Resource-efficient quantum matrix processing with commutator scaling

Rhik Mazumder, James Watson, Samson Wang

Introduction. A growing body of work explores an *early fault-tolerant* regime where one wishes to run an algorithm of real-world relevance on a qubit and depth-constrained quantum computer [Cam21; LT20; LT22; WBC22; Wan+23; Wan+25; WMB24; Kis+25]. Here, the motivation is not necessarily to achieve optimal algorithmic complexity overall; but instead to minimize the complexity of individual quantum circuit runs, in terms of quantum gate depth and number of qubits. We note that whilst nominally these algorithms are conceived to be run on error-corrected quantum computers, there is no a priori reason why they also cannot be run on noisy hardware with error mitigation [Blu+23].

Within the realm of quantum simulation, product formulae are already very amenable to running on current hardware and small-scale have been a mainstay of quantum computing experiments for many years. In addition to their conceptual and practical simplicity which allows this, Product formulae exhibit commutator scaling [Chi+21] which other (otherwise asymptotically superior) methods do not. We seek algorithms outside of dynamics which can incorporate the useful properties of product formulae, whilst simultaneously maintaining good asymptotic runtime guarantees and being amenable to running on early fault tolerant quantum computers.

Result 1: algorithm framework. In our work we present a matrix processing algorithm framework which can return the following quantities

$$\text{estimate } \mathrm{Tr}[\rho f(H)] \text{ and } \mathrm{Tr}[f(H)\rho f(H)^\dagger O] \text{ to additive precision } \varepsilon\,, \tag{1}$$

$$\text{return } \vec{v} \text{ such that } \|\vec{v}-\vec{p}\|_2 \leq \varepsilon \text{ for } p_i := |\langle i|\, f(H)\, |\psi\rangle|^2. \tag{2}$$

The quantum circuits that need to be run are very simple; we simply run Hadamard tests on product formulae unitaries corresponding to different time steps and total evolution times.

Instructive to understand our algorithm is to investigate the compilation of a time signal of the form $\mathrm{tr}[\rho e^{iHT}]$. We consider the matrix to be expressed as the sum of Γ terms as $H = \sum_{\gamma=1}^{\Gamma} H_\gamma$ such that time evolution of each H_γ is available in $O(1)$ time. It will be useful to reference the 1-norm of H which we denote as $\Gamma := \sum_{\gamma=1}^{\Gamma} \|H_\gamma\|$. Our analysis applies to any general staged product formula $\mathcal{P}(t)$, and we say that a formula is pth-order if it satisfies $\mathcal{P}(t) = e^{iHt} + O(t^{p+1})$. To simplify presentation we presume the number of stages is fixed. For instance, we can consider Trotter-Suzuki formulas of some fixed order p. We present a complexity comparison in Table 1.

We see that product formulae are unique that their dependence on matrix parameters consists of a quantity $(\alpha_{\mathrm{comm}}^{(p+1)})^{1/p}$. It is known that $(\alpha_{\mathrm{comm}}^{(p+1)})^{1/p} << \Lambda$ in many settings [Chi+21]. However, product formula have substantially worse dependence on the error parameter than all other settings. Qubitization instantiates time evolution in optimal circuit depth in the time parameter, but it is costly to implement in practice and requires substantially more ancillary qubits than other approaches.

Our gate complexities exponentially improve the error dependence of product formulae, at very modest increase in sample overhead. The complexities depend on a new quantity λ_{comm} which is a set of nested commutators which also appears in the algorithmic complexity of [WW24a;

Cha+25], and its explicit form is presented in our technical manuscript. Previously, its behavior was understood only in limited regimes. In particular, it is divergent when using best known bounds for k-local systems.

Result 2: refined gate complexities and commutator scaling. Our second contribution is to demonstrate refined gate complexities in a number of settings that exploit commutator scaling. These were all previously unknown for extrapolated product formulae circuits or not explicitly instantiated. We give results for

- k-local systems, subsuming with it models of power-law interactions.
- Scenarios where commutator scaling for standard product formula is only understood for a fixed order. Previous studies could only exploit commutator scaling when scaling is understood at every order.
- Interpolation of our scaling with randomized formula, exploiting the ideas of [Gün+25], which can be particularly advantageous for matrix ensembles with long tails.
- Settings where the input state is of well-defined Fermion number.

