Tradeoffs Between Compactness and Complexity in Quantum Simulations of Chemistry

Ryan Babbush *et al.*

Google AI Quantum

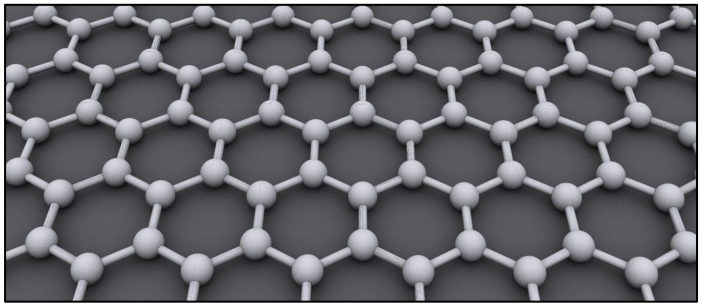

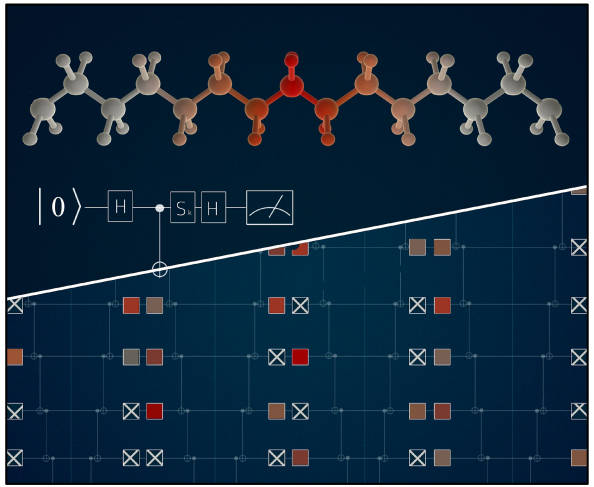

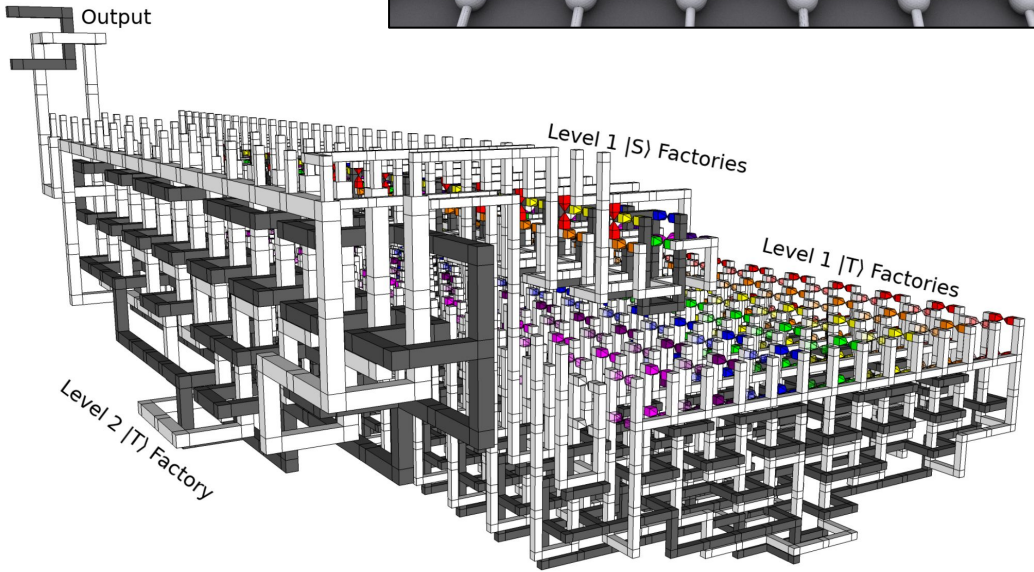

Chemistry is interacting electrons

Google AI Quantum

Systems of interacting electrons are near ideal application for quantum simulation
They're classically hard to simulate but extremely important

CPU seconds

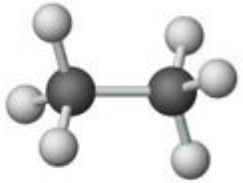

CPU minutes

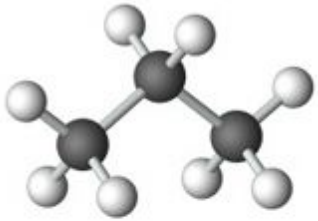

CPU days

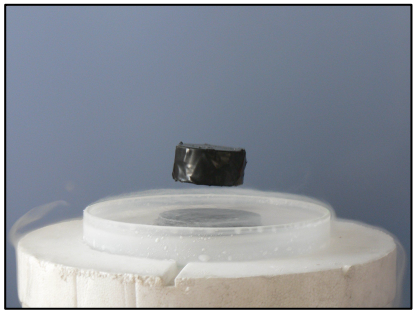

The prospect of more efficient simulations is both scientifically exciting and valuable!

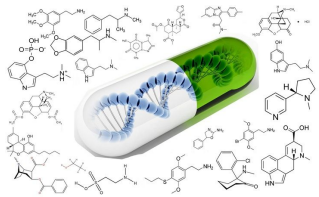

The molecular electronic structure problem

Google AI Quantum

Goal is to solve for the energy of molecule

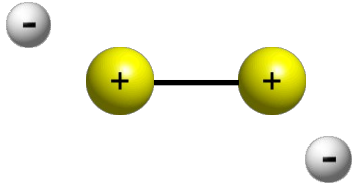

$$H = \hat{T}_{\text{nuc}} + \hat{T}_{\text{elec}} + \hat{V}_{\text{nuc-nuc}} + \hat{V}_{\text{nuc-elec}} + \hat{V}_{\text{elec-elec}}$$

Clamp nuclei under Born-Oppenheimer approx and solve for energy surface

$$H(R) = \hat{T}_{\text{elec}} + \hat{V}_{\text{nuc-elec}}(R) + \hat{V}_{\text{elec-elec}}$$

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

Should only use quant comp for "the hard part"

