

CAN ONE DO
QUANTUM
CHEMISTRY WITH A
SMALL QUANTUM
COMPUTER?

Bryan Clark, UIUC

with Dave Wecker, Bela Bauer, Matt Hastings, Matthias Troyer

Boston CCP 2014

The take-home message:

Suppose you have today's quantum algorithms but tomorrow's quantum computers.

Can you do highly accurate simulations of molecules which are twice as large as those that can be done today?

Answer: no

Interestingly, quantum computers are 'slow' for a different reason than classical computers.

We will need better algorithms!

Note: not the scaling limit

One day we will have quantum computers...

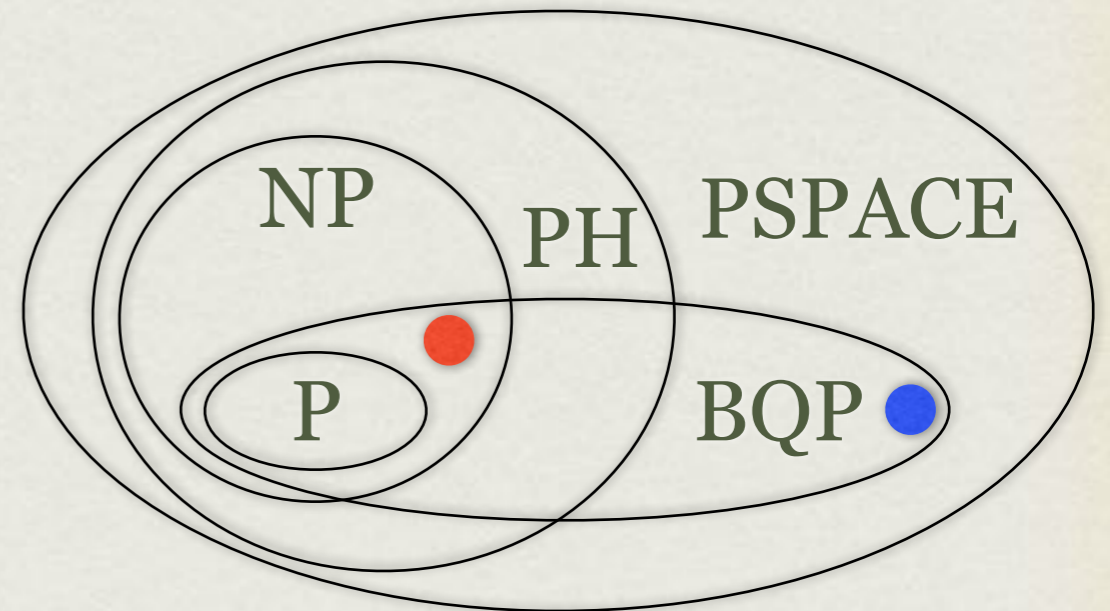


What will we do with them?

Quantum computers don't scale significantly better than classical computers for almost anything.

- factoring
- quantum simulation

But certainly they should be useful for simulating quantum mechanics...



Nature isn't classical and if you want to make a simulation of nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy.

--- Feynman "[Simulating Physics with Computers](#)", *International Journal of Theoretical Physics*, volume 21, 1982, p. 467-488, at p. 486 (final words)

Quantum Simulations

- Take a Hamiltonian H
- Compute properties (say energy) of ground state Ψ_0

We are willing to tolerate some small error.

For chemistry: significantly below room temperature.

For physics: smaller than other energy scales

$$H = \sum_{pq} t_{pq} c_p^\dagger c_q + \frac{1}{2} \sum_{pqrs} \underbrace{V_{pqrs}}_{\substack{\text{Calculated once per molecule} \\ N^4 \text{ numbers}}} c_p^\dagger c_q^\dagger c_r c_s$$

Sparse and structure-full....

What are classical computers good at?

Linear Algebra

Exact diagonalization

Quantum Chemistry (CCSD(T))


DMRG/Tensor Networks

Sampling

Quantum Monte Carlo

Linear Algebra

Just 'write' down $|\Psi_T\rangle$



Arbitrary states in the Hilbert space can't be represented by a small number of parameters.

Represent state by n 4-tensors
Polynomial parameters

Ground states aren't arbitrary. They are very special!

- sub-volume law
- simpler sign structure

Current quantum algorithms don't seem to use this fact.

Exponential parameters

Ground states



Equation of State Calculations by Fast Computing Machines

THE purpose of this paper is to describe a general method, suitable for fast electronic computing machines, of calculating the properties of any substance which may be considered as composed of interacting individual molecules. Classical statistics is assumed,

(Received March 6, 1953)


$$|\Psi\rangle = \alpha_0 |\Psi_0\rangle + e^{-\beta(E_1 - E_0)} \alpha_1 |\Psi_1\rangle + e^{-\beta(E_2 - E_0)} \alpha_2 |\Psi_2\rangle \quad e^{-\beta H}$$

Approach ground state exponentially fast!

Just sample.

Variance is exponential so need exponential number of samples

But depending on what system size you're doing, you might be ok...



Even for sign-problem highly-entangled problems there are often classical algorithms which achieve a 'good enough' approximation.