To showcase our refined gate complexities along with an application of our general framework in one go, we present complexity of our approach for a ground state energy estimation task in Table 2. Here our goal is to find the ground state energy of H, given efficient preparation of an ansatz state ρ with overlap with the ground space $\eta := \mathrm{tr}[\rho\Pi_0]$, where Π_0 denotes the projector to the ground space of H. We see here that commutator scaling is highly advantageous, and we inherit this scaling with a minor logarithmic overhead. Moreover, we exponentially improve the dependence on the gate complexity on η^{-1} compared to the approach with vanilla Trotter formulae.

method	ancillary qubits	max depth per sample $\widetilde{O}(\cdot)$	sample overhead $O(\cdot)$	commutator scaling?
Product formulae [Chi+21]	1	$\frac{\Gamma(\alpha_{\mathrm{comm}}^{(p+1)})^{1/p}\, T^{1+1/p}}{\varepsilon^{1/p}}$	$1/\varepsilon^2$	✓
Qubitization [LC19]	$\lceil\log(\Gamma)\rceil + 3$	$\Gamma\left(\Lambda T + \frac{\log(1/\epsilon)}{\log\log(1/\epsilon)}\right)$	$1/\varepsilon^2$	
Random compiler [WBC22]	1	$\Lambda^2 T^2$	$1/\varepsilon^2$	
Our work	1	$\Gamma\,(\lambda_{\mathrm{comm}} T)^{(1+1/p)}$	$\frac{(\log\log(1/\varepsilon))^2}{\varepsilon^2}$	✓

Table 1: Complexity comparison for the task of approximating a time signal $\mathrm{tr}[\rho e^{iHT}]$ via a Hadamard test given efficient preparation of ρ. In order to allow comparison with the random compiler of [WBC22] we have assumed that the Hamiltonian is given as a linear combination of Γ Pauli matrices. The quantity $\alpha_{\mathrm{comm}}^{(p+1)} \leq \frac{1}{2}(2\Lambda)^{p+1}$ is a sum of nested commutators.

method	ancillary qubits	max depth per sample $\widetilde{O}(\cdot)$	samples $\widetilde{O}(\cdot)$
QET-U + Qubitization [DLT22; LC17]	$\lceil \log(\Gamma) \rceil + 3$	$\Gamma\Lambda\varepsilon^{-1}$	$1/\eta$
QET-U + Trotter [DLT22]	1	$\Gamma\Lambda^{1/p}\eta^{1/p}\varepsilon^{-1-1/p}$	$1/\eta$
Random compiler [WBC22]	1	$\Lambda^2\varepsilon^{-2}$	$1/\eta^2$
Our work	1	$\Gamma(\Lambda^{1/p} + \log(\Lambda\varepsilon^{-1}))\varepsilon^{-1-1/p}$	$1/\eta^2$

Table 2: **Phase estimation complexity comparisonm k-local systems.** For Hamiltonian $H = \sum_{\ell=1}^{\Gamma} h_\ell$, $\Lambda = \sum_{\ell=1}^{\Gamma} \|h_\ell\|$, $\eta = \mathrm{tr}[\rho\Pi_0]$ the initial state overlap, defined in Lemma . We have assumed constant order p and constant per-site energy.

Comparison to prior art. Extrapolation was first introduced in [WW24b] for application to estimation time-evolved observables. Recent work [Cha+25] considers a general framework which can also instantiate Task 1 with the same gate complexity as ours. However, we give a full exposition of this problem in the matrix function picture in our work for completeness. Our innovation beyond prior art is two fold: first, we use randomization to give refined sample overheads that only incur polylog-logarithmic multiplicative overhead, and allows a novel task as defined in Task 2. Second, we provide novel understanding of commutator scaling in extrapolated circuits which applies also settings considered in prior art – this opens up gate complexities with commutator scaling for k-local systems and alternative refined scalings.