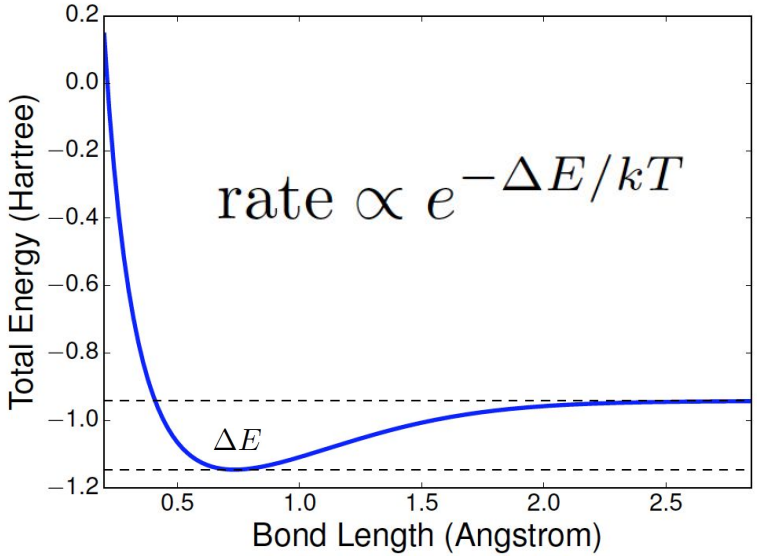

The second quantized Galerkin discretization

Google AI Quantum

To represent wavefunctions on computer one must discretize space (confine to basis)

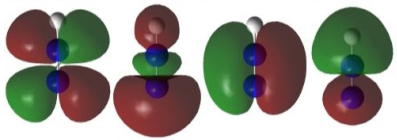

$$|\psi\rangle = a_1 |0011\rangle + a_2 |0101\rangle + a_3 |1001\rangle + a_4 |0110\rangle + a_5 |1010\rangle + a_6 |1100\rangle$$

If η electrons confined to N locations, there are (N choose η) configurations!

Galerkin discretization in MOs leads to O(N^4) Hamiltonian terms at all sizes

$$H = \frac{1}{2} \sum_{pqrs} \langle pq| \hat{V}_{\text{elec-elec}} |sr\rangle \, a_p^\dagger a_q^\dagger a_r a_s + \sum_{pq} \langle p| \hat{T}_{\text{elec}} + \hat{V}_{\text{nuc-elec}}(R) |q\rangle \, a_p^\dagger a_q$$

2nd quantization = antisymmetry in operators

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

Jordan-Wigner transformation

$$a_p^\dagger = (X_p - iY_p) \, Z_{p-1} Z_{p-2} \cdots Z_0$$
$$a_p = (X_p + iY_p) \, Z_{p-1} Z_{p-2} \cdots Z_0$$

Hamiltonian representation is tied to cost

Google AI Quantum

Physical Review X 8 (1), 011044 (2018)

Using molecular orbitals leads to Hamiltonian with O(N^4) terms

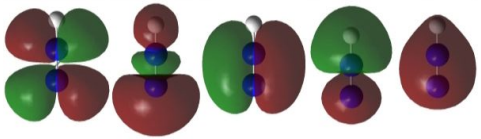

$$H = \sum_{p,q} h_{pq}\, a_p^\dagger a_q + \frac{1}{2} \sum_{p,q,r,s} h_{pqrs}\, a_p^\dagger a_q^\dagger a_r a_s$$

By carefully changing the basis, we obtain Hamiltonian with O(N^2) terms

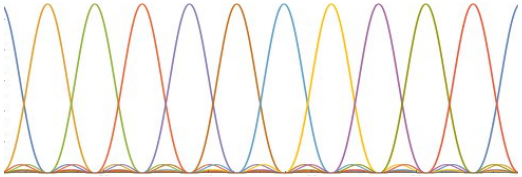

$$H = \sum_{pq} T_{pq} a_p^\dagger a_q + \sum_{p} U_p n_p + \sum_{p \neq q} V_{pq} n_p n_q$$

Plane waves are the natural basis for periodic systems

Molecules need ~100x more plane waves than Gaussians

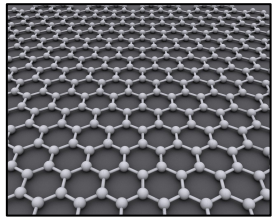

Physical Review Letters 120 (11), 110501 (2018) shows N depth Trotter step on linear chip

Can also perform rotations of single-particle basis in N/2 depth on linear array

Preparing initial states with strong support

Google AI Quantum

Tubman *et al.*, arXiv:1809.05523

The chemistry you learn in high school is a mean-field description, and it kind of works!

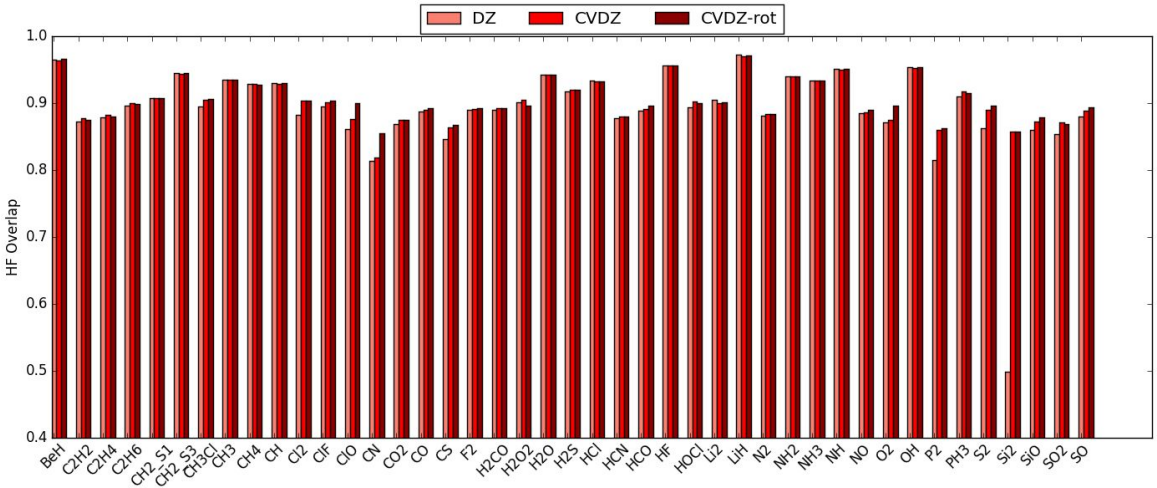

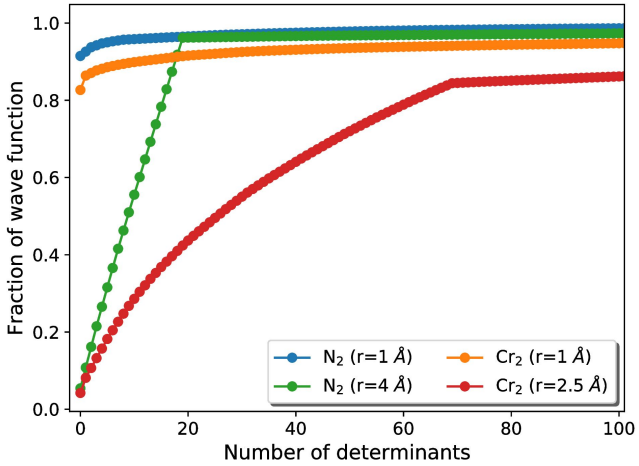

Most methods capture *either* static correlation (sparse) or dynamic correlation (low rank)

Some molecules, especially out of equilibrium, cannot be described by one configuration

Can prepare superposition of L arbitrary configurations at O(L) T gate complexity

We want a fault-tolerant quantum computer!

