Entanglement accelerates quantum simulation

Andrew Childs
University of Maryland

COMPUTER SCIENCE
UNIVERSITY OF MARYLAND

UMIACS
University of Maryland
Institute for Advanced
Computer Studies

JOINT CENTER FOR
QUANTUM INFORMATION
AND COMPUTER SCIENCE

Institute for Robust
Quantum Simulation

Qi Zhao
University of Hong Kong

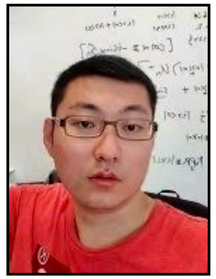

You Zhou
Fudan University

arXiv:2406.02379

Quantum speedup

Quantum computers allow for *interference between computational paths* in an *exponentially large state space*

Hype: Quantum computers will revolutionize all areas of computing (logistics, financial modeling, weather prediction, natural language processing, climate change, machine learning…).

Reality: Problems must have special structure to take advantage of quantum computing. Very few problems have convincing evidence for significant quantum speedup (cryptanalysis, simulating quantum mechanics, …).

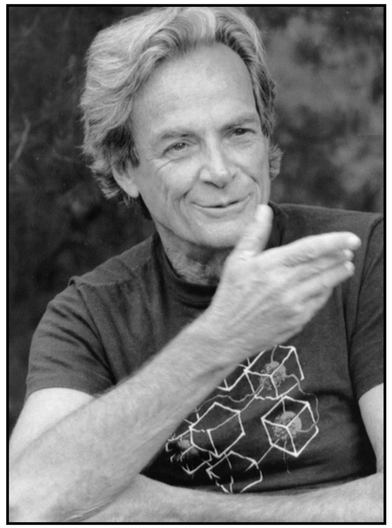

"… nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy."

Richard Feynman
Simulating physics with computers (1981)

