Quantum Computing for Modeling of Molecules and Materials

Canada Research Chairs
Chaires de recherche du Canada

Canada

What do recent collapses of quantum algorithms mean for quantum computing?

Thomas E. Baker
Canada Research Chair in Quantum Computing for Modelling of Molecules and Materials

Department of Physics & Astronomy
Department of Chemistry
Centre for Advanced Materials and Related Technology
University of Victoria

CANADA RESEARCH CHAIRS
CHAIRES DE RECHERCHE DU CANADA

UVIC

What is quantum computing?

Classical computing:

- Move electrons around
- Copy objects around
- Python, Java, C++, Fortan

Quantum Computing

- Special rules
- Can we get something?

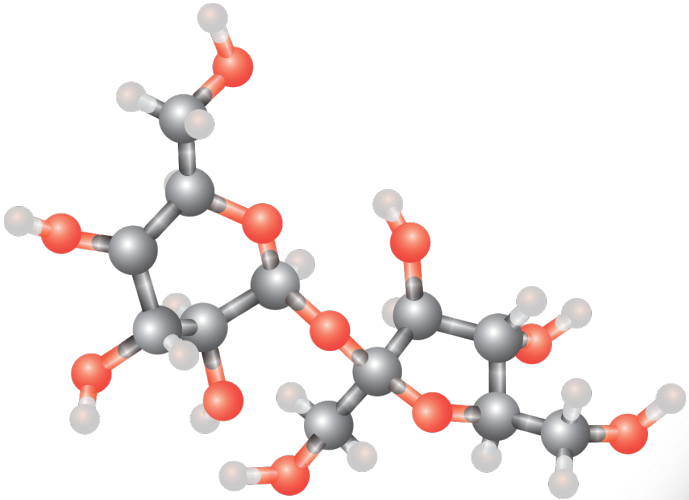

https://en.wikipedia.org/wiki/Disaccharide

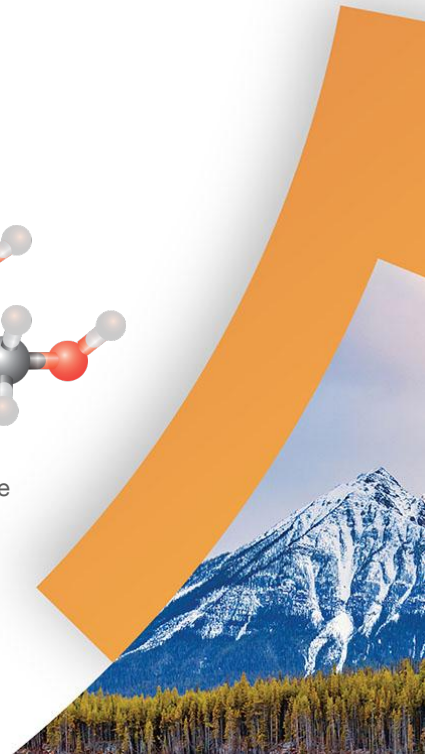

CANADA RESEARCH CHAIRS
