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

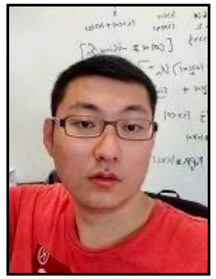

You Zhou
Fudan University

arXiv:2406.02379

Main idea

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

Quantum simulation

Quantum simulation provided the original application of quantum computers…

… and remains their most compelling application, since it is

- relevant to many practical problems (chemistry, materials, high-energy physics, etc.) and
- close to the native behavior of quantum devices.

But quantum simulation remains challenging to implement. Qubit counts for classically hard instances are modest, but gate counts and circuit depths seem daunting.

We should try to to reduce resource requirements as much as possible…

… while retaining strong theoretical guarantees about quality of simulation.

The Hamiltonian simulation problem

Input: Hamiltonian H
Evolution time t
Allowed error ϵ
Unknown initial state $|\psi(0)\rangle$

Output: Final state $|\psi(t)\rangle = e^{-iHt}|\psi(0)\rangle$

Standard measure of error: worst-case ℓ_2 distance

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

simulation output state

simulated unitary operation

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

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 error at most ϵ.

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

Other methods (e.g., quantum signal processing) are better asymptotically, but product formulas are simple, require no ancilla qubits, and may perform better for modest-size problems.

[Childs, Su, Tran, Wiebe, Zhu 21; QIP 2020]

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; QIP 2022]

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 may 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: $U - \mathscr{U} = \sum_j E_j + \cdots$

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

Simulation error:

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

($\widetilde{\ }$ ← restricted to j,j'; $d_{j,j'}$ ← dimension of the j,j' subsystem)

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 product formula approximation* $\mathscr{U}$*,*

$$\|(U - \mathscr{U})\,|\psi\rangle\| = \mathcal{O}\left(\sqrt{\sum_{j,j'} \|E_j^\dagger E_{j'}\|\,\mathrm{tr}\,|\rho_{j,j'} - \mathbb{I}/d_{\rho_{j,j'}}|}\; t^{p+1} + \|E\|_F t^{p+1} + \alpha_{p+2} t^{p+2}\right).$$

where E_j *are the leading-order local error terms in the Trotter approximation,* $E = \sum_j E_j$*,* $\rho_{j,j'}$ *is the reduced state of* $|\psi\rangle$ *on the support of* $E_j^\dagger E_{j'}$*, and* α_{p+2} *is the next-leading-order spectral-norm commutator sum.*