Computational quantum physics

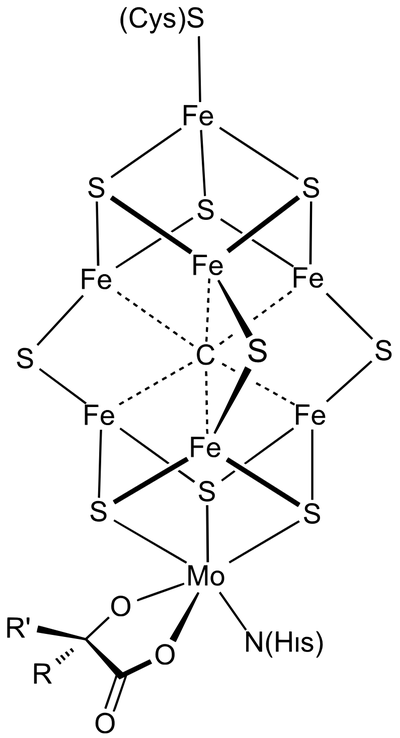

quantum chemistry

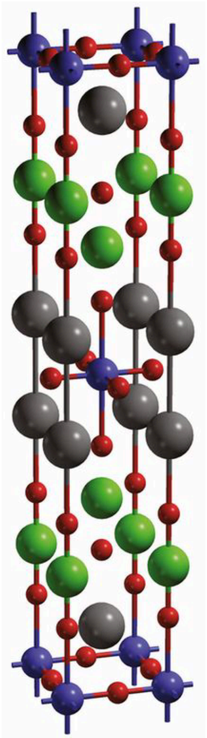

condensed matter physics/
materials science

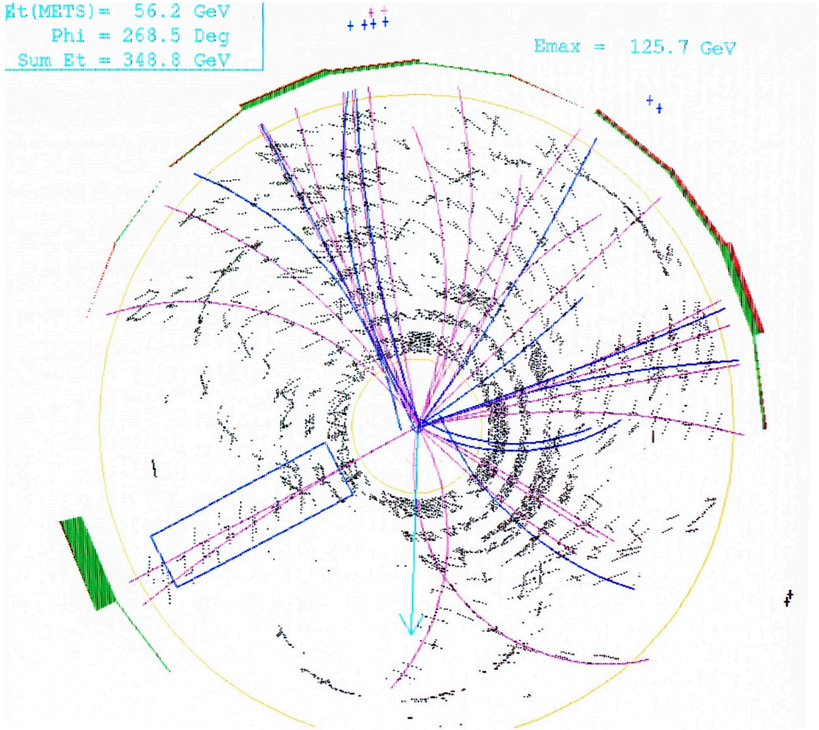

nuclear/particle
physics

Implementing quantum algorithms

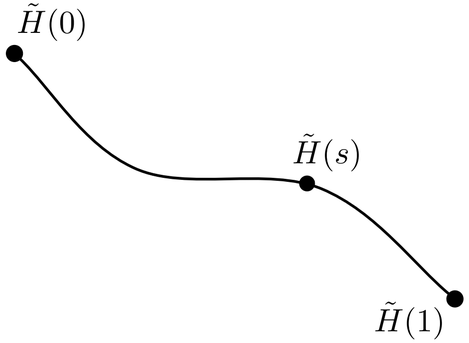

adiabatic optimization

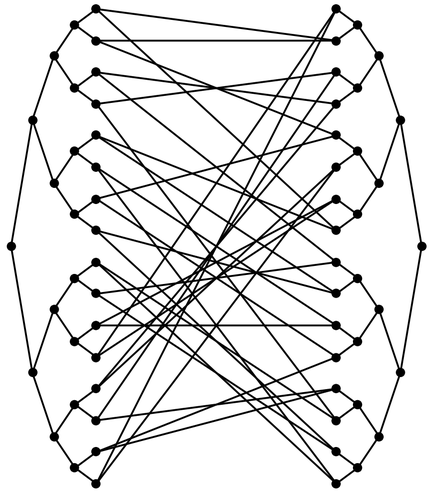

exponential speedup by quantum walk

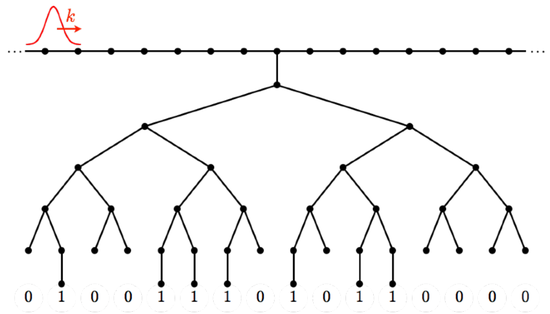

evaluating Boolean formulas

$$A|x\rangle = |b\rangle$$

linear/ differential equations, convex optimization

Simulation as a theoretical tool

Quantum simulation algorithms provide convenient descriptions of the time evolution operator.

Can we use them as theoretical tools to reason about physics?