Metric of 'good enough': chemical accuracy

1 milliHartree (out of 100 Hartree)

To chemical accuracy:

N=50 spin orbitals - Lanczos

N=70 spin orbitals - Tensor Networks

N >> 100 spin orbitals - CCSD(T) on weakly correlated

Electrons ~ 54 - QMC Variants


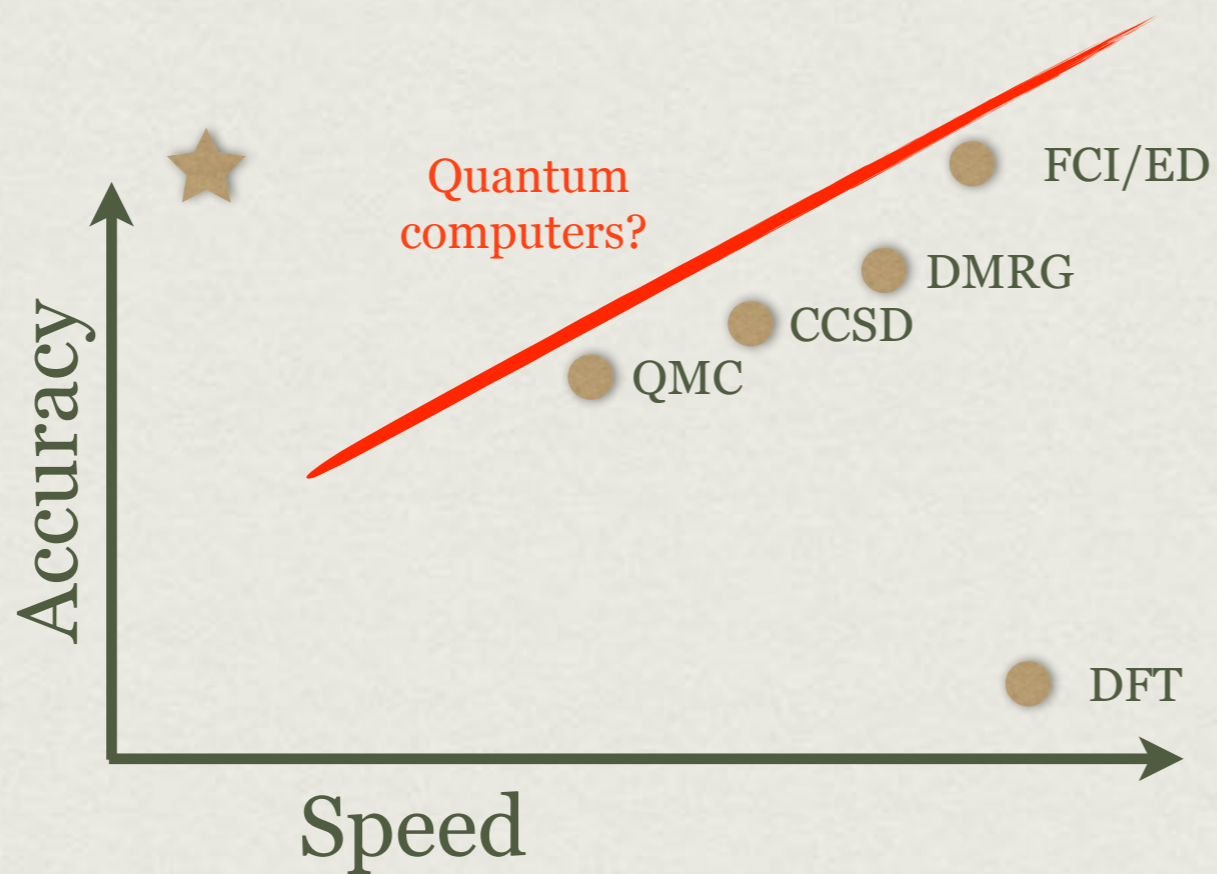
ED: 1 site per year

Tensor Networks: ~10 site per year

Aside: Any QC algorithm ~30-50 qubits simulatable
38 qubits -> 10 minutes per step



Why aren't classical computers good enough?



The three simple steps to quantum simulations on a quantum computer.

1. Build a quantum computer

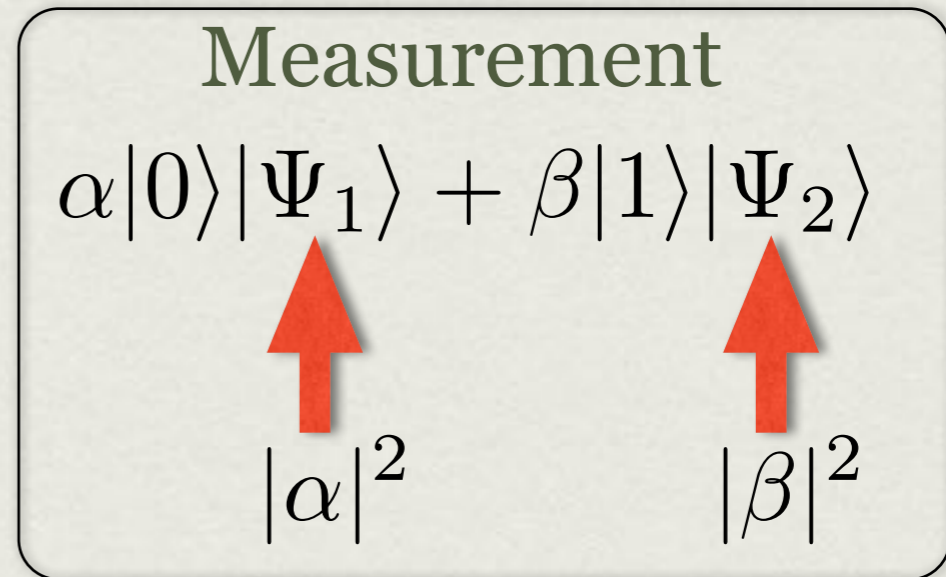


Error correction...


2. Develop quantum algorithms

3. Profit/Do great science

What are quantum computers good at?



We need to move our starting wavefunction to our final one.