CHAIRES DE RECHERCHE DU CANADA

Old debate: Noise

Classical error correction

- Makes cellphone calls better

Quantum error correction

- Implementation of logical qubits
- Only early implementations

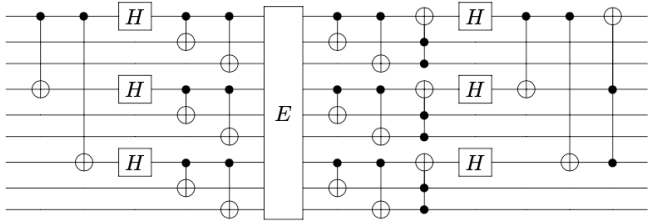

P. W. Shor, PRA **52**, R2493 (1995)

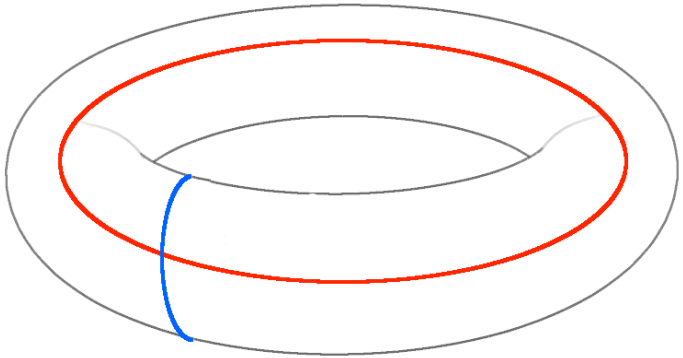

Toric codes

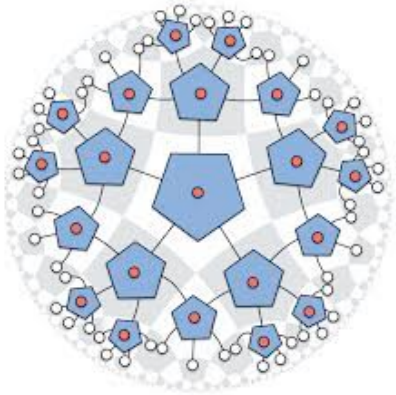

Holographic codes

CANADA RESEARCH CHAIRS
CHAIRES DE RECHERCHE DU CANADA

Quantum algorithms

Grover's search algorithm

L. K. Grover, "Quantum Mechanics Helps in Searching for a Needle in a Haystack," *Phys. Rev. Lett.* **79**, 325 (1997)

- Classical analogy with coupled oscillators

L. K. Grover, "From Schrödinger's equation to the quantum search algorithm," *Pramana* **56**, 333 (2001)

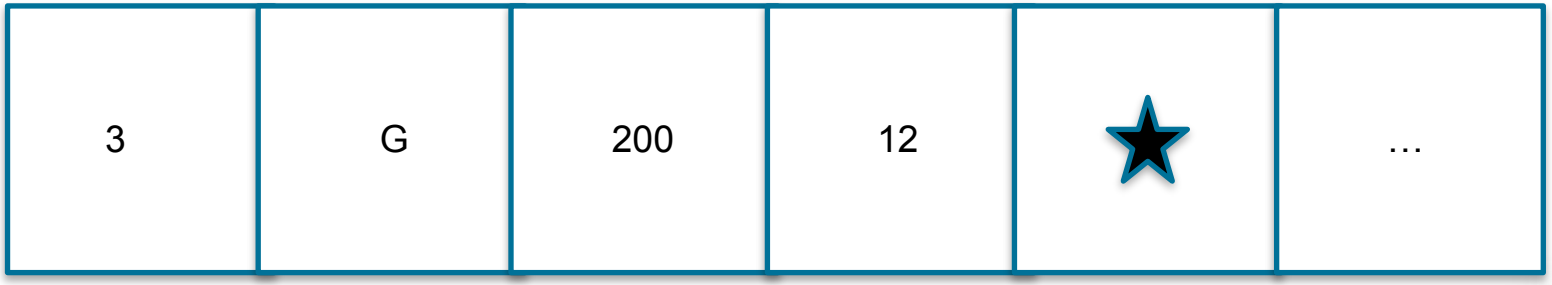

- Classical: O(N); Quantum: O($\sqrt{N}$)