Example: Lieb-Robinson bounds from quantum simulation algorithms

PHYSICAL REVIEW X **9,** 031006 (2019)

Locality and Digital Quantum Simulation of Power-Law Interactions

Minh C. Tran,[1,2,3,*] Andrew Y. Guo,[1,2] Yuan Su,[1,4,5] James R. Garrison,[1,2] Zachary Eldredge,[1,2] Michael Foss-Feig,[6,1,2] Andrew M. Childs,[1,4,5] and Alexey V. Gorshkov[1,2]

[1]*Joint Center for Quantum Information and Computer Science, NIST/University of Maryland, College Park, Maryland 20742, USA*
[2]*Joint Quantum Institute, NIST/University of Maryland, College Park, Maryland 20742, USA*
[3]*Kavli Institute for Theoretical Physics, University of California, Santa Barbara, California 93106, USA*
[4]*Department of Computer Science, University of Maryland, College Park, Maryland 20742, USA*
[5]*Institute for Advanced Computer Studies, University of Maryland, College Park, Maryland 20742, USA*
[6]*United States Army Research Laboratory, Adelphi, Maryland 20783, USA*

(Received 25 September 2018; revised manuscript received 21 February 2019; published 10 July 2019)

The propagation of information in nonrelativistic quantum systems obeys a speed limit known as a Lieb-Robinson bound. We derive a new Lieb-Robinson bound for systems with interactions that decay with distance r as a power law, $1/r^{\alpha}$. The bound implies an effective light cone tighter than all previous bounds. Our approach is based on a technique for approximating the time evolution of a system, which was first introduced as part of a quantum simulation algorithm by Haah *et al.*, FOCS'18. To bound the error of the approximation, we use a known Lieb-Robinson bound that is weaker than the bound we establish. This result brings the analysis full circle, suggesting a deep connection between Lieb-Robinson bounds and digital quantum simulation. In addition to the new Lieb-Robinson bound, our analysis also gives an error bound for the Haah *et al.* quantum simulation algorithm when used to simulate power-law decaying interactions. In particular, we show that the gate count of the algorithm scales with the system size better than existing algorithms when $\alpha > 3D$ (where D is the number of dimensions).

DOI: 10.1103/PhysRevX.9.031006

Subject Areas: Atomic and Molecular Physics, Condensed Matter Physics, Quantum Information

I. INTRODUCTION

Lieb-Robinson bounds limit the rate at which information can propagate in systems that obey the laws of Lieb-Robinson bound. Thus, understanding the fundamental limit on the speed of information propagation in these systems holds serious physical implications, including for the applications mentioned above. Despite many efforts in

Main idea of this talk

Quantum simulation algorithms can be more efficient when the state being simulated is entangled—exactly the case where classical simulation is hard!

Quantum dynamics

The dynamics of a quantum system are determined by its *Hamiltonian* H, a Hermitian operator whose spectrum describes the energy levels of the system.

For time-independent H: $$i\frac{\mathrm{d}}{\mathrm{d}t}|\psi(t)\rangle = H|\psi(t)\rangle \qquad \Rightarrow \qquad |\psi(t)\rangle = e^{-iHt}|\psi(0)\rangle$$

For time-dependent $H(t)$: $$i\frac{\mathrm{d}}{\mathrm{d}t}|\psi(t)\rangle = H(t)|\psi(t)\rangle \quad \Rightarrow \quad |\psi(t)\rangle = \mathcal{T}e^{-i\int_0^t H(\tau)\,\mathrm{d}\tau}|\psi(0)\rangle$$

The Hamiltonian simulation problem

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

Standard measure of error: worst-case ℓ_2 distance

$$\max_{|\psi(0)\rangle} \| |\psi(t)\rangle - |\psi_{\text{sim}}\rangle \| = \| e^{-iHt} - U_{\text{sim}} \|$$

initial state: $|\psi(0)\rangle$

$e^{-iHt}|\psi(0)\rangle$

simulation output state: $|\psi_{\text{sim}}\rangle$

simulated unitary operation: U_{sim}

Goal: Understand the cost of simulation as a function of t, ϵ, and parameters of H.

(cost could be circuit size, circuit depth, queries, …)

Also important for applications but not covered in this talk: state preparation, measurement, representing fermions/bosons using qubits

Hamiltonian simulation is BQP-complete

How do the abilities of classical and quantum computers to simulate Hamiltonian dynamics compare?

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.