Google AI Quantum

100K - 1M superconducting qubits on a planar array
10^{-3} error rates and roughly 150 μs state distillation

How much is that going to cost?

More than we paid for YouTube, but less than we paid for Motorola?

What interesting, classically intractable problems can we solve with that? How quickly?

Most promising application is small but difficult (~100 qubit) chemistry problems

PNAS, 114: 7555-7560 (2017): ~10^{14} T gates
state distillation alone would takes at least a gigaqubitmonth

Phys. Rev. X 8, 041015 (2018): ~10^{8} T gates
we can do with 700k qubits in 4 hours

Cost of fault-tolerant implementations

Google AI Quantum

Year	arXiv Reference	Algorithm	Representation	Space Complexity	Gate Complexity
2005	quant-ph/0604193	Trotter	2nd Quantized Gaussians	$\mathcal{O}(N)$	$\mathcal{O}(\text{poly}(N/\epsilon))$
2010	1001.3855	Trotter	2nd Quantized Gaussians	$\mathcal{O}(N)$	$\mathcal{O}(N^{11}/\epsilon^{3/2})$
2012	1208.5986	Trotter	2nd Quantized Gaussians	$\mathcal{O}(N)$	$\mathcal{O}(N^{10}\log(N)/\epsilon^{3/2})$
2013	1312.1695	Trotter	2nd Quantized Gaussians	$\mathcal{O}(N)$	$\mathcal{O}(N^{10}/\epsilon^{3/2})$
2013	1312.2579	Trotter	1st Quantized Gaussians	$\mathcal{O}(\eta \log N)$	$\mathcal{O}(\eta^2 N^8/\epsilon^{3/2})$
2014	1403.1539	Trotter	2nd Quantized Gaussians	$\mathcal{O}(N)$	$\mathcal{O}(N^9/\epsilon^{3/2})$
2014	1406.4920	Trotter	2nd Quantized Gaussians	$\mathcal{O}(N)$	$\mathcal{O}(N^6/\epsilon^{3/2})$
2015	1506.01020	Taylor Series	2nd Quantized Gaussians	$\mathcal{O}(N)$	$\mathcal{O}(N^5\text{polylog}(N/\epsilon)/\epsilon)$
2015	1506.01029	Taylor Series	1st Quantized Gaussians	$\mathcal{O}(\eta \log N)$	$\mathcal{O}(\eta^2 N^3\text{polylog}(N/\epsilon)/\epsilon)$
2017	1706.00023	Taylor Series	2nd Quantized Plane Waves	$\mathcal{O}(N)$	$\mathcal{O}(N^{11/3}\text{polylog}(N/\epsilon)/\epsilon)$
2018	1805.03662	Qubitization	2nd Quantized Plane Waves	$\mathcal{O}(N)$	$\mathcal{O}(N^3/\epsilon + N^2\log(1/\epsilon)/\epsilon)$
2018	1805.00675	Interaction Picture	2nd Quantized Plane Waves	$\mathcal{O}(N \log N)$	$\mathcal{O}(N^2\text{polylog}(N/\epsilon)/\epsilon)$
2018	1807.09802	Interaction Picture	1st Quantized Plane Waves	$\mathcal{O}(\eta \log N)$	$\mathcal{O}(\eta^{8/3}N^{1/3}\text{polylog}(N/\epsilon)/\epsilon)$

TABLE I. Lowest gate complexity quantum algorithms for sampling in the eigenbasis of the electronic structure Hamiltonian. N is number of spin-orbitals, $\eta \ll N$ is number of electrons and ϵ is target precision.

We'll talk about these algorithms, then discuss how to overcome plane wave basis error

Phase estimation to sample in eigenbasis of H

Google AI Quantum

Goal is to sample eigenvalues of *H* on which initial state ψ has support, to precision ϵ

$$H\,|n\rangle = E_n\,|n\rangle \qquad \mathrm{Prob}\,(E_n) = |\langle\psi|n\rangle|^2$$

Can accomplish with phase estimation

$$\mathcal{W}\,|n\rangle = e^{-if(E_n)}\,|n\rangle \qquad \mathcal{O}\left(\frac{C_{\mathcal{W}}}{\epsilon}\left|\frac{\mathrm{d}f}{\mathrm{d}E}\right|^{-1}\right)$$

Time evolution is typical

$$\mathcal{W} \approx e^{-iHt} \qquad \left|\frac{\mathrm{d}f}{\mathrm{d}E}\right|^{-1} = \frac{1}{t} \qquad \mathcal{O}\left(\frac{C_{\mathcal{W}}\,(\epsilon, t)}{\epsilon\, t}\right)$$

Can use a quantum walk

$$\mathcal{W} = e^{i\arccos(H/\lambda)} \qquad \left|\frac{\mathrm{d}f}{\mathrm{d}E}\right|^{-1} \leq \lambda \qquad \mathcal{O}\left(\frac{C_{\mathcal{W}}\,\lambda}{\epsilon}\right)$$

Focus of Phys. Rev. X 8, 041015 (2018) is compiling circuits for this *W* to surface code

$$C_{\mathcal{W}} = \mathcal{O}\left(N + \log(1/\epsilon)\right) \qquad \lambda_{\mathrm{chem}} = \mathcal{O}\left(N^2\right) \qquad \lambda_{\mathrm{hub}} = \mathcal{O}\left(N\right)$$

Oracular framework for LCU quantum walks

Google AI Quantum

Simulation based on Hamiltonian as linear combinations of unitaries (1202.5822) $H = \sum_{\ell=0}^{L-1} w_\ell U_\ell$

High level framework in terms of queries to two “qubitization” oracles (1610.06546)

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

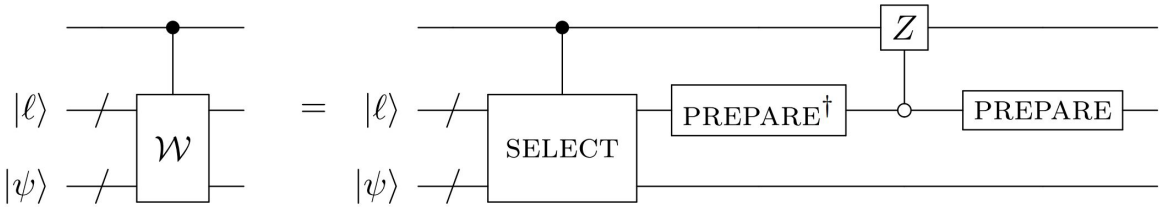

We can perform phase estimation directly on qubitized quantum walk (1711.10460, 1711.11025)

Avoids overhead to synthesize $e^{-i H t}$ via quantum signal processing (1606.02685)

We’re not finished! How to realize SELECT and PREPARE?