References

[Blu+23] Nick S Blunt, Laura Caune, Róbert Izsák, Earl T Campbell, and Nicole Holzmann. "Statistical phase estimation and error mitigation on a superconducting quantum processor". In: *PRX Quantum* 4.4 (2023). arXiv:2304.05126, p. 040341.

[Cam21] Earl T Campbell. "Early fault-tolerant simulations of the Hubbard model". In: *Quantum Science and Technology* 7.1 (2021). arXiv:2012.09238, p. 015007.

[Cha+25] Shantanav Chakraborty, Soumyabrata Hazra, Tongyang Li, Changpeng Shao, Xinzhao Wang, and Yuxin Zhang. "Quantum singular value transformation without block encodings: Near-optimal complexity with minimal ancilla". arXiv:2504.02385. 2025.

[Chi+21] Andrew M. Childs, Yuan Su, Minh C. Tran, Nathan Wiebe, and Shuchen Zhu. "Theory of Trotter Error with Commutator Scaling". In: *Phys. Rev. X* 11.1 (Feb. 2021). arXiv:1912.08854.

[DLT22] Yulong Dong, Lin Lin, and Yu Tong. "Ground-state preparation and energy estimation on early fault-tolerant quantum computers via quantum eigenvalue transformation of unitary matrices". In: *PRX Quantum* 3.4 (2022). arXiv:2204.05955, p. 040305.

[Gün+25] Jakob Günther, Freek Witteveen, Alexander Schmidhuber, Marek Miller, Matthias Christandl, and Aram Harrow. "Phase estimation with partially randomized time evolution". arXiv:2503.05647. 2025.

[Kis+25] Oriel Kiss, Utkarsh Azad, Borja Requena, Alessandro Roggero, David Wakeham, and Juan Miguel Arrazola. "Early fault-tolerant quantum algorithms in practice: Application to ground-state energy estimation". In: *Quantum* 9 (2025). arXiv:2405.03754, p. 1682.

[LC17] Guang Hao Low and Isaac L. Chuang. "Optimal Hamiltonian Simulation by Quantum Signal Processing". In: *Phys. Rev. Lett.* 118.1 (2017). arXiv:1606.02685, p. 010501.

[LC19] Guang Hao Low and Isaac L. Chuang. "Hamiltonian Simulation by Qubitization". In: *Quantum* 3 (2019). arXiv:1610.06546, p. 163.

[LT20] Lin Lin and Yu Tong. "Near-optimal ground state preparation". In: *Quantum* 4 (2020). arXiv:2002.12508, p. 372.

[LT22] Lin Lin and Yu Tong. "Heisenberg-Limited Ground-State Energy Estimation for Early Fault-Tolerant Quantum Computers". In: *PRX Quantum* 3.1 (2022). arXiv:2102.11340, p. 010318.

[Wan+23] Guoming Wang, Daniel Stilck França, Ruizhe Zhang, Shuchen Zhu, and Peter D Johnson. "Quantum algorithm for ground state energy estimation using circuit depth with exponentially improved dependence on precision". In: *Quantum* 7 (2023). arXiv:2209.06811, p. 1167.

[Wan+25] Guoming Wang, Daniel Stilck França, Gumaro Rendon, and Peter D Johnson. "Efficient ground-state-energy estimation and certification on early fault-tolerant quantum computers". In: *Physical Review A* 111.1 (2025). arXiv:2304.09827, p. 012426.

[WBC22] Kianna Wan, Mario Berta, and Earl T. Campbell. "Randomized Quantum Algorithm for Statistical Phase Estimation". In: *Phys. Rev. Lett.* 129.3 (2022). arXiv:2110.12071, p. 030503.

[WMB24] Samson Wang, Sam McArdle, and Mario Berta. "Qubit-efficient randomized quantum algorithms for linear algebra". In: *PRX Quantum* 5.2 (2024). arXiv:2302.01873, p. 020324.

[WW24a] James D Watson and Jacob Watkins. "Exponentially Reduced Circuit Depths Using Trotter Error Mitigation". arXiv:2408.14385. 2024.

[WW24b] James D Watson and Jacob Watkins. "Exponentially Reduced Circuit Depths Using Trotter Error Mitigation". In: *arXiv preprint arXiv:2408.14385* (2024).