Quantum algorithms for the early fault-tolerance regime

Prof. Dr. Mario Berta

Physics Colloquium May 29th, 2023
Technology Innovation Institute Abu Dhabi

Quantum information science

- Theory of information processing: Mathematical foundations by Turing, Shannon, etc.
- Abstract theory independent of underlying physics?
- Physics changes at different length scales (energies), notion of information for microscopical systems described by quantum physics?
- Deep finding: Quantum information $\neq$ classical information!
- Led to whole new research area of quantum technologies and quantum computing

RWTH Aachen University: Berta group

- Theory of quantum information science
- Institute for Quantum Information RWTH Aachen University
- Members:

Mario Berta
Professor of Physics

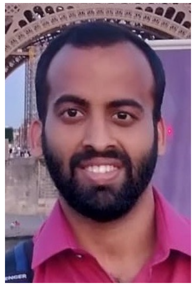

Sreejith Sreekumar
Postdoc

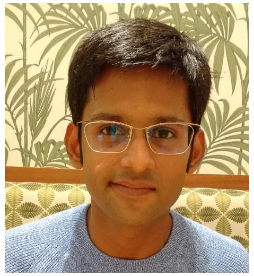

Aditya Nema
Postdoc

Tobias Rippchen
PhD student

Julius Zeiss
PhD student

Gereon Koßmann
PhD student

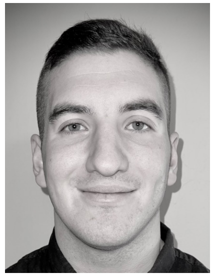

Brandon Augustino
Postdoc (Imperial)

Navneeth Ramakrishnan
PhD student (Imperial)

Samson Wang
PhD student (Imperial)

+ 4x Master students Physics / Computer Science

+3 more postdocs incoming

Theory of quantum information science

- Our focus areas:
  1. Mathematical foundations of quantum information
  2. Quantum algorithm development
- Visiting Reader at Department of Computing Imperial College London
- Industry ties with Amazon Web Services Center for Quantum Computing

JARA INSTITUTE QUANTUM INFORMATION
AN INITIATIVE OF
RWTH AACHEN UNIVERSITY
JÜLICH Forschungszentrum
Imperial College London
EPSRC
aws
erc European Research Council
ML4Q MATTER AND LIGHT FOR QUANTUM COMPUTING

This talk: Quantum algorithm development

Quantum algorithms

- Early ideas by Feynman and others on quantum simulation in the 1980s
- Query complexity separation results in the circuit model in the 1990s
- Peter Shor (1999) breakthrough result:

  n-bit integer factorization in quantum complexity $O(n^2 \log n)$ versus the classical complexity $O\left(\exp\left(1.9 \cdot n^{\frac{1}{3}} (\log n)^{\frac{2}{3}}\right)\right)$
- Steady progress on quantum algorithm development since, recent flurry of activities and results
- Goal: Quantify classical-quantum complexity boundary

Classical versus quantum technologies

- Basic question from complexity-theoretic viewpoint:

  Do algorithms based on quantum components, such as quantum processing units (QPU) / quantum memory / quantum random access memory (QRAM) etc., provide computational advantages compared to leading methods based on classical components?

- Goal is to identify use cases / areas of applications with
  - large (super-quadratic) quantum speed-up
  - minimal quantum footprint, i.e., use classical whenever possible

Regimes for quantum algorithm design

- Nascent state of quantum technologies gives noisy and intermediate scale quantum (NISQ) computing, i.e.,
  - NISQ analogue simulators, not universal, not fully programmable
  - NISQ digital quantum circuits, inbuilt noise resilience, error mitigation, severe scaling limitations, etc.
- For NISQ regime rigorous guarantees and scaling questions are challenging
- In contrast, fully quantum error-corrected and scaling quantum computer
- Any intermediate regimes of interest?

Early fault-tolerant regime

- Run many small circuits
- Leverage some classical power
- Usually no runtime/success guarantees

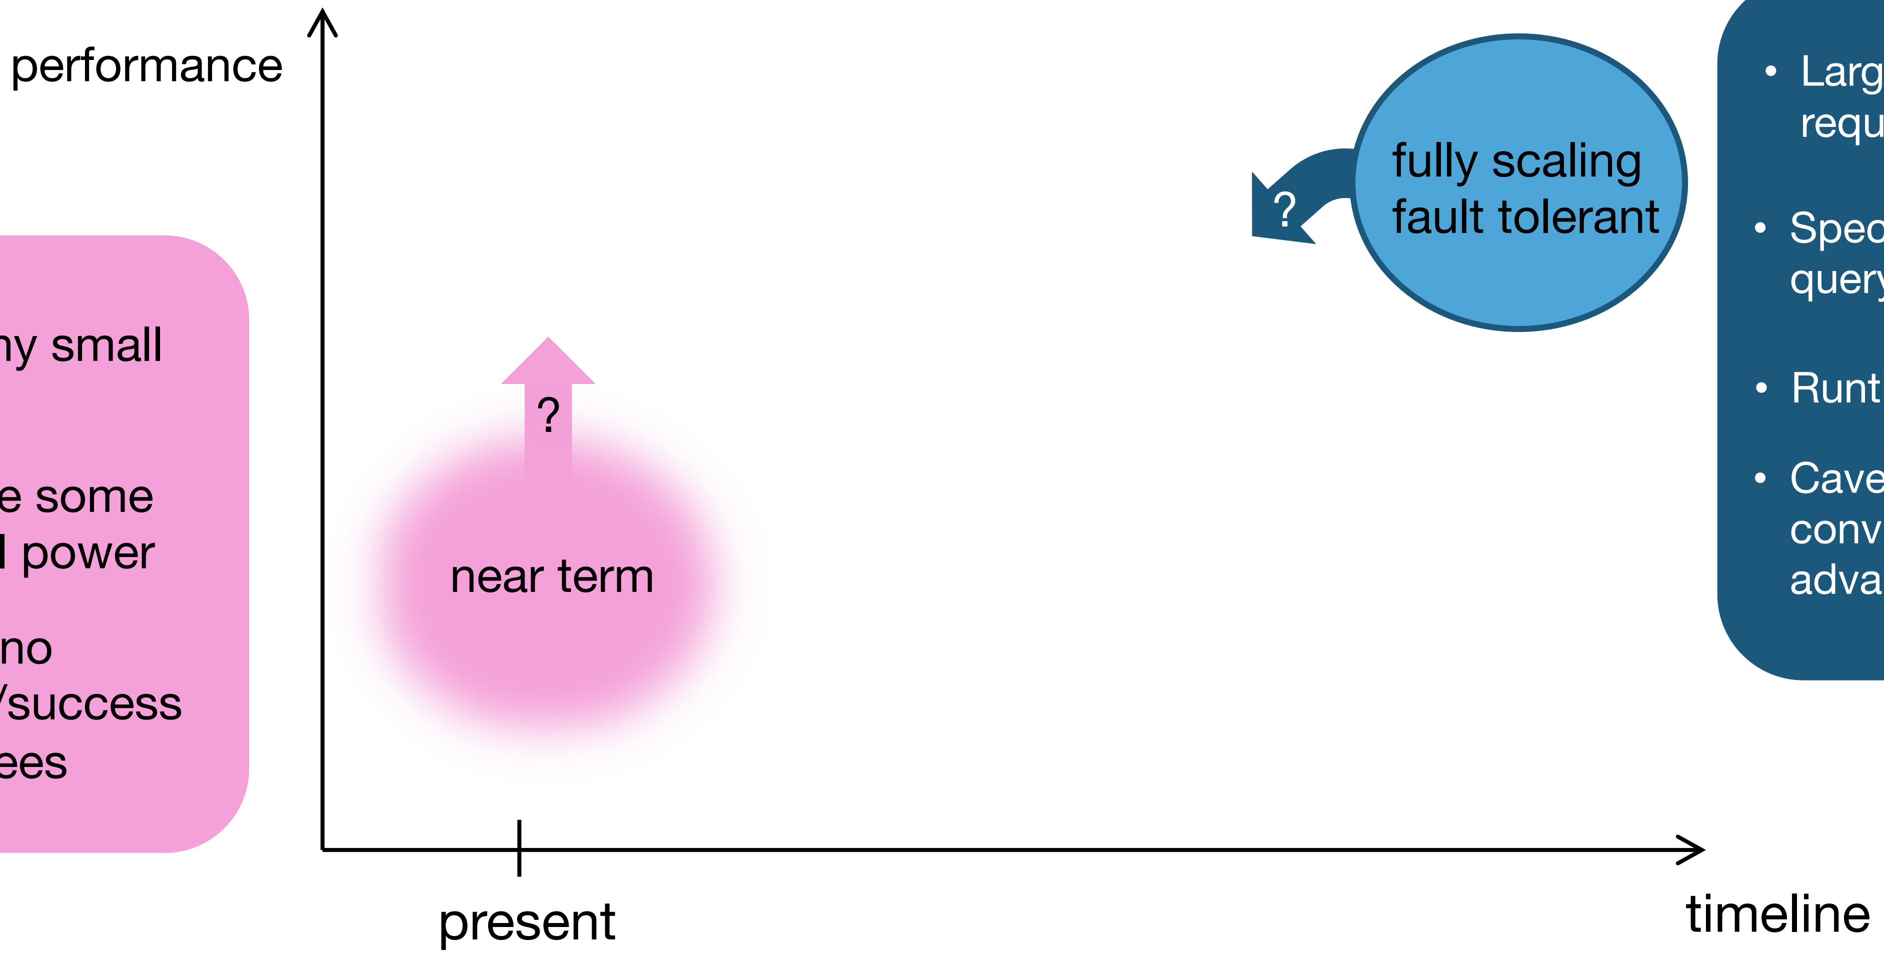

- Large hardware requirements
- Specified in terms of query complexity
- Runtime guarantees
- Caveats can cloud convincing practical advantage