Product formula simulation

Suppose we want to simulate $H = \sum_{j=1}^{L} H_j$

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

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

$$e^{-i(A+B)t} = 1 - i(A+B)t - \tfrac{1}{2}(A^2 + AB + BA + B^2)t^2 + O(t^3)$$

$$\begin{aligned} e^{-iAt}e^{-iBt} &= (1 - iAt - \tfrac{1}{2}A^2t^2 + O(t^3))(1 - iBt - \tfrac{1}{2}B^2t^2 + O(t^3)) \\ &= 1 - i(A+B)t - \tfrac{1}{2}(A^2 + 2AB + B^2)t^2 + O(t^3) \end{aligned}$$

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

To ensure error at most ϵ, take $r = O(t^2/\epsilon)$

[Lloyd 96]

Higher-order product formulas

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

E.g., second order: $S_2(t) := e^{-iAt/2}e^{-iBt}e^{-iAt/2} = e^{-i(A+B)t} + O(t^3)$

Can construct expansions to arbitrarily high order [Suzuki 92]:

$$\text{let } S_{2k}(t) := S_{2k-2}(p_k t)^2 S_{2k-2}(q_k t) S_{2k-2}(p_k t)^2$$

Then you can find numbers p_k and q_k (as functions of k) so that

$$S_{2k}(t) = e^{-i(A+B)t} + O(t^{2k+1})$$

Using the order-$2k$ expansion for an L-term Hamiltonian, the number of exponentials required for an approximation with error at most ϵ is at most [Berry, Ahokas, Cleve, Sanders 07]

$$5^{2k} L^2 ht \left(\frac{Lht}{\epsilon}\right)^{1/2k} \qquad h := \max_j \|H_j\|$$

Improved error analysis

Numerics suggest that product formula error bounds based on Taylor series expansion can be very loose in practice.

Alternative: *local error analysis* provides convenient integral representations of the error

[Descombes, Thalhammer 10]

Example (first order):

$$e^{-iBt}e^{-iAt} - e^{-i(A+B)t} = \int_0^t \mathrm{d}\tau_1 \int_0^{\tau_1} \mathrm{d}\tau_2 \, e^{-i(A+B)(t-\tau_1)} e^{i(\tau_2-\tau_1)B} [A,B] e^{-i\tau_2 B} e^{-i\tau_1 A}$$

Therefore $\|e^{-iBt}e^{-iAt} - e^{-i(A+B)t}\| \leq \|[A,B]\| t^2$

Advantages:

- no explicit sum over higher-order terms
- commutator scaling
- can offer much better performance, both asymptotically and in practice

[Childs, Su 19]

A theory of Trotter error

Local error analysis can be generalized to give tight bounds on the error of arbitrary-order product formula approximations with any number of terms.

Theorem. A pth-order product formula approximates the evolution of $H = \sum_{j=1}^{L} H_j$ with additive error $O(\alpha t^{p+1})$ where

$$\alpha := \sum_{j_1,\ldots,j_{p+1}} \left\| [H_{j_{p+1}}, [\cdots, [H_{j_2}, H_{j_1}] \cdots]] \right\|.$$

Therefore $O\left(Lt(\alpha t/\epsilon)^{1/p}\right)$ gates suffice to simulate H for time t with constant accuracy.

This gives much tighter rigorous analysis of product formulation simulations (among other applications).

[Childs, Su, Tran, Wiebe, Zhu 19]

Linear-time simulation

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

Proof is based on the hardness of computing parity: given a black box for the bits of a string $x \in \{0,1\}^n$, computing $x_1 \oplus \cdots \oplus x_n$ requires $\Omega(n)$ quantum queries.

[Beals, Buhrman, Cleve, Mosca, de Wolf 98; Farhi, Goldstone, Gutmann, Sipser 98]

Construct a sparse Hamiltonian that encodes the parity:

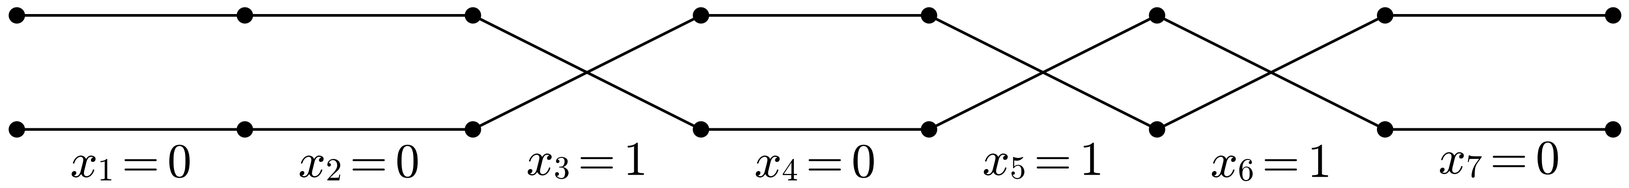

Simulation for time $t = O(n)$ can be used to determine the parity (choosing weights to ensure perfect transfer), so a general-purpose simulation method must use $\Omega(t)$ gates.