Quantum algorithms run faster, but…

Tomography

- Copenhagen interpretation of measurement
- Many operations to measure

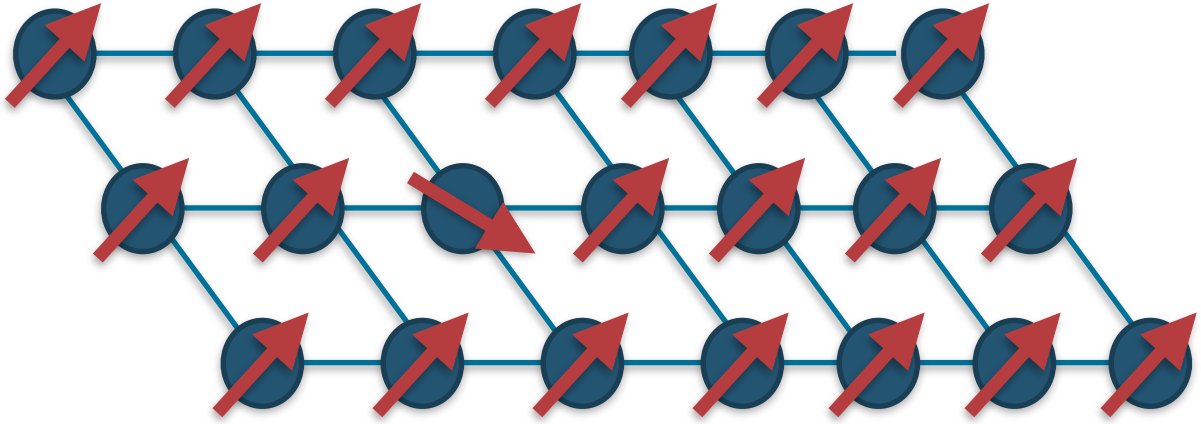

Some other algorithms

Shor's algorithm

- Breaks encryption
- Prime factorization of an integer

Deutsch-Josza algorithm

- Good mathematical example

Broad categories

- Quantum machine learning
- Quantum chemistry

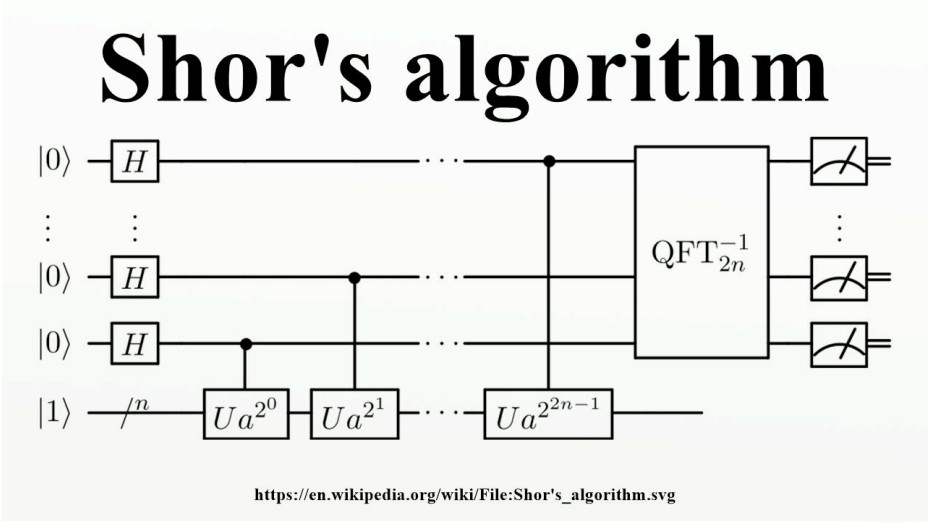

CANADA RESEARCH CHAIRS
CHAIRES DE RECHERCHE DU CANADA

Quantum chemistry

$$e^{-iH\Delta t} = e^{-iH_A\Delta t}e^{-iH_B\Delta t} + O(\Delta t^2)$$

$O(N^2)$ But a large prefactor!

