Randomized Quantum Algorithm for Statistical Phase Estimation

Physical Review Letters 129, 030503 (2022)

with Kianna Wan (Stanford) and Earl Campbell (Riverlane)

Mario Berta

RWTH AACHEN UNIVERSITY

Problem: Ground state energy estimation

- Given n-qubit Hamiltonian

$$\mathrm{H} := \sum_{l=1}^{L} \alpha_l P_l \text{ with } P_l \text{ } n\text{-qubit Paulis}$$

and one-norm $\lambda := \sum_{l=1}^{L} |\alpha_l|$, together with efficiently preparable n-qubit ansatz state ρ with overlap

$$\langle \phi_0 | \rho | \phi_0 \rangle \geq \eta > 0$$

for ground state $|\phi_0\rangle\langle\phi_0|$ with energy E_0

- Goal: Compute estimate $\tilde{E}_0$ with precision $\left|\tilde{E}_0 - E_0\right| \leq \Delta$

Goals for early fault-tolerance scheme

1. Minimize number of qubits needed – only one ancilla

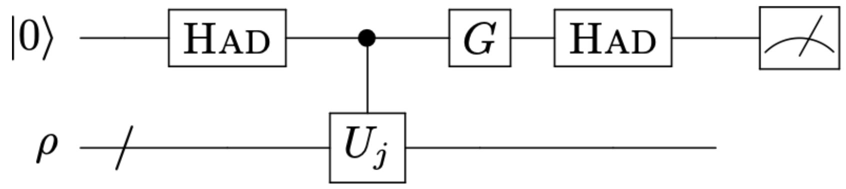

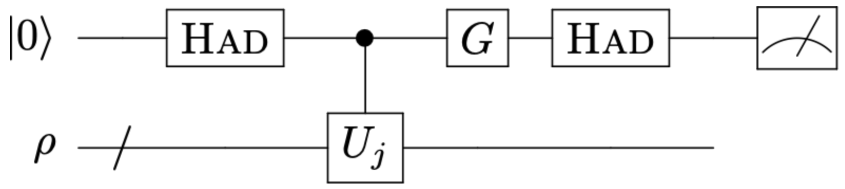

2. Independent of the number L of Pauli terms in H – instead, depending on one-norm $\lambda \leq L$
3. Trade-off gate versus sample complexity
4. Decrease error by solely taking more samples

Main Result

Algorithm ground state energy estimation

- Output $\tilde{E}_0$ with $|\tilde{E}_0 - E_0| \leq \Delta$ with probability $1 - \xi$ by employing

$$C_{sample} = \tilde{O}(\eta^{-2}) \quad \left[= O\left(\eta^{-2}\log^2(\lambda\Delta^{-1}\log(\eta^{-1}))\log(\xi^{-1}\log(\lambda\Delta^{-1}))\right)\right]$$

quantum circuits on $n + 1$ qubits, each using one copy of ρ and

$$C_{gate} = \tilde{O}(\lambda^2\Delta^{-2}) \quad \left[= O(\lambda^2\Delta^{-2}\log^2(\eta^{-1}))\right]$$

single-qubit Pauli rotations $\exp(i\theta P_l)$

- Plus: Clifford gates – generated by CNOT, H, and S (Paulis)

Complexity ground state energy estimation

- n qubit Hamiltonian, $n+1$ qubits with quantum complexities independent of L:

$$C_{gate} = \tilde{O}(\lambda^2\Delta^{-2}) \text{ for } C_{sample} = \tilde{O}(\eta^{-2})$$

- Randomized algorithm with classical pre- and post-processing
- Comparison state-of-the-art qubitization based approach:

Gate complexity $\tilde{O}\left(\sqrt{L}\lambda\Delta^{-1}\right)$ for $\tilde{O}\left(\sqrt{L}\right)$ qubits → total $\tilde{O}(L\lambda\Delta^{-1})$

[Lee et al., PRX Quantum (2021)]

Basic idea

- Cumulative distribution function (CDF) relative to ρ:

$$C(x) \coloneqq Tr[\rho\Pi_{\leq x}]$$

- Evaluate $C(x)$ from quantum routine?
- Eigenvalue thresholding
- Give ground state energy estimate $\tilde{E}_0$ via binary search