Implementing SELECT

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

Google AI Quantum

$$H = \sum_{p\neq q,\sigma} \frac{T(p-q)}{2}\left(X_{p,\sigma}\overrightarrow{Z}\,X_{q,\sigma} + Y_{p,\sigma}\overrightarrow{Z}\,Y_{q,\sigma}\right) + \sum_{(p,\alpha)\neq(q,\beta)} \frac{V(p-q)}{4} Z_{p,\alpha}Z_{q,\beta} - \sum_{p,\sigma}\left(\frac{T(0)+U(p)+\sum_q V(p-q)}{2}\right) Z_{p,\sigma}$$

Next step is to write specifications for the circuit:

$$\text{SELECT}_{\text{CHEM}}\,|\theta,U,V,p,\alpha,q,\beta\rangle\,|\psi\rangle = (-1)^\theta\,|\theta,U,V,p,\alpha,q,\beta\rangle \otimes \begin{cases} Z_{p,\alpha}\,|\psi\rangle & U \wedge \neg V \wedge ((p,\alpha)=(q,\beta)) \\ Z_{p,\alpha}Z_{q,\beta}\,|\psi\rangle & \neg U \wedge \quad V \wedge ((p,\alpha)\neq(q,\beta)) \\ X_{p,\alpha}\overrightarrow{Z}\,X_{q,\alpha}\,|\psi\rangle & \neg U \wedge \neg V \wedge (p<q) \wedge (\alpha=\beta) \\ Y_{q,\alpha}\overrightarrow{Z}\,Y_{p,\alpha}\,|\psi\rangle & \neg U \wedge \neg V \wedge (p>q) \wedge (\alpha=\beta) \\ \text{UNDEFINED} & \text{otherwise.} \end{cases}$$

Brute-force solution has $O(N^3 \log N)$ complexity

Our circuits

T count: 12N + 8 log N + O(1)

Ancilla qubit count: log N + O(1)

$$|p\rangle\,|\psi\rangle \mapsto |p\rangle\,X_p\,|\psi\rangle$$

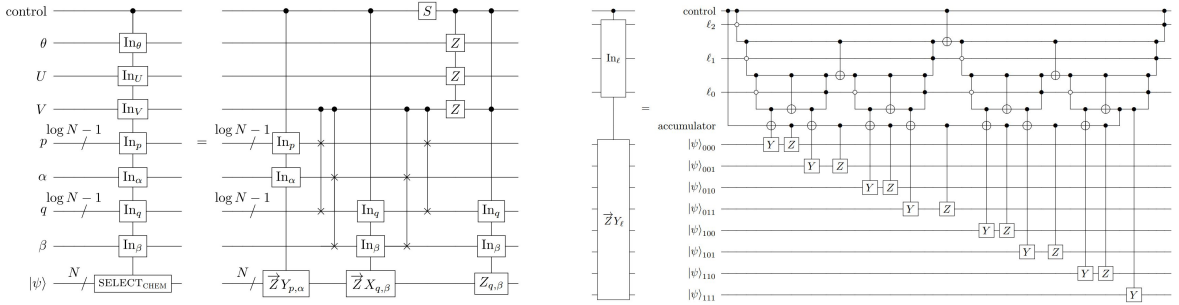

Control NOT on single bit of “streamed” unary register that is ON iff streamed index = |p⟩

Implementing PREPARE

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

Google AI Quantum

$$H = \sum_{p\neq q,\sigma} \frac{T(p-q)}{2}\left(X_{p,\sigma}\vec{Z}\,X_{q,\sigma} + Y_{p,\sigma}\vec{Z}\,Y_{q,\sigma}\right) + \sum_{(p,\alpha)\neq(q,\beta)} \frac{V(p-q)}{4} Z_{p,\alpha}Z_{q,\beta} - \sum_{p,\sigma}\left(\frac{T(0)+U(p)+\sum_q V(p-q)}{2}\right) Z_{p,\sigma}$$

Brute force: $O(N^2 \log (N/\epsilon))$ gates

Translational invariance: $O(N \log (N/\epsilon))$ gates

Our T count: $6N + O(\log(N/\epsilon))$

"QROM" for read-only data lookup

$$\text{QROM}_d \cdot \sum_{\ell=0}^{L-1} \alpha_\ell |\ell\rangle|0\rangle = \sum_{\ell=0}^{L-1} \alpha_\ell |\ell\rangle|d_\ell\rangle$$

T complexity for QROM to load L entry database is 4L

$$\sqrt{\frac{1}{L}}\sum_{\ell=0}^{L-1} |\ell\rangle\,|0\rangle \mapsto \sqrt{\frac{1}{L}}\sum_{\ell=0}^{L-1} |\ell\rangle\,|\text{alt}_\ell\rangle\,|\text{keep}_\ell\rangle \mapsto \sum_{\ell=0}^{L-1}\sqrt{\frac{\omega_\ell}{\lambda}}\,|\ell\rangle\,|\text{junk}_\ell\rangle$$

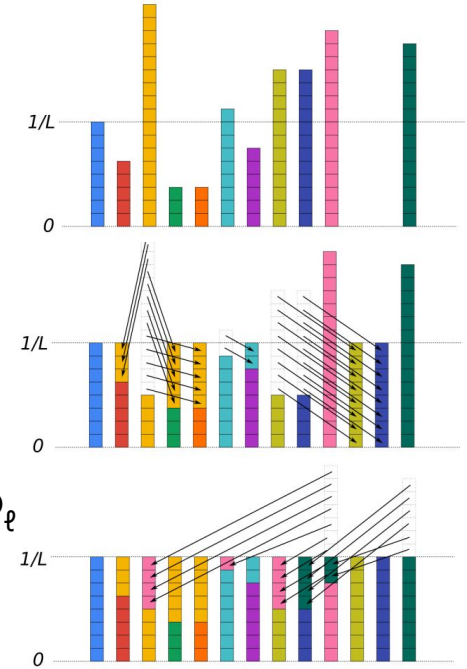

Coherent alias sampling: swap $|\ell\rangle$ with $|\text{alt}_\ell\rangle$ with probability = keep_ℓ