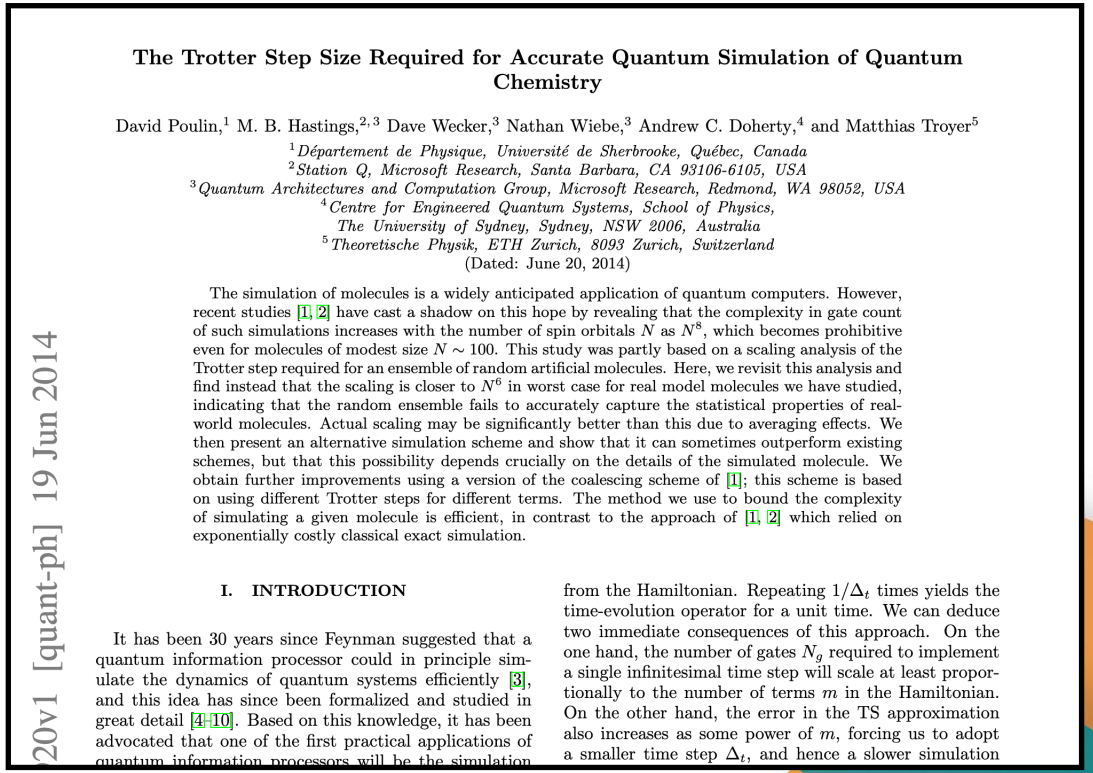

The Trotter Step Size Required for Accurate Quantum Simulation of Quantum Chemistry

David Poulin,[1] M. B. Hastings,[2,3] Dave Wecker,[3] Nathan Wiebe,[3] Andrew C. Doherty,[4] and Matthias Troyer[5]

[1] *Département de Physique, Université de Sherbrooke, Québec, Canada*
[2] *Station Q, Microsoft Research, Santa Barbara, CA 93106-6105, USA*
[3] *Quantum Architectures and Computation Group, Microsoft Research, Redmond, WA 98052, USA*
[4] *Centre for Engineered Quantum Systems, School of Physics, The University of Sydney, Sydney, NSW 2006, Australia*
[5] *Theoretische Physik, ETH Zurich, 8093 Zurich, Switzerland*

(Dated: June 20, 2014)

The simulation of molecules is a widely anticipated application of quantum computers. However, recent studies [1, 2] have cast a shadow on this hope by revealing that the complexity in gate count of such simulations increases with the number of spin orbitals N as N^8, which becomes prohibitive even for molecules of modest size $N \sim 100$. This study was partly based on a scaling analysis of the Trotter step required for an ensemble of random artificial molecules. Here, we revisit this analysis and find instead that the scaling is closer to N^6 in worst case for real model molecules we have studied, indicating that the random ensemble fails to accurately capture the statistical properties of real-world molecules. Actual scaling may be significantly better than this due to averaging effects. We then present an alternative simulation scheme and show that it can sometimes outperform existing schemes, but that this possibility depends crucially on the details of the simulated molecule. We obtain further improvements using a version of the coalescing scheme of [1]; this scheme is based on using different Trotter steps for different terms. The method we use to bound the complexity of simulating a given molecule is efficient, in contrast to the approach of [1, 2] which relied on exponentially costly classical exact simulation.

I. INTRODUCTION

It has been 30 years since Feynman suggested that a quantum information processor could in principle simulate the dynamics of quantum systems efficiently [3], and this idea has since been formalized and studied in great detail [4–10]. Based on this knowledge, it has been advocated that one of the first practical applications of quantum information processors will be the simulation

from the Hamiltonian. Repeating $1/\Delta_t$ times yields the time-evolution operator for a unit time. We can deduce two immediate consequences of this approach. On the one hand, the number of gates N_g required to implement a single infinitesimal time step will scale at least proportionally to the number of terms m in the Hamiltonian. On the other hand, the error in the TS approximation also increases as some power of m, forcing us to adopt a smaller time step Δ_t, and hence a slower simulation