Early fault-tolerant regime

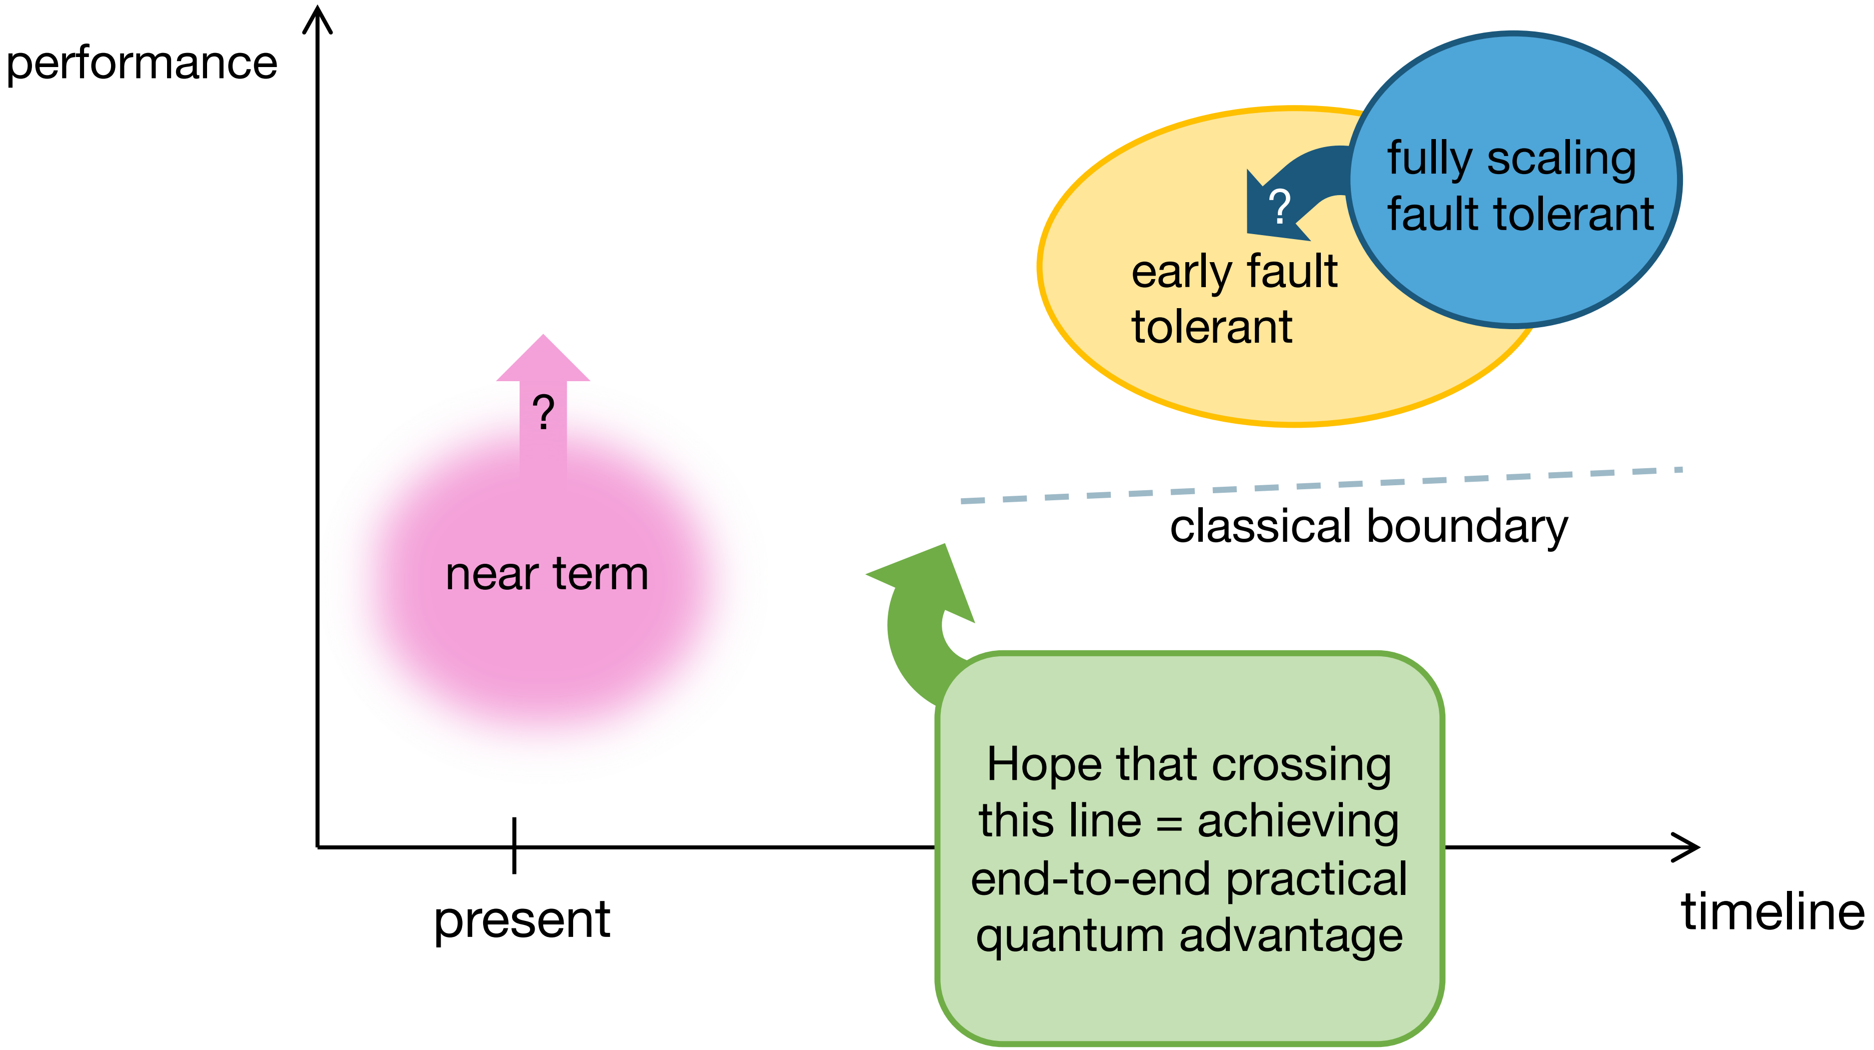

Early fault-tolerance characteristics

- Limited number of logical qubits, slow quantum clock speed from error correction overhead
- Price of resources from most expensive to cheap:
    1. Number of qubits
    2. Depth of quantum circuits
    3. Sample complexity
    4. Classical pre- and post-processing
- Goal is flexible trade-off between different resources
- Stay with provable worst-case guarantees + add strong heuristic about average case performances

Our work on early fault-tolerance

- Hybrid classical-quantum schemes with end-to-end complexity analysis

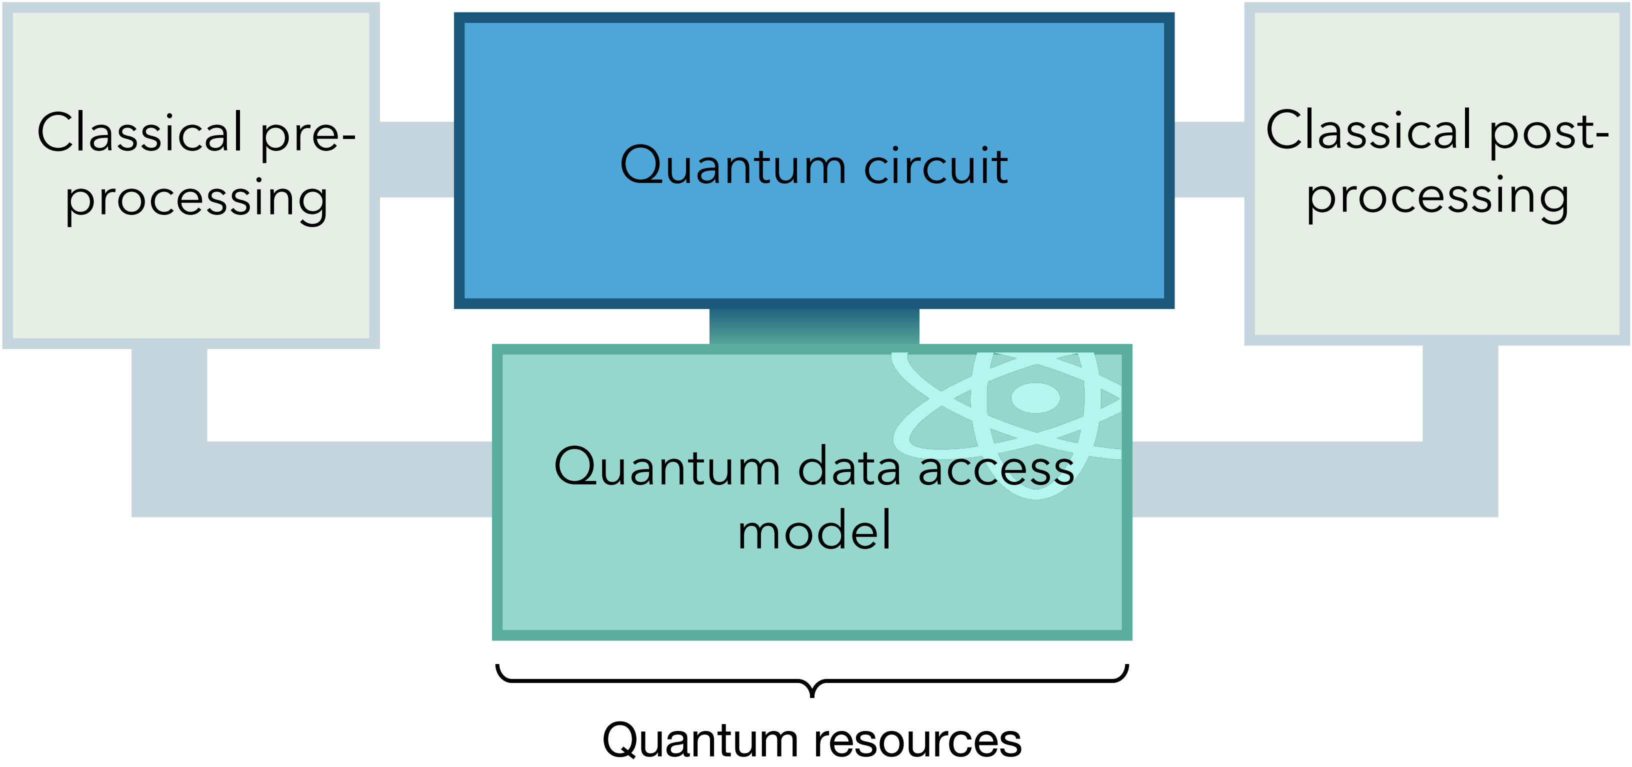

- Resource estimates for comparison with state-of-the-art classical methods

Example I: Ground state energy estimation

Randomized quantum algorithm for statistical phase estimation
QIP21, Physical Review Letters (2022) with Campbell and Wan

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 $|\tilde{E}_0 - E_0| \leq \Delta$

Early fault-tolerance approach

1. Minimize number of qubits needed – only one ancilla

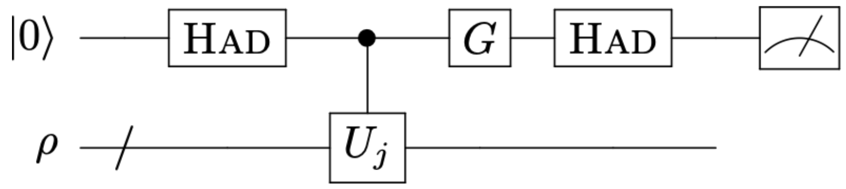

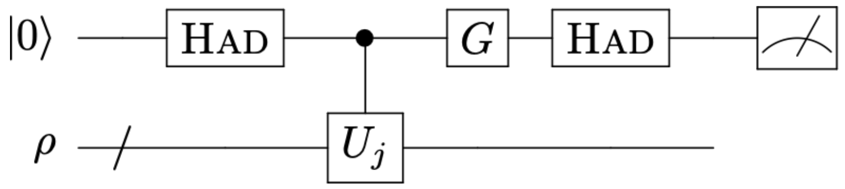

2. Trade-off gate versus sample complexity
3. Decrease error by solely taking more samples
4. Independent of the number L of Pauli terms in H – instead, depending on one-norm $\lambda \leq L$

Algorithmic result: Quantum phase estimation

- Output $\tilde{E}_0$ with $\left|\tilde{E}_0 - E_0\right| \leq \Delta$ 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 (Pauli gates)

Complexity quantum phase 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 $\tilde{O}(L\lambda\Delta^{-1})$

Basic idea

- Cumulative distribution function (CDF) relative to ρ is

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

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

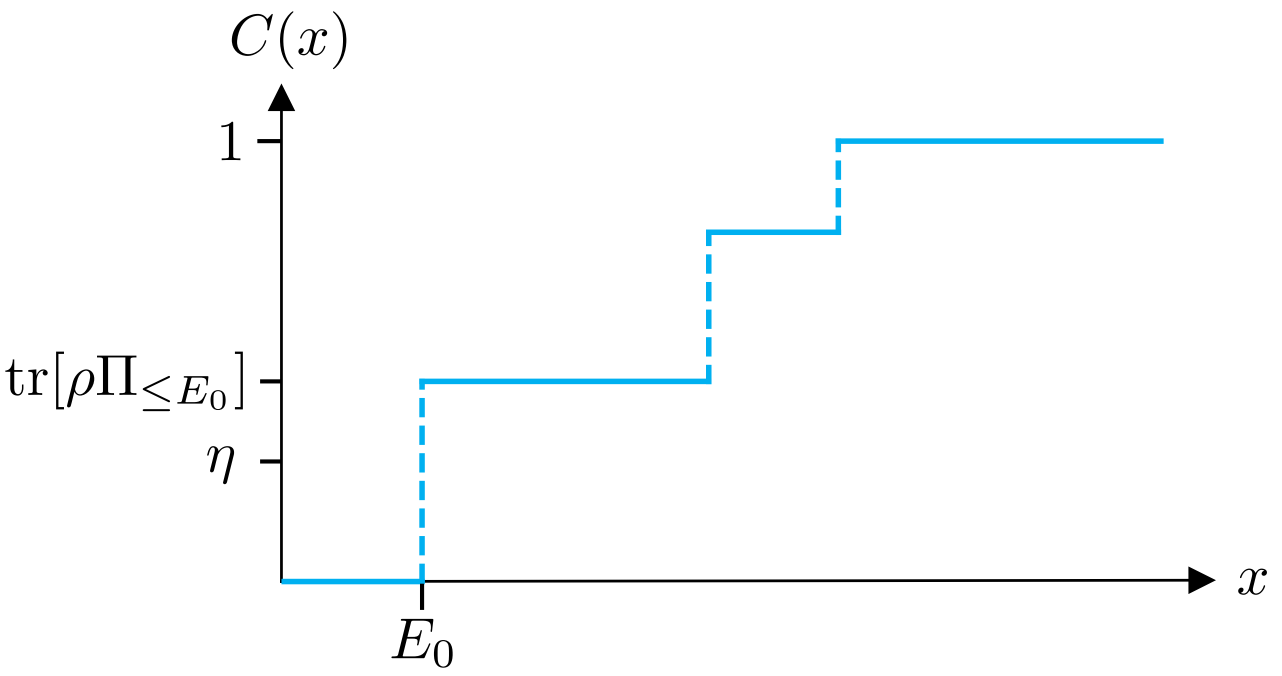

Workhorse A: Hadamard test

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

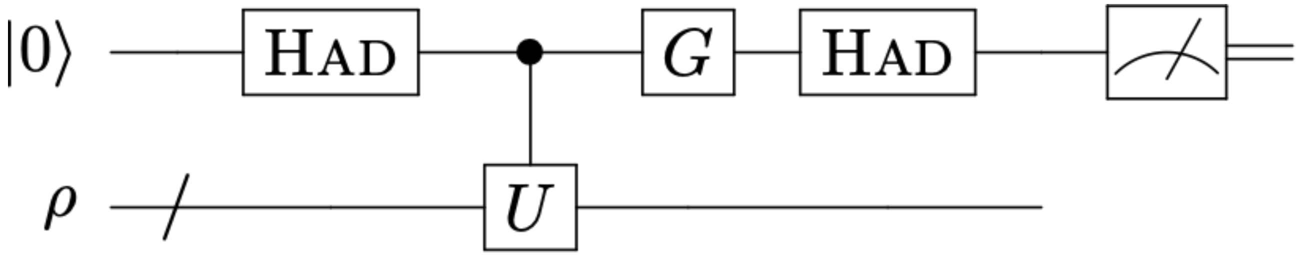

$$X = \begin{cases} +1 & \text{if } |0\rangle \\ -1 & \text{if } |1\rangle \end{cases}$$

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

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

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):

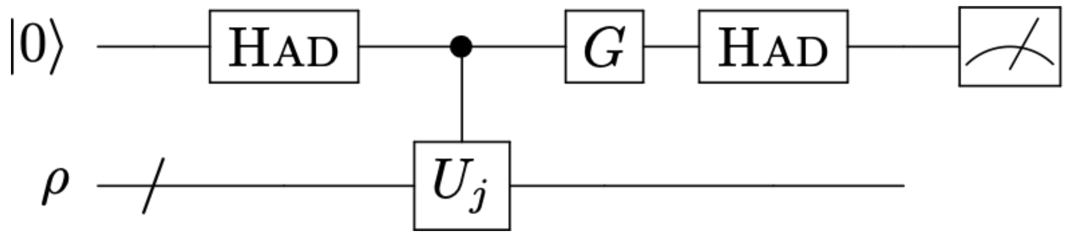

$\Rightarrow$ estimate of $\mathrm{tr}[\rho U_j]$

- Take average of samples, number 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 $\Theta(x)$:

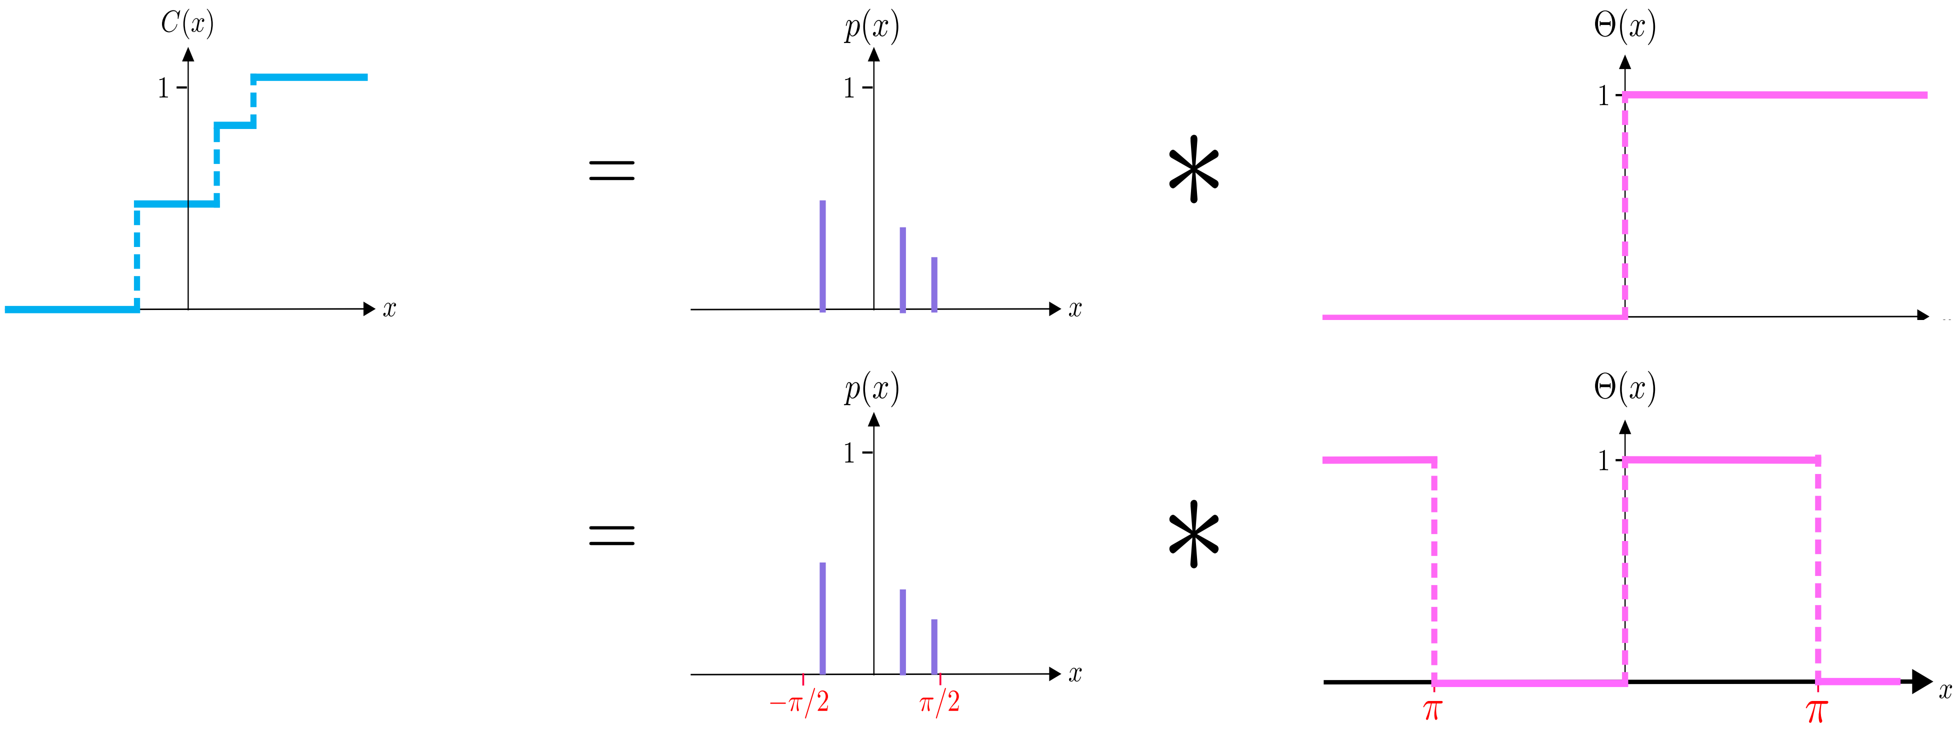

CDF via Fourier series

- 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:

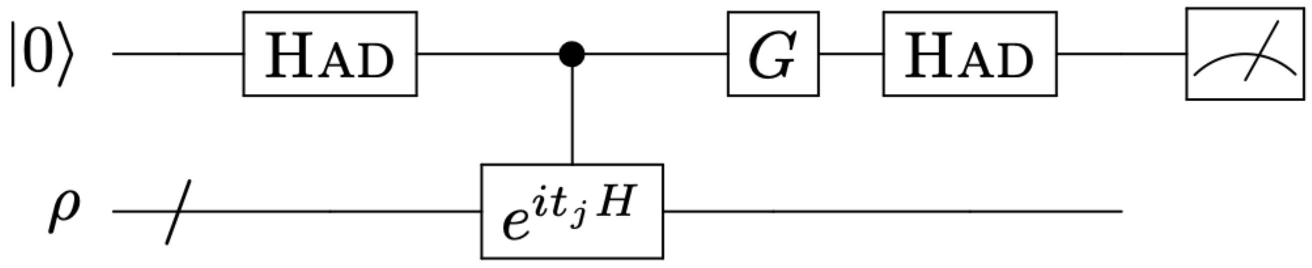

Hadamard test on Fourier series

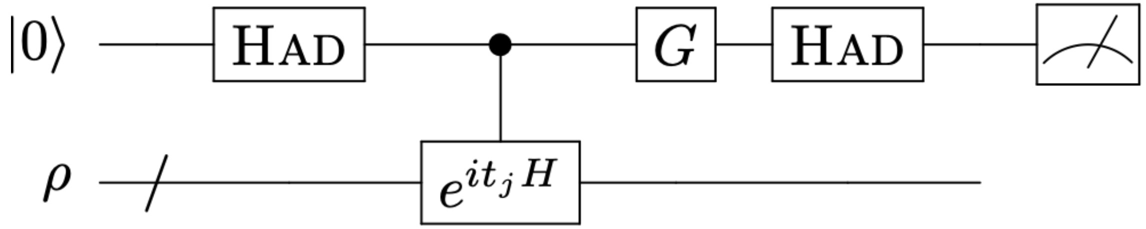

$$C(x) \approx \sum_{j \in S} \hat{F}_j e^{ijx} \cdot Tr\left[\rho e^{it_j H}\right]$$

- 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? Novel random compiler lemma for Hamiltonian simulation:

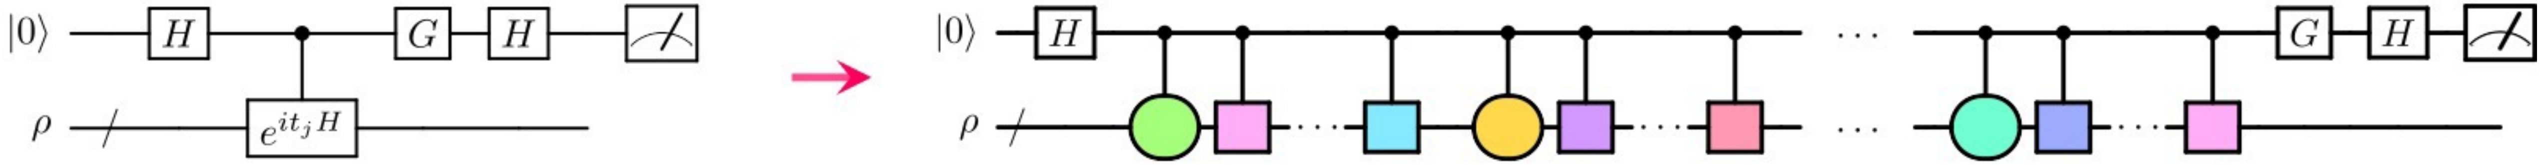

Versus previous random qDRIFT compiler: [Campbell, PRL (2019)]

FeMoco benchmark – resource trade-offs

- 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 $e^{i\theta P_l}$
- Comparison: Qubitization with heuristic truncations

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

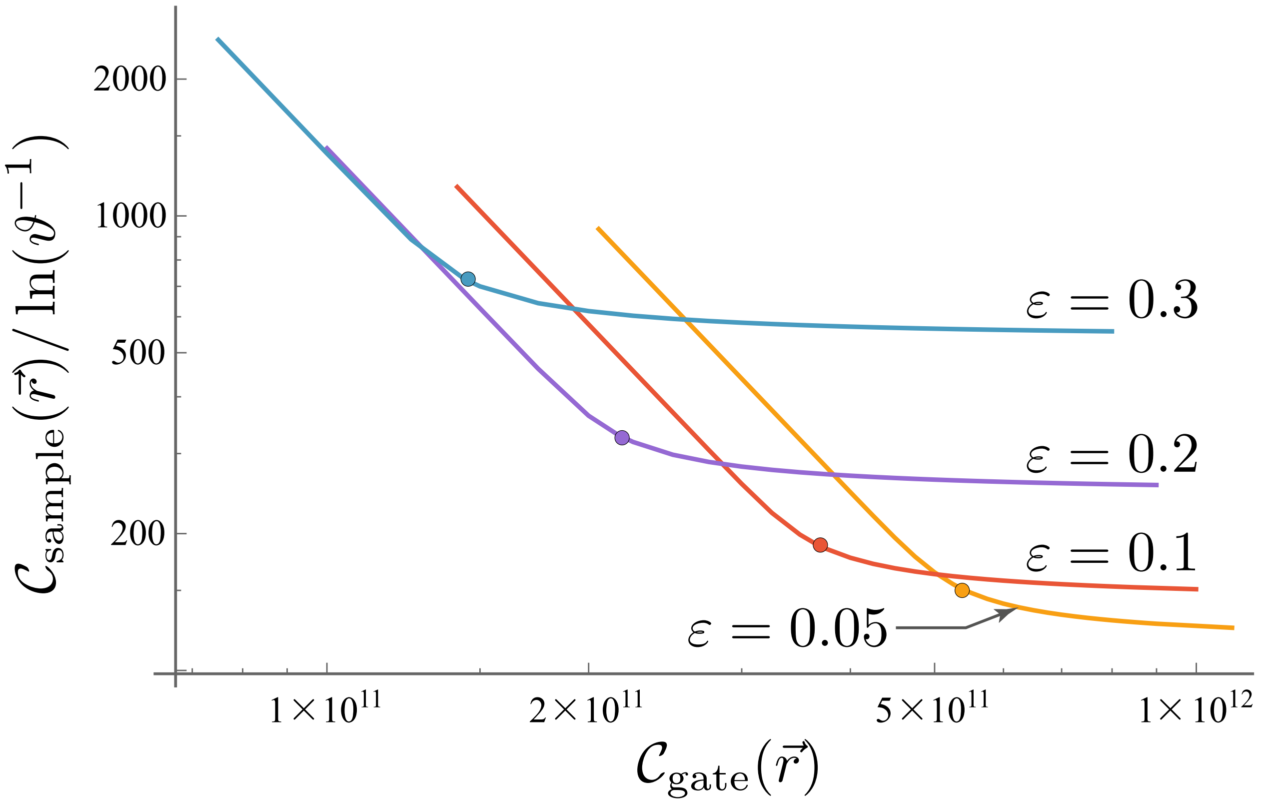

Hydrogen chain benchmark – scaling

- For length N chain, one-norm estimate $\lambda \approx O(N^{1.34})$
- 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 $C_{gate} = \tilde{O}(N^{2.3}\Delta^{-1})$
  C. tensor hypercontraction method $C_{gate} = \tilde{O}(N^{2.1}\Delta^{-1})$
- Extensive properties $\Delta \propto N$ interesting for our methods: $C_{gate} = \tilde{O}(N^{0.68})$

Example II: Linear algebra on classical data

Qubit-efficient randomized quantum algorithms for linear algebra
QCTIP23, TQC23, arXiv:2302.01873 (2023) with McArdle and Wang

Data comes via classical description

"Early fault-tolerant algorithms for classical data"

