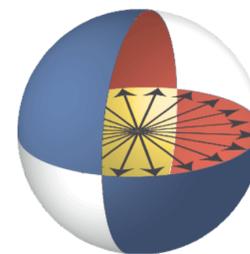


Simulating quantum mechanics with quantum computers

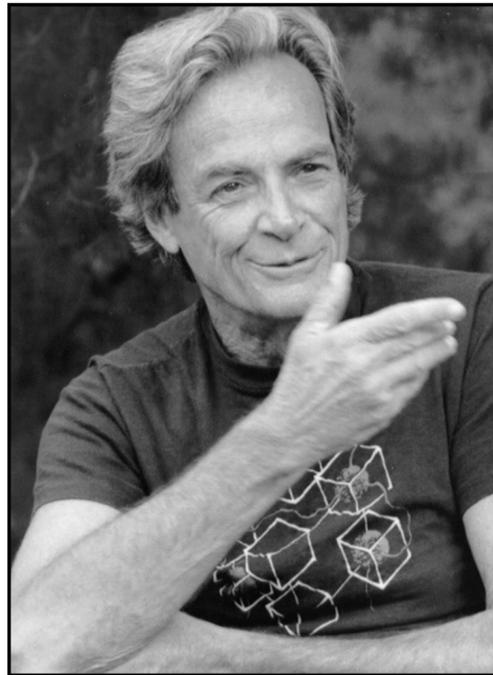
Andrew Childs
University of Maryland



UMIACS
University of Maryland
Institute for Advanced
Computer Studies



JOINT CENTER FOR
QUANTUM INFORMATION
AND COMPUTER SCIENCE



“... nature isn’t classical, dammit, 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.”

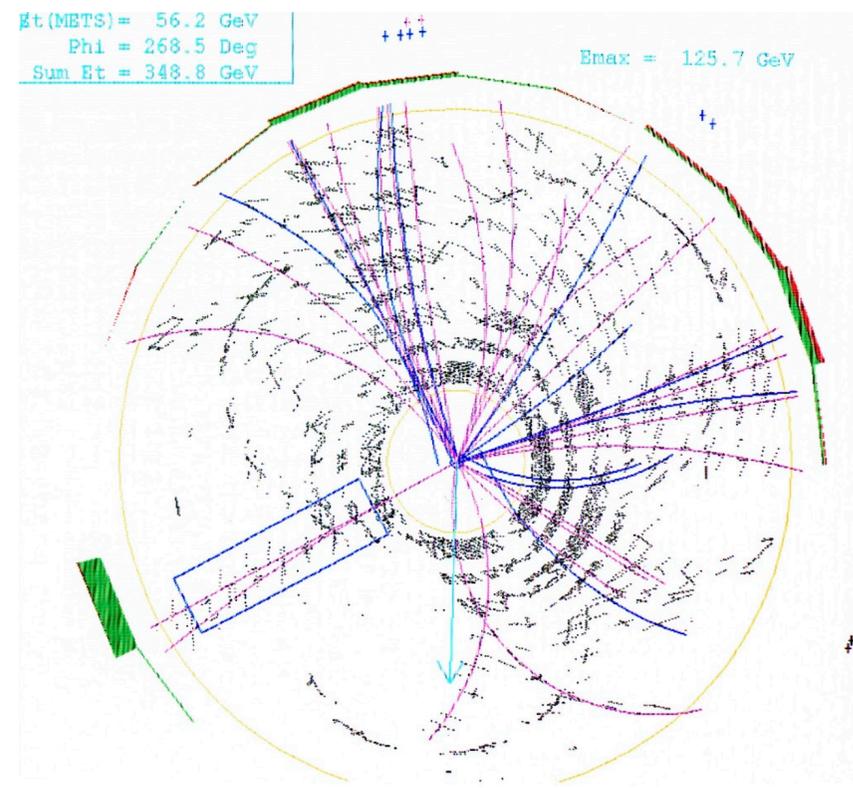
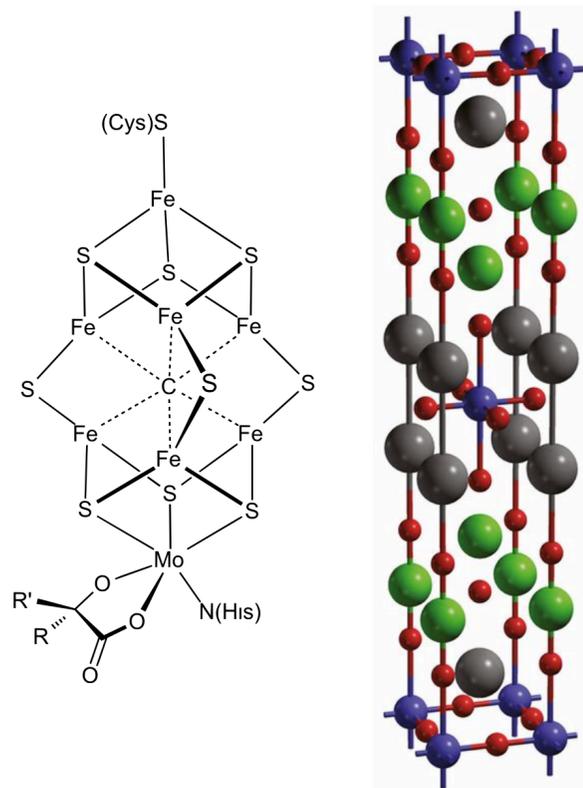
Richard Feynman

Simulating physics with computers (1981)

Why simulate quantum mechanics?