[Berry, Ahokas, Cleve, Sanders 05]

Applying phase estimation to a quantum walk constructed from H gives an $O(t)$ simulation.

[Childs10; Berry, Childs 12]

High-precision simulation

Can we improve the dependence on ϵ?

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

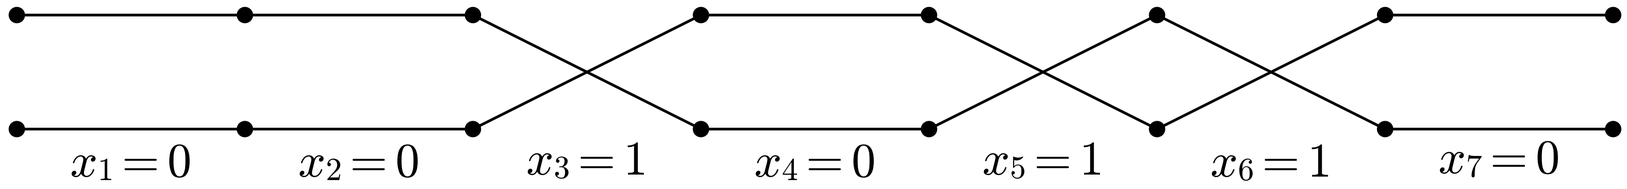

Quantum walk simulation: $O(1/\sqrt{\epsilon})$

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

Improved quantum algorithm:

Directly implement the truncated Taylor series of $\exp(-iHt)$, cost $O\big(t\frac{\log(t/\epsilon)}{\log\log(t/\epsilon)}\big)$

LCU Lemma: Implement $U = \sum_j \beta_j V_j$ with complexity $O(\sum_j |\beta_j|)$

[Berry, Childs, Cleve, Kothari, Somma 14 & 15]

Optimal tradeoff between t and ϵ

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

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

An alternative method for implementing a linear combination of unitary operations, *quantum signal processing*, gives an optimal tradeoff. [Low, Chuang 16, 17]

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.

Quantum signal processing (and more general *quantum singular value transformation*) gives a powerful framework for designing other quantum algorithms [Gilyén, Su, Low, Wiebe 19].

Lattice Hamiltonians

We've focused on the complexity as a function of t (evolution time) and ϵ (allowed error). What about the dependence on system size?

Consider n spins with nearest-neighbor interactions on a grid of fixed dimension. To simulate evolution for constant time, previous methods (LCU, QSP, high-order PF with standard analysis) use $O(n^2)$ gates, with circuit depth (execution time with parallel gates) $O(n)$.

Execution time should not have to be extensive!

Can give a simulation with $\tilde{O}(n)$ gates, $\tilde{O}(1)$ depth (nearly optimal!) [Haah, Hastings, Kothari, Low 18]

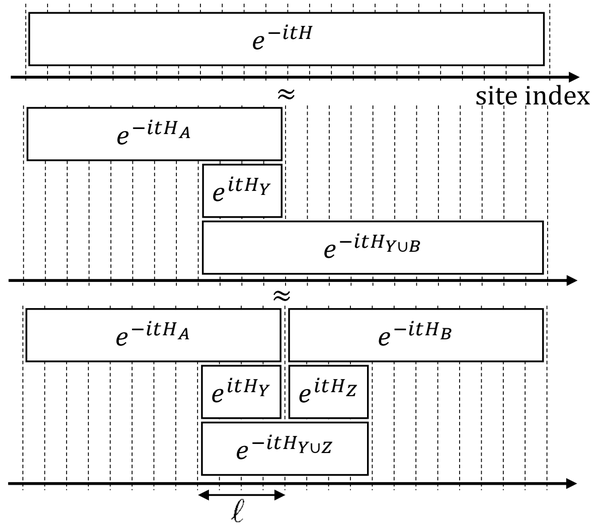

- Lieb-Robinson bound limits the speed of propagation
- Simulate small regions with negative-time evolutions to correct the boundaries

Commutator bounds on high-order product formulas give nearly the same scaling.