Hardware efficient
&
Provable guarantees

Quantum algorithms for classical data

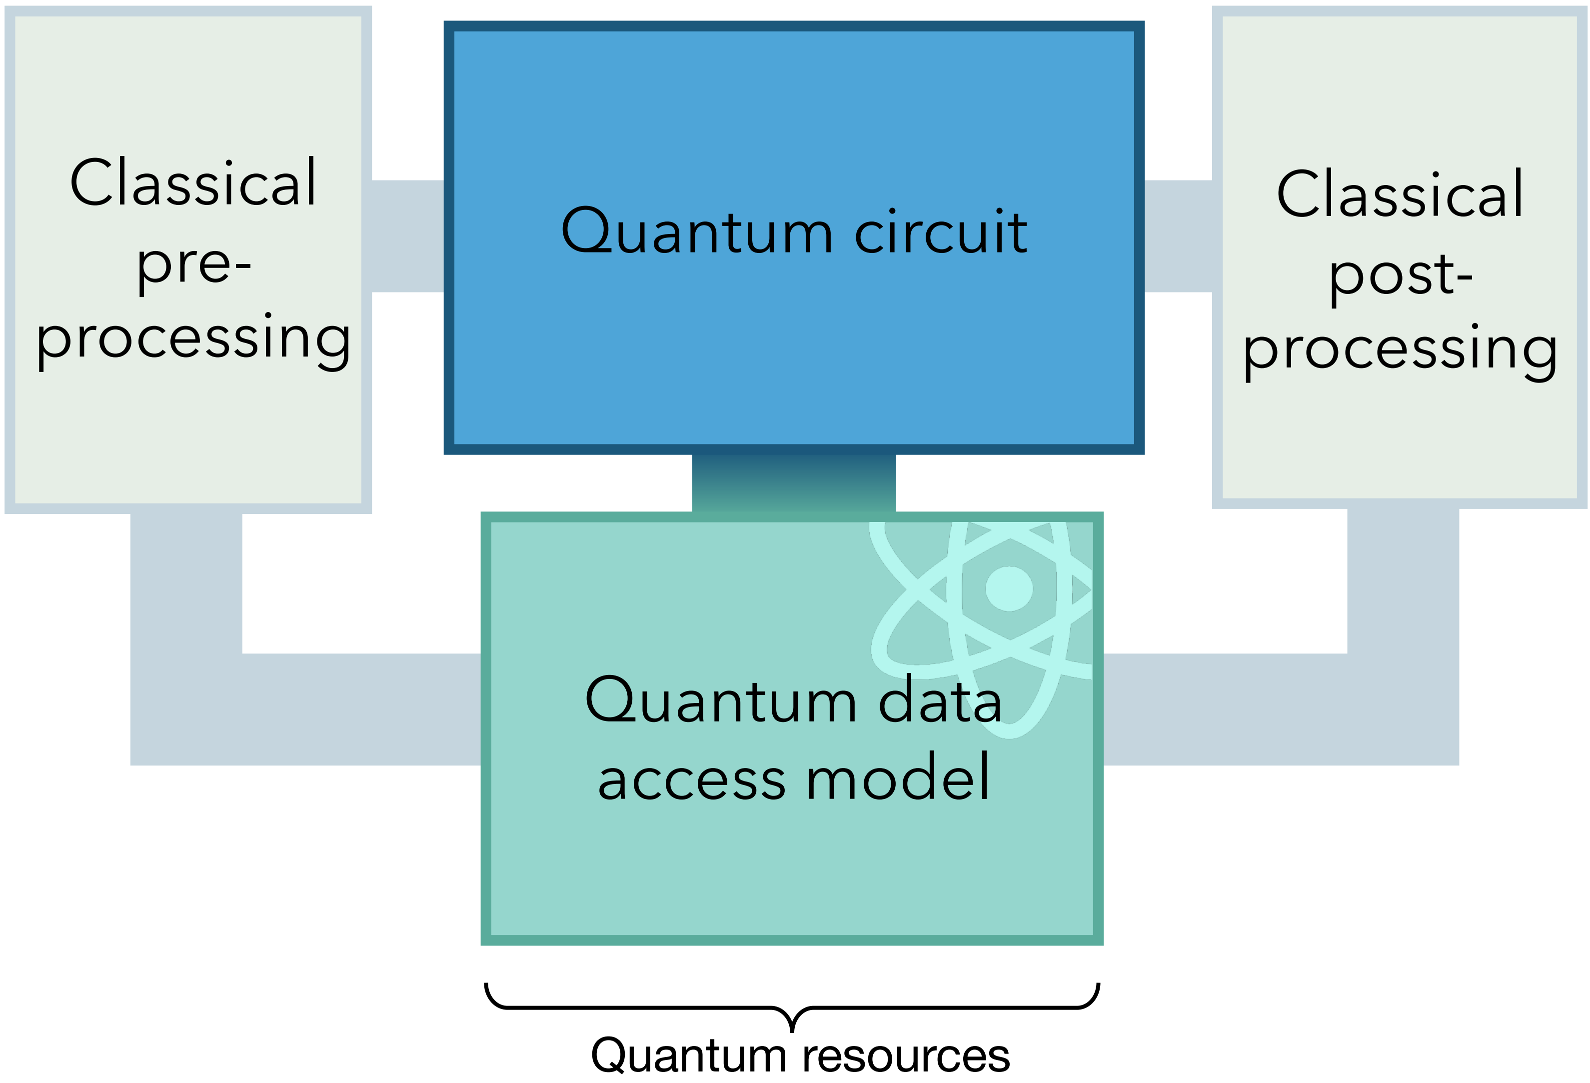

Example task: Linear system of equations

- Task: Given $N \times N$ complex matrix A and length N vector b, sample properties of length N solution vector $x = A^{-1}b$
- Classical algorithms:
  - Gaussian elimination $O(N^{\omega})$ with $\omega < 2.373$
  - Randomized $O\left(s\kappa_F(A)\log\frac{1}{\varepsilon}\right)$ for ε-approximation with $s =$ row sparsity and condition number $\kappa_F(A) = ||A||_F \cdot ||A^{-1}||$
  - Dequantized $\tilde{O}\left(\frac{\kappa_F^6(A)\kappa^2(A)}{\varepsilon^2}\right)$ for ε-approximation with $\kappa(A) = ||A|| \cdot ||A^{-1}||$
- Disclaimer: Condition number dependence $\kappa(A), \kappa_F(A)$? Input model?

Quantum linear system solver

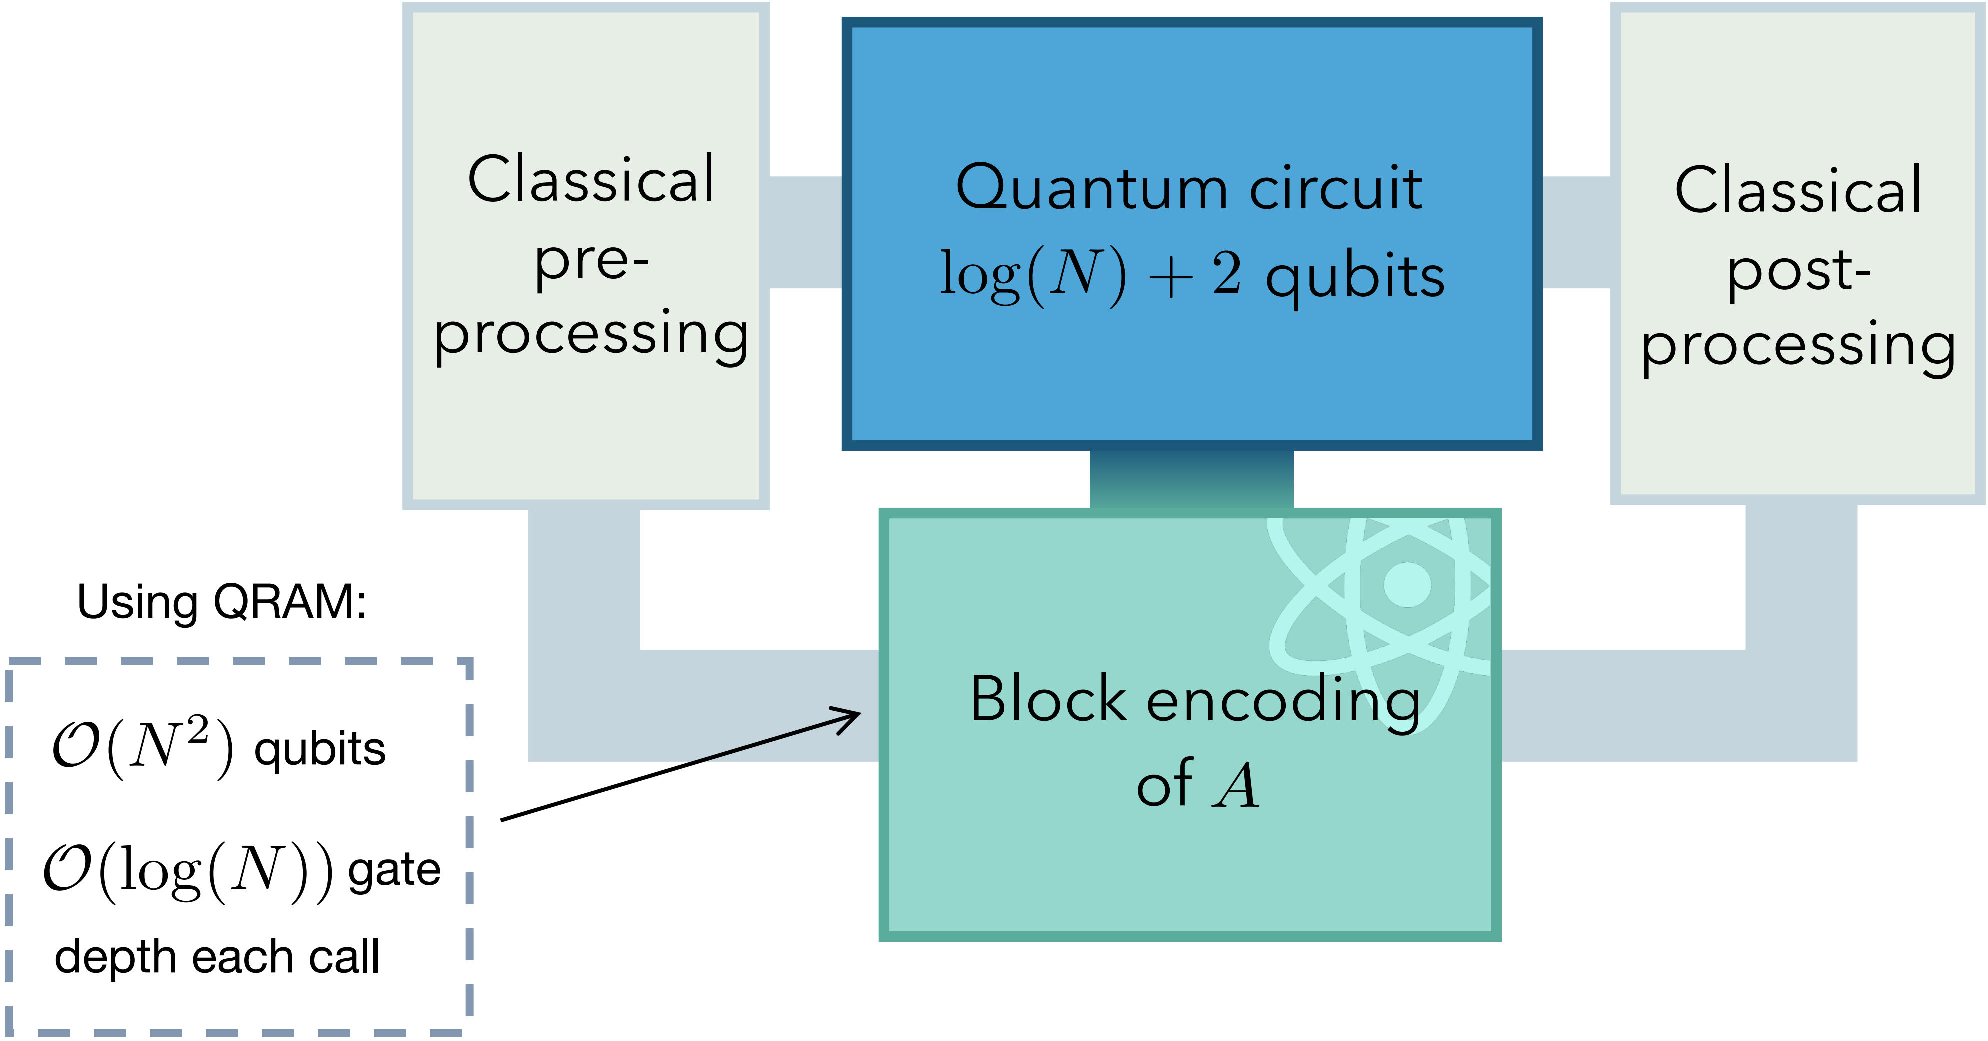

Quantum linear algebra setting

- Task (i): Given $N \times N$ complex matrix A, a function f, and preparations for $|\phi\rangle, |\psi\rangle$, sample from

$$\langle\psi|f(A)|\phi\rangle$$

- Task (ii): Given $N \times N$ complex matrix A, a function f, a preparation for $|\phi\rangle$, and an observable O, sample from

$$Tr[f(A)|\phi\rangle\langle\phi|f(A)^* O]$$

- Linear system solver corresponds to function $f(x) = x^{-1}$
- Other functions of interest: $\exp(ix), \exp(-x^2), \exp(x), \Theta(x), \cdots$