The options:

1. Project from nearby
2. Adiabatically evolve

Project from nearby: $|\alpha_0|^2 \approx 1$

- Start close
- Do a lot of quantum work to get closer



$e^{-iT H}$ This is your computational bottleneck.

Phase Estimation

- Start close to ground state: $|\Psi_T\rangle = \sum_i \alpha_i |0\rangle |\Psi_i\rangle$
- Apply Phase Estimation: $\sum_i \alpha_i |E_i\rangle |\Psi_i\rangle$
- Measure E_i with probability α_i^2 getting



What T do we need?

Set by required accuracy: $T \approx 6000 E_h^{-1}$

Interesting note: What matters is absolute accuracy not relative accuracy.

Phase estimation is the core of factoring too.

3. Project from nearby: $|\alpha_0|^2 \approx 1$



$e^{-iT H}$ This is your computational bottleneck.

Phase Estimation

- Start close to ground state: $|\Psi_T\rangle = \sum_i \alpha_i |0\rangle |\Psi_i\rangle$
- Apply Phase Estimation: $\sum_i \alpha_i |E_i\rangle |\Psi_i\rangle$
- Measure E_i with probability α_i^2 getting

What T do we need?


Set by required accuracy: $T \approx 6000 E_h^{-1}$

Interesting note: What matters is absolute accuracy not relative accuracy.

Adiabatic Evolution

- Start in Ψ_i as ground state of H_i
- $\exp[-itH(t)]$

This is your computational bottleneck.



Easy to make

$$H_i |\Psi_i\rangle = E_i |\Psi_i\rangle$$

$$H_{\text{chemistry}} |\Psi_0\rangle = E_0 |\Psi_0\rangle$$


The answer

What T do we need?


Set by minimal gap

Quantum simulations on quantum computers relies essentially on one trick: Apply e^{-iTH} quickly.

Trotter Breakup



Quantum Walks



To get a ground state, we need to be able to do time evolution quickly.