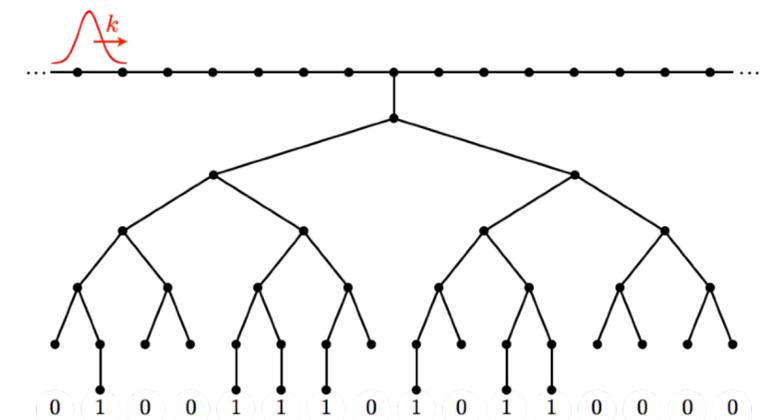
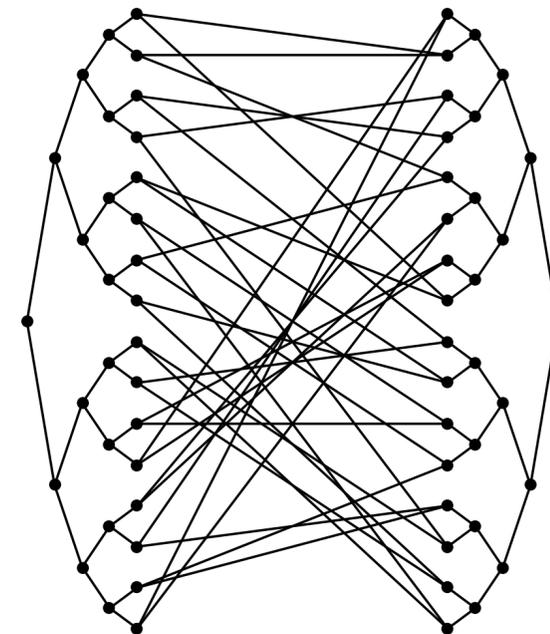
Computational chemistry/physics

- chemical reactions (e.g., nitrogen fixation)
- properties of materials
- condensed matter physics
- particle physics



Implementing quantum algorithms

- continuous-time quantum walk (e.g., for formula evaluation, search, ...)
- adiabatic quantum computation (e.g., for optimization or state generation)
- linear/differential equations



Simulating quantum mechanics with quantum computers

Algorithms

- Can we give an efficient algorithm?
- What is the best possible complexity as a function of various parameters?

Implementation

- What classically-hard simulations are easiest for a quantum computer?
- What simulation algorithm is best in practice for medium-scale problems?
- What optimizations can be applied to improve the implementation of algorithms?
- Are there resource tradeoffs (e.g., time vs. space)?
- Can we reliably do a classically-hard simulation without fault tolerance?
- How do the details of an experimental system (connectivity of qubits, timescales for different gates, etc.) interact with algorithmic issues?

Algorithms

Quantum dynamics

The dynamics of a quantum system are determined by its *Hamiltonian* H .

$$i \frac{d}{dt} |\psi(t)\rangle = H |\psi(t)\rangle \quad \Rightarrow \quad |\psi(t)\rangle = e^{-iHt} |\psi(0)\rangle$$

Quantum simulation problem: Given a description of the Hamiltonian H , an evolution time t , and an initial state $|\psi(0)\rangle$, produce the final state $|\psi(t)\rangle$ (to within some error tolerance ϵ)

A classical computer cannot even represent the state efficiently.

A quantum computer cannot produce a complete description of the state.

But given succinct descriptions of

- the initial state (suitable for a quantum computer to prepare it efficiently) and
 - a final measurement (say, measurements of the individual qubits in some basis),
- a quantum computer can efficiently answer questions that (apparently) a classical one cannot.

Local and sparse Hamiltonians

Local Hamiltonians [Lloyd 96]

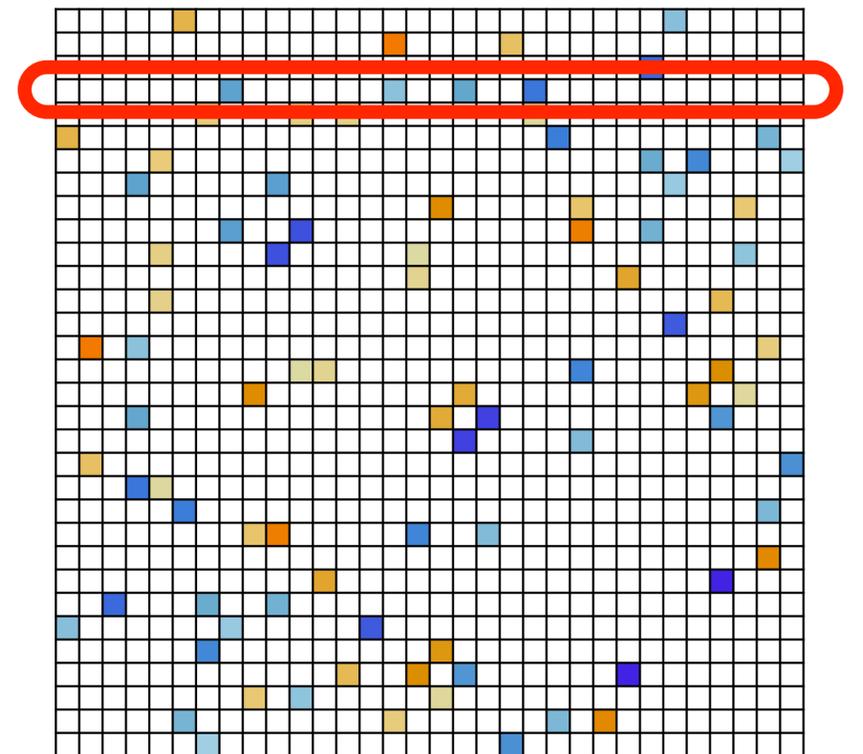
$$H = \sum_{\ell=1}^L H_{\ell} \text{ where each } H_{\ell} \text{ acts on } k = O(1) \text{ qubits}$$

Sparse Hamiltonians [Aharonov, Ta-Shma 03]

At most d nonzero entries per row, $d = \text{poly}(\log N)$
(where H is $N \times N$)

In any given row, the location of the j th nonzero entry and its value can be computed efficiently

$H =$



Note: A k -local Hamiltonian with L terms is d -sparse with $d = 2^k L$

Product formula simulation

Suppose we want to simulate $H = \sum_{\ell=1}^L H_{\ell}$

Combine individual simulations with the Lie product formula. E.g., with two terms:

$$\lim_{r \rightarrow \infty} \left(e^{-iAt/r} e^{-iBt/r} \right)^r = e^{-i(A+B)t}$$

$$\left(e^{-iAt/r} e^{-iBt/r} \right)^r = e^{-i(A+B)t} + O(t^2/r)$$

To ensure error at most ϵ , take

$$r = O\left(\frac{\|H\|t^2}{\epsilon}\right)$$

[Lloyd 96]

To get a better approximation, use higher-order formulas.