[Childs, Su, Tran, Wiebe, Zhu 19]

State of the art

We have algorithms with optimal performance as a function of

evolution time t: $\Theta(t)$

allowed error ϵ: $\Theta\left(\frac{\log(1/\epsilon)}{\log\log(1/\epsilon)}\right)$

system size n: $\tilde{\Theta}(n)$

Furthermore, we can achieve the optimal tradeoff between t and ϵ: $\Theta\left(t + \frac{\log(1/\epsilon)}{\log\log(1/\epsilon)}\right)$

Can we do better?

Of course not!
We already have optimal algorithms.

Of course we can!
Improve performance in practice,
change the problem specification,
⋮

Random inputs

Worst-case input states could be atypical. Can we give a more efficient simulation with extra information about the input state?

Recall that for a pth-order formula, the worst-case error is $O(\alpha t^{p+1})$ where

$$\alpha := \sum_{j_1,\ldots,j_{p+1}} \left\| [H_{j_{p+1}}, [\cdots, [H_{j_2}, H_{j_1}] \cdots]] \right\|.$$

We can give a tighter bound if the input state is chosen at random from some distribution. For a 1-design ensemble on a d-dimensional system, the average error is $O(\alpha_F t^{p+1})$ where

$$\alpha_F := \frac{1}{\sqrt{d}} \sum_{j_1,\ldots,j_{p+1}} \left\| [H_{j_{p+1}}, [\cdots, [H_{j_2}, H_{j_1}] \cdots]] \right\|_F.$$

This is no worse, but can sometimes give much better performance (for the same algorithm!).

[Zhao, Zhou, Shaw, Li, Childs 22]

Are random inputs practical?

The random-input analysis suggests that product-formula simulation typically performs much better than the worst case.

The setting is somewhat flexible (only requires a 1-design input distribution), but it does require the input to be drawn from some well-behaved distribution.

In applications, we may not have a nice distribution, or care about a particular input state.

Can we identify specific, typical features that make simulation faster?

State-dependent error bound

The leading-order contribution to the error only depends on the marginals of the state corresponding to the leading-order error terms of the product formula approximation.

Leading-order PF error: $U - \mathscr{U} = \sum_j E_j + \cdots$

← ideal evolution; ← product formula; ← local terms; ← higher-order terms

Simulation error:

$$\begin{aligned} \|(U-\mathscr{U})\,|\psi\rangle\|^2 &\approx \sum_{j,j'} \langle\psi|\, E_j^\dagger E_{j'} |\psi\rangle \\ &\approx \sum_{j,j'} \operatorname{tr}\Big(\widetilde{E_j^\dagger E_{j'}} (\rho_{j,j'} - \mathbb{I}/d_{j,j'}) \Big) + \sum_{j,j'} \operatorname{tr}\Big(E_j^\dagger E_{j'}\Big)/2^n \\ &\lessapprox \sum_{j,j'} \|E_j^\dagger E_{j'}\| \operatorname{tr}|\rho_{j,j'} - \mathbb{I}/d_{j,j'}| + \|E\|_F^2 \end{aligned}$$

← restricted to j,j'

If the marginals are maximally mixed, then the error is given by the Frobenius norm of the product formula error (as for a random input state).

More generally, there is a contribution from the spectral norm of the product formula error, but it is small if the marginals are close to maximally mixed.

Entanglement accelerates quantum simulation

Theorem. *For a given pure quantum state* $|\psi\rangle$ *and quantum evolution* $U = e^{-iHt}$ *with* p*th-order Trotter approximation* $\mathscr{U}$*,*

$$\|(U - \mathscr{U})\,|\psi\rangle\| = \mathcal{O}\left(t^{p+1}\sqrt{\sum_{j,j'} \|E_j^\dagger E_{j'}\|\,\mathrm{tr}\,|\rho_{j,j'} - \mathbb{I}/d_{\rho_{j,j'}}|} + t^{p+1}\|E\|_F\right).$$

where E_j *are the leading-order local error terms in the Trotter approximation,* $E = \sum_j E_j$*, and* $\rho_{j,j'}$ *is the reduced state of* $|\psi\rangle$ *on the support of* $E_j^\dagger E_{j'}$*.*

Highly entangled states have marginals that are close to maximally mixed. (Could recast the bound in terms of entropies of the marginals.)