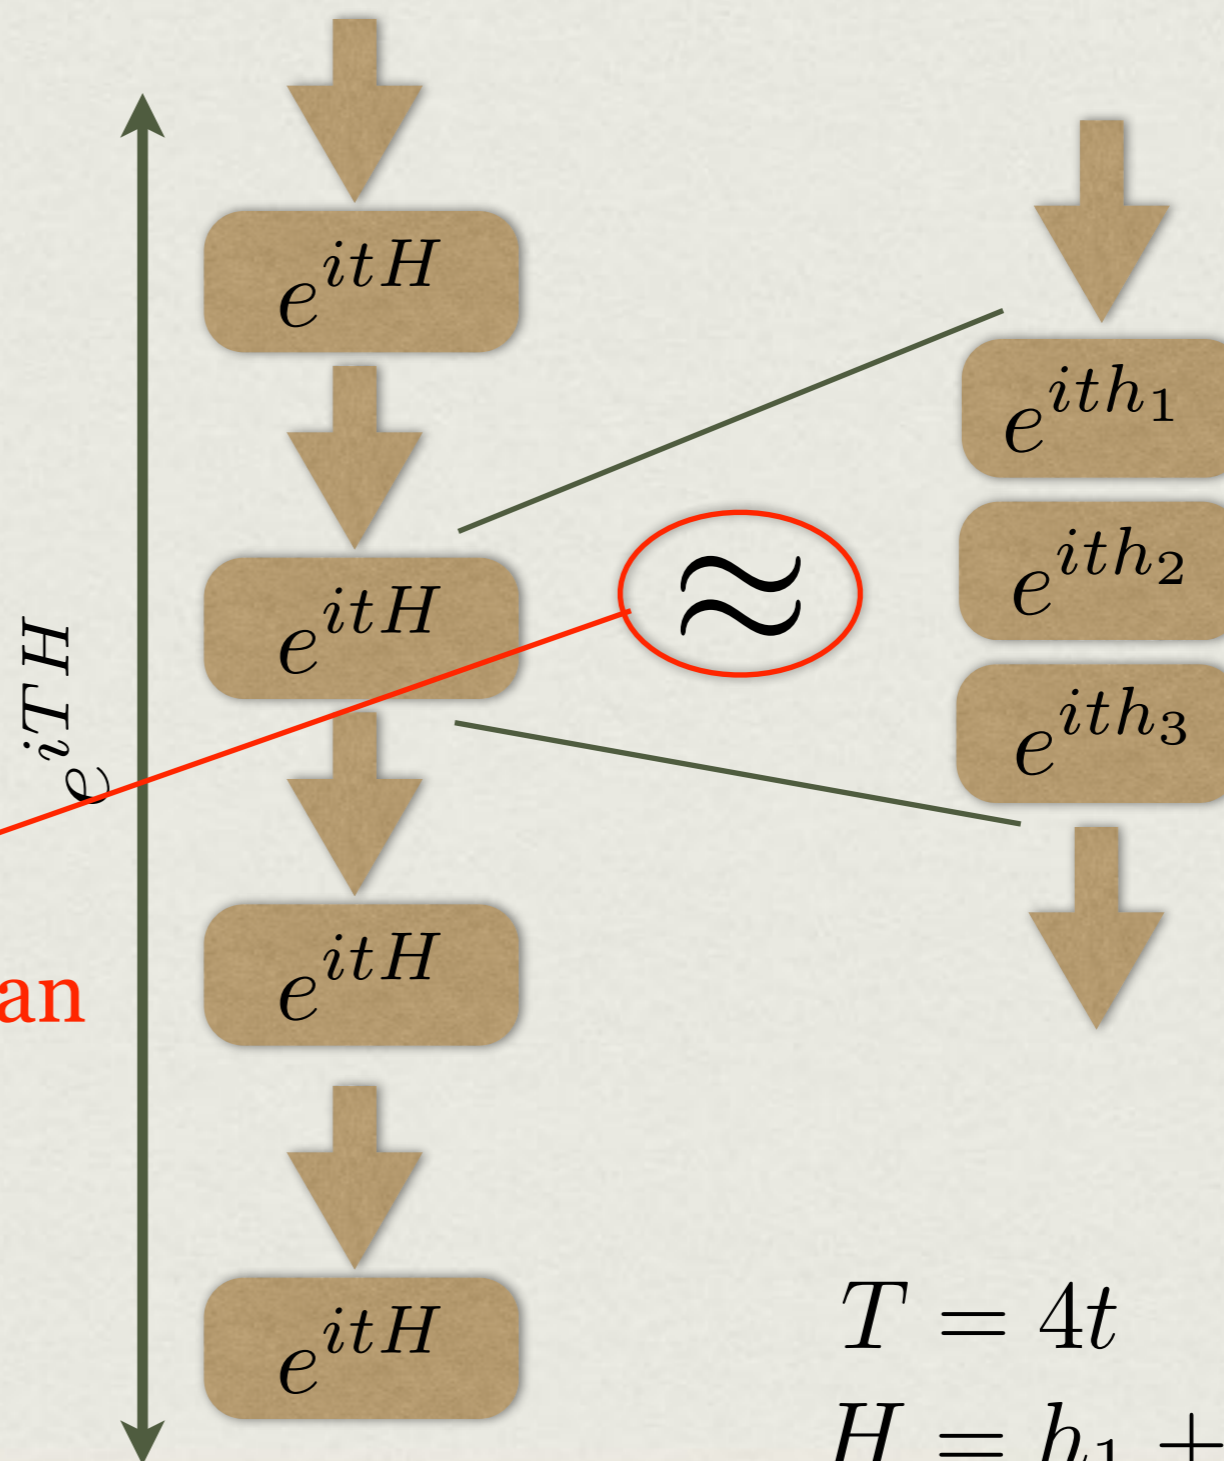
Trotter Breakup

What T ?

What t ?

How many h_i ?

How costly h_i ?




This approximation can be painful.

Only need a few coherent qubits.

$$T = 4t$$
$$H = h_1 + h_2 + h_3$$

Basis: Real Space $e^{-H} = e^{-K} e^{-V}$



What T? $T \approx 6000 E_h^{-1}$

What t? $t \approx 0.01$

How many h_i ? 2

How costly h_i ? N^2

1 million grid points \longleftrightarrow 100 points per dimension

6×10^{17} gates

Basis: Quantum chemistry

$$H = \sum_{pq} t_{pq} c_p^\dagger c_q + \frac{1}{2} \sum_{pqrs} V_{pqrs} c_p^\dagger c_q^\dagger c_r c_s$$

Calculated once per molecule
N⁴ numbers

Can localized orbitals pull this N⁴ down*

Yes, for really big systems.

Yes, for long linear systems (say of Carbon) by factor of ~5 for 10 carbons

Not likely for moderate sizes, non-linear systems.

* Exploiting locality in quantum computation for quantum chemistry: McClean, Babbush, Love, Guzik


Computing number of gates: $O(N^4 \times N) = O(N^5)$

terms in Hamiltonian

Jordan-Wigner strings for sign

Some of this back from parallelization
or smart tricks (Bravyi-Kitaev)

Matches empirically



Computing $1/\tau$

* Maybe not so bad: arXiv:1406.4920

Theory: # trotter steps for fixed time (for fixed trace norm distance)


$$O(\underline{m}^{1+1/\underline{2k}}) \quad \text{m: terms in Hamiltonian} \quad \text{For } k=1: m^{3/2}$$

k: trotter order

O(n) terms mutually commute. New theoretical bound:

$$\text{For } k=1: Km^{1/2} \sim m^{3/4}m^{1/2} \sim m^{1.25}$$

Empirically: # trotter steps for fixed time (for fixed energy error)



□ Computed with imaginary molecules

□ For $k=1$, $m^{1.08} - m^{1.27}$

terms in trotter series $m \sim N^4$

Scaling: $N^4 - N^5$

Putting it together ...

(Gates per trotter step) x (Steps per fixed time) x (time)

$$N_g$$

$$1/\tau$$

$$T$$

$$N^4 - N^5$$

$$N^2 - N^4$$


$$6000E_h^{-1}$$

$$6000N^6 - 6000N^9$$

Water (STO-3G): 10^{10} serial gates (441 x 441 matrix - 14 s.o.)
(by counting)

Parallelization saves factor of 20

Quantum Walks*



$$O(d^{2/3} ((\log \log d)t \|H\|)^{4/3})$$

$$O(N^{8/3} ((\log \log N^4) \|H\|)^{4/3})$$

$$\|H\| \rightarrow O(N) \quad (\text{operator norm})$$

$$O(N^{8/3} N^{4/3} ((\log \log N^4))^{4/3}) = O(N^4) \quad \text{oracle queries}$$



$$O(d\Lambda_{max}t)$$

$$O(N^4) \quad \text{oracle queries}$$

This looks quite promising.....

until you think about the 'oracle'...

$$\text{Oracle: } U_f |x, i\rangle |0\rangle = |\phi_{x,i}\rangle |y_i, H_{x,y_i}\rangle$$



N^4 numbers in the box
space x time = N^4

Want only N qubits



time: $N^3 \times N$ (Jordan-Wigner) = N^4

With quantum walks: N^8 time.