E.g., second order:

$$\left(e^{-iAt/2r} e^{-iBt} e^{-iAt/2r} \right)^r = e^{-i(A+B)t} + O(t^3/r^2)$$

Systematic expansions to arbitrary order are known [Suzuki 92]

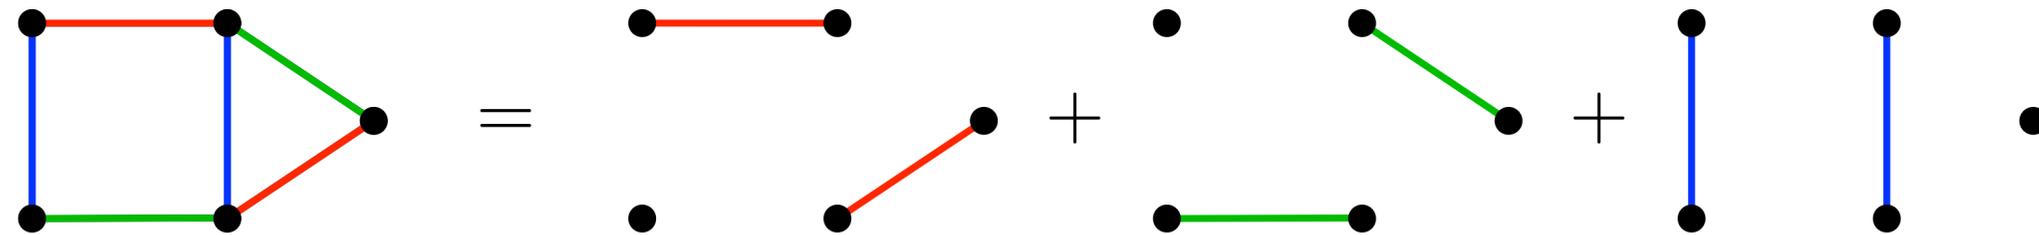
Using the $2k$ th order expansion, the number of exponentials required for an approximation with error at most ϵ is at most

$$5^{2k} L^2 \|H\| t \left(\frac{L \|H\| t}{\epsilon} \right)^{1/2k}$$

[Berry, Ahokas, Cleve, Sanders 07]

Sparse Hamiltonians and coloring

Strategy: Color the edges of the graph of H . Then the simulation breaks into small pieces that are easy to handle.

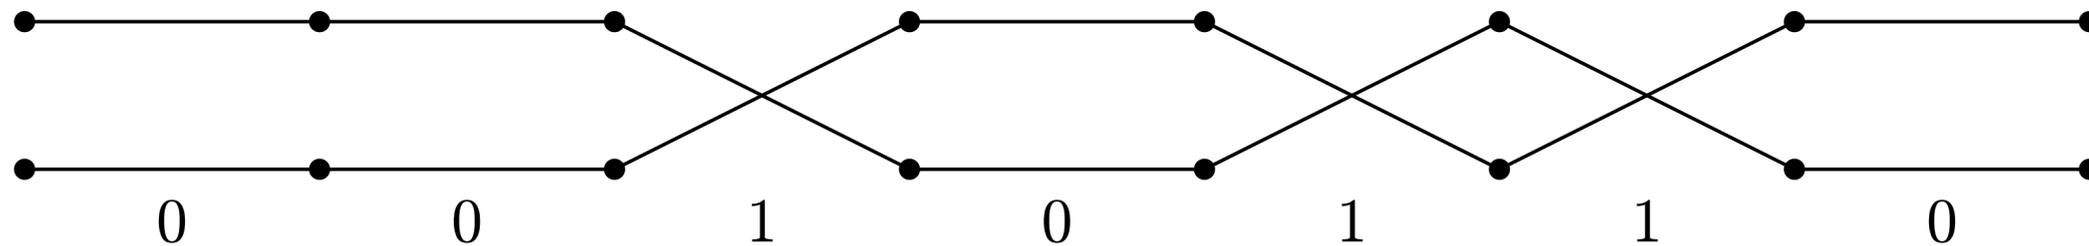


A sparse graph can be efficiently colored using only local information, so this gives efficient simulations.

Sometimes we can do better with different graph decompositions.

Real-time simulation?

No fast-forwarding theorem: Simulating Hamiltonian dynamics for time t requires $\Omega(t)$ gates.



Complexity of k th order product formula simulation is $O(5^{2k} t^{1+1/2k})$.

Can we give an algorithm with complexity precisely $O(t)$?

Systems simulate their own dynamics in real time!

Hamiltonian simulation by quantum walk

Quantum walk corresponding to H

Alternately reflect about $\text{span}\{|\psi_j\rangle\}_{j=1}^N$,

$$|\psi_j\rangle := |j\rangle \otimes \left(\nu \sum_{k=1}^N \sqrt{H_{jk}^*} |k\rangle + \nu_j |N+1\rangle \right),$$

and swap the two registers.

If H is sparse, this walk is easy to implement.

Spectral theorem: Each eigenvalue λ of H corresponds to two eigenvalues $\pm e^{\pm i \arcsin \lambda}$ of the walk operator (with eigenvectors closely related to those of H).

Simulation by phase estimation

$$|\lambda\rangle \mapsto |\lambda\rangle |\widetilde{\arcsin \lambda}\rangle \quad (\text{phase estimation})$$

$$\mapsto e^{-i\lambda t} |\lambda\rangle |\widetilde{\arcsin \lambda}\rangle$$

$$\mapsto e^{-i\lambda t} |\lambda\rangle \quad (\text{inverse phase est})$$

Theorem: $O(t/\sqrt{\epsilon})$ steps of this walk suffice to simulate H for time t with error at most ϵ .