Idea I: Parallelize quantum sub-routines

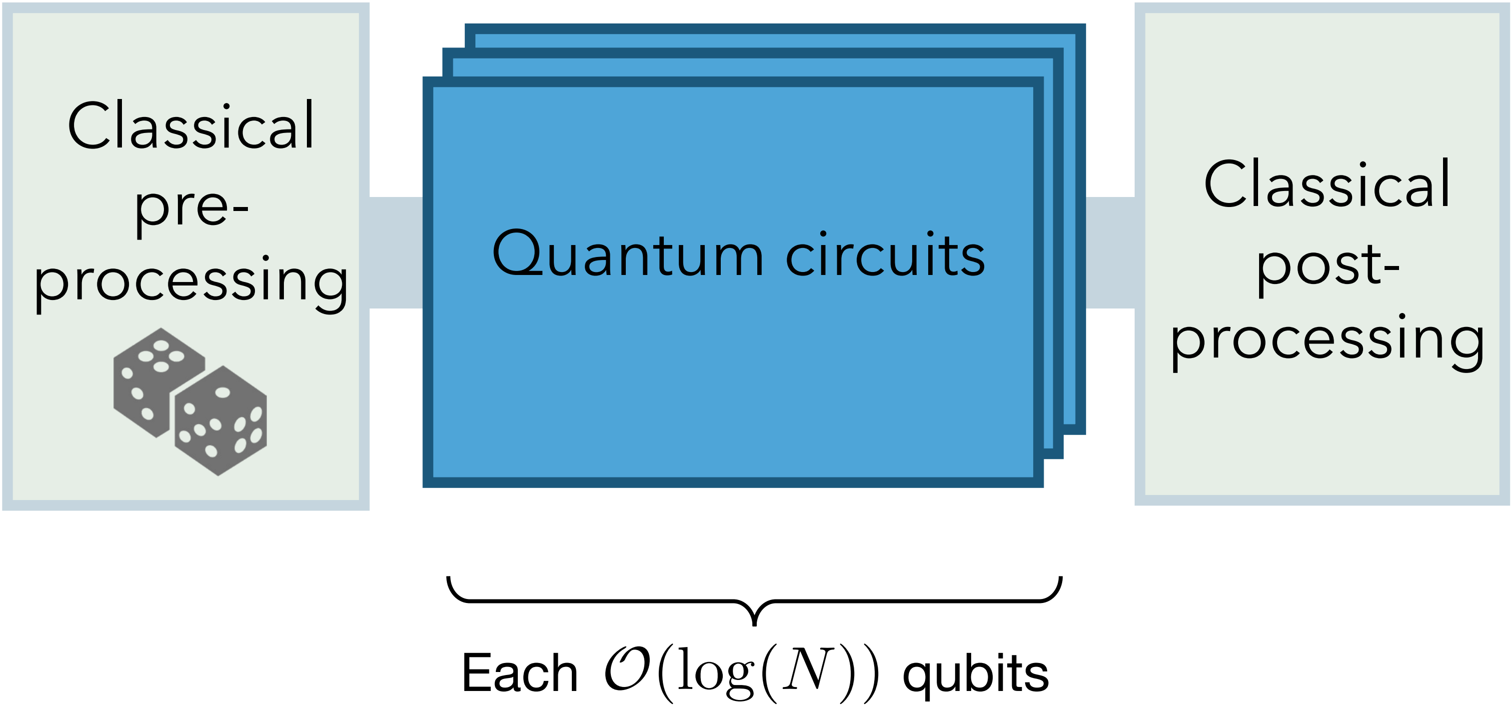

Idea II: Classical instead of quantum access

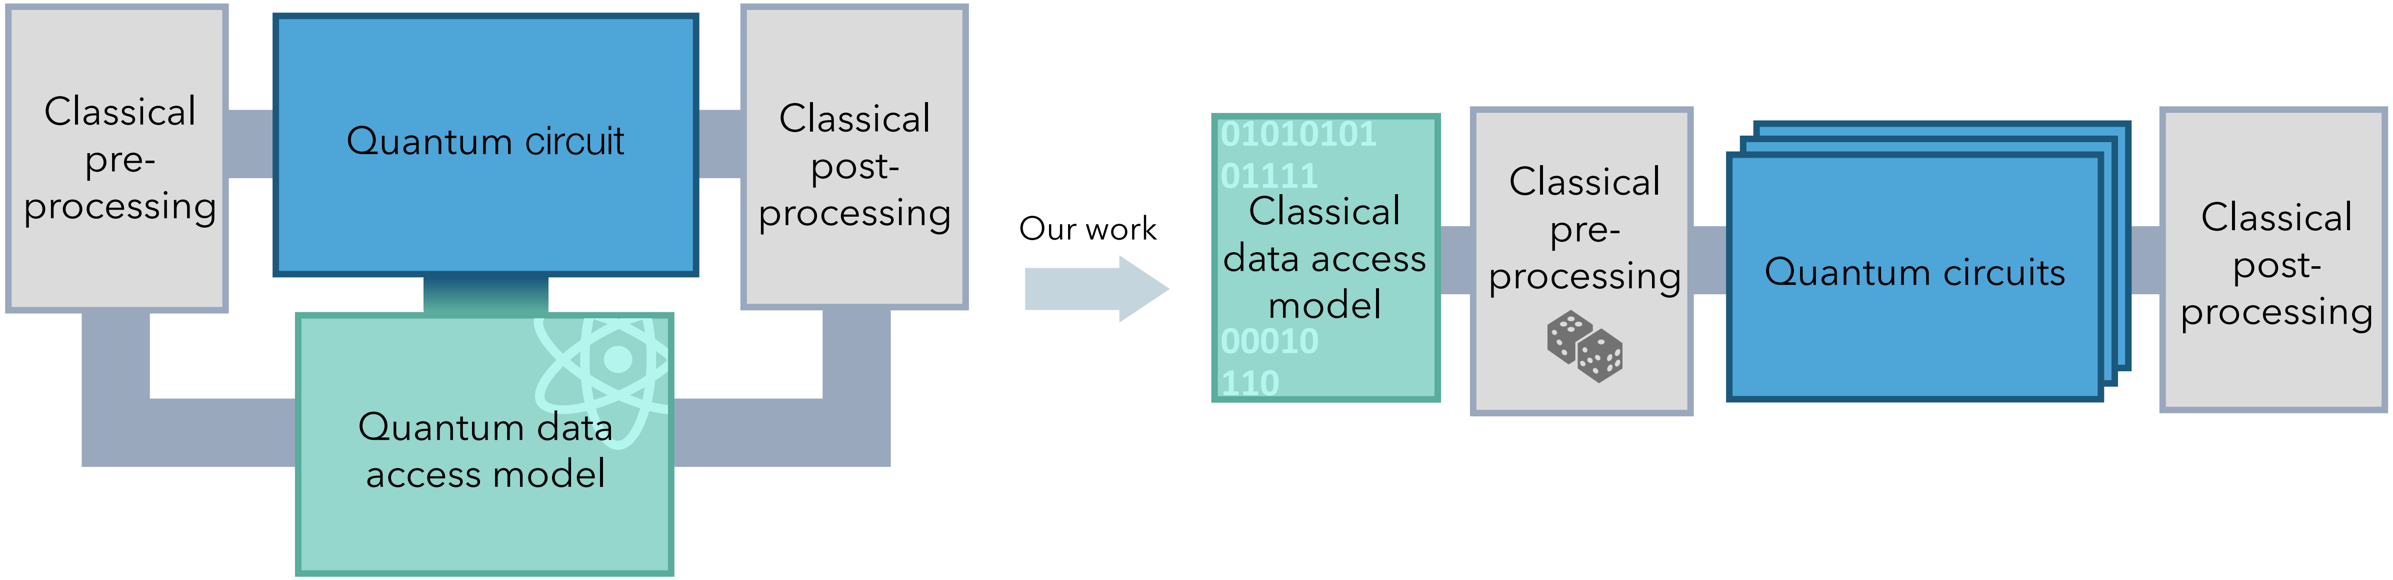

Classical access model

- Focus on $N \times N$ matrices with given Pauli decomposition $A = \sum_{l=1}^{L} a_l P_l$ and function some function f, sample properties of $f(A)$
- Classical access model = ability to sample from $\{a_l\}_l$ and known Pauli weight $\lambda(A) = \sum_l |a_l|$
- Example: A = Hamiltonian NB: $\lambda(A) \leq L \leq N^2$ but often even $\lambda(A) = O(\log N)$
- Use Hadamard test and importance sampling, just now with Fourier series of function f!

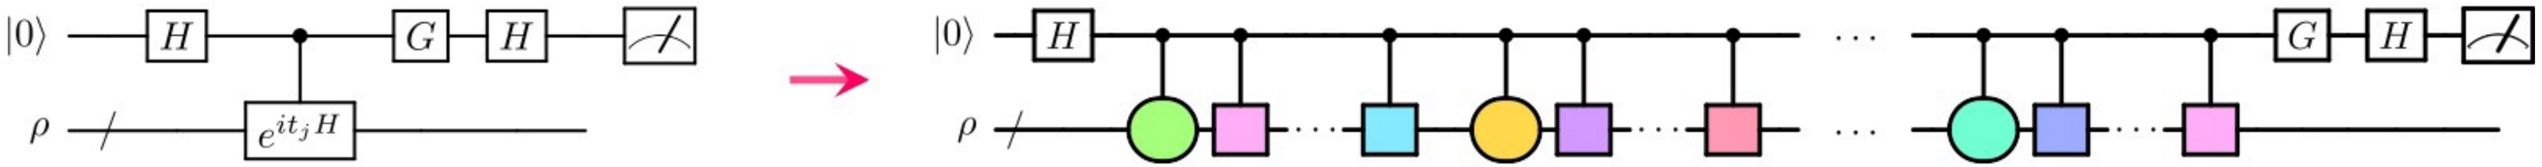

Quantum linear algebra result

Fourier approximation of f

$$|f(x) - s(\varepsilon, D_A, x)| \leq \varepsilon\,, \quad \forall x \in D_A$$

$$s(\varepsilon, D_A, x) = \sum_{k\in F} \alpha_k(\varepsilon, D_A) \exp\left(it_k(\varepsilon, D_A)x\right)$$

Pauli decomposition of A

$$\{a_l\}_l \;\; s.t. \; A = \sum_l^L a_l P_l$$

$$\lambda = \sum_l^L |a_l|$$

Given $|\phi\rangle, |\psi\rangle, O$

```
Sampling Alg.
```

1. ---------------
2. ---------------
3. ---------------
4. ---------------
5. ---------------

…

(i) sample $\langle\psi|f(A)|\phi\rangle$

(ii) sample $Tr[f(A)|\phi\rangle\langle\phi|f(A)^*O]$

Quantum linear algebra result

Fourier approximation of *f*

$$|f(x) - s(\varepsilon, D_A, x)| \leq \varepsilon\,, \quad \forall x \in D_A$$

$$s(\varepsilon, D_A, x) = \sum_{k \in F} \alpha_k(\varepsilon, D_A) \exp\left(i t_k(\varepsilon, D_A) x\right)$$

Pauli decomposition of *A*

$$\{a_l\}_l \;\; s.t. \; A = \sum_l^L a_l P_l$$

$$\lambda = \sum_l^L |a_l|$$

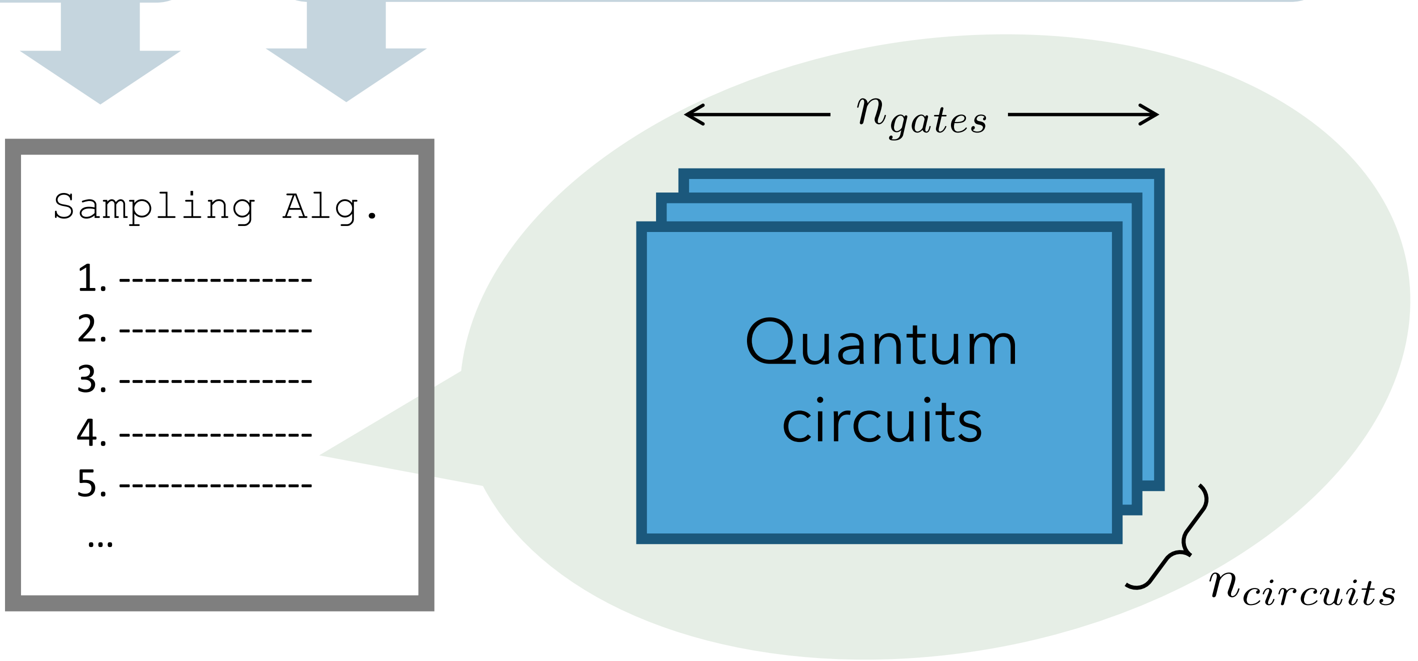

Complexities quantum linear system solver

- Task (i): Given $N \times N$ complex matrix $A = \sum_{l=1}^{L} a_l P_l$ in Pauli input model, and preparations for $|\phi\rangle, |b\rangle$, sample from $\langle\phi|A^{-1}|b\rangle$
- Only use $\log N + 2$ qubits in total!
- Flexible trade-off possible, one choice

$$n_{gates} = \tilde{O}(\lambda^2(A) \cdot ||A^{-1}||^2),\ n_{circuits} = \tilde{O}\left(\frac{||A^{-1}||^2}{\varepsilon^2}\right)$$ NB: often $\lambda(A) = O(\log N)$