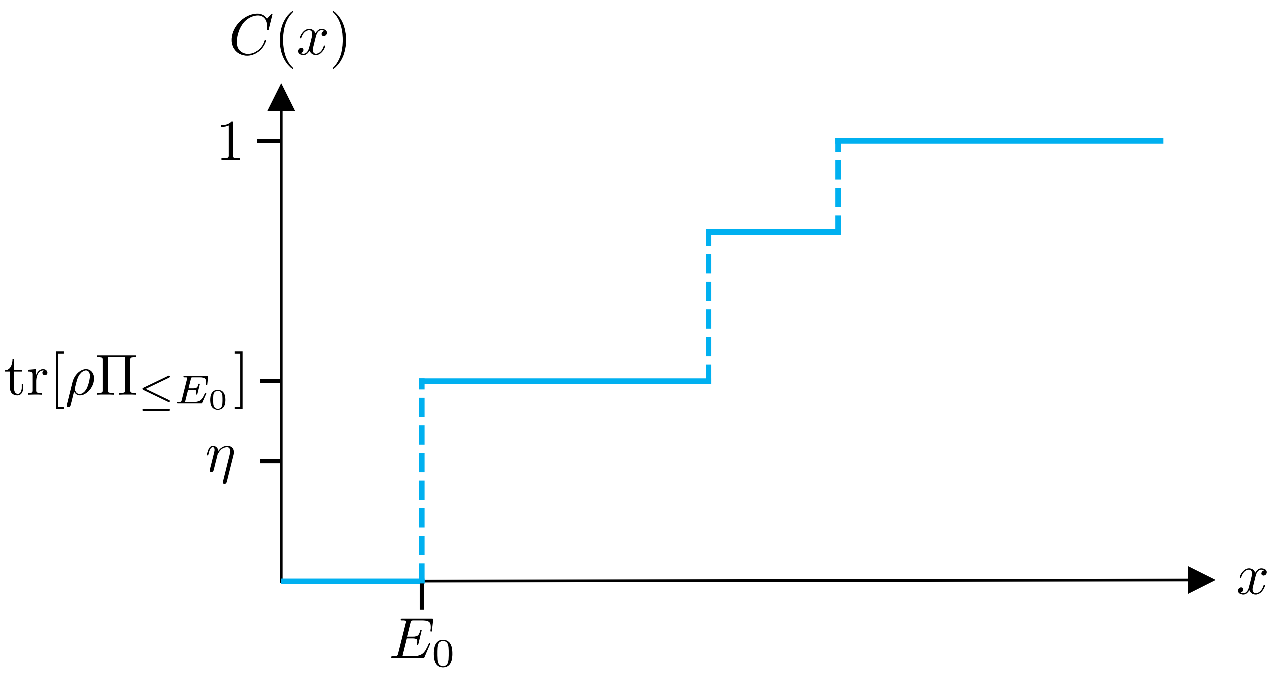

[Lin & Tong, PRX Quantum (2022)]
[Martyn et al., PRX Quantum (2021)]

Quantum routine to evaluate CDF

Workhorse A: Hadamard test

- Input: n-qubit state ρ together with n-qubit unitary U
- Circuit:

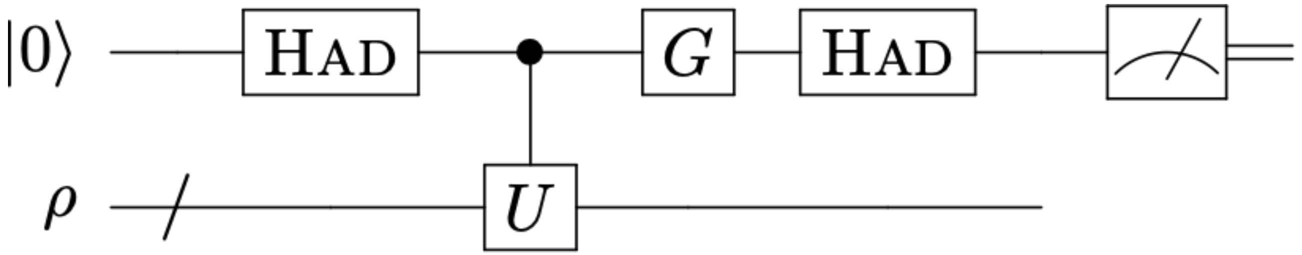

$$X = \begin{cases} +1 & \text{if } |0\rangle \\ -1 & \text{if } |1\rangle \end{cases}$$

- Output: unbiased estimate of $Tr[\rho U]$ from

$$G = I \quad \Rightarrow \quad \mathbb{E}[X] = \mathrm{Re}(\mathrm{tr}[\rho U])$$
$$G = S^\dagger \quad \Rightarrow \quad \mathbb{E}[X] = \mathrm{Im}(\mathrm{tr}[\rho U])$$

Workhorse B: Importance sampling

- Estimate linear combination $\sum_j a_j Tr[\rho U_j]$ for unitaries U_j with $a_j > 0$ and normalization $A := \sum_j a_j$
- Sample j with probability $a_j \cdot A^{-1}$ and perform Hadamard test on (ρ, U_j):

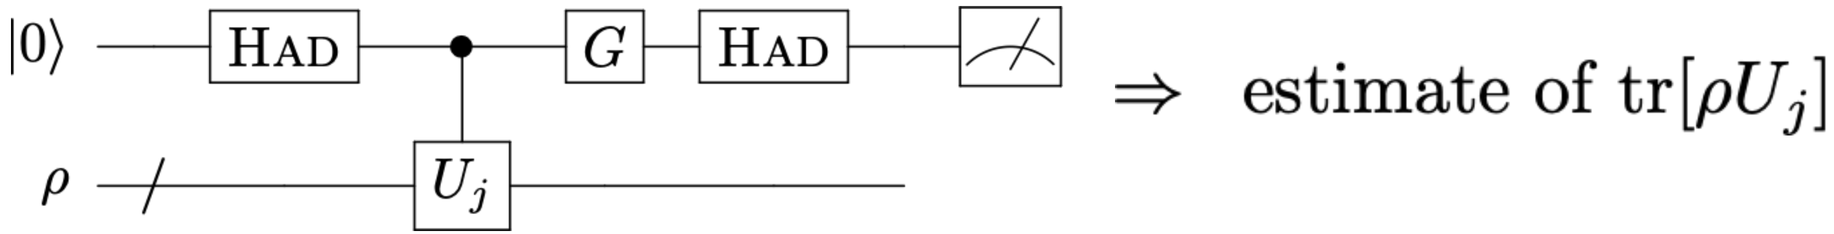

- Take average of samples, number of required is $[A^2\sigma^{-2}]$ for variance $\sigma > 0$
- Expected gate complexity becomes $A^{-1} \cdot \sum_j a_j COST(C - U_j)$

Towards quantum implementation of CDF

- Normalize Hamiltonian with $c \cdot ||H||_\infty \leq c \cdot \lambda$ to put spectrum in $\left[-\frac{\pi}{2}, +\frac{\pi}{2}\right]$
- CDF $C(x) \equiv Tr[\rho \Pi_{\leq x}] = (\Theta * p)(x)$ from convolution with Heaviside $\Theta(x)$:

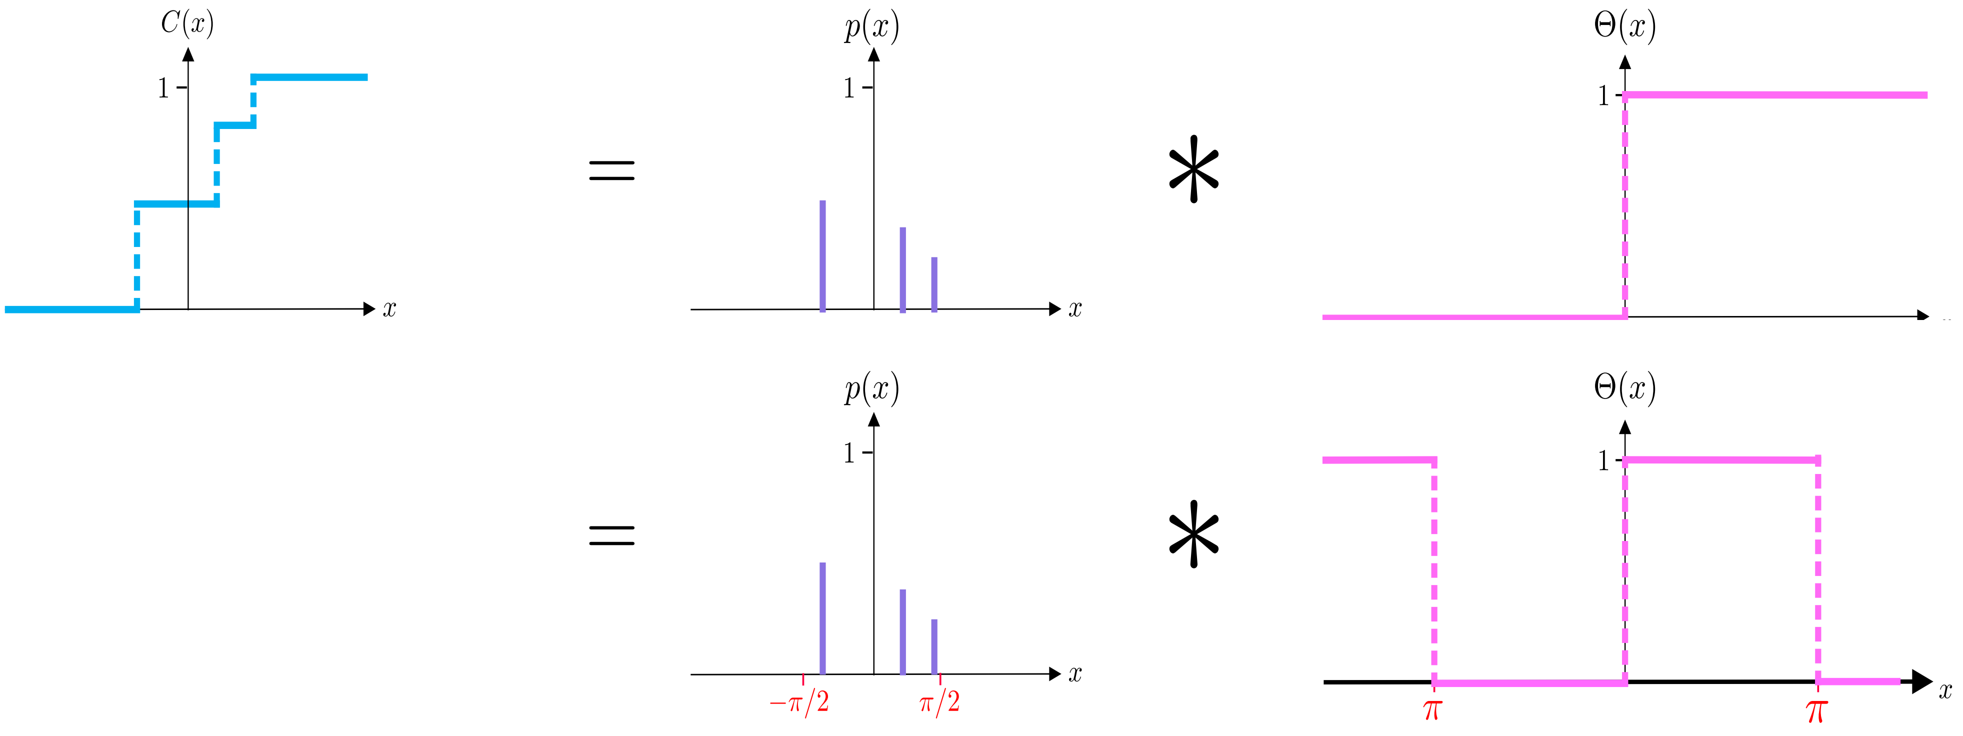

CDF via Fourier series

- Replace Heaviside $\Theta(x)$ by finite Fourier series $F(x) \coloneqq \sum_{j\in S} \hat{F}_j e^{ijx}$
- Approximate CDF:

$$C(x) \approx (p * F)(x) = \sum_{j\in S} \hat{F}_j e^{ijx} \cdot Tr\left[\rho e^{it_j H}\right]$$

with runtimes $t_j = j \times$ normalization

- Hadamard test + importance sampling + Hamiltonian simulation:

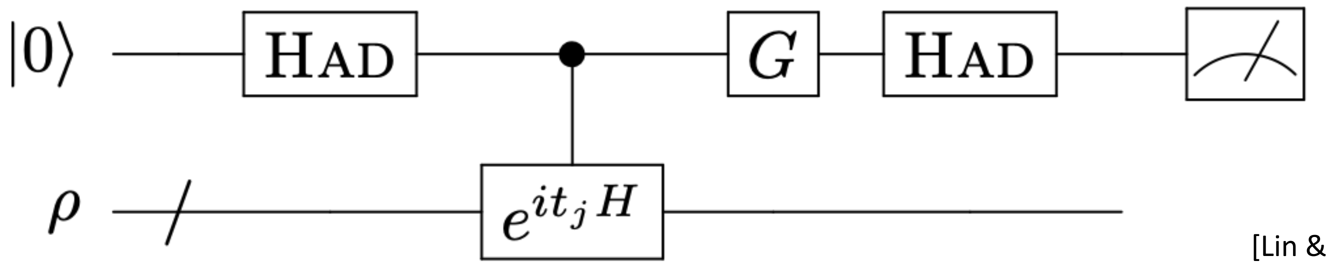

[Lin & Tong, PRX Quantum (2022)]

Fourier series lemma (Heaviside function)

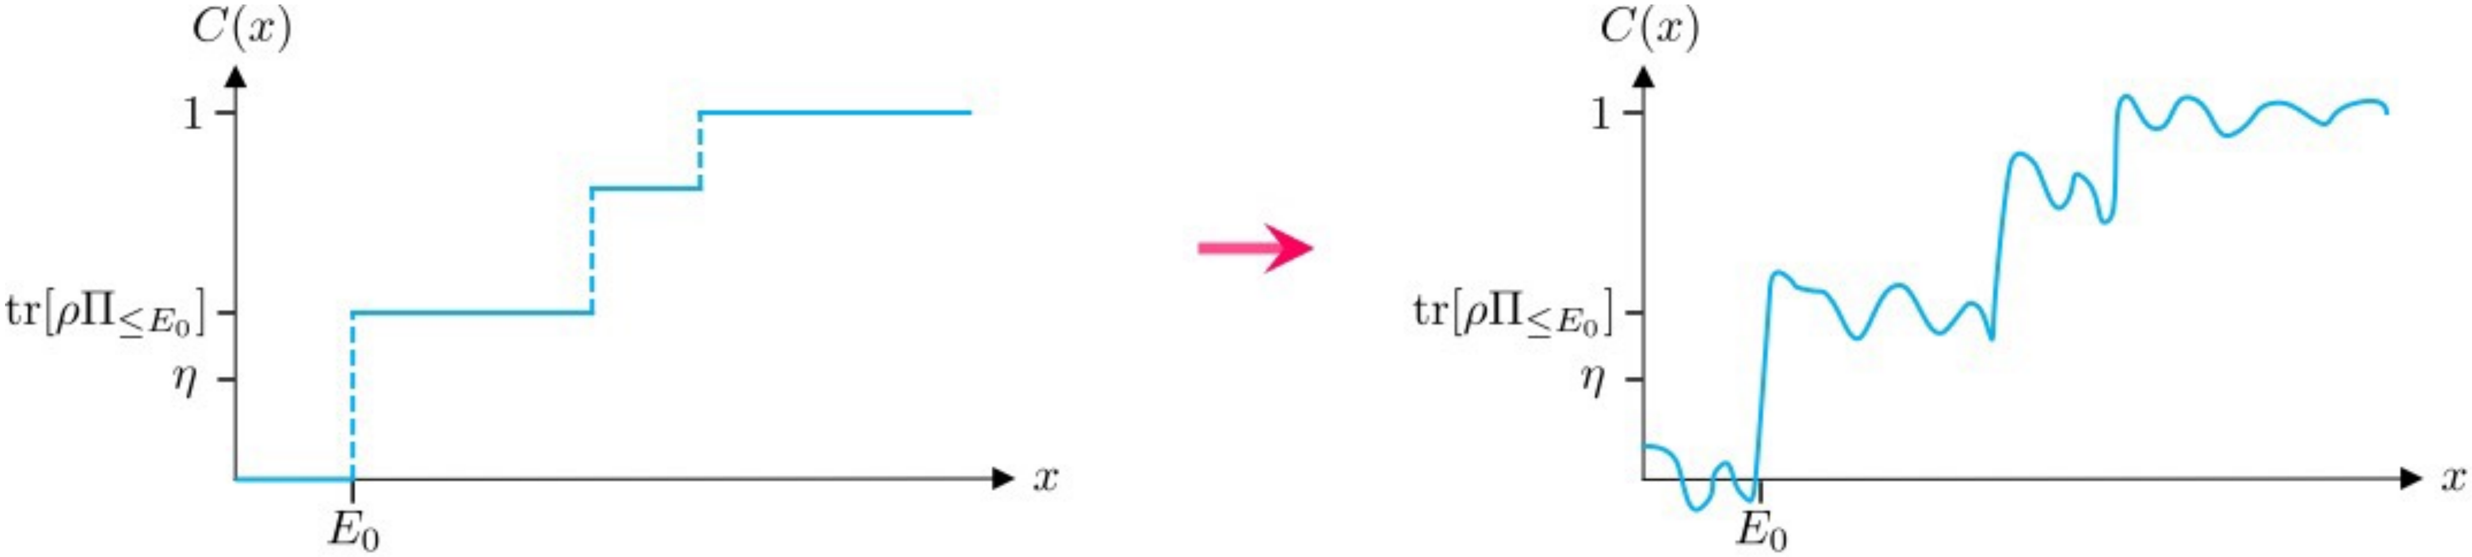

- Improved Fourier series approximation of Heaviside function
- Technical contribution:

Gate complexity for precision $\Delta > 0$ from $O\left(\Delta^{-2}\log^2(\Delta^{-1})\right)$ to $O(\Delta^{-2})$

[Lin & Tong, PRX Quantum (2022)]