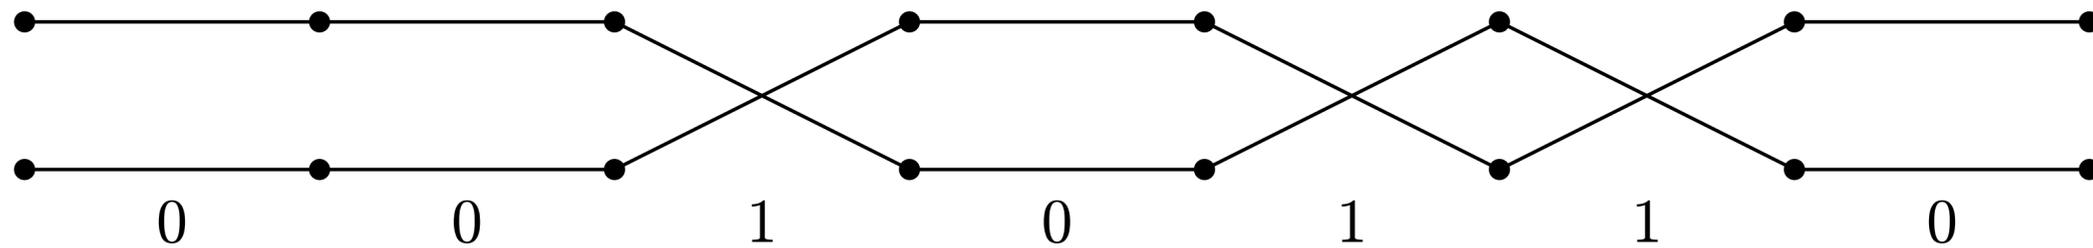
High-precision simulation?

Can we improve the dependence on ϵ ?

Many approximate computations can be done with complexity $\text{poly}(\log(1/\epsilon))$:

- computing π
- boosting a bounded-error subroutine
- Solovay-Kitaev circuit synthesis
- and more...

Lower bound (based on the *unbounded-error* query complexity of parity): $\Omega\left(\frac{\log(1/\epsilon)}{\log \log(1/\epsilon)}\right)$



Quantum walk simulation: $O(1/\sqrt{\epsilon})$

Product formulas ($2k$ th order): $O(5^{2k} \epsilon^{-2k})$

Can we do better?

Hamiltonian simulation by linear combinations of unitaries

Main idea: Directly implement the series

$$e^{-iHt} = \sum_{k=0}^{\infty} \frac{(-iHt)^k}{k!} \\ \approx \sum_{k=0}^K \frac{(-iHt)^k}{k!}$$

Write $H = \sum_{\ell} \alpha_{\ell} H_{\ell}$ with H_{ℓ} unitary.

Then

$$\sum_{k=0}^K \sum_{\ell_1, \dots, \ell_k} \frac{(-it)^k}{k!} \alpha_{\ell_1} \cdots \alpha_{\ell_k} H_{\ell_1} \cdots H_{\ell_k}$$

is a linear combination of unitaries.

LCU Lemma: Given the ability to perform unitaries V_j with unit complexity, one can perform the operation $U = \sum_j \beta_j V_j$ with complexity $O(\sum_j |\beta_j|)$. Furthermore, if U is (nearly) unitary then this implementation can be made (nearly) deterministic.

Main ideas:

- Implement U with some amplitude:

$$|0\rangle|\psi\rangle \mapsto \sin \theta |0\rangle U|\psi\rangle + \cos \theta |\Phi\rangle$$

- Boost the amplitude for success by *oblivious amplitude amplification*

Query complexity: $O\left(t \frac{\log(t/\epsilon)}{\log \log(t/\epsilon)}\right)$

Tradeoff between t and ϵ

Combining known lower bounds on the complexity of simulation as a function of t and ϵ gives

$$\Omega\left(t + \frac{\log \frac{1}{\epsilon}}{\log \log \frac{1}{\epsilon}}\right) \quad \text{vs. upper bound of} \quad O\left(t \frac{\log \frac{t}{\epsilon}}{\log \log \frac{t}{\epsilon}}\right)$$

Very recent work, using an alternative method for implementing a linear combination of quantum walk steps, gives an optimal tradeoff.

Main idea: Encode the eigenvalues of H in a two-dimensional subspace; use a carefully-chosen sequence of single-qubit rotations to manipulate those eigenvalues.

To compute the rotation angles, we must find the roots of a high-degree polynomial to high precision. This can be done in polynomial time (classically), but it's expensive in practice.