Okay for this control register to be entangled with some junk

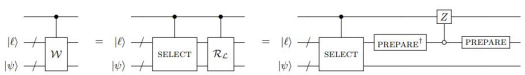
SELECT
SELECT
PREPARE†
PREPARE

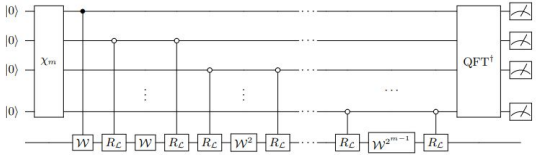
QFT†

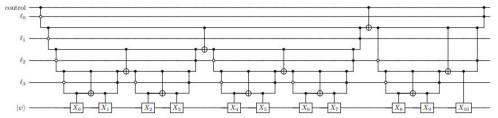

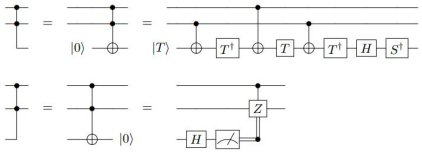

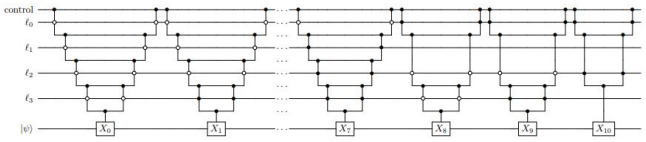
control

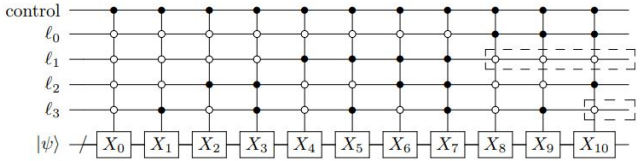
control

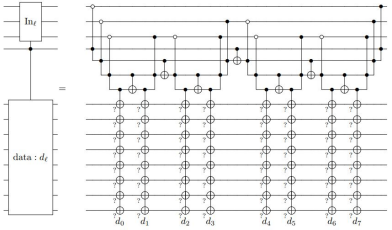

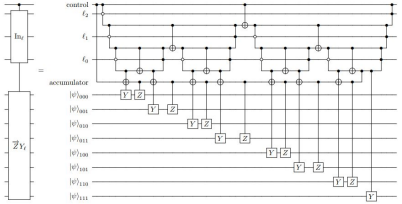
control
accumulator

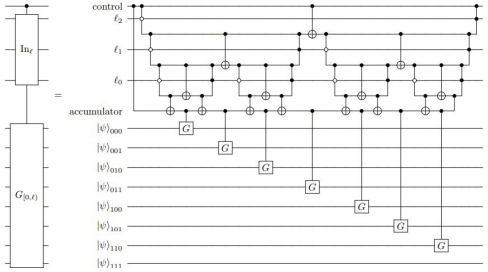
control
accumulator

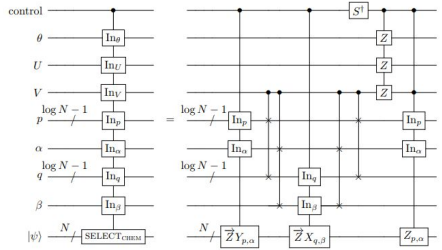
control
SELECT_CHEM

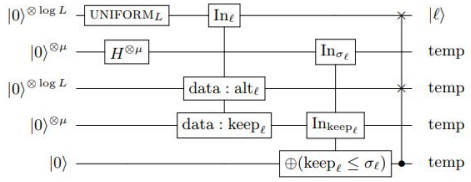
UNIFORM_L
temp
temp
temp
temp
temp

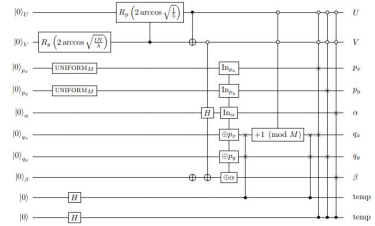

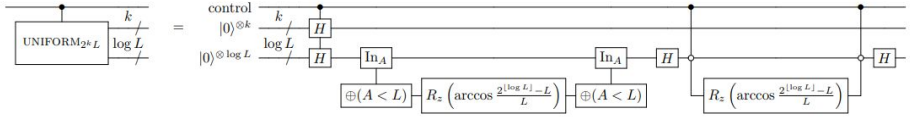
control

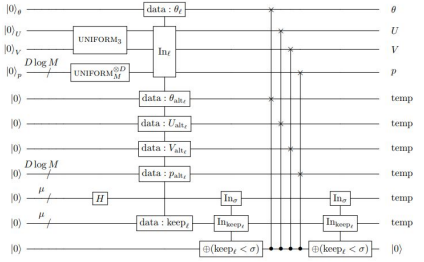
UNIFORM_3
temp
temp
temp
temp
temp
temp
temp

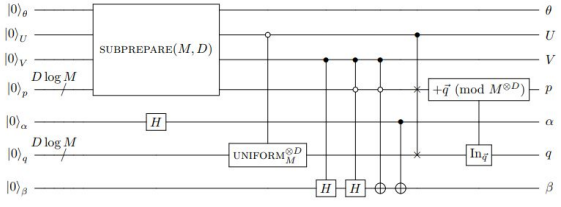
SUBPREPARE(M, D)
UNIFORM_M^⊗D

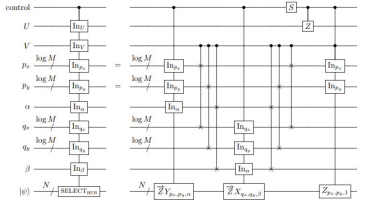
control
SELECT_UCCSD

Compiling to surface code fault-tolerant gates

Google AI Quantum

Topological diagrams indicate stabilizer measurements

T gates require state distillation, very slow! $|\mathrm{T}\rangle \equiv \mathrm{T}\,|+\rangle$