Hadamard test on Fourier series

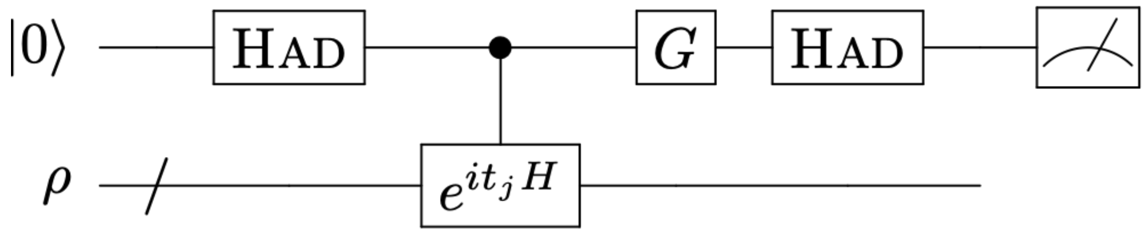

$$C(x) \approx \sum_{j\in S} \hat{F}_j e^{ijx} \cdot Tr\left[\rho e^{it_j H}\right]$$

- Implement Hamiltonian simulation unitary $\mathrm{U_j} = e^{it_j H}$ for $\mathrm{H} = \sum_{l=1}^{L} \alpha_l P_l$
- Independent of L? Technical contribution:

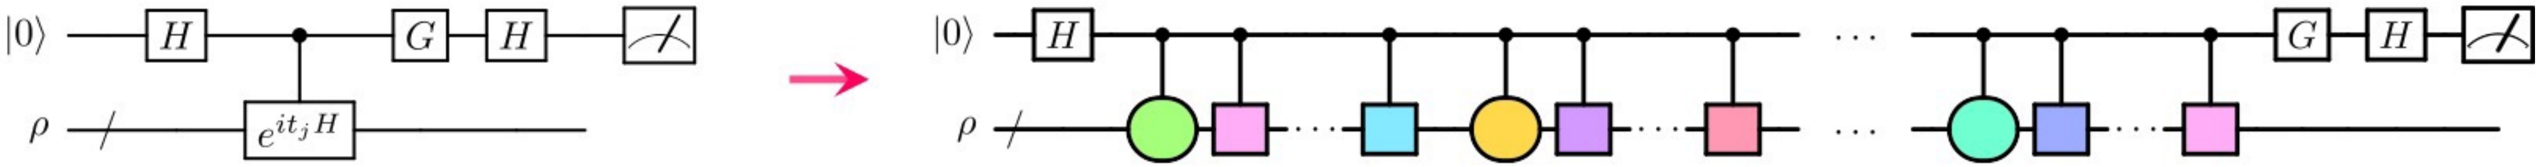

novel random compiler lemma (Hamiltonian simulation)

Versus previous random compiler:
[Campbell, PRL (2019)]

Random compiler lemma (Hamiltonian simulation)

- For e^{itH} with $\mathrm{H} = \sum_{l=1}^{L} \alpha_l P_l$, we give linear combination of unitaries (LCU) $e^{itH} = \sum_k b_k U_k$ such that:
  - I. $\mu(r) := \sum_k b_k \leq \exp(t^2 r^{-1})$
  - II. $COST(C - U_k) = r$ controlled single qubit Pauli rotations $\forall k$

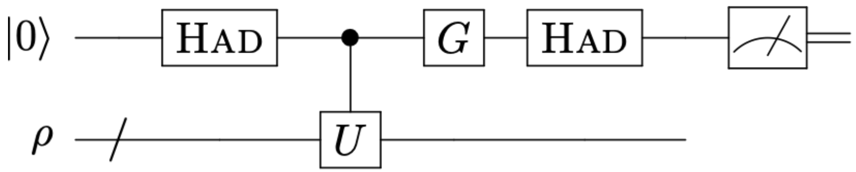

- Gate complexity r versus sample complexity $\exp(t^2 r^{-1})$

Versus previous LCU methods:
[Berry et al., PRL (2015)]

- Example: $r = 2t^2 \rightarrow \mu \leq \sqrt{e}$ and $COST(C - U_k) = 2t^2$
- Use this on: $C(x) \approx \sum_{j \in S} \hat{F}_j e^{ijx} \cdot Tr\left[\rho e^{it_j H}\right]$

Random compiler for CDF

- CDF $C(x) \approx \sum_j \hat{F}_j e^{ijx} \cdot Tr\left[\rho e^{it_j H}\right]$ becomes $C(x) \approx \sum_j \sum_k \hat{F}_j e^{ijx} b_k^{(j)} Tr\left[\rho U_k^{(j)}\right]$

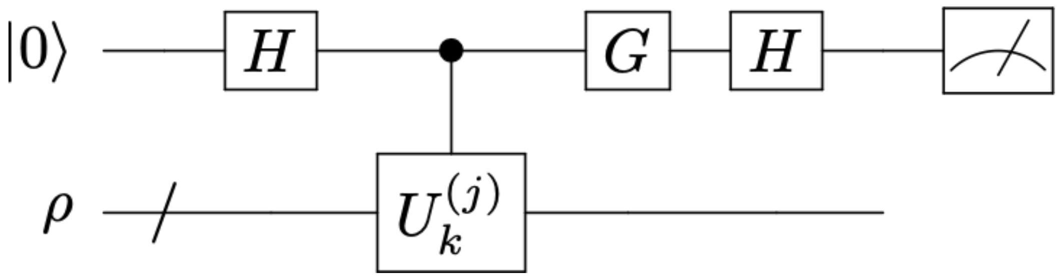

- $e^{it_j H} = \sum_k b_k^{(j)} U_k^{(j)}$ decomposition for runtime vector $\vec{r} = (r_j)_j \in \mathbb{N}^{|S|}$ as:

  I. $\mu_j \coloneqq \mu_j(r) \coloneqq \sum_k b_k^{(j)} \leq \exp\left(t_j^2 r_j^{-1}\right)$

  II. $COST\left(C - U_k^{(j)}\right) = r_j$

Putting things together

- CDF decomposition $C(x) \approx \sum_j \sum_k \hat{F}_j e^{ijx} b_k^{(j)} Tr\left[\rho U_k^{(j)}\right]$
- $C_{gate} = \left(\sum_{i\in S} |\hat{F}_i|\mu_i\right)^{-1} \cdot \left(\sum_{j\in S} |\hat{F}_j|\, \mu_j r_j\right)$
- $C_{sample} \propto \left(\sum_{j\in s} |\hat{F}_j|\mu_j\right)^2$

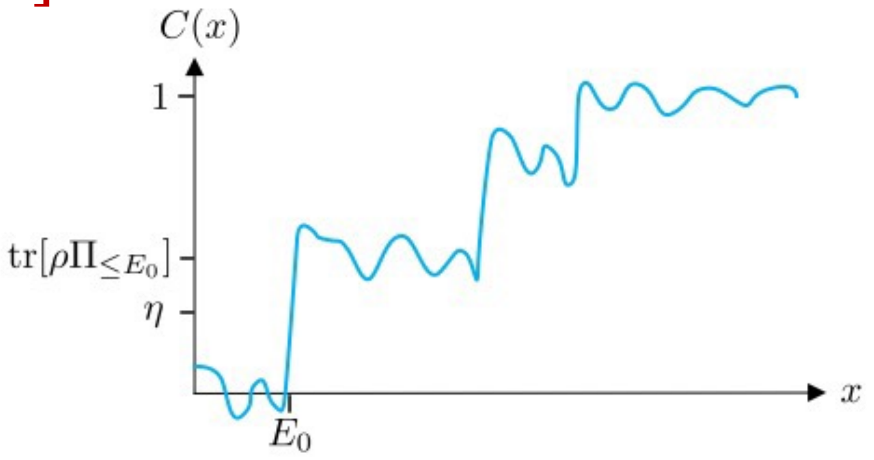

- As $\mu_j \leq e^{t_j^2 r_j^{-1}}$ choosing $r_j = 2t_j^2$ $\forall j$ gives $\mu_j \leq \sqrt{e}$:

$$C_{gate} \propto \left(\sum_{i\in S} |\hat{F}_i|\right)^{-1} \left(\sum_{j\in S} |\hat{F}_j|\, j^2\right) \rightarrow C_{gate} = \tilde{O}(\lambda^2 \Delta^{-2})$$

$$C_{sample} \propto \left(\sum_{j\in s} |\hat{F}_j|\right)^2 \rightarrow C_{sample} = \tilde{O}(\eta^{-2})$$

Example systems

Finite size numerical analysis

- Asymptotic complexity from fixed runtime vector $\vec{r}$ with $r_j = 2t_j^2 \; \forall j \in S$
- Optimize $\vec{r}$ to minimize C_{gate}, C_{sample}, or $C_{gate} \cdot C_{sample}$ for different settings?
- High-dimensional optimization problem, technical contribution: approximate dimension reduction that allows for efficient classical pre-processing
- Leads to flexible resource trade-offs:

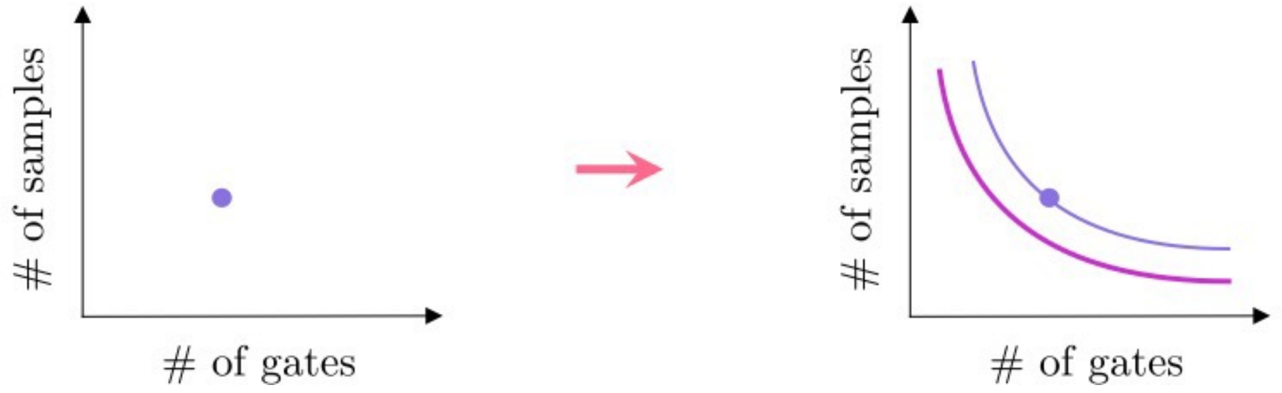

FeMoco benchmark

- Li et al. FeMoco Hamiltonian with 152 spin orbitals: 152+1=153 qubits
- Chemical accuracy $\Delta = 0.0016$ Hartree, one-norm $\lambda = 1511$
- Gate complexity in single-qubit Pauli rotations $e^{i\theta P_l}$
- T gate or Toffoli-gate complexity similar
- Qubitization using heuristic truncations: [Lee et al., PRX Quantum (2021)]