To simulate evolution for long times, divide into short segments such that the total error is small.

Improvement over worst-case simulation

Consider simulating an n-site nearest-neighbor lattice Hamiltonian in constant dimensions with a pth-order product formula.

Worst case: $\alpha := \sum_{j_1,\ldots,j_{p+1}} \left\| [H_{j_{p+1}}, [\cdots, [H_{j_2}, H_{j_1}] \cdots]] \right\| = O(n)$

Number of Trotter steps to evolve for time t with error at most ϵ: $O(t(nt/\epsilon)^{1/p})$

Average case: $\alpha_F := \frac{1}{\sqrt{d}} \sum_{j_1,\ldots,j_{p+1}} \left\| [H_{j_{p+1}}, [\cdots, [H_{j_2}, H_{j_1}] \cdots]] \right\|_F = O(\sqrt{n})$

Number of Trotter steps to evolve for time t with error at most ϵ: $O(t(\sqrt{n}t/\epsilon)^{1/p})$

This analysis sets strictly better parameters for *the same algorithm*.

Example

Ising model with mixed fields: $H = h_x \sum_{j=1}^{N} X_j + h_y \sum_{j=1}^{N} Y_j + J \sum_{j=1}^{N-1} X_j X_{j+1}$

$t = N$
$\epsilon = 10^{-5}$

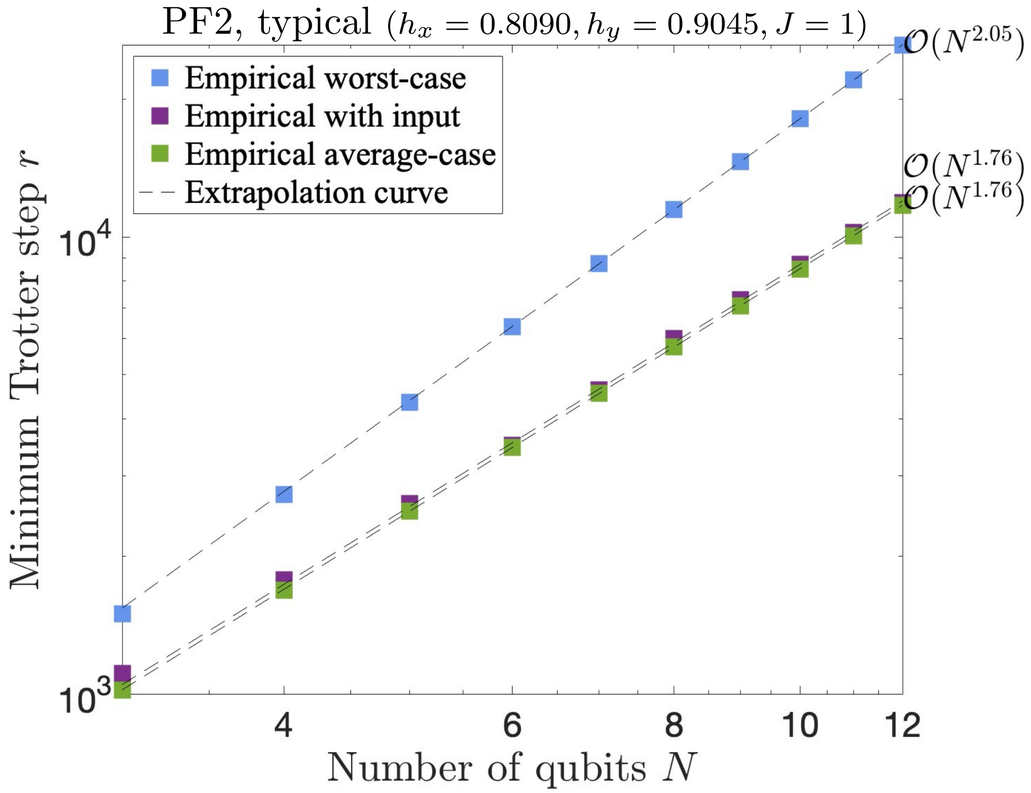

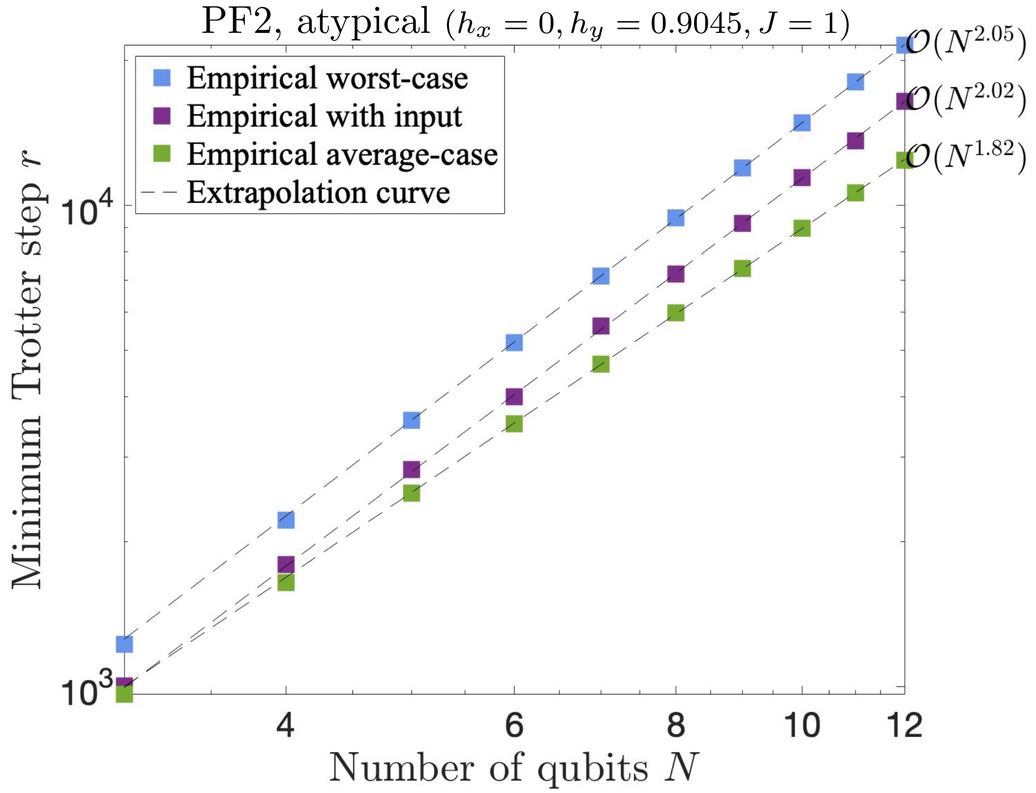

Theory prediction: worst case $O(N^2)$; average case $O(N^{1.75})$.

Adaptive algorithm

Typical Hamiltonians rapidly produce entangled states, even when the state is initially unentangled.

How can we tell that this is the case, and see how entangled the state has become, so that simulation for longer times can be more efficient?

Periodically measure the state, using shadow tomography to estimate the entanglement of its marginals. Depending on the parameters, this can be advantageous even though we have to restart the simulation when measuring a state.

Example: Consider first-order simulation of an n-site lattice Hamiltonian that makes the marginals highly mixed in constant time. Worst-case simulation: $O(n^4 t^2)$. With $O(n^2)$ measurements after constant evolution time, cost is $O(n^4 t + n^{3.5} t^2)$.

Can also consider more sophisticated strategies with multiple intermediate measurements.

Conclusions

- Simulation with product formulas performs better if the state is entangled.
  (If the state is weakly entangled, we can give an efficient classical simulation!)
- Achieves the improved performance of average-case simulation under a more natural assumption.
- No change to the algorithm; just a better analysis of its parameters.
- If necessary, entanglement can be estimated adaptively with reasonable overhead.

Future directions

- Find explicit, readily computable error bounds with tight constant factors.
- Explore the extent of the improvement for practical simulations.
- Are there similar improvements for other quantum evolutions, e.g., algorithms that implement imaginary time evolution?
- Under what other conditions can simulation bounds be improved?