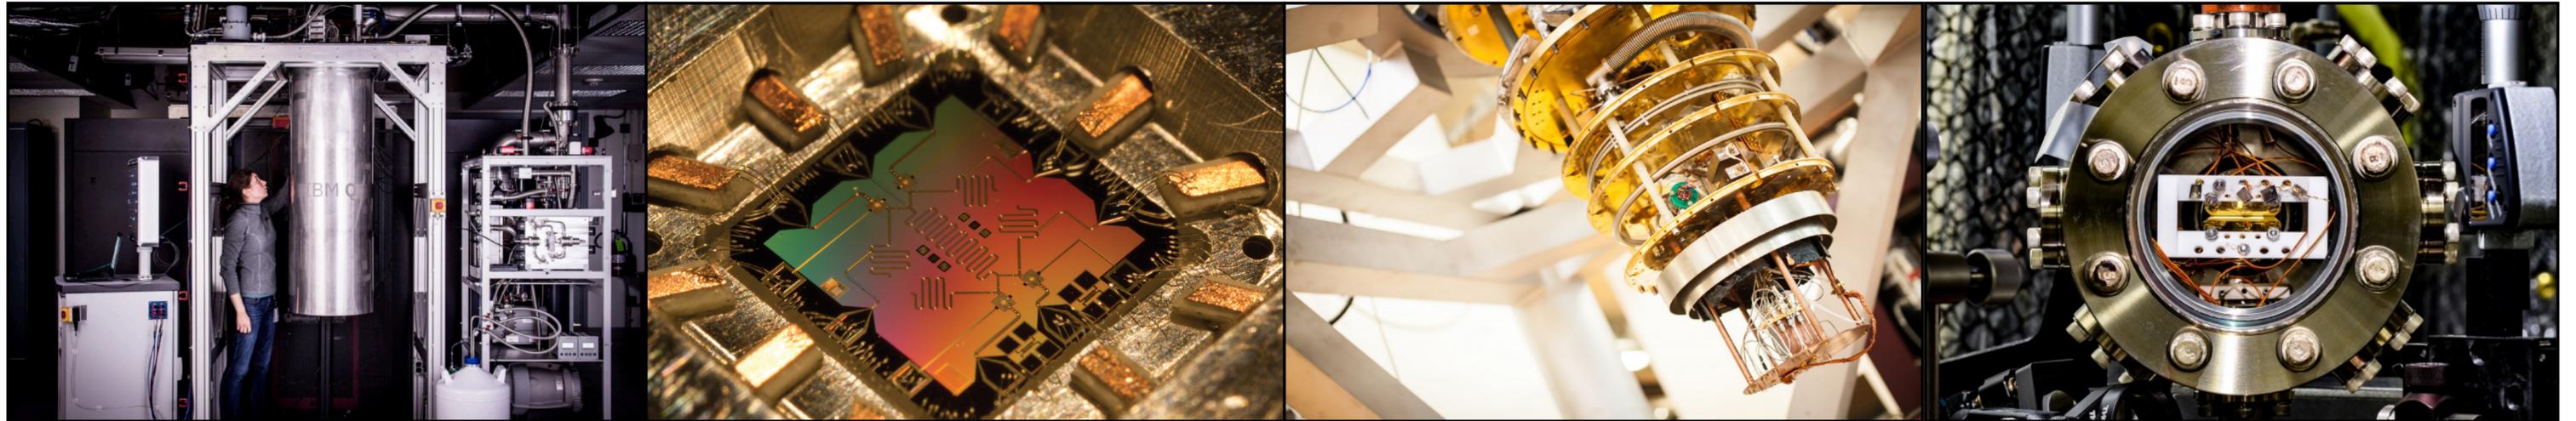
Algorithm comparison

Algorithm	Query complexity	Gate complexity
Product formula, 1st order	$O(d^4 t^2 / \epsilon)$	$O(d^4 t^2 / \epsilon)$
Product formula, (2k)th order	$O(5^{2k} d^3 t (\frac{dt}{\epsilon})^{1/2k})$	$O(5^{2k} d^3 t (\frac{dt}{\epsilon})^{1/2k})$
Quantum walk	$O(dt / \sqrt{\epsilon})$	$O(dt / \sqrt{\epsilon})$
Fractional-query simulation	$O(d^2 t \frac{\log(dt/\epsilon)}{\log \log(dt/\epsilon)})$	$O(d^2 t \frac{\log^2(dt/\epsilon)}{\log \log(dt/\epsilon)})$
Taylor series	$O(d^2 t \frac{\log(dt/\epsilon)}{\log \log(dt/\epsilon)})$	$O(d^2 t \frac{\log^2(dt/\epsilon)}{\log \log(dt/\epsilon)})$
Linear combination of q. walk steps	$O(dt \frac{\log(dt/\epsilon)}{\log \log(dt/\epsilon)})$	$O(dt \frac{\log^{3.5}(dt/\epsilon)}{\log \log(dt/\epsilon)})$
Quantum signal processing	$O(dt + \frac{\log(1/\epsilon)}{\log \log(1/\epsilon)})$	$O(dt + \frac{\log(1/\epsilon)}{\log \log(1/\epsilon)})$

OPTIMAL!

Implementation

Toward practical quantum speedup



IBM

Google/UCSB

Delft

Maryland

Important early goal: demonstrate quantum computational advantage
... but can we find a *practical* application of near-term devices?

Challenges

- Improve experimental systems
- Improve algorithms and their implementation, making the best use of available hardware

Our goal: Produce concrete resource estimates for the simplest possible practical application of quantum computers

What to simulate?

~~Quantum chemistry?~~ Spin systems!

Heisenberg model on a ring:
$$H = \sum_{j=1}^n (\vec{\sigma}_j \cdot \vec{\sigma}_{j+1} + h_j \sigma_j^z) \quad h_j \in [-h, h] \text{ uniformly random}$$

This provides a model of *self-thermalization* and *many-body localization*.

The transition between thermalized and localized phases (as a function of h) is poorly understood. Most extensive numerical study: fewer than 25 spins. [Luitz, Laflorencie, Alet 15]

Could explore the transition by preparing a simple initial state, evolving, and performing a simple final measurement. Focus on the cost of simulating dynamics.

For concreteness: $h = 1, \quad t = n, \quad \epsilon = 10^{-3}, \quad 20 \leq n \leq 100$