$C_{gate} = 3.2 \cdot 10^{10}$ on 2196 qubits

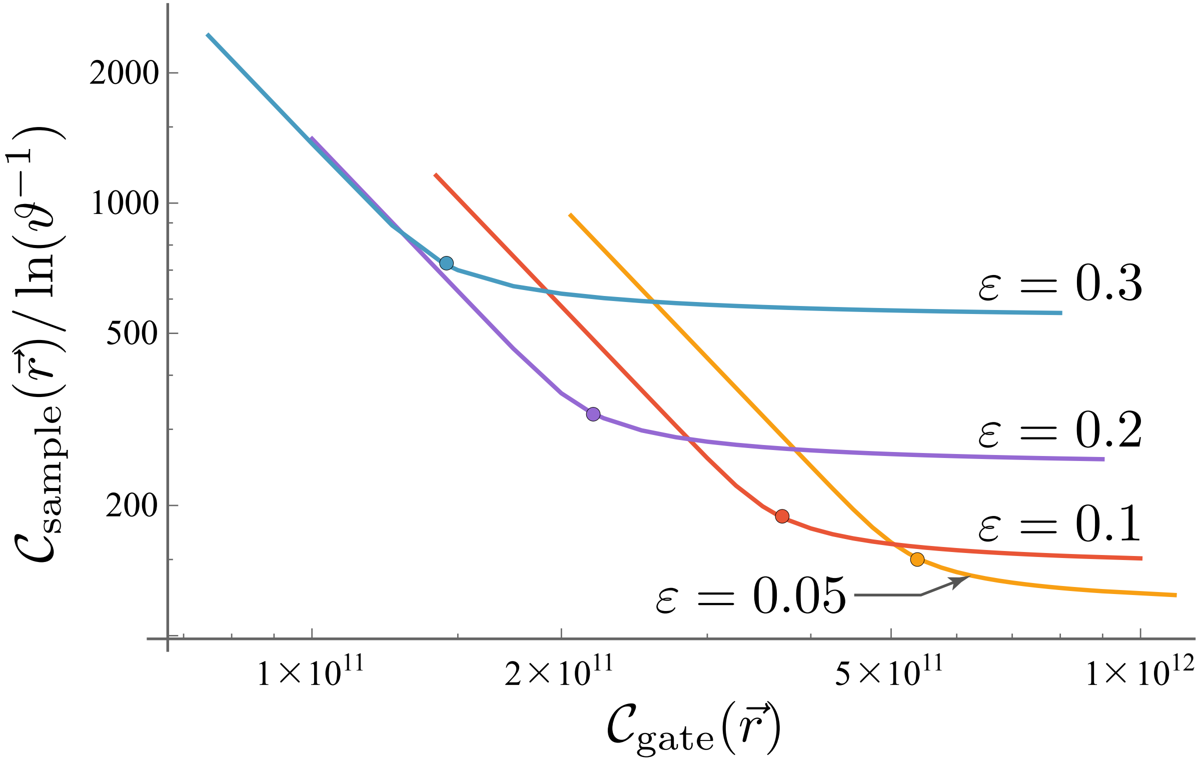

[Koridon et al., PRR (2021)]

Hydrogen chains benchmark

- For length N chain, one-norm estimate $\lambda \approx O(N^{1.34})$ [Koridon et al., PRR (2021)]
- Our work $C_{gate} = \tilde{O}(N^{2.68}\Delta^{-2})$
- Qubitization based approaches:

  A. rigorous $C_{gate} = \tilde{O}(N^{3.34}\Delta^{-1})$

  B. sparse method $C_{gate} = \tilde{O}(N^{2.3}\Delta^{-1})$ [Berry et al., Quantum (2019)]

  C. tensor hypercontraction method $C_{gate} = \tilde{O}(N^{2.1}\Delta^{-1})$ [Lee et al., PRX Quantum (2021)]

- Extensive properties $\Delta \propto N$ interesting for our methods: $C_{gate} = \tilde{O}(N^{0.68})$

Conclusion

Recap main result

- Given: n-qubit Hamiltonian $\mathrm{H} = \sum_{\mathrm{l}=1}^{\mathrm{L}} \alpha_{\mathrm{l}} \mathrm{P}_{\mathrm{l}}$ with $\lambda = \sum_{\mathrm{l}=1}^{\mathrm{L}} |\alpha_{\mathrm{l}}|$, plus ansatz state ρ with ground state overlap $\langle \phi_0 | \rho | \phi_0 \rangle \geq \eta > 0$
- Output: ground state energy estimate $\tilde{E}_0$ with $\left|\tilde{E}_0 - E_0\right| \leq \Delta$
- Result: $n + 1$ qubits, $C_{gate} = \tilde{O}(\lambda^2 \Delta^{-2})$, $C_{sample} = \tilde{O}(\eta^{-2})$
- Advantages:

I. rigorous estimates

II. only depends on $\lambda \leq L$

III. only uses one ancilla

IV. flexible trade-off gate versus sample complexity

V. decrease error by solely taking more samples → still state preparation bottleneck!

Extension: General matrix arithmetic

- General matrices A, instead of Hamiltonians H
- General functions $f(x)$ such as, e.g., x^{-1}, instead of Heaviside $\theta(x)$
- Goal to outperform (probabilistic) classical methods with early fault-tolerance
- Quantum singular value transformation (QSVT): $||A||_F$ or $s(A) \cdot ||A||_{max}$ [Gilyen et al., STOC (2019)]

- Qubit-efficient randomized quantum algorithms for linear algebra, Wang, McArdle, B., arXiv:2302.01873 (2023)
  - $\mathrm{A} = \sum_{l=1}^{L} \alpha_l P_l$ Paulis with $\lambda = \sum_{l=1}^{L} |\alpha_l|$, gives λ^2 complexity (input model!)
  - no QRAM needed