You can trade-off time for space here.

until you think about the 'oracle'...

$$\text{Oracle: } U_f |x, i\rangle |0\rangle = |\phi_{x,i}\rangle |y_i, H_{x,y_i}\rangle$$



N^3

Put games in the box

time: At least N^3 but probably more

With quantum walks: At least N^7 time.

Where are we?

Want ground state

Adiabatic

Phase estimation

Need time-evolution

Quantum walks

Trotter Decomp

Quantum chemistry

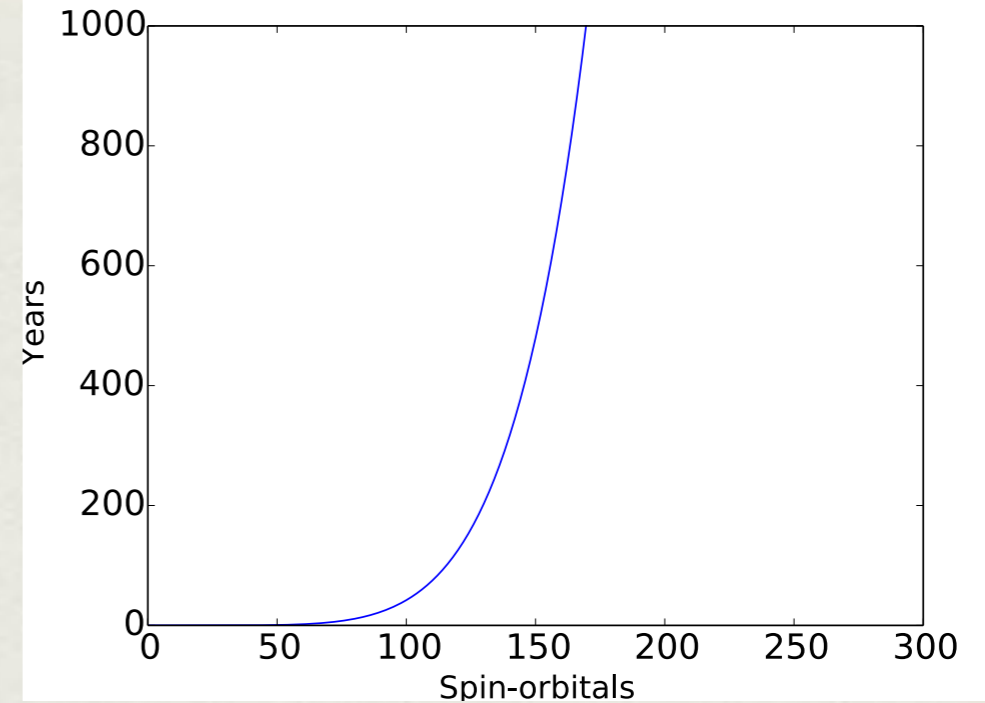
Real space

The runtime...

If you believe in a classical Moore's law, if you want to do 125 spin-orbitals you would be better to wait for a better classical computer

A hard problem


Chromium dimer with a really good basis (Dunning cc-pV5Z) you need 300 spin-orbitals



Logical qubit time: 1 microsecond

The way forward:

- Faster QC: Gigahertz logical qubit time. (fantasy?)
- Much better error correction
- Better algorithms: Even in the last few months, the community gained a few orders of magnitude.
- We already have massively parallel classical computers. Can this help us?



A comment about parallelization....

Computers aren't getting much faster.

But there are many cores (GPU's, clusters, etc.)

Without coherence...



No network



With coherence...



Network



Keeping qubits
coherent is hard

Comparing quantum and classical results....

Quantum

Polynomial

'Bad' time-step error

Cost in absolute not relative accuracy

Costly error correction

No checkpointing

Hard to parallelize

Classical

Exponential


No time step error

Cost often in relative accuracy

No error correction

Checkpointing

Easy to parallelize




Maybe what we need is a chimera....



Suppose you have, a massively parallel classical computer and as smaller quantum computer. What would you do?



A quantum computer...



hands samples to a classical computer.



A (slightly imaginary) limit:

megahertz quantum computer

gigahertz classical computer with arbitrary number of processors.

This is much better then running the quantum computer all the way.

When there are not an arbitrary number of processors, how this trades off depends on the details.



Conclusions

- ❑ Today's algorithms, tomorrow's quantum computer:
 - ❑ 150 spin orbitals -> order of a millenium
- ❑ Classical computers get hit by an exponential variance.
- ❑ Quantum computers get hit by bad time step errors and having to act with a lot of gates.
- ❑ Algorithms will get better
- ❑ Maybe using quantum computers to see a classical computer is a good start.

What's the runtime?

	Logical Qubit time	Computation time
'Fast'	1 micro-second	3000 years
'Fantasy'	1 ns	3 years

Plus...no checkpointing!

This is (no matter how good hardware gets) unrealistic.

- Can we do better?
- Is there anything better already in the literature?



I want to simulate quantum systems.

I have a big classical computer.



Soon we'll have ~100 qubit quantum computer that can run for a day (hey, I'm optimistic)

When should I use my classical computer, my small quantum computer, and when is it hopeless?



Algorithms


Quantum Algorithms

Exact classical algorithms

Classically such algorithms exist that don't depend on your starting wave-function. Because everything is unitary that can't be the case here.