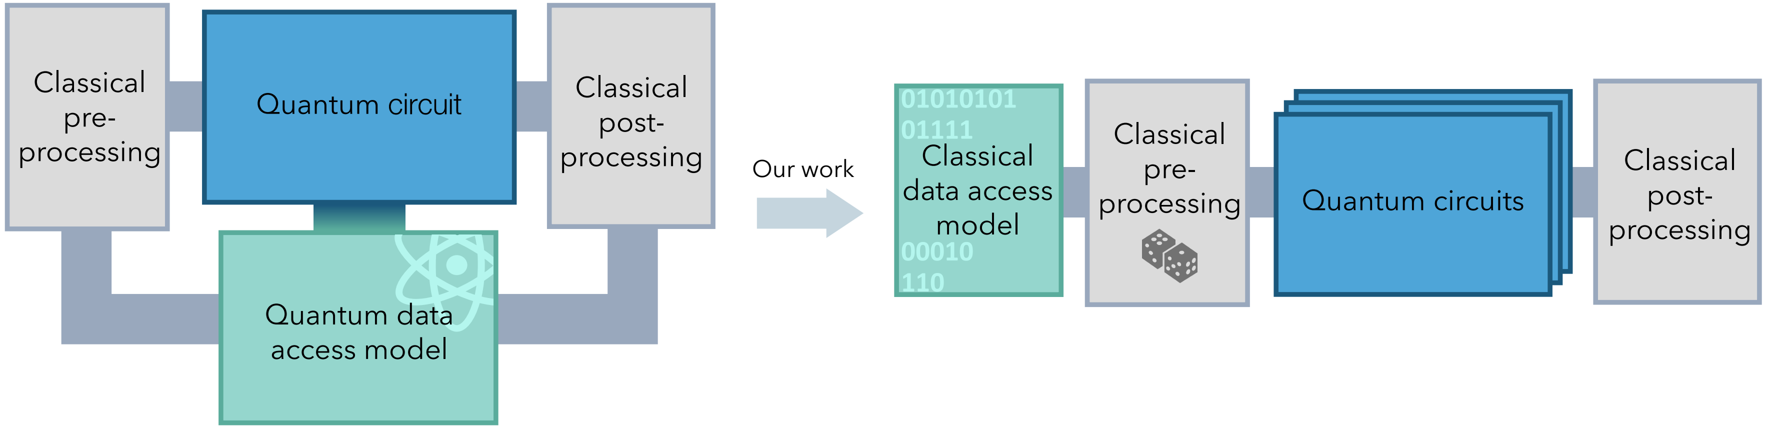

Example III: Quantum state preparation

Quantum state preparation without coherent arithmetic
arXiv:2210.14892 (2022) with McArdle, Gilyen

Quantum state preparation problem

- Classical data not from table but generated via functions
- Given the function $f : [a, b] \to \mathbb{R}$, prepare the n-qubit quantum state

$$|\Psi_f\rangle := \frac{1}{\mathcal{N}_f} \cdot \sum_{x=0}^{2^n-1} f(\bar{x})|x\rangle$$

with uniform grid $\bar{x} := a + \frac{x(b-a)}{2^n}$ and normalization $\mathcal{N}_f := \sqrt{\sum_{\bar{x}} f(\bar{x})}$

- Important sub-routine in a variety of quantum algorithms, for different functions of interest
- Minimize number of ancilla qubits and quantum gates

Standard approach(es)

- Amplitude oracle $|x\rangle|0\rangle \mapsto |x\rangle|f(\bar{x})\rangle$ that prepares g-bit approximation of the values $f(\bar{x})$
- Implemented via reversible computation, using piecewise polynomial approximation of the function $f(x)$
- Alternatively, reading values stored in a quantum memory
- Downsides:
  - Handcrafted for every function + discretization of values of function
  - Large ancilla cost – not suited for early fault-tolerant regime
- Other approaches: Grover-Rudolph, adiabatic, repeat until success, matrix product states, etc. (similar bottlenecks)

Quantum eigenvalue transformation (QET)

- Apply functions to the eigenvalues of a Hermitian matrix
- An (α, m)-block encoding of an n-qubit Hermitian A is an $(n+m)$-qubit unitary U_A with

$$\mathrm{A} = \alpha \cdot \left(\langle 0|^{\otimes m} \otimes 1_n\right) U_A \left(|0\rangle^{\otimes m} \otimes 1_n\right)$$

- Base functions are even degree d polynomials

→ QET circuit output is block encoding U_{A^d} of the matrix A^d

- Implementation cost:

$\frac{d}{2}$ applications of U and U^* each + $\mathrm{O(d)}$ other gates in between

QET continued

- Example circuit for even degree d polynomial and $m = 1$:

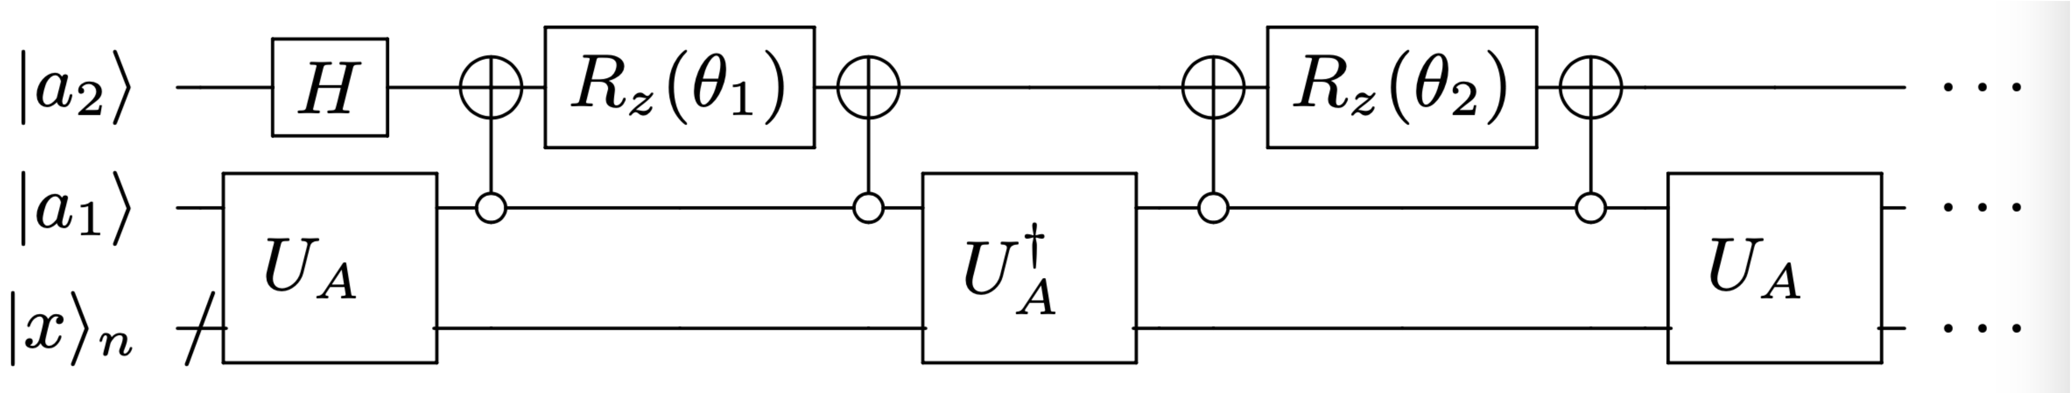

- Efficient classical pre-computation of angle set $\{\theta_1, \theta_2, \cdots, \theta_d\}$
- Odd polynomials, general functions via polynomial approximation, complexity given by degree of polynomial

Main idea: State preparation via QET

- Create low-cost block encoding of $A \coloneqq \sum_{x=0}^{2^n-1} \sin\left(\frac{x}{2^n}\right) |x\rangle\langle x|$ via

  = (1,1) block encoding

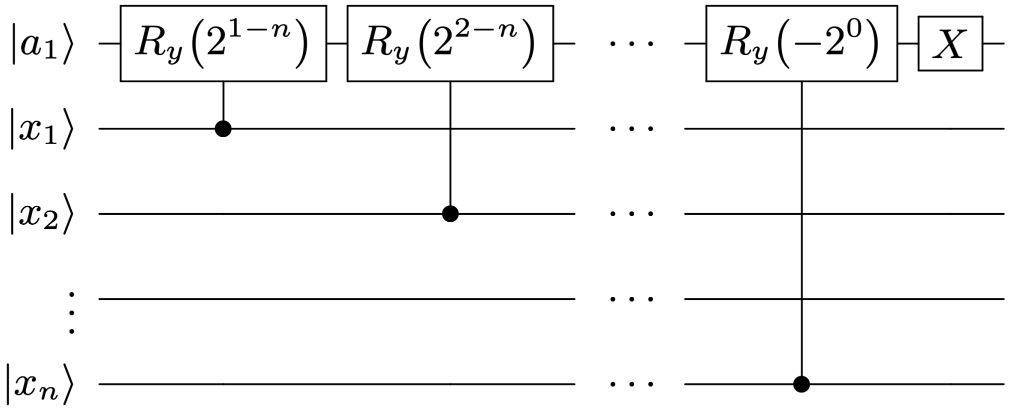

$R_Y(\alpha) \coloneqq \exp(-i\alpha Y)$

- Idea: Applying QET, convert this into block encoding of $\sum_{x=0}^{2^n-1} f(\bar{x})|x\rangle\langle x|$ using polynomial approximation of $f\big((b-a)\arcsin(\cdot) + a\big)$

- Run circuits on input $|x_1 \cdots x_n\rangle \otimes |0 \cdots 0\rangle_a = |+\rangle^{\otimes n}|0 \cdots 0\rangle_a$ and use amplitude amplification to maximize probability of $|\Psi_f\rangle \otimes |0 \cdots 0\rangle_a$

Quantum complexities

- For sufficiently smooth functions f, we can prepare a quantum state $|\Psi_{\tilde{f}}\rangle$ that is ε-close in trace distance to $|\Psi_f\rangle$ using

$\tilde{\mathcal{O}}\left(\frac{n\log(\varepsilon^{-1})}{\mathcal{F}_{\tilde{f}}^{[N]}}\right)$ gates + 4 ancilla qubits with discretized L_2-norm filling-fraction ($N \coloneqq 2^n$)

$$\mathcal{F}_f^{[N]} \coloneqq \frac{\sqrt{\frac{(b-a)}{N}\sum_{x=0}^{N-1}|f(\bar{x})|^2}}{\sqrt{(b-a)|f|_{\max}^2}} \approx \frac{\sqrt{\int_a^b |f(\bar{x})|^2 d\bar{x}}}{\sqrt{(b-a)|f|_{\max}^2}} =: \mathcal{F}_f^{[\infty]}$$

- Show analytical results via minimax polynomial
- In practice use instead (works even better):
  - Remez approximation or even just Local Taylor series
  - L_2-approximation on grid

Analytical performance: Gaussians

- Example function $f_\beta(x) := \exp\left(-\frac{\beta}{2}x^2\right)$

- For $\varepsilon \in \left(0, \frac{1}{2}\right)$ and $0 \leq \beta \leq 2^n$ we can prepare the $[-1,1]$ uniform grid Gaussian state on n qubits up to ε-precision with gate complexity

$$O\left(n \cdot \log^{\frac{5}{4}}\left(\frac{1}{\varepsilon}\right)\right) + 3 \text{ ancilla qubits}$$

  for $\beta \geq \log\left(\frac{1}{\varepsilon}\right)$.

- Note: All other approaches use hundreds of ancilla qubits

Numerical benchmarking: tanh(x)

- Example function $\tanh(x)$ in the range $x \in [0,1]$ on $n = 32$ gives

Method	# Ancilla qubits	# Toffoli gates
QET (this work)	3	9.7×10^4
Black-box state amplitude oracle	216	6.9×10^4
Grover-Rudolph amplitude oracle	> 959	$> 2.0 \times 10^5$

- Cost are lower bounds minimizing gate count, based on state-of-the-art amplitude oracles (which could potentially be improved)
- Other methods give even higher costs

Conclusion / Outlook

Quantum algorithms for early fault-tolerance

- Motto: Classical whenever possible, use as few qubits as possible
- Finding: Early fault-tolerant methods can even be competitive with state-of-the-art (non-qubit aware) schemes in terms of asymptotic complexities
- Needed: More quantum resource counts for different applications, end-to-end complexity analyses
- Guiding questions:
  - What quantum algorithms do we eventually want to run?
  - For what applications is the quantum footprint the smallest to become competitive with classical methods?

Thank you!

RWTH AACHEN UNIVERSITY

Some references

Paper references of our work

- A randomized quantum algorithm for statistical phase estimation
  QIP21, Physical Review Letters (2022) with Campbell, Wan
- Qubit-efficient randomized quantum algorithms for linear algebra
  QCTIP23, TQC23, arXiv:2302.01873 (2023) with McArdle, Wang
- Quantum state preparation without coherent arithmetic
  arXiv:2210.14892 (2022) with McArdle, Gilyen
- Quantum resources required to block-encode a matrix of classical data
  IEEE Transactions on Quantum Engineering (2022) with Clader, Dalzell, Stamatopoulos, Salton, Zeng
- A streamlined quantum algorithm for topological data analysis with exponentially fewer qubits
  QIP22, arXiv:2209.12887 (2022) with McArdle, Gilyen
- Sparse random Hamiltonians are quantumly easy
  QIP22, arXiv:2302.03394 (2023) with Chen, Dalzell, Brandão, Tropp

Extra content ground state energy

Fourier series lemma (Heaviside function)