CANADA RESEARCH CHAIRS
CHAIRES DE RECHERCHE DU CANADA

Quantum machine learning

Quantum versus Classical Learnability

Rocco A. Servedio* and Steven J. Gortler†

Division of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138
{rocco,sjg}@cs.harvard.edu

Abstract

This paper studies fundamental questions in computational learning theory from a quantum computation perspective. We consider quantum versions of two well-studied classical learning models: Angluin's model of exact learning from membership queries and Valiant's Probably Approximately Correct (PAC) model of learning from random examples. We give positive and negative results for quantum versus classical learnability. For each of the two learning models described above, we show that any concept class is information-theoretically learnable from polynomially many quantum examples if and only if it is information-theoretically learnable from polynomially many classical examples. In contrast to this information-theoretic equivalence betwen quantum and classical learnability, though, we observe that a separation does exist between *efficient* quantum and classical learnability. For both the model of exact learning from membership queries and the PAC model, we show that under a widely held computational hardness assumption for classical computation (the intractability of factoring), there is a concept class which is polynomial-time learnable in the quantum version but not in the classical version of the model.

Quantum machine learning: HHL algorithm

A quantum-inspired classical algorithm for recommendation systems

Ewin Tang

May 10, 2019

Abstract

We give a classical analogue to Kerenidis and Prakash's quantum recommendation system, previously believed to be one of the strongest candidates for provably exponential speedups in quantum machine learning. Our main result is an algorithm that, given an $m \times n$ matrix in a data structure supporting certain ℓ^2-norm sampling operations, outputs an ℓ^2-norm sample from a rank-k approximation of that matrix in time $O(\text{poly}(k)\log(mn))$, only polynomially slower than the quantum algorithm. As a consequence, Kerenidis and Prakash's algorithm does not in fact give an exponential speedup over classical algorithms. Further, under strong input assumptions, the classical recommendation system resulting from our algorithm produces recommendations exponentially faster than previous classical systems, which run in time linear in m and

Quantum machine learning: Shadow Tomography

Shadow Tomography of Quantum States*

Scott Aaronson
University of Texas
Austin, TX, USA
aaronson@cs.utexas.edu

ABSTRACT

We introduce the problem of *shadow tomography*: given an unknown quantum mixed state ρ of dimension D, as well as known two-outcome measurements $E_1, \ldots, E_M$, estimate the probability that E_i accepts ρ, to within additive error ε, for each of the M measurements. How many copies of ρ are needed to achieve this, with high probability? Surprisingly, we give a procedure that solves the problem by measuring only $\widetilde{O}\left(\varepsilon^{-5} \cdot \log^4 M \cdot \log D\right)$ copies. This means, for example, that we can learn the behavior of an arbitrary n-qubit state, on *all* accepting/rejecting circuits of some fixed polynomial size, by measuring only $n^{O(1)}$ copies of the state. This resolves an open problem of the author, which arose from his work on private-key quantum money schemes, but which also has applications to quantum copy-protected software, quantum advice, and quantum one-way communication. Recently, building on this work, Brandão et al. have given a different approach to shadow tomography using semidefinite programming, which achieves a

cleverness will ever let us recover a classical description of ρ, even approximately, by measuring ρ. Of course, the destructive nature of measurement is what opens up many of the cryptographic possibilities of quantum information, including quantum key distribution and quantum money.

In general, the task of recovering a description of a D-dimensional quantum mixed state ρ, given many copies of ρ, is called *quantum state tomography*. This task can be shown for information-theoretic reasons to require $\Omega\left(D^2\right)$ copies of ρ, while a recent breakthrough of O'Donnell and Wright [24] and Haah et al. [17] showed that $O\left(D^2\right)$ copies also suffice.[1] Unfortunately, this number can be astronomically infeasible: recall that, if ρ is a state of n entangled qubits, then $D = 2^n$. No wonder that the world record, for full[2] quantum state tomography, is 10-qubit states, for which millions of measurements were needed [25].