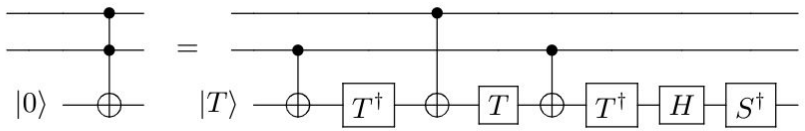

Braided T factory in our context is ~150,000 qubits

There have been recent improvements

Lattice surgery CCZ 3X fewer qubits, 4X faster

About one week to distill 10^9 CCZ states in serial

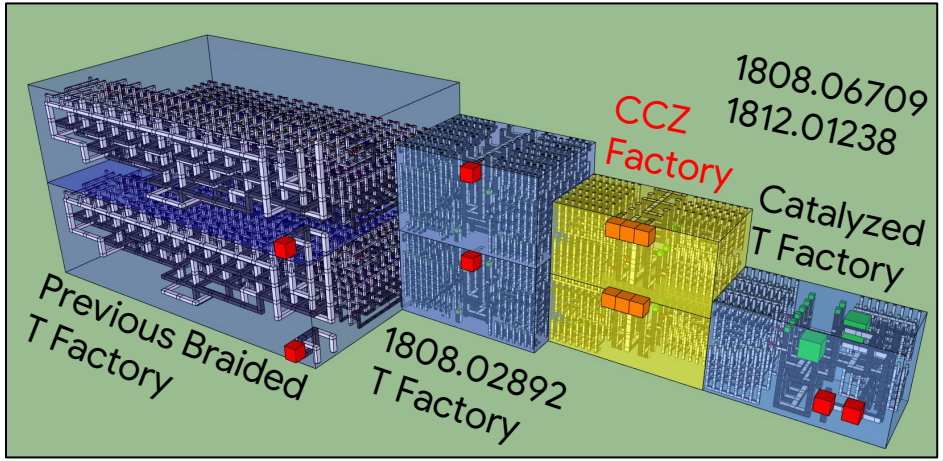

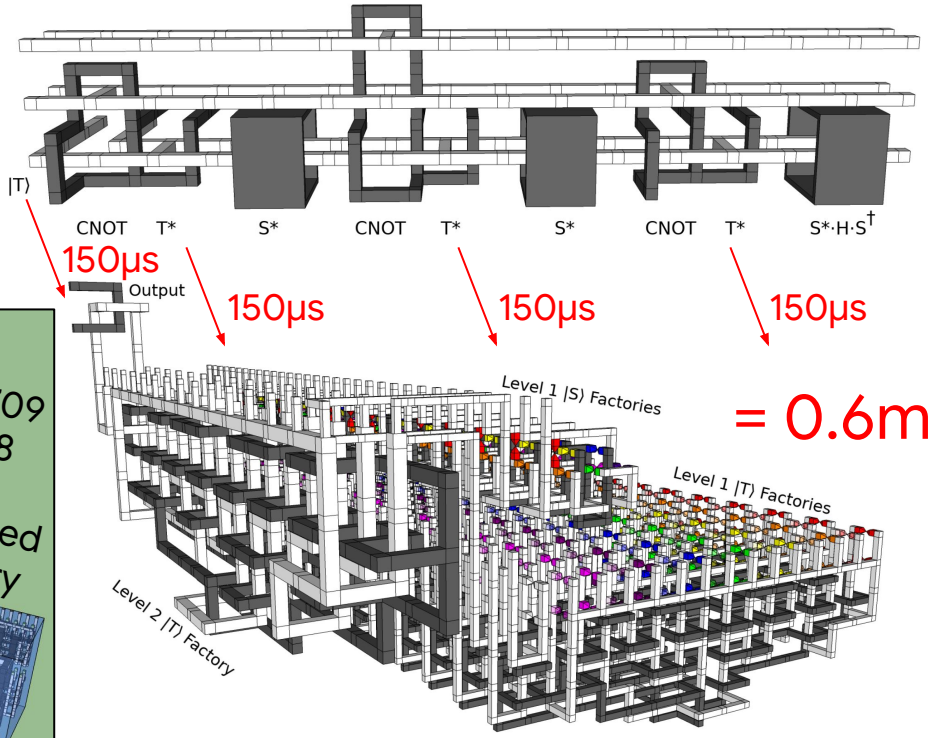

Resource estimates - Phys. Rev. X 8, 041015

Google AI Quantum

Compiled all bottlenecks to surface code gates by hand and with automated system

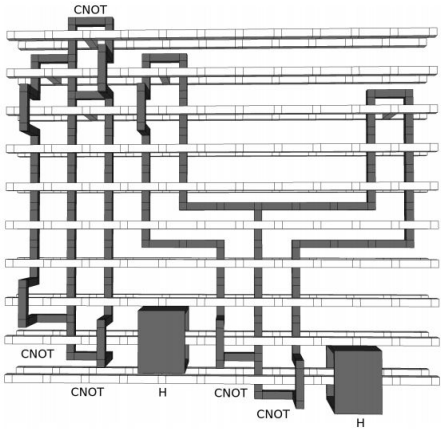

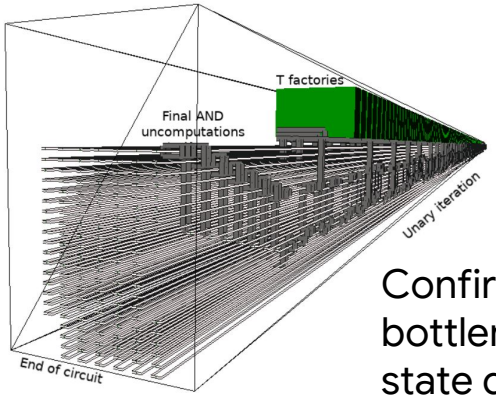

Confirmed bottleneck is state distillation

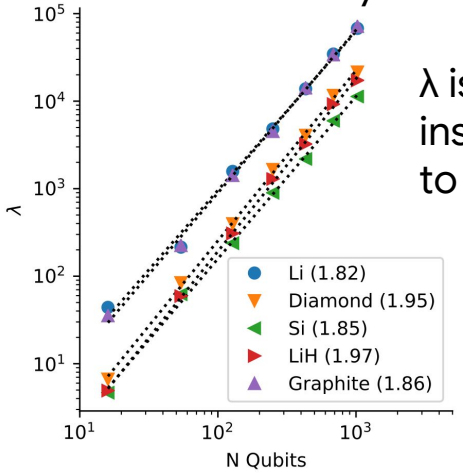

λ is fairly insensitive to instance

Estimated resources for crystalline solids (e.g. diamond, graphite, LiH metal, silicon, jellium)

Estimated resources for Hubbard with ϵ = t / 100, u = 4 t, λ = 4 N t

Can error-correct real applications with < 1M superconducting qubits!

problem		physical qubits		run time (hours)	
system	spin-orbitals N	$p = 10^{-3}$	$p = 10^{-4}$	$p = 10^{-3}$	$p = 10^{-4}$
Hubbard model	128	7.0×10^5	2.2×10^5	3.8	2.1
Hubbard model	200	1.1×10^6	3.0×10^5	10	5.3
Hubbard model	800	4.7×10^6	1.2×10^6	170	9.3
Electronic structure	128	8.0×10^5	2.7×10^5	2.5	1.4
Electronic structure	250	1.5×10^6	4.0×10^5	15	7.5
Electronic structure	1024	6.7×10^6	1.6×10^6	680	350

Exploiting low rankness of Coulomb operator

Google AI Quantum

Motta *et al.* arXiv:1808.02625

$$V = \sum_{pqrs} V_{pqrs} a_p^\dagger a_q a_r^\dagger a_s \qquad W_{pq|rs} = V_{pqrs} \qquad W \in \mathbb{R}^{N^2 \times N^2} \qquad W g_\ell = \lambda_\ell g_\ell$$