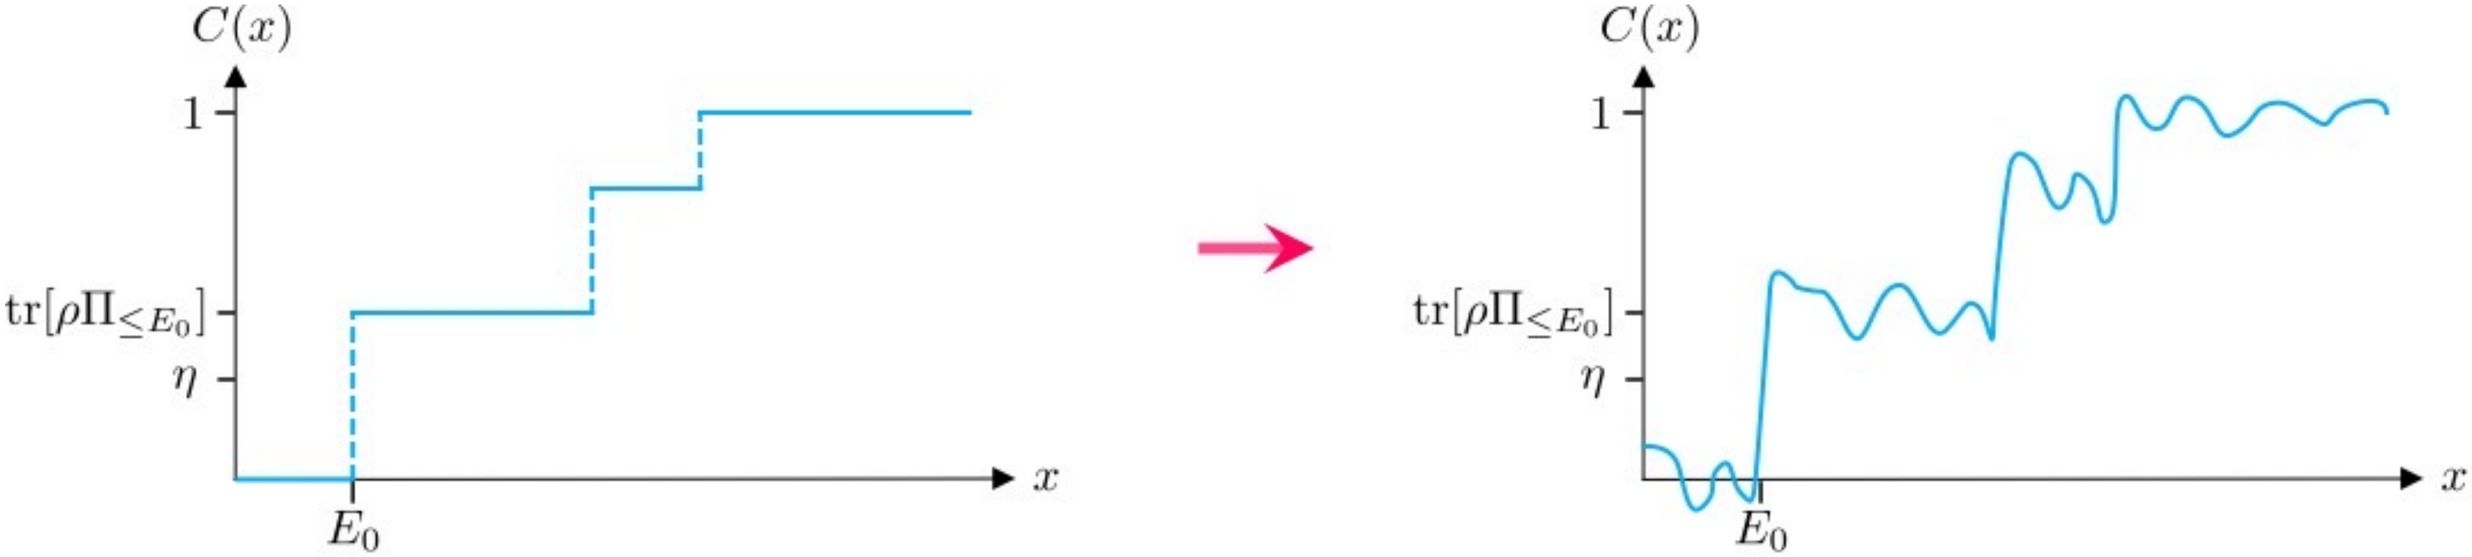

- Improved Fourier series approximation of Heaviside function
- Technical contribution:

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)]

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:

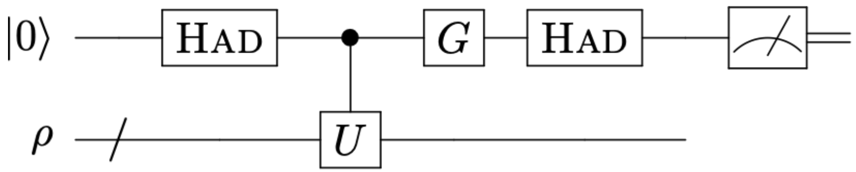

  I. $\mu(r) \coloneqq \sum_k b_k \leq \exp(t^2 r^{-1})$
  II. $COST(C - U_k) = r$ controlled single qubit Pauli rotations $\forall k$

- Gate complexity r versus sample complexity $\exp(t^2 r^{-1})$
- 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 $C(x) \approx \sum_j \sum_k \hat{F}_j e^{ijx} b_k^{(j)} Tr\left[\rho U_k^{(j)}\right]$

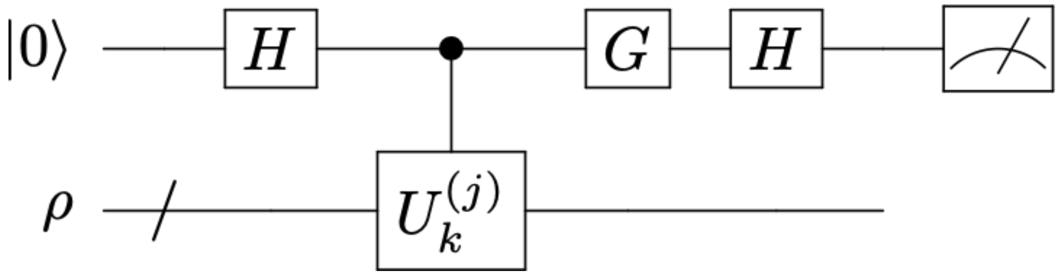

- $e^{it_j H} = \sum_k b_k^{(j)} U_k^{(j)}$ decomposition for runtime vector $\vec{r} = (r_j)_j \in \mathbb{N}^{|S|}$ as:

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

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

Putting things together

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

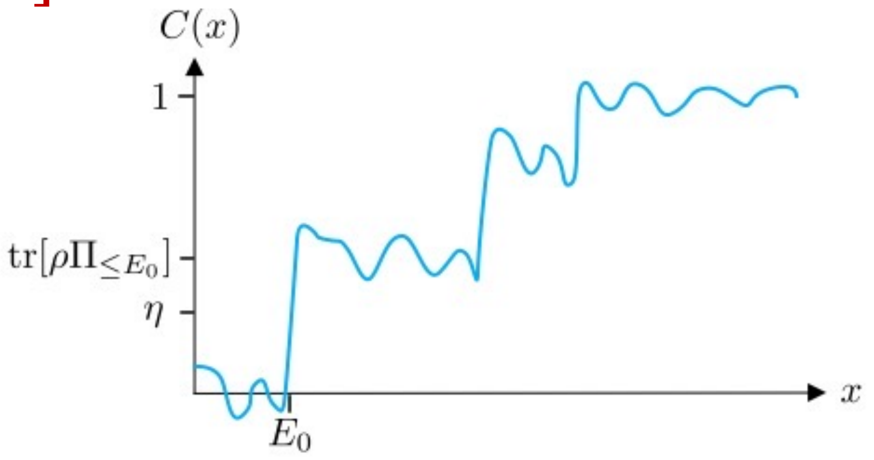

- As $\mu_j \leq e^{t_j^2 r_j^{-1}}$ choosing $r_j = 2t_j^2$ $\forall j$ gives $\mu_j \leq \sqrt{e}$:

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

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

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:

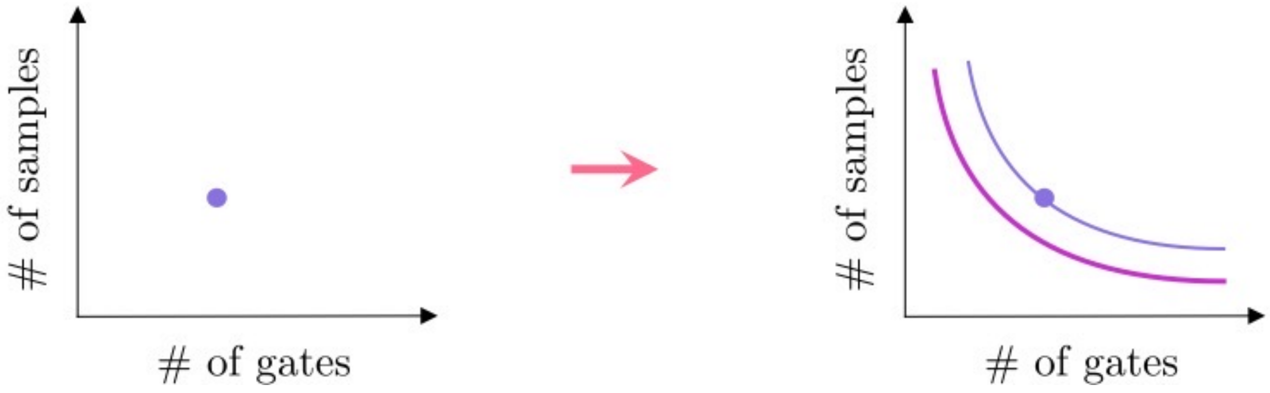

Extra: Proof Fourier series lemma

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

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

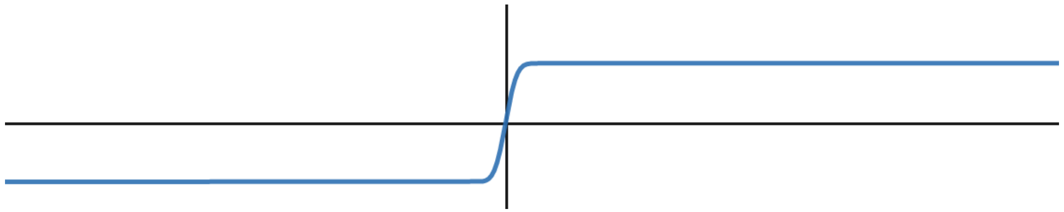

- Fourier series: $\Theta(x) \approx \mathrm{erf}\big(\beta \sin(x)\big) \approx \sum_k c_k T_k\left(\cos\left(\frac{\pi}{2} - x\right)\right)$

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

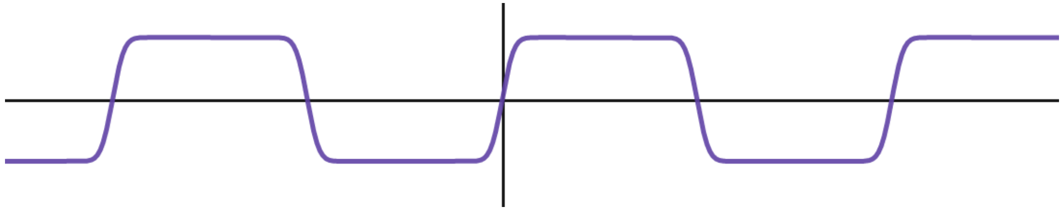

Extra: Proof random compiler lemma

- 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

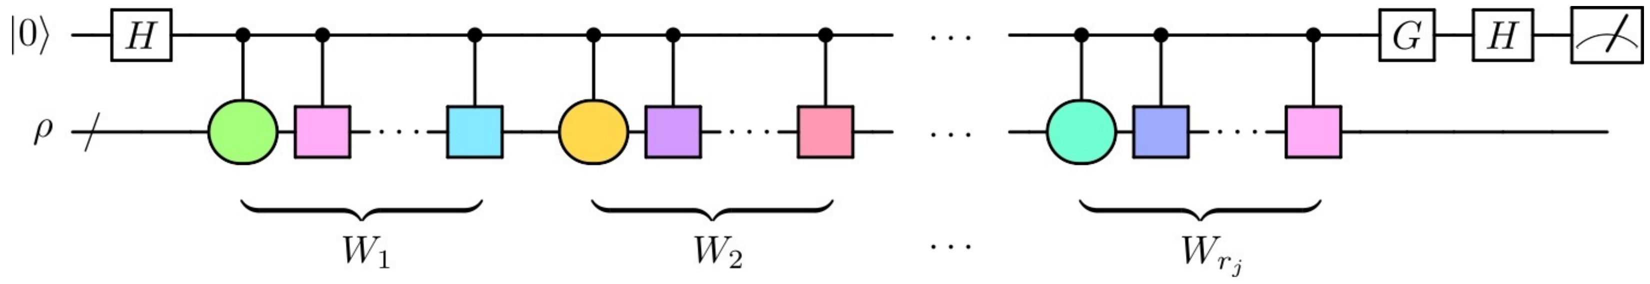

Extra: qDRIFT comparison

[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 $P_{l_1}, \dots, P_{l_r}$ 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

Extra content linear algebra

Sampling algorithm for (i)

$s(A) \approx f(A) \longrightarrow$ Decompose into linear combination of strings of gates composed of n_{gates} gates

For $m = 1, 2, ..., n_{circuits}$:

Sample one string of gates : (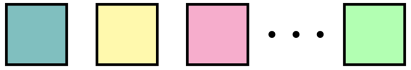)

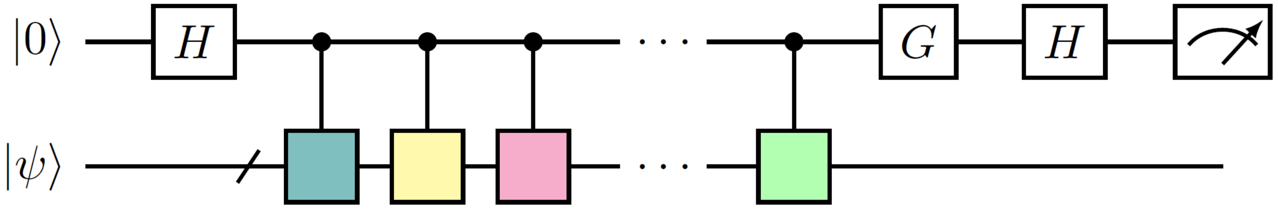

Run circuit, obtain single measurement statistic o_m

Multiply o_m by “weight” of linear combination. Record.