Besides the practical issue, this state of affairs could be viewed as an epistemic problem for quantum mechanics itself. If learning a full description of an n-qubit state ρ requires measuring $\exp(n)$

Quantum machine learning: A way out

PHYSICAL REVIEW LETTERS **126,** 190505 (2021)

Editors' Suggestion

Information-Theoretic Bounds on Quantum Advantage in Machine Learning

Hsin-Yuan Huang,[1,2] Richard Kueng,[3] and John Preskill[1,2,4,5]

[1]*Institute for Quantum Information and Matter, Caltech, Pasadena, California 91125, USA*
[2]*Department of Computing and Mathematical Sciences, Caltech, Pasadena, California 91125, USA*
[3]*Institute for Integrated Circuits, Johannes Kepler University Linz, Linz 4040, Austria*
[4]*Walter Burke Institute for Theoretical Physics, Caltech, Pasadena, California 91125, USA*
[5]*AWS Center for Quantum Computing, Pasadena, California 91125, USA*

(Received 12 January 2021; revised 17 March 2021; accepted 2 April 2021; published 14 May 2021)

We study the performance of classical and quantum machine learning (ML) models in predicting outcomes of physical experiments. The experiments depend on an input parameter x and involve execution of a (possibly unknown) quantum process $\mathcal{E}$. Our figure of merit is the number of runs of $\mathcal{E}$ required to achieve a desired prediction performance. We consider classical ML models that perform a measurement and record the classical outcome after each run of $\mathcal{E}$, and quantum ML models that can access $\mathcal{E}$ coherently to acquire quantum data; the classical or quantum data are then used to predict the outcomes of future experiments. We prove that for any input distribution $\mathcal{D}(x)$, a classical ML model can provide accurate predictions *on average* by accessing $\mathcal{E}$ a number of times comparable to the optimal quantum ML model. In contrast, for achieving an accurate prediction on *all* inputs, we prove that the exponential quantum advantage is possible. For example, to predict the expectations of all Pauli observables in an n-qubit system ρ, classical ML models require $2^{\Omega(n)}$ copies of ρ, but we present a quantum ML model using only $\mathcal{O}(n)$ copies. Our results clarify where the quantum advantage is possible and highlight the potential for classical ML models to address challenging quantum problems in physics and chemistry.

Another recent success: Qubitization

Hamiltonian Simulation by Qubitization

Guang Hao Low[1] and Isaac L. Chuang[2]

[1]Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

[2]Department of Electrical Engineering and Computer Science, Department of Physics, Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

July 12, 2019

We present the problem of approximating the time-evolution operator $e^{-i\hat{H}t}$ to error ϵ, where the Hamiltonian $\hat{H} = (\langle G|\otimes\hat{\mathcal{I}})\hat{U}(|G\rangle\otimes\hat{\mathcal{I}})$ is the projection of a unitary oracle $\hat{U}$ onto the state $|G\rangle$ created by another unitary oracle. Our algorithm solves this with a query complexity $\mathcal{O}\big(t + \log(1/\epsilon)\big)$ to both oracles that is optimal with respect to all parameters in both the asymptotic and non-asymptotic regime, and also with low overhead, using at most two additional ancilla qubits. This approach to Hamiltonian simulation subsumes important prior art considering Hamiltonians which are d-sparse or a linear combination of unitaries, leading to significant improvements in space and gate complexity, such as a quadratic speed-up for precision simulations. It also motivates useful new instances, such as where $\hat{H}$ is a density matrix. A key technical result is 'qubitization', which uses the controlled version of these oracles to embed any $\hat{H}$ in an invariant SU(2) subspace. A large class of operator functions of $\hat{H}$ can then be computed with optimal query complexity, of which $e^{-i\hat{H}t}$ is a special case.

Contents

1 Introduction 2

$$H|\psi\rangle = E|\psi\rangle$$

$$A|\psi\rangle = \alpha|\psi\rangle + \beta|\psi^{\perp}\rangle$$

Conclusion

- Quantum algorithms can be good
  - New perspective on algorithm development
  - Some algorithms have fallen…
- Quantum machine learning
  - How to learn data faster
- Some hope going forward
  - Qubitization
  - Shadow tomography
- Quantum chemistry
  - Some efficient algorithms, but they require error correction

CANADA RESEARCH CHAIRS
CHAIRES DE RECHERCHE DU CANADA

UVIC