Approximate classical
algorithms

Real Space $e^{-H} = e^{-K} e^{-V}$



What T? $T \approx 6000 E_h^{-1}$

What t? $t \approx 0.01$

How many h_i ? **2**

How costly h_i ? $N \log(N)$

The quantum chemistry Hamiltonian:


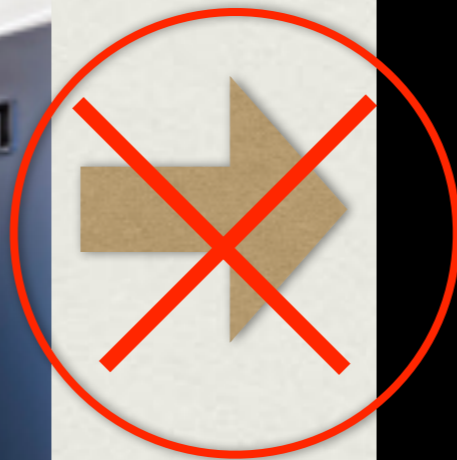


$$H = \sum_{pq} t_{pq} c_p^\dagger c_q + \frac{1}{2} \sum_{pqrs} V_{pqrs} c_p^\dagger c_q^\dagger c_r c_s$$

Calculated once per molecule
 N^4 numbers

- Rows and columns indexed by N bit binary numbers with N_e 1's
- $|H_{ab}|$ equal to $f(i,j,k,l)$ where $a \rightarrow b$ by turning off (i,j) and on (k,l)
- $\text{Sign}[H_{ab}]$ depends on parity of electrons between (i,j,k,l)


Only N^4 non-zeros per row

Sparse and structure-full....




Q. Why can't classical computers simulate quantum system?
(i.e. what special about quantum mechanics)


Zero point motion and l_2 norm



Interference



Statistics



We are weighing classical computers against (small) quantum computers. What can classical computers do?

'Lore': Sign-free ground states

Use quantum Monte Carlo....

This lore relies on fast mixing times of Markov chains (for path integral Monte Carlo) or a well controlled population (diffusion Monte Carlo).

* 'known' to diverge exponentially for typical systems

* Energy 'straightforward'; sampling $|\Psi|^2$ less-so

* No evidence of long mixing times if you start in correct phase

Low entanglement ground states


Use tensor networks....

Questions about finding low-entanglement ansatz

Harder in molecules because tensor networks rely on locality (except for CGTNS*)

* Bauer, et al






The quantum algorithm for getting ground state (energies)?

Two generic approaches

Adiabatic

- Start in Ψ_i as ground state of H_i
- $\exp[-itH(t)]$



Phase Estimation

- Start close to ground state: $|\Psi_T\rangle = \sum_i \alpha_i |0\rangle |\Psi_i\rangle$
- Apply Phase Estimation: $\sum_i \alpha_i |E_i\rangle |\Psi_i\rangle$
- Measure E_i with probability α_i^2 getting Ψ_i


Let's just worry about this for the moment...

Quantum Phase Estimation

How quickly can this be done?

Algorithms

- Trotter Decomposition
- Sparse Hamiltonian Problem
 - Quantum Walks
 - Trotter



e^{-iTH} This is your computational bottleneck.

What T do we need?

Set by required accuracy: $T \approx 6000E_h^{-1}$


Interesting note: What matters is absolute accuracy not relative accuracy.

Quantum Phase Estimation

How quickly can this be done?

Algorithms

- Trotter Decomposition
- Sparse Hamiltonian Problem
 - Quantum Walks
 - Trotter



$e^{-iT H}$ This is your computational bottleneck.

What T do we need?

Set by required accuracy: $T \approx 6000 E_h^{-1}$

Interesting note: What matters is absolute accuracy not relative accuracy.

Trotter

$$\exp[-itH]^{T/t}$$

$$\exp \left[-it \sum_{pqrs} V_{pqrs} c_p^\dagger c_q^\dagger c_r c_s \right]^{T/t}$$

\approx

of trotter steps

$$\left(\prod_{pqrs} \exp \left[-it V_{pqrs} c_p^\dagger c_q^\dagger c_r c_s \right] \right)^{T/t}$$

Cost per term

How many terms?

Parallel Circuit	Global R_z	H, Y, Y^\dagger	CNOT	CR_z	BSM	Total
H_{pp}				1		1
H_{pq}		8	2	4	4	18
H_{pqqp}	1		2	3		1+5
H_{pqqr}		4	8	4	4	24
H_{pqrs}		$8 \cdot 2$	$8 \cdot 2$	$8 \cdot 1$	$8 \cdot 2$	$8 \cdot 7$


Computing number of gates: $O(N^4 \times N) = O(N^5)$

terms in Hamiltonian

Jordan-Wigner strings for sign

Some of this back from parallelization

Matches empirically



Computing $1/\tau$

Theory: # trotter steps for fixed time (for fixed trace norm distance)


$$O(\underline{m}^{1+1/\underline{2k}}) \quad \text{m: terms in Hamiltonian} \quad \text{For k=1: } m^{3/2}$$

k: trotter order

O(n) terms mutually commute. New theoretical bound:

$$\text{For k=1: } Km^{1/2} \sim m^{3/4}m^{1/2} \sim m^{1.25}$$

Empirically: # trotter steps for fixed time (for fixed energy error)



□ Computed with imaginary molecules

□ For k=1, $m^{1.08} - m^{1.27}$

terms in trotter series $m \sim N^4$

Scaling: $N^4 - N^5$

Putting it together ...

