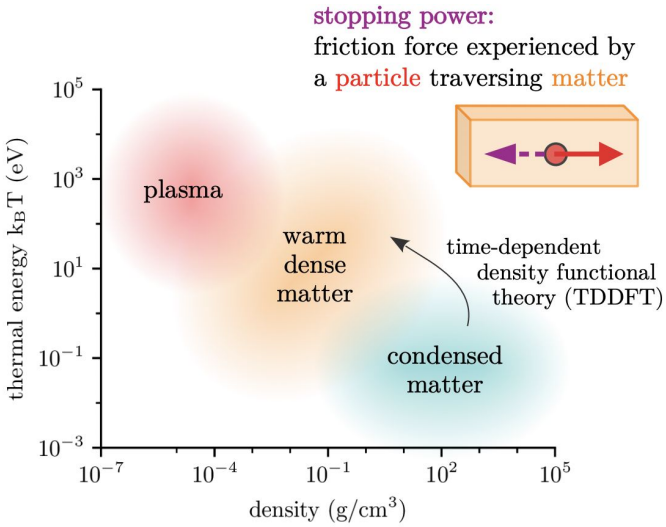

Quantum AI

Classically computing stopping power

arXiv:2308.12352

Quantum AI

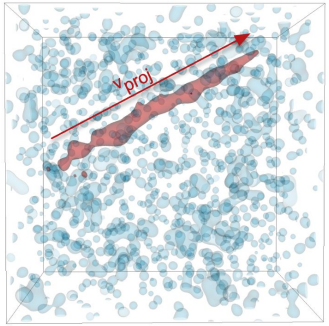

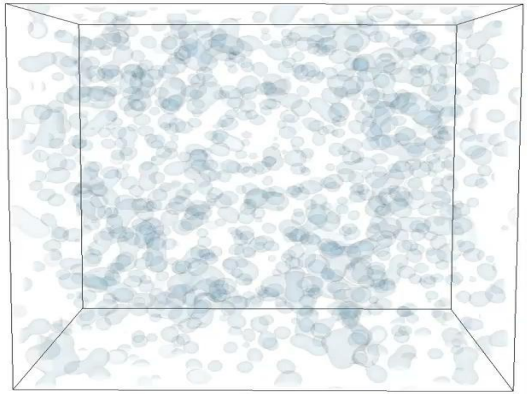

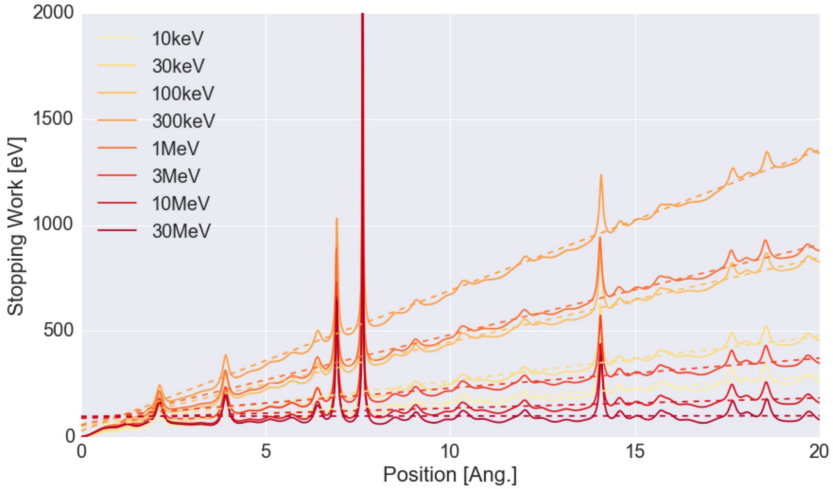

- DOE spends billions of CPU-hours computing stopping power with TDDFT
- TDDFT not very accurate but no other choice - difficult to measure via experiment
- Our collaborators alone use 10% of Trinity Supercomputer (~$200 mil) for this, annually