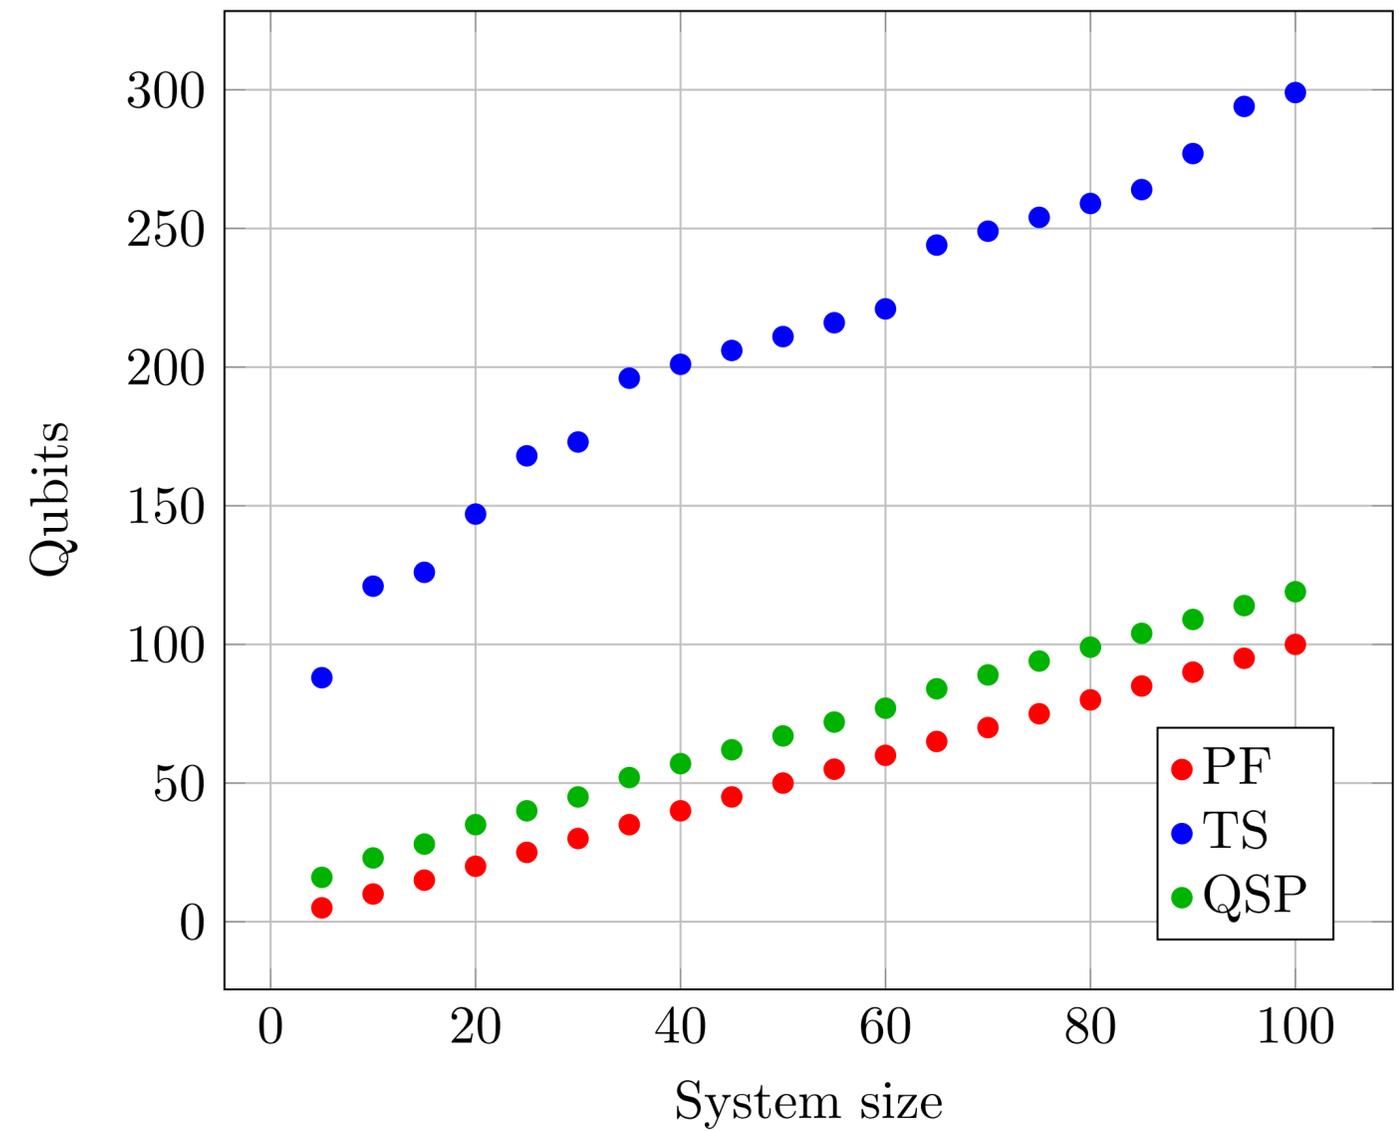
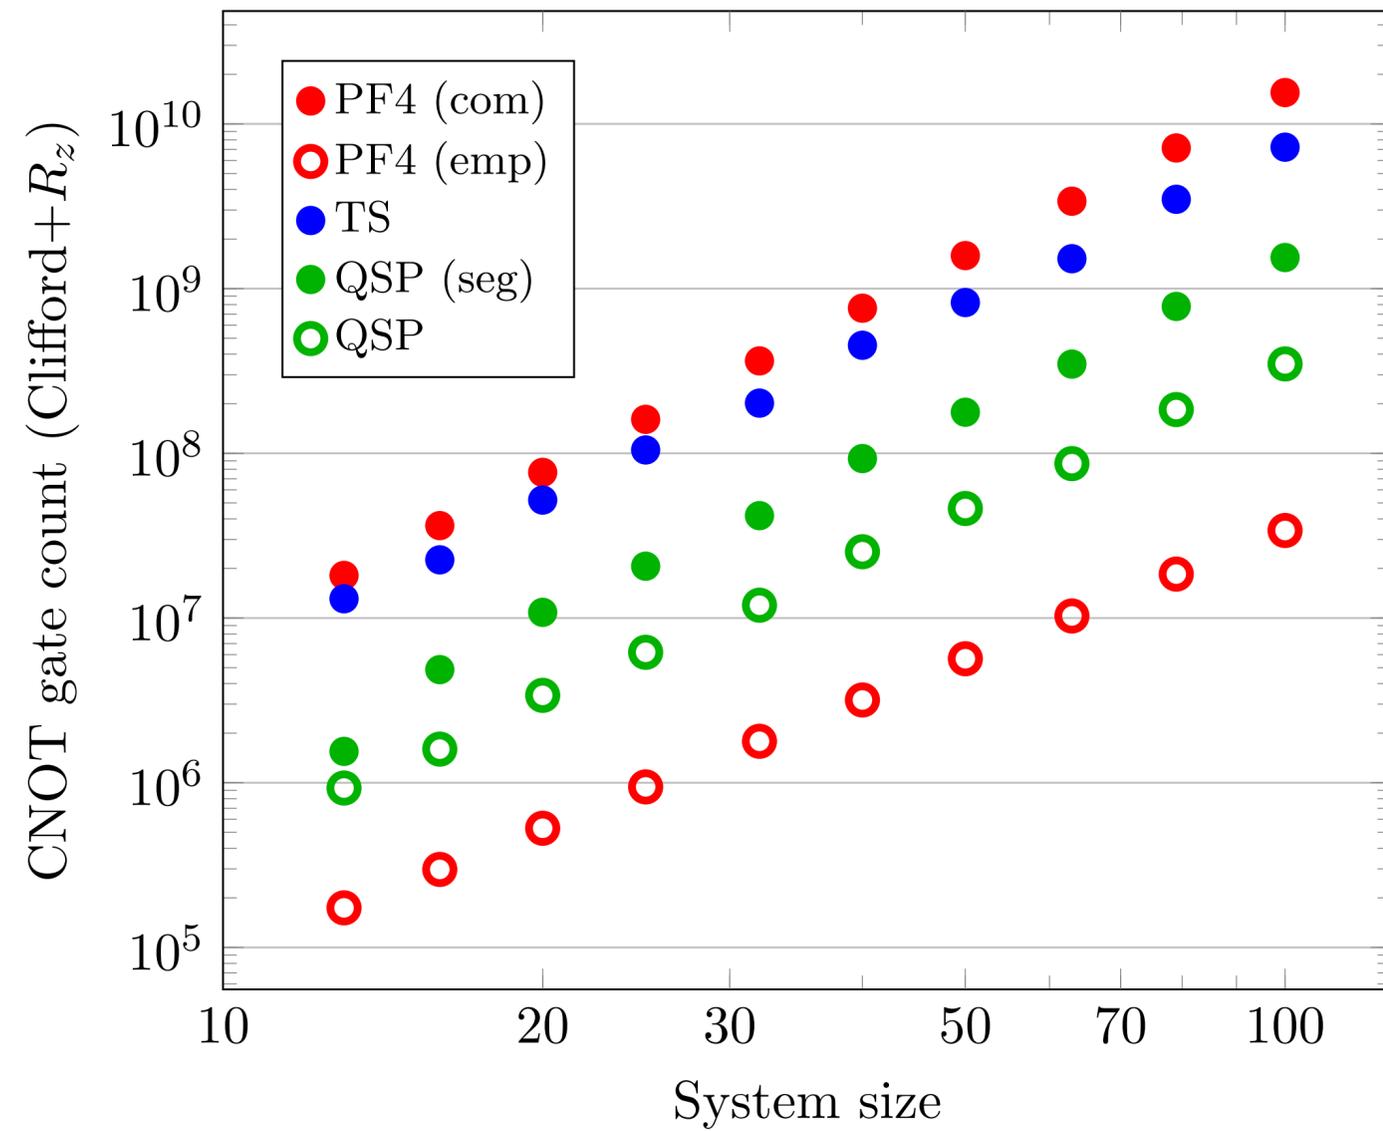
Algorithms