(Gates per trotter step) x (Steps per fixed time) x (time)

$$N_g$$

$$1/\tau$$

$$T$$

$$N^5$$

$$N^4$$

$$6000 E_h^{-1}$$

$$6000 N^9$$

Water (STO-3G): 10^{10} serial gates (441 x 441 matrix - 14 s.o.)
(by counting)

Fe₂S₂ (STO-3G): 10^{18} serial gates (112 s.o.)
(by extrapolation)

75 years of quantum
Moore's law

Parallelization saves factor of 20

With 100 qubits, can never save more than factor of 100

In parallel: $6000 N^8$

Localized orbitals



(Gates per trotter step) x (Steps per fixed time) = Total

$$N_g \quad 1/\tau$$

quartic: N^5 N^4 N^9

quadratic: N^3 $N^{3/2}$ $N^{4.5}$


10,000 localized orbitals <---> 100 delocalized orbitals

Other Approaches

Different breakup

$$e^{-H} = e^{-h_1} e^{-h_2} e^{-h_3} \dots e^{-h_n}$$

Time step extrapolation:




Quantum Phase Estimation

How quickly can this be done?

Algorithms

- Trotter Decomposition
- Sparse Hamiltonian Problem
 - Trotter
 - Quantum Walks



e^{-iTH} This is your computational bottleneck.

What T do we need?

Set by required accuracy: $T \approx 6000E_h^{-1}$

Interesting note: What matters is absolute accuracy not relative accuracy.

We also need to evaluate $1/\tau$ and number of gates.

Sparse Hamiltonian Problem

Given an oracle to elements of $H = \sum_{j=1}^m H_j$, compute $\exp[-iT H]$
d non-zeros per row

$$\text{Oracle: } U_f |x, i\rangle |0\rangle = |\phi_{x,i}\rangle |y_i, H_{x,y_i}\rangle$$

Quantum chemistry Hamiltonian: $d = N^4$


Two 'current' winners:

D. W. Berry, R. Cleve, and R. D. Somma, Preprint (2013), [arXiv:1308.5424](https://arxiv.org/abs/1308.5424).

'Trotter Approach'

D. W. Berry and A. M. Childs, Quantum Information & Computation **12**, 29 (2012).

'Quantum Walks'



Trotter Steps*

$$O(d^2 T \log^3(Td))$$

$$O(N^8 T \log^3(TN^4)) \quad \text{oracle queries.}$$


Not much better than our current results.

Crossover $N \lesssim 100$

* D. W. Berry, R. Cleve, and R. D. Somma, Preprint (2013), arXiv:1308.5424.

What we've ignored....


- ❑ Error correction
~ factor of 100



- ❑ Adiabatically evolving, measurements, ...

Similar difficulties

Speed depends on gap



Simulating quantum mechanical ground states:

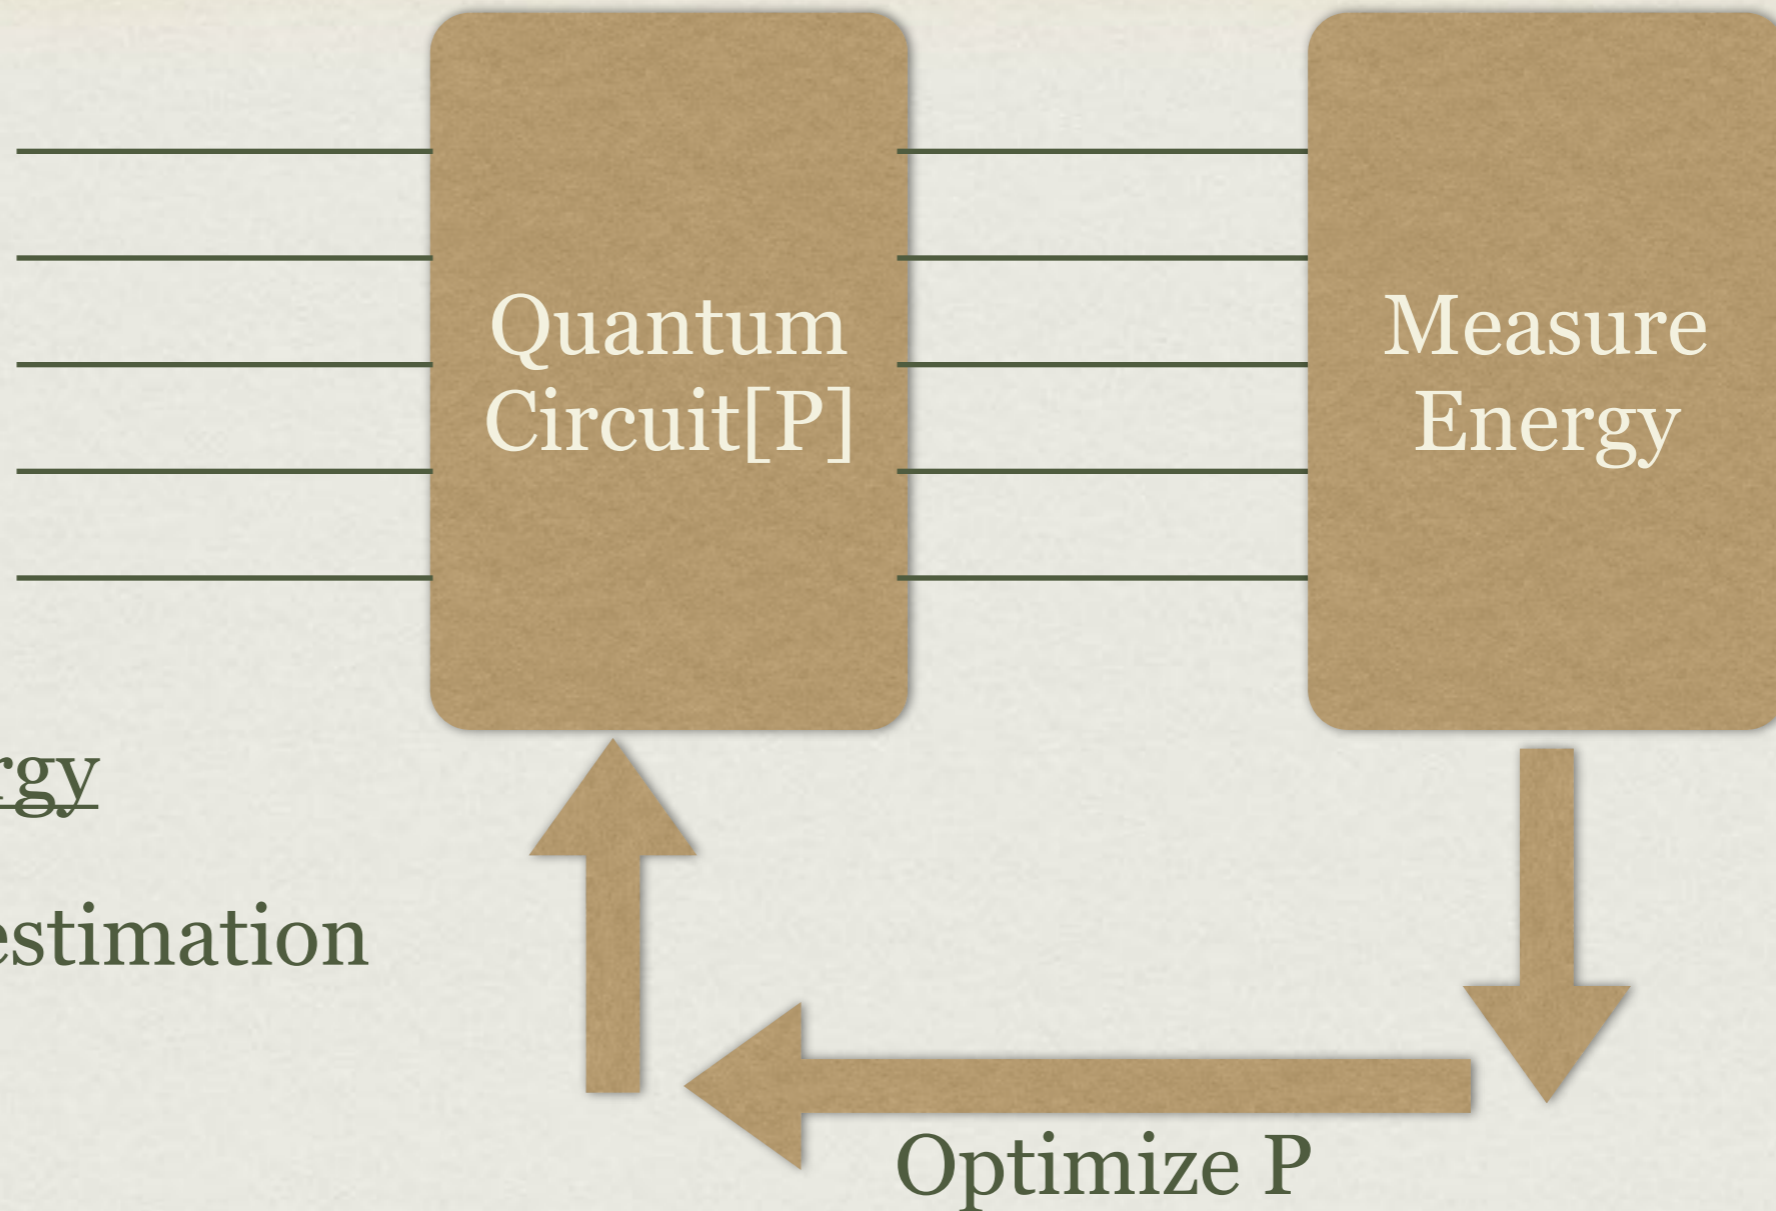
- Hermitian Matrix H
- has some ground state $|\Psi_0\rangle$
- with some ground state energy E_0
- which we want to approximate to a (chemical) accuracy.

$$e^{iT H} |\Psi_i\rangle = |\Psi[T]\rangle$$

$$e^{-T H} |\Psi_i\rangle = |\Psi_0\rangle \text{ dissipative.}$$

Beyond Phase Estimation

Variational*



Measure energy

Phase estimation

Collapse

$$N^4 \text{ terms} \times N^2 \text{ scaling of Monte-Carlo Error} = N^8$$


+

horrible non-linear optimization

* Apsuru Guzik

Real Space

$$e^{-H} = e^{-K} e^{-V}$$



(Gates per trotter step) x (Steps per fixed time) = Total

$$N_g$$

$$1/\tau$$

quartic:

$$N^5$$

$$N^4$$

$$N^9$$

$$N^2$$

$$2$$

$$N^2$$

1 million grid points <---> 100 delocalized orbitals

100 points per dimension

Other problems: Antisymmetrization, etc.

Adiabatic*


- $O(N^4)$ terms $> 10^8$
- $O(N^3)$ terms coupled to single qubit $> 10^8$
- 4 orders of magnitude range in accuracy
- Error correction?

one must resort to the indignity of numerical simulations.


* Aspuru Guzik

Summary

- ❑ Quantum mechanics in nature is full of structure ...
 - ❑ classical computers are powerful ...
 - ❑ and we don't need arbitrarily low error...
 - ❑ This sets a high bar for quantum computation.
- ❑ Quantum chemistry on quantum computers scale around N^8
 - ❑ and it looks like we need 10^{18} gates to do something useful.
- ❑ Classical approximate algorithms scale better than exact quantum algorithms. When the approximation is good enough, go classical.



1. Quantum walks



2. Adiabatic Evolution

3. 'Measure' from nearby

GARBAGE