Resource estimates for stopping power

arXiv:2308.12352

Quantum AI

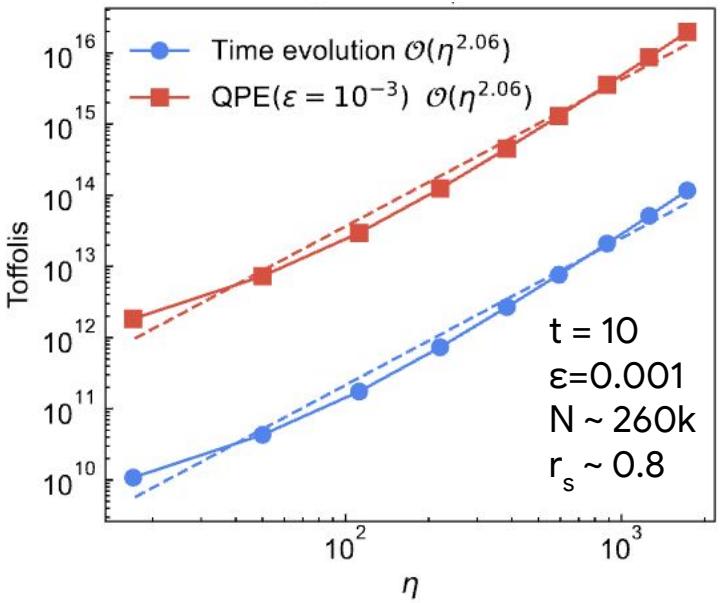

reference benchmarks

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

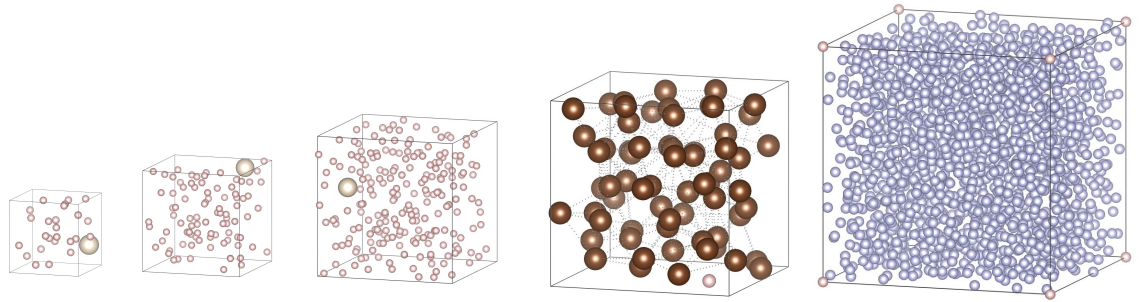

Projectile + Host	η	QSP Toffoli	Trotter Toffoli	QSP Qubits	Trotter Qubits
Helium + Hydrogen (50%)	28	1.166×10^{14}	2.143×10^{12}	1749	444
Helium + Hydrogen (75%)	92	4.239×10^{15}	5.852×10^{13}	3309	1680
Helium + Hydrogen	218	4.154×10^{16}	2.667×10^{14}	5650	3948
Hydrogen+Deuterium	1729	4.424×10^{19}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

Quantum AI

Part III: various interesting non-electronic structure but chemistry adjacent things to explore with quantum computers

Google

Exponential speedup in simulating classical oscillators

Quantum AI

arXiv:2303.13012, accepted *PRX + FOCS*

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:

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

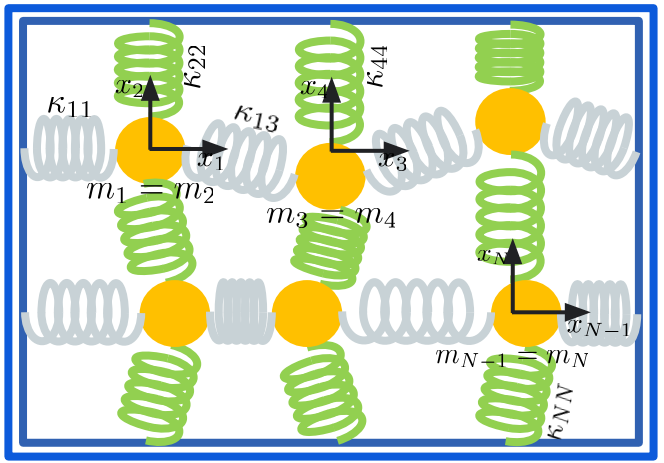

- If spring constants/masses efficiently computable, can simulate N oscillators in O(polylog(N))
- We show problem is BQP-Complete and prove relativized exponential speedup
- Many potential applications: modeling electrical grids, mechanical engineering, classical wave equation, molecular vibrations, statistical mechanics of fields, etc.

Quantum algorithm for difficult NMR spectra

Quantum AI

PRX Quantum 3, 030345 (2022)

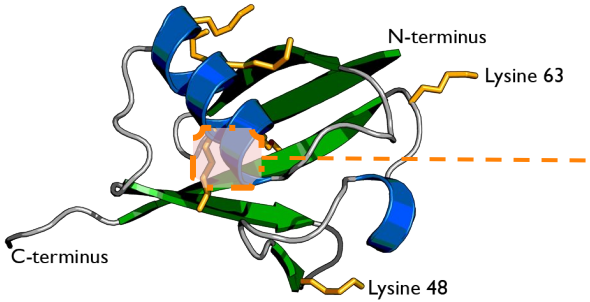

protein

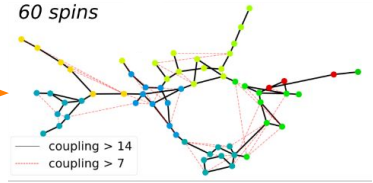

quantum spin model

Time-dependent magnetic field

Response

NMR spectrometer

Inverse problem: infer molecular structure (spin Hamiltonian) from nuclear magnetic resonance data

Hard to solve classically when dipolar coupling is strong (e.g. when protein confined to surface)

Easy to solve on quantum device potentially enabling new and more powerful forms of NMR experiments

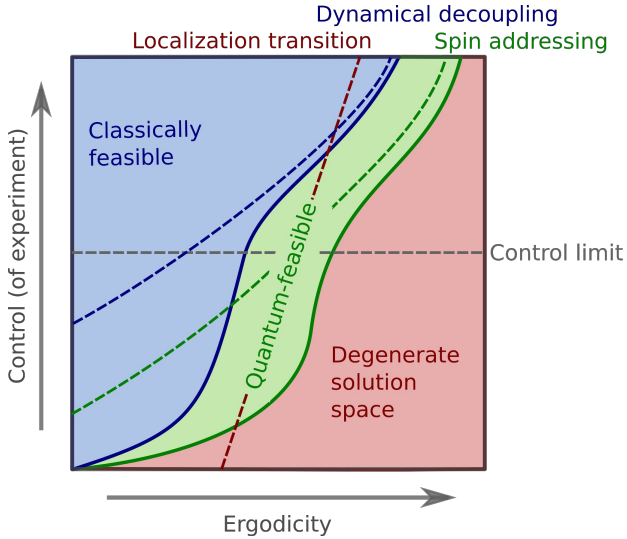

Quantum enhanced experiments

Quantum AI

Science 376, 1182-1186 (2022)