 = Pauli gate + Pauli rotation

Average over $m \longrightarrow$ Get answer $\approx \langle\psi| \, s(A) \, |\psi\rangle \approx \langle\psi| \, f(A) \, |\psi\rangle$

Sampling algorithm for (ii)

For $m = 1, 2, ..., n_{circuits}$:

Sample **two** strings of gates $(\ \cdots\ )$, $(\ \cdots\ )$

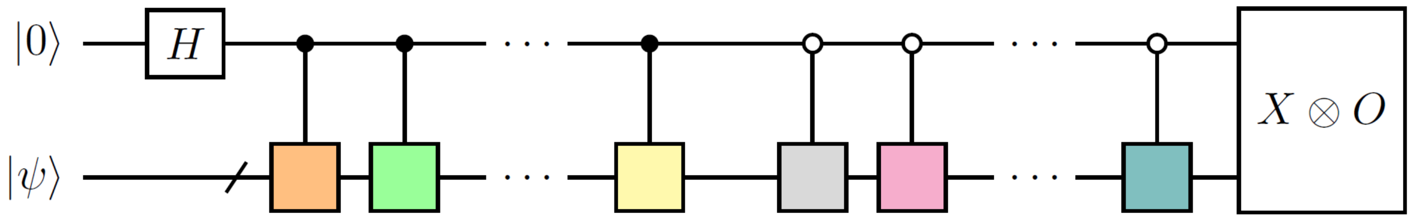

Run circuit, obtain single measurement statistic o_m

Multiply o_m by “weight” of linear combination. Record.

Average over m $\longrightarrow$ Get answer $\approx \langle\psi| s(A)^\dagger O s(A) |\psi\rangle \approx \langle\psi| f(A)^\dagger O f(A) |\psi\rangle$

Linear Systems

Given $N \times N$ matrix A and state $|\vec{b}\rangle$, with probability at least $1-\delta$

prepare $\langle\psi| A^{-1} |\vec{b}\rangle$

prepare $\langle\vec{b}| (A^{-1})^\dagger O A^{-1} |\vec{b}\rangle$

- Qubit count: $\log(N)+1$ (Hermitian), $\log(N)+2$ (general)
- Number of quantum circuits:

$$n_{circuits} = \widetilde{\mathcal{O}}\left(\log\left(\tfrac{2}{\delta}\right)\frac{\kappa^2}{\|A\|^2}\frac{1}{\varepsilon^2}\right), \quad \widetilde{\mathcal{O}}\left(\log\left(\tfrac{2}{\delta}\right)\frac{\kappa^4}{\|A\|^4}\frac{1}{\varepsilon^2}\right)$$

- Number of (non-Clifford) gates: $n_{gates} = \widetilde{\mathcal{O}}\left(\lambda^2\frac{\kappa^2}{\|A\|^2}\log^2(1/\varepsilon)\right)$

Extra content state preparation

Main result complexities

- Discretized L_2-norm filling-fraction ($N := 2^n$) as

$$\mathcal{F}_f^{[N]} := \frac{\sqrt{\frac{(b-a)}{N}\sum_{x=0}^{N-1}|f(\bar{x})|^2}}{\sqrt{(b-a)|f|_{\max}^2}} \approx \frac{\sqrt{\int_a^b |f(\bar{x})|^2 d\bar{x}}}{\sqrt{(b-a)|f|_{\max}^2}} =: \mathcal{F}_f^{[\infty]}$$

- **Theorem I**: Given a degree d_δ polynomial approximation $\tilde{f}$ of f,$^{(*)}$ we can prepare a quantum state $|\Psi_{\tilde{f}}\rangle$ that is ε-close in trace distance to $|\Psi_f\rangle$ using $O\left(\frac{n d_\delta}{\mathcal{F}_{\tilde{f}}^{[N]}}\right)$ gates + 4 ancilla qubits, for $\delta = \varepsilon \min\{\mathcal{F}_f^{[N]}, \mathcal{F}_{\tilde{f}}^{[N]}\}$.

$(*)$ when $\tilde{f}(\cdot)$ applied to $\sin\left(\frac{x}{N}\right)$ approximates $\frac{f(\bar{x})}{|f|_{\max}}$ to L_∞-error on $[a, b]$

Main result complexities simplified

- **Theorem II**: For sufficiently smooth functions f,$^{(*)}$we can prepare a quantum state $|\Psi_{\tilde{f}}\rangle$ that is ε-close in trace distance to $|\Psi_f\rangle$ using

$$\widetilde{\boldsymbol{O}}\left(\frac{n\log(\varepsilon^{-1})}{\mathcal{F}_{\tilde{f}}^{[N]}}\right) \text{ gates + 4 ancilla qubits.}$$

$(*)$ need L_∞-approximation $\delta \propto \exp(-d_\delta)$ for degree d_δ polynomial

- Show analytical results via minimax polynomial
- In practice use instead (works even better):
  - Remez approximation or even just Local Taylor series
  - L_2-approximation on grid

Complexity comparison literature

	# Non-Clifford gates	# Ancilla qubits	Rigorous error bounds	Function
QET (this work)	$\mathcal{O}\left(\frac{n d_\epsilon}{\mathcal{F}_{\tilde{f}}^{[N]}}\right)$	4	✓	Polynomial/Fourier approximation
Black-box amplitude oracle	$\mathcal{O}\left(\frac{g_\epsilon^2 \tilde{d}_\epsilon}{\mathcal{F}_f^{[N]}}\right)$	$\mathcal{O}(g_\epsilon \tilde{d}_\epsilon)$	✓	General
Grover-Rudolph amplitude oracle	$\mathcal{O}\left(n g_\epsilon^2 \tilde{d}_\epsilon\right)$	$\mathcal{O}(g_\epsilon \tilde{d}_\epsilon)$	✓	Efficiently integrable probability distribution
Adiabatic amplitude oracle	$\mathcal{O}\left(\frac{g_\epsilon^2 \tilde{d}_\epsilon}{\left(\mathcal{F}_f^{[N]}\right)^4 \epsilon^2}\right)$	$\mathcal{O}(g_\epsilon \tilde{d}_\epsilon)$	✓	General
Matrix product state	$\mathcal{O}(n)$	0	✗	Matrix product state $d = 2$ approximation

Note: g_ε-bit amplitude oracles with degree $\tilde{d}_\varepsilon$ piecewise polynomial approximation ($\tilde{d}_\varepsilon \neq d_\varepsilon$ in general)

Outlook

- Introduced versatile method for preparing a quantum state whose amplitudes are given by some known function
- Based on the QET, orders of magnitude savings in ancilla qubits
- Needed: More detailed practical resource estimates, more functions
- Open questions:
  - Example square root function $\sqrt{\bar{x}}$ for $\bar{x} \in [0,1]$, non-differentiable at $\bar{x} = 0$

    → use $\sqrt{\bar{x} + a}$ instead?
  - Multivariate functions via multivariate QET?

Thank you.

Algorithm: Setup

- Treat special case: $a = -1, b = 1$, with function $f(x) = f(-x)$
- Goal: Prepare the n-qubit quantum state

$$|\Psi_f\rangle = \frac{1}{\mathcal{N}_f} \cdot \sum_{x=-N/2}^{N/2-1} f(\bar{x})|x\rangle \text{ with } \bar{x} = \frac{2x}{N}, \text{ and } \mathcal{N}_f = \sqrt{\sum_{\bar{x}} f(\bar{x})}$$

1. Start with block encoding of $A = \sum_{x=-N/2}^{N/2-1} \sin\left(\frac{2x}{N}\right)|x\rangle\langle x|$
2. QET to convert into block encoding of $\sum_{x=-N/2}^{N/2-1} f(\bar{x})|x\rangle\langle x|$
3. $O\left(1/\mathcal{F}_{\tilde{f}}^{[N]}\right)$ rounds of exact amplitude amplification (extra ancilla)

- Need to start with (extensive) classical pre-processing!

Algorithm: Quantum circuits

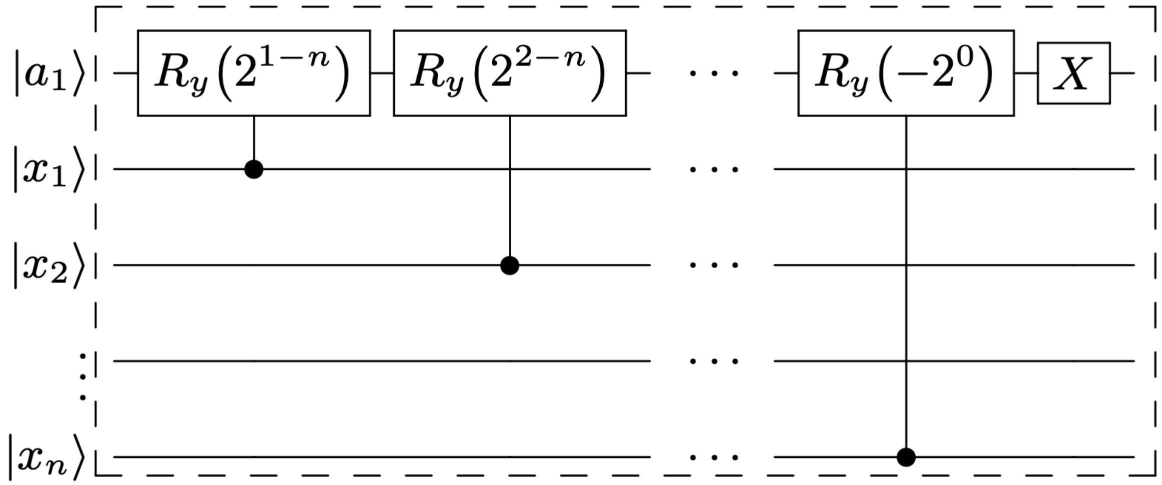

1. $U_{\sin}$ block encoding circuit

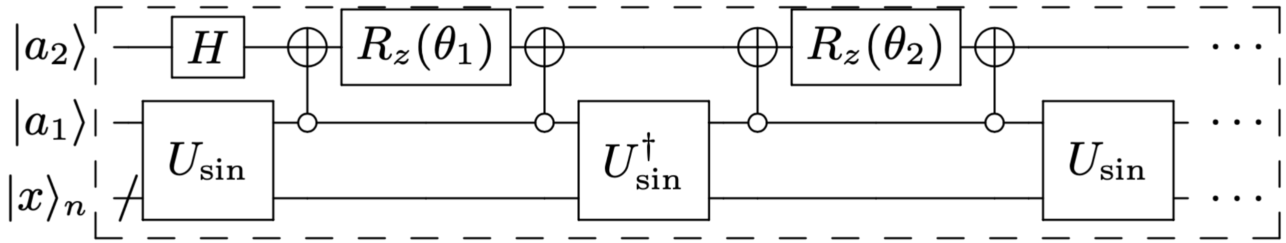

2. $U_{\tilde{f}}$ block encoding circuit

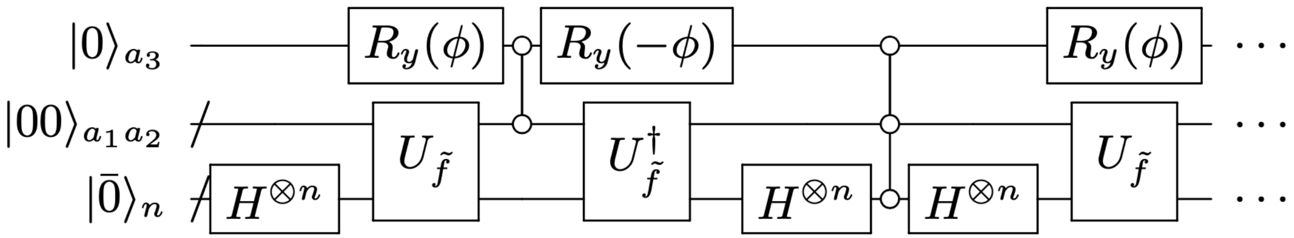

3. Amplitude amplification (exact) circuit

Algorithm: Classical pre-computation

- Compute polynomial $h(y)$ such that

  $|h(y)|_{\max}^{y\in[-1,1]} \leq 1$ and $\left|h\big(\sin(y)\big) - \frac{f(y)}{|f(y)|_{\max}^{y\in[-1,1]}}\right|_{\max}^{y\in[-1,1]} \leq \delta$

  leading to approximation $\tilde{f}(x) \coloneqq \mathrm{h}\big(\sin(\bar{x})\big)$

  (Remez algorithm / local Taylor series / L_2-approximation on grid / …)

- Compute discretized L_2-norm filling-fraction $\mathcal{F}_{\tilde{f}}^{[N]} \approx \mathcal{F}_{\tilde{f}}^{[\infty]}$ of $\tilde{f}(x)$

  (choose depending on how large $N = 2^n$ is)

- Compute QET angle set $\{\theta_1, \theta_2, \cdots, \theta_d\}$ of polynomial $\tilde{f}(x)$

  (analytically good Haah method or numerically good Dong *et al.* method)

Extension: Non-smooth functions

- First approach: Use coherent inequality test with flag qubit for piecewise QET polynomial implementation

→ for k discontinuities this requires $(k+n)$ ancilla qubits and $2kn$ Toffoli gates for the inequality comparison

- Second approach: Example triangle function for $\bar{x} \in [0,1]$

$$f(\bar{x}) = \begin{cases} \bar{x} & 0 \leq \bar{x} \leq 1/3 \\ \frac{1}{2}(1-\bar{x}) & 1/3 < \bar{x} \leq 1 \end{cases}$$

instead use

$$\bar{f}(\bar{x}) = \begin{cases} \bar{x} & 0 \leq \bar{x} \leq \frac{1}{3} \\ \text{Unspecified} & \frac{1}{3} < \bar{x} < 2 \\ \frac{1}{2}\left(\frac{7}{3} - \bar{x}\right) & 2 \leq \bar{x} \leq \frac{7}{3} \end{cases}$$

→use coherent inequality test to flip for $\bar{x} > \frac{1}{3}$ and in the end reverse this inequality check

Extension: Fourier based QET

- Block-encoding of A is replaced by controlled time evolution

$$V(A) := |0\rangle\langle 0| \otimes 1 + |1\rangle\langle 1| \otimes \exp(iAt)$$

- Fourier-based QET uses calls to $V(A)$, together with single-qubit-rotations, to apply a function $f(\cdot)$ in Fourier series form to A
- We can implement $V(A)$ for diagonal $A = \sum_x \bar{x}|x\rangle\langle x|$ using n controlled Z-rotations
- Example with compact Fourier series: Cycloid function

→ $n = 32$ for $\bar{x} \in [0,2\pi]$, gives 7.35×10^5 Toffoli gates
+ 3 ancillas qubits

From wikipedia