Algorithm	Gate complexity (t, ϵ)	Gate complexity (n)
Product formula (PF), 1st order	$O(t^2 / \epsilon)$	$O(n^5)$
Product formula (PF), $(2k)$ th order	$O(5^{2k} t^{1+1/2k} / \epsilon^{1/2k})$	$O(5^{2k} n^{3+1/k})$
Quantum walk	$O(t / \sqrt{\epsilon})$	$O(n^4 \log n)$
Fractional-query simulation	$O\left(t \frac{\log^2(t/\epsilon)}{\log \log(t/\epsilon)}\right)$	$O\left(n^4 \frac{\log^2 n}{\log \log n}\right)$
Taylor series (TS)	$O\left(t \frac{\log^2(t/\epsilon)}{\log \log(t/\epsilon)}\right)$	$O\left(n^3 \frac{\log^2 n}{\log \log n}\right)$
Linear combination of q. walk steps	$O\left(t \frac{\log^{3.5}(t/\epsilon)}{\log \log(t/\epsilon)}\right)$	$O\left(n^4 \frac{\log^2 n}{\log \log n}\right)$
Quantum signal processing (QSP)	$O(t + \log(1/\epsilon))$	$O(n^3)$
Segmented QSP, q iterates per segment	$O(t^{1+2/q} / \epsilon^{2/q})$	$O(n^{3+4/q})$