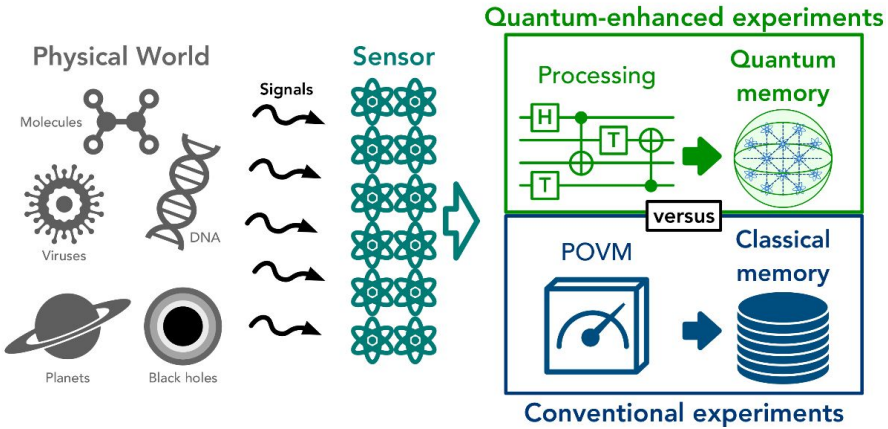

By entangling two copies of N qubit state (e.g. from a sensor), we can learn properties with 2^N fewer queries vs single copy

With limited data, one can achieve quantum advantage with very few qubits; achieved on our chips!

Many possibilities here for more sensitive probes of chemistry

Talk summary

- Critical to identify chemical problems where quantum advantage would be impactful
- We have a few examples in strongly correlated molecules but could use more; many examples in ab initio materials but quantum algorithms must improve
- Quantum algorithms for electron dynamics are sometimes even faster than classical mean-field methods, especially when needing to resolve correlation or temperature
- 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 frontiers outside of electronic structure in helping to interpret NMR spectra (maybe even NISQ-viable), molecular vibrations, machine learning of quantum data

Quantum AI

Thank you!

Dominic Berry (Macquarie)
Nicholas Rubin (Google)
Joonho Lee (Harvard)
Andrew Baczewski (Sandia)

and many others!

Google

Appendix: Galerkin discretizations

Quantum AI

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

$$U_{pq}\left(\vec{R}\right) = \sum_{\ell=1}^{L} \int dr\, \phi_p^*(r) \left(\frac{\zeta_\ell}{|r - R_\ell|}\right) \phi_q(r)$$

$$V_{pqrs}^{(\alpha,\beta)} = \int dr_1\, dr_2\, \phi_p^*(r_1)\, \phi_q^*(r_2) \left(\frac{\alpha\beta}{|r_1 - r_2|}\right) \phi_r(r_2) \phi_s(r_1)$$

Second quantization space = O(N)

$$T = \sum_{p,q=1}^{N} T_{pq}^{(1)} a_p^\dagger a_q$$ sparsity =O(N^2)

$$U = -\sum_{p,q=1}^{N} U_{pq}\left(\vec{R}\right) a_p^\dagger a_q$$ sparsity =O(N^2)

$$V = \frac{1}{2} \sum_{p,q,r,s=1}^{N} V_{pqrs}^{(1,1)} a_p^\dagger a_q^\dagger a_r a_s$$ sparsity =O(N^4)

First quantization space = O(η log N)

$$T = \sum_{i=1}^{\eta} \sum_{p,q=1}^{N} T_{pq}^{(1)} |p\rangle\langle q|_i$$ sparsity =O(ηN)

$$U = -\sum_{i=1}^{\eta} \sum_{p,q=1}^{N} U_{pq}\left(\vec{R}\right) |p\rangle\langle q|_i$$ sparsity =O(ηN)

$$V = \frac{1}{2} \sum_{i,j=1}^{\eta} \sum_{p,q,r,s=1}^{N} V_{pqrs}^{(1,1)} |p\rangle\langle s|_i |q\rangle\langle r|_j$$ sparsity =O($\eta^2 N^2$)

$$|p\rangle\langle q|_i \equiv \mathbb{1}_1 \otimes \mathbb{1}_2 \cdots \otimes |p\rangle\langle q|_i \cdots \otimes \mathbb{1}_\eta$$