Thank you

RWTH AACHEN UNIVERSITY

- Randomized quantum algorithm for statistical phase estimation, Wan, B., Campbell, Physical Review Letters 129, 030503 (2022)
- Efficient randomized quantum algorithms for linear algebra, Wang, McArdle, B., arXiv:2302.01873 (2023)

Extra slides

Extra: Proof Fourier series lemma

- Rigorous argument via truncated Chebyshev series of rescaled error function:

$$\mathrm{erf}(\beta y) = 2\pi^{-\frac{1}{2}} \int_0^{\beta y} e^{-t^2} dt \approx \sum_k c_k T_k(y)$$

[Low & Chuang, arxiv:1707.05391 (2017)]

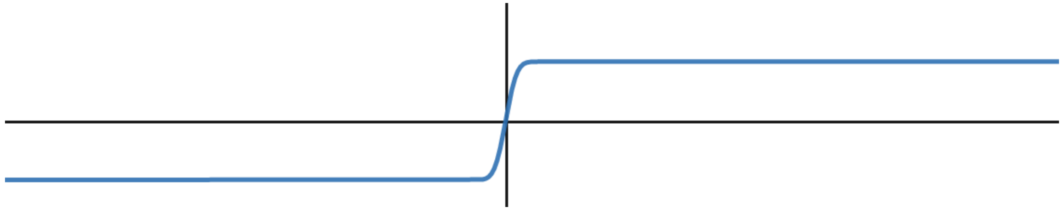

- Fourier series: $\Theta(x) \approx \mathrm{erf}(\beta \sin(x)) \approx \sum_k c_k T_k\left(\cos\left(\frac{\pi}{2} - x\right)\right)$

using $T_k(\cos(\cdot)) = \cos(k(\cdot))$

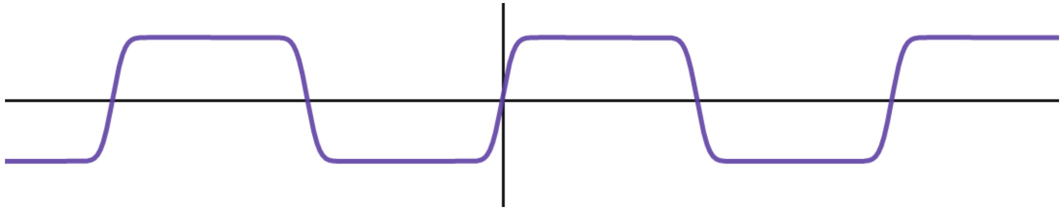

Extra: Proof random compiler lemma

- For $\mathrm{H} = \sum_{l=1}^{L} \alpha_l P_l$ and $r \in \mathbb{N}$: $e^{iHt} = \left(e^{iHtr^{-1}}\right)^r = (1 + itr^{-1}H + \cdots)^r$

$$1 + itr^{-1}H = \sum_{l=1}^{L} p_l(1 + itr^{-1}P_l) \propto \sum_{l=1}^{L} p_l e^{i\theta P_l} \text{ for } \theta = \arccos\left(\sqrt{1 + t^2 r^{-2}}\right)$$

- Similarly handle higher order terms – contain Paulis as well

- To sample U_k from $e^{iHt} = \sum_k b_k U_k$: independently sample r unitaries $W_1, \ldots, W_r$ from decomposition of $e^{iHtr^{-1}}$ and implement product

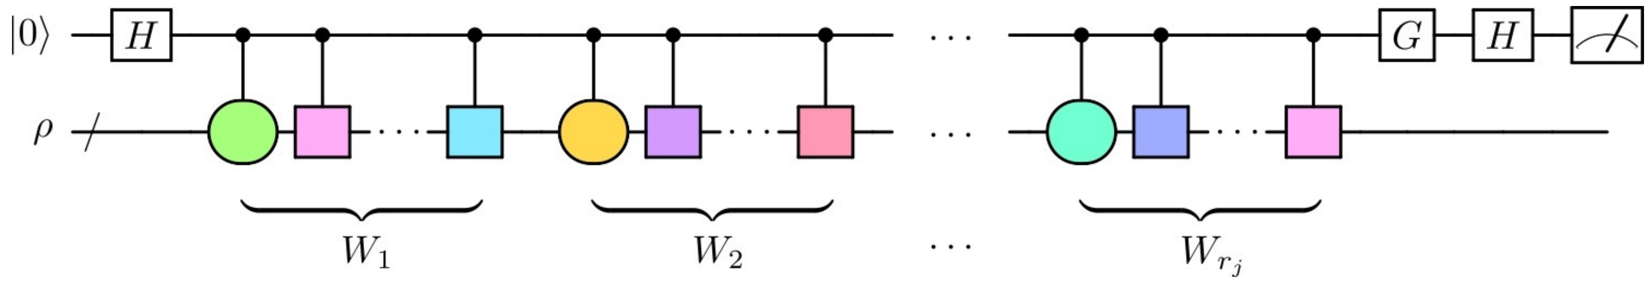

Extra: qDRIFT comparison

[Campbell, PRL (2019)]

- qDRIFT approximates quantum channel

$$\rho \mapsto e^{iHt} \rho e^{-iHt} \text{ for } H = \sum_{l=1}^{L} p_l P_l \text{ (normalized)}$$

by sampling r Paulis $P_{l_1}, \ldots, P_{l_r}$ independently with $\Pr[P_l] = p_l$ and putting

$$\mathrm{V} := e^{itr^{-1} P_{l_1}} \cdots e^{itr^{-1} P_{l_r}}$$

- qDRIFT compilation error can only be suppressed by increasing gate count r
- Our random compiler: approximates unitary $\mathrm{U} = e^{iHt}$ and compilation error can be suppressed arbitrarily by simply taking more samples