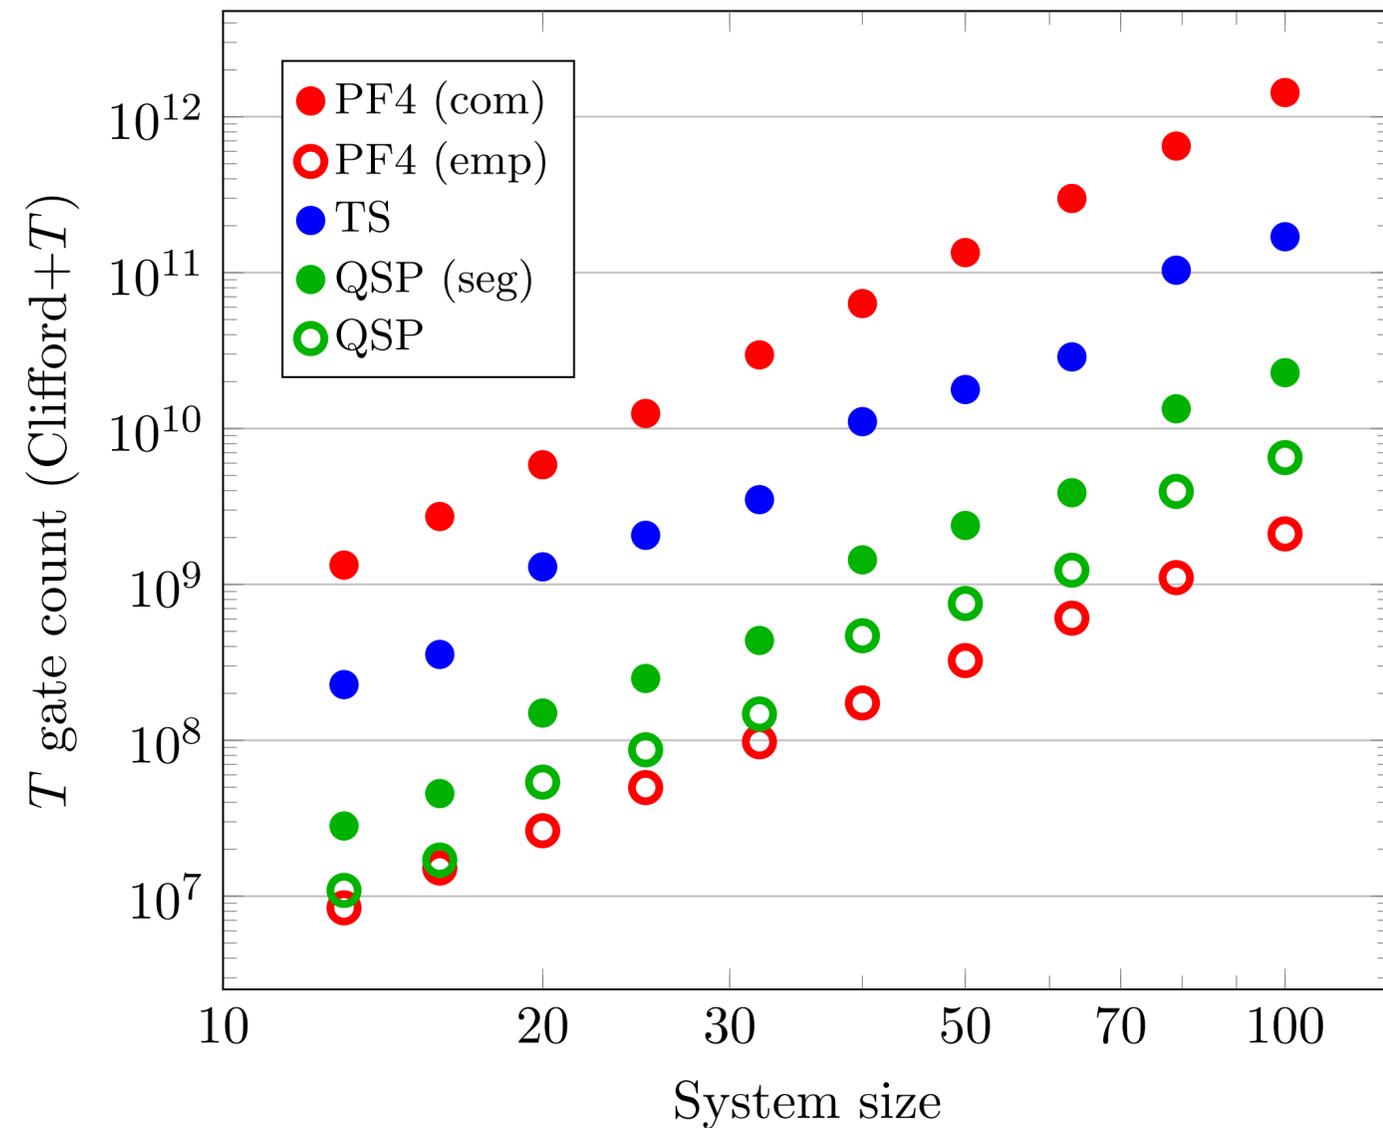
Algorithms

Algorithm	Gate complexity (t, ϵ)	Gate complexity (n)
Product formula (PF), 1st order	$O(t^2 / \epsilon)$	$O(n^5)$
Product formula (PF), (2k)th order	$O(5^{2k} t^{1+1/2k} / \epsilon^{1/2k})$	$O(5^{2k} n^{3+1/k})$
Quantum walk	$O(t / \sqrt{\epsilon})$	$O(n^4 \log n)$
Fractional-query simulation	$O\left(t \frac{\log^2(t/\epsilon)}{\log \log(t/\epsilon)}\right)$	$O\left(n^4 \frac{\log^2 n}{\log \log n}\right)$
Taylor series (TS)	$O\left(t \frac{\log^2(t/\epsilon)}{\log \log(t/\epsilon)}\right)$	$O\left(n^3 \frac{\log^2 n}{\log \log n}\right)$
Linear combination of q. walk steps	$O\left(t \frac{\log^{3.5}(t/\epsilon)}{\log \log(t/\epsilon)}\right)$	$O\left(n^4 \frac{\log^2 n}{\log \log n}\right)$
Quantum signal processing (QSP)	$O(t + \log(1/\epsilon))$	$O(n^3)$
Segmented QSP, q iterates per segment	$O(t^{1+2/q} / \epsilon^{2/q})$	$O(n^{3+4/q})$

Resource estimates



Resource estimates



Simulating a 50-qubit system (PF4, empirical):

- 50 qubits
- 3.2×10^8 T gates

Factoring a 1024-bit number [Kutin 06]:

- 3132 qubits
- 5.7×10^9 T gates

Simulating FeMoco [Reiher et al. 16]:

- 111 qubits
- 1.0×10^{14} T gates

Outlook

Super-classical quantum simulation without invoking fault tolerance?

- Improved error bounds (e.g., empirical error bound for QSP algorithm?)
- Optimized implementations
- Alternative target systems
- New simulation algorithms
- Experiments!

Resource estimates for more practical models

- Architectural constraints, parallelism
- Fault-tolerant implementations

Better provable bounds for simulation algorithms

- Product formula error bounds beyond the triangle inequality
- Efficient synthesis of the QSP circuit

Resource estimates for simulating quantum field theory

Algorithms



Dominic Berry



Richard Cleve



Robin Kothari

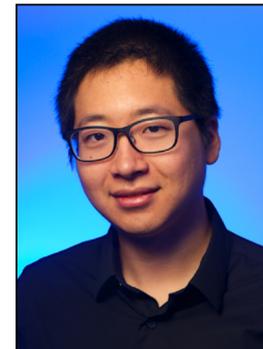


Rolando Somma

Implementation



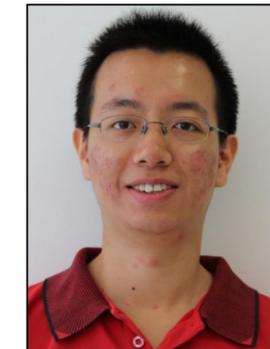
Dmitri Maslov



Yunseong Nam



Neil Julien Ross



Yuan Su



Product formula comparisons