Well known in chemistry that W is low rank so that L = O(N)

$$V = \sum_{\ell=0}^{L-1} \lambda_\ell \left(\sum_{pq} [g_\ell]_{pq} a_p^\dagger a_q \right)^2 \qquad R_\ell \left(\sum_{pq} [g_\ell]_{pq} a_p^\dagger a_q \right) R_\ell^\dagger = \sum_p f_{\ell p} n_p$$

g matrix is of rank O(log N) when N grows towards thermodynamic limit

$$V = \sum_{\ell=0}^{L-1} \lambda_\ell \left(R_\ell \left(\sum_p f_{\ell p} n_p \right) R_\ell^\dagger \right)^2 = \sum_{\ell=0}^{L-1} \lambda_\ell R_\ell \left(\sum_p f_{\ell p} n_p \right) R_\ell^\dagger R_\ell \left(\sum_p f_{\ell p} n_p \right) R_\ell^\dagger = \sum_{\ell=0}^{L-1} \lambda_\ell R_\ell \left(\sum_{pq} f_{\ell p} f_{\ell q} n_p n_q \right) R_\ell^\dagger$$

This is sum of L diagonal Hamiltonians, each in different single-particle basis

$$\prod_{\ell=0}^{L-1} R_\ell \exp\left(-i\lambda_\ell \sum_{pq} f_{\ell p} f_{\ell q} n_p n_q \right) R_\ell^\dagger$$

Simulating diagonal Hamiltonian and R(u) as before gives O(L N) = O(N^2) depth circuit

But is it really low rank?

Google AI Quantum

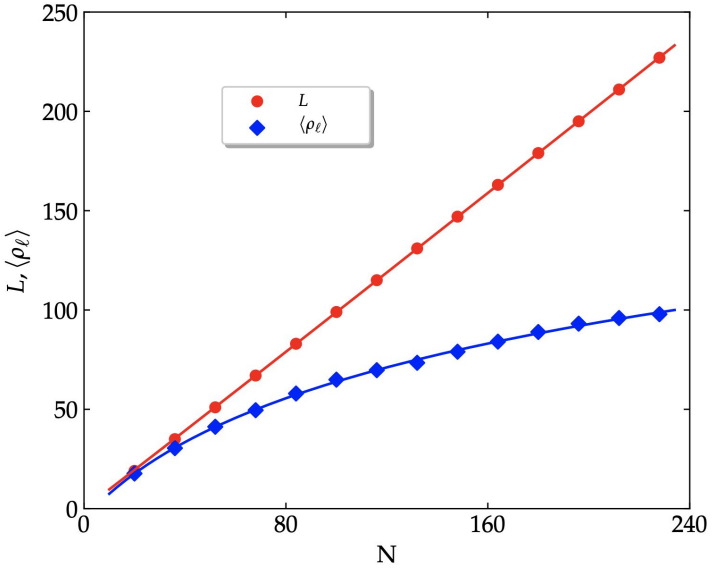

Rank of both truncations for a hydrogen ring in STO-6G basis

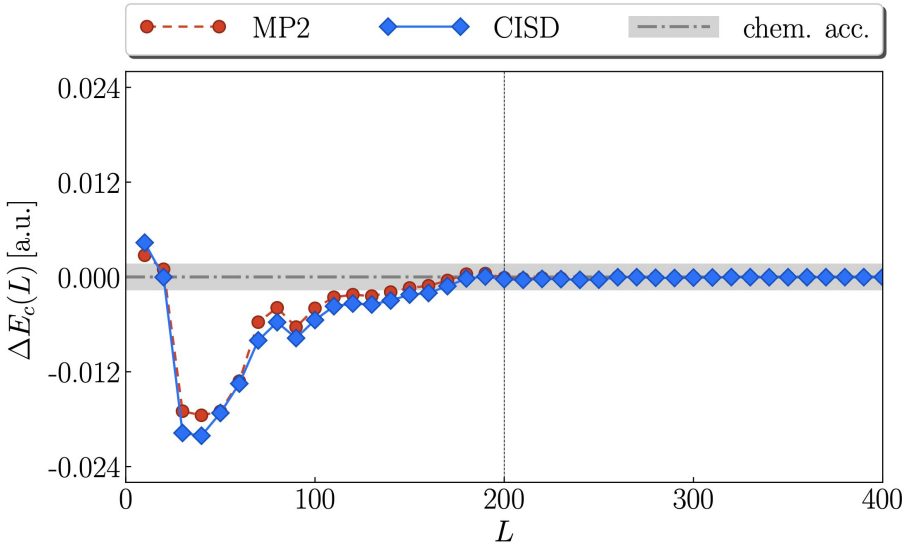

Determination of L for 152 qubit active space of FeMoco molecule

Reduces Trotter step cost for arbitrary basis simulation from $O(N^4)$ to $O(N^3)$ or $O(N^2 \log N)$

Hamiltonians that are both compact and sparse

Google AI Quantum

$$V = \sum_{pqrs} V_{pqrs} a_p^\dagger a_q a_r^\dagger a_s \qquad W_{pq|rs} = V_{pqrs} \qquad W \in \mathbb{R}^{N^2 \times N^2}$$

Molecular orbitals

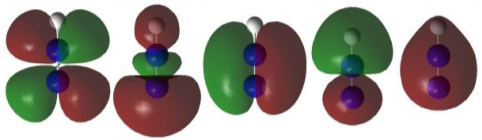

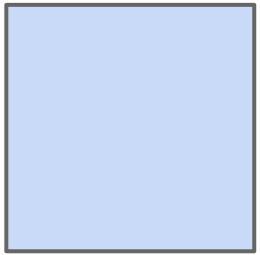

W is dense
Few qubits
$O(N^4)$ terms

Sincs

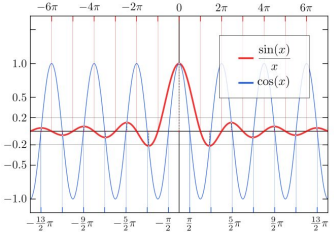

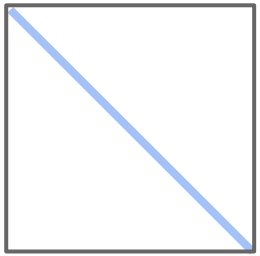

W is diagonal
Many qubits
$O(N^2)$ terms

Gausslets

JCP 147, 244102 (2017)