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

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

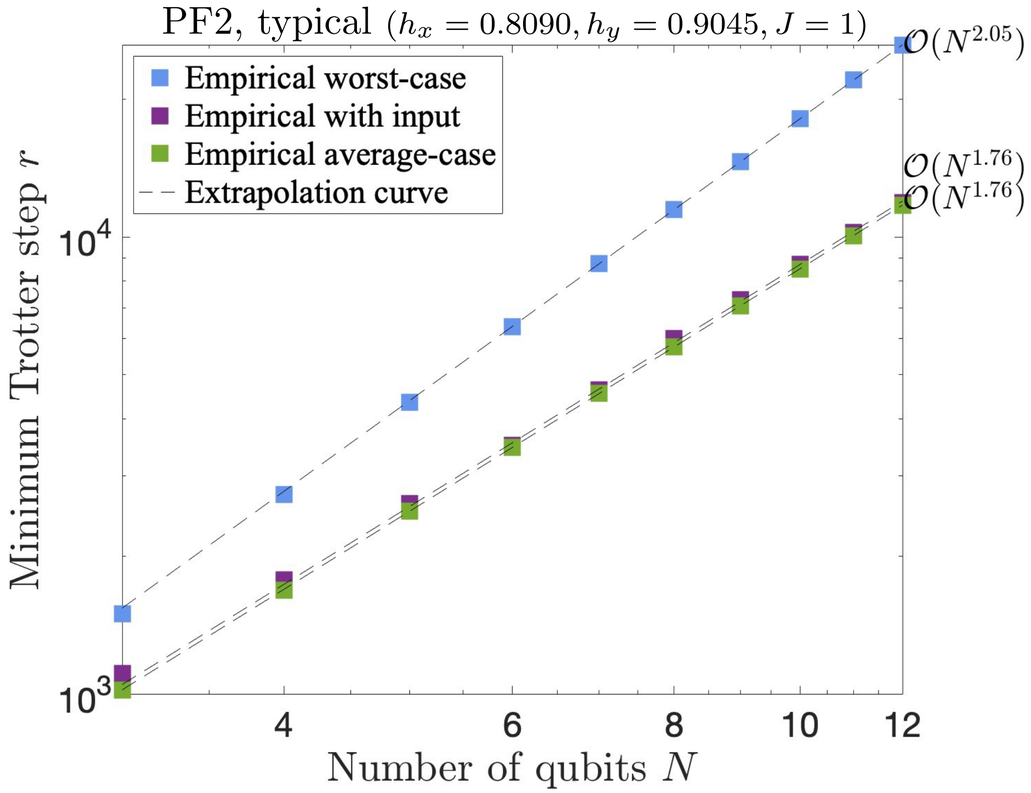

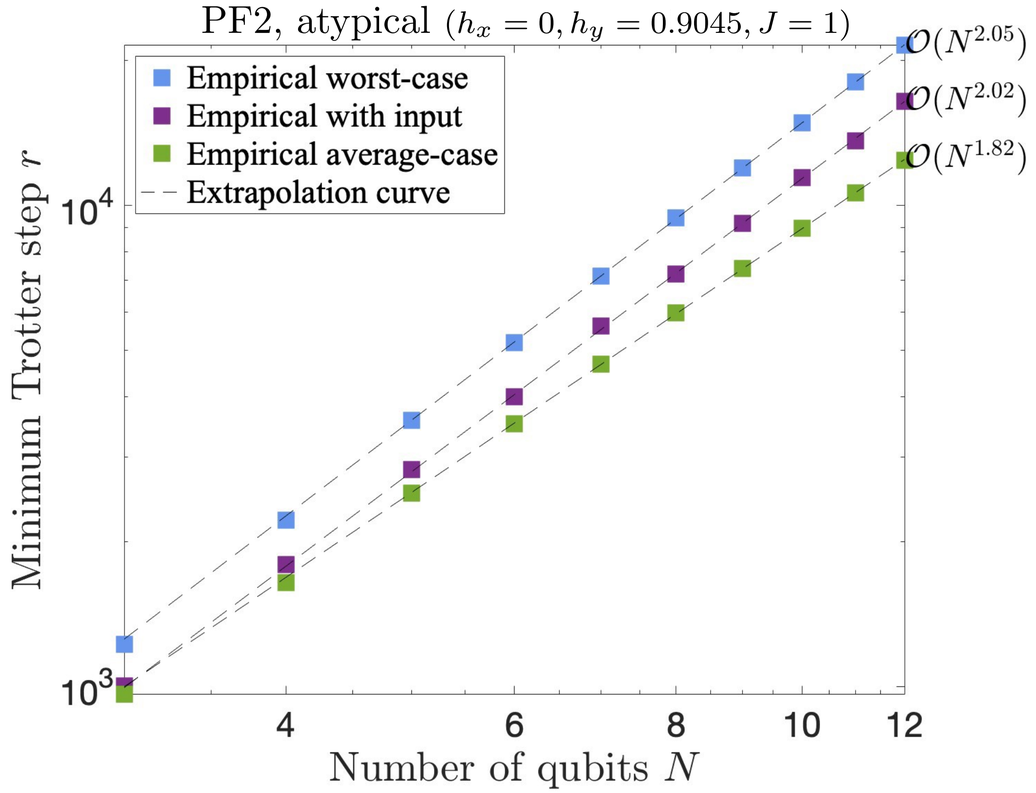

Theory prediction: worst case $O(N^2)$; average case $O(N^{1.75})$.

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 isn't 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?