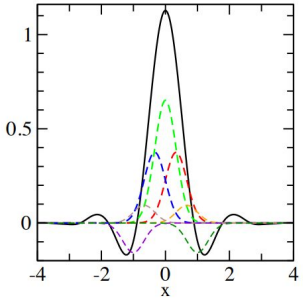

W is diagonal
Fewer qubits
$O(N^2)$ terms

Challenging to construct

Adaptive local basis by discontinuous Galerkin

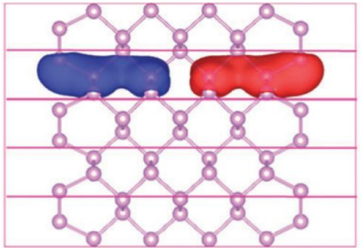

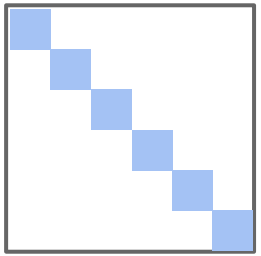

W is sparse
Fewer qubits
$O(b^2 N^2)$ terms

Sublinear scaling with first-quantization

arXiv:1807.09802

Google AI Quantum

η registers store "which plane wave" in log N bits wavefunction is explicitly antisymmetric

$$|\psi\rangle = \sum_{p_\ell \in G} \alpha_{p_1 \cdots p_\eta} |p_1 \cdots p_i \cdots p_j \cdots p_\eta\rangle = -\sum_{p_\ell \in G} \alpha_{p_1 \cdots p_\eta} |p_1 \cdots p_j \cdots p_i \cdots p_\eta\rangle$$

Hamiltonian is sparse, but not local $||T||_1 = O(\eta^{1/3} N^{2/3})$ $||V||_1 = O(\eta^{5/3} N^{1/3})$

$$T = \sum_{j=1}^{\eta} \sum_{p,q,\nu \in G} \frac{\|k_p\|^2}{2} |p\rangle\langle p|_j \qquad V = \frac{2\pi}{\Omega} \sum_{i,j=1}^{\eta} \sum_{p,q,\nu \in G} \frac{1}{\|k_\nu\|^2} |p+\nu\rangle\langle p|_i \cdot |q-\nu\rangle\langle q|_j$$

Using methods from (1805.00675), simulate in frame of T, obtain complexity $O(\eta^{8/3} N^{1/3})$

Consider $\eta=50$, $N=10^6$ then $\eta^{8/3} N^{1/3} \sim$ 3 million gates and $\eta \log N \sim$ one thousand qubits

Open question: can we exponentially suppress basis error? i.e. achieve O(polylog(N))

Back-of-the-envelop FeMoco

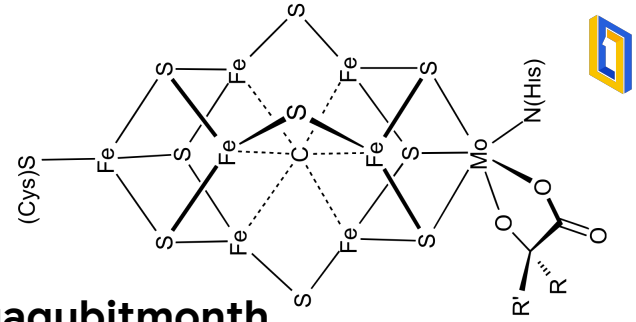

PNAS, 114: 7555-7560 (2017): ~10^{14} T gates > 1 gigaqubitmonth

ΔE = 0.0016 Ha N = 108 η = 54 L = 200 $\lambda \sim 10^4$ Ha M (trotter steps) = ~ 10^6

Randomized compiled Trotter (1811.08017) > $\lambda^2 / \Delta E^2$ > 10^{14} T gates

Trotter with low rank technique (1808.02625) ~ 10^6 M > 10^{12} T gates

Qubitization with low rank factorization > N $L^{1/2}$ $\lambda / \Delta E \sim 10^{11}$ T gates > 10^7 qubitdays

Interaction picture approach (1807.09802) > $\eta^{8/3}$ $N^{1/3}$ / $\Delta E \sim 10^9$ T gates > 10^5 qubitdays

Gausslet and discontinuous Galerkin methods mixed with some of the above: ???

Can we do molecules like FeMoco with 1M physical qubits at 10^{-3} error rates?

Google AI Quantum

Thank you!

Craig Gidney (Google)

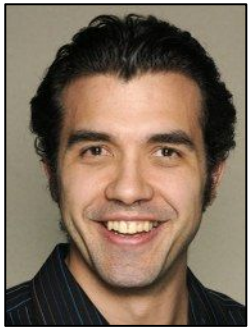

Dominic Berry (Macquarrie)

Also, does anybody want to go surfing this afternoon?

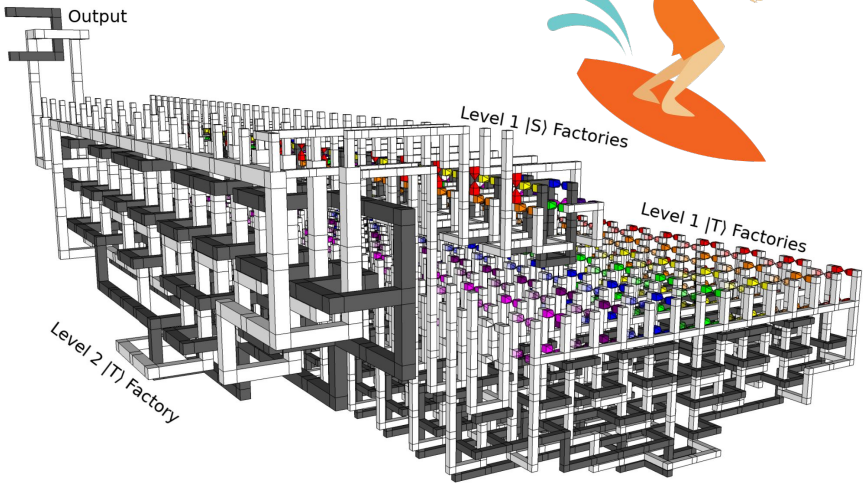

Jarrod McClean (Google)
Austin Fowler (Google)
Garnet Chan (Caltech)
Mario Motta (Caltech)
Nathan Wiebe (Microsoft)
Steven White (UC Irvine)
Lin Lin (UC Berkeley)

OpenFermion

www.openfermion.org

arXiv:1710.07629

Google AI Quantum

OpenFermion is an Apache 2 open source project for quantum simulation:

- Generate Hamiltonians for arbitrary molecules and materials in arbitrary basis sets
- Automatically compiles quantum algorithms to circuits
- Google software standards; ~60K lines of code at 99.9% test coverage

OpenFermion is a community! Over 3 dozen contributors from 20 institutions

130 active (visible) forks and use in nearly all new papers suggests that entire field is using

Framework and platform agnostic

- Google Cirq, Microsoft LIQUID, IBM QISKit, Xanadu Strawberry, Rigetti Forest, etc.
- Runs on Linux, Mac, and Windows with optional Docker installation