Quantum Algorithm Development

Mario Berta

Physics Colloquium October 30th, 2023

Information science

- Theory of information processing: Mathematical foundations in 1940s
- Abstract theory independent of implementation

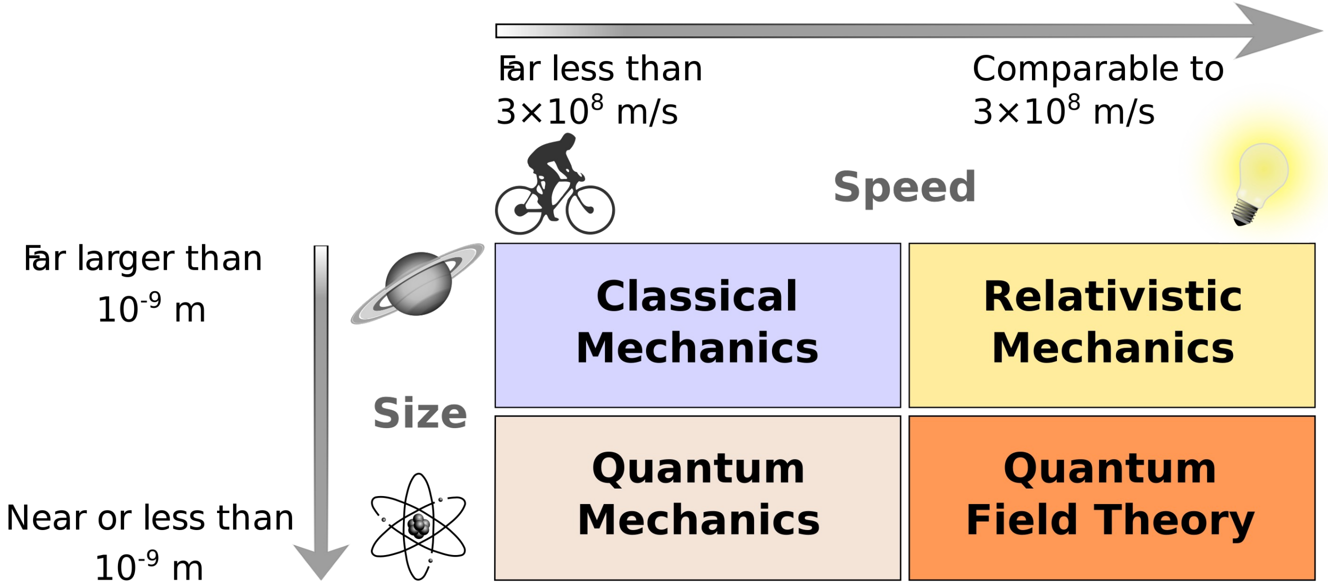

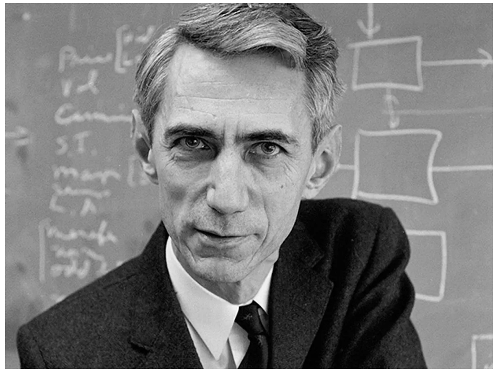

Claude Shannon

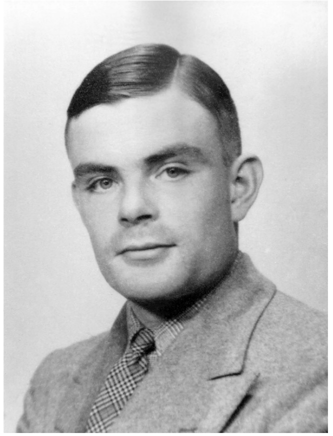

Alan Turing

- Question: Also independent of underlying physics?

Quantum information science

- Notion of information for microscopical systems described by quantum mechanics? Quantum information $\neq$ classical information

- Bell's theorem (1964): Quantum mechanics is incompatible with local hidden-variable theories

John Stewart Bell

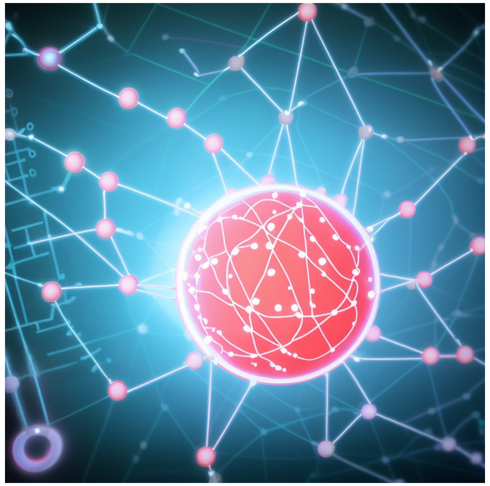

- New research area based on quantum technologies: Computing, communication, cryptography, sensing, …

New group @RWTH Aachen

JARA INSTITUTE QUANTUM INFORMATION

AN INITIATIVE OF

RWTH AACHEN UNIVERSITY

JÜLICH Forschungszentrum

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


Mario Berta
Professor of Physics

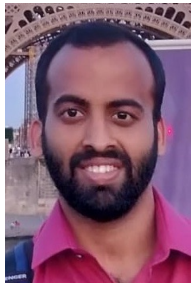
Sreejith Sreekumar
Postdoc

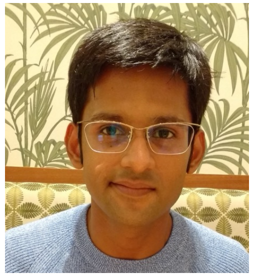
Aditya Nema
Postdoc

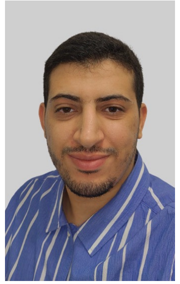
Aadil Oufkir
Postdoc

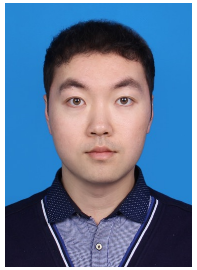
Yongsheng Yao
Postdoc (incoming)


Michael X. Cao
Postdoc (incoming)


Tobias Rippchen
PhD student


Julius Zeiss
PhD student


Gereon Koßmann
PhD student


Navneeth Ramakrishnan
PhD student (Imperial)


Samson Wang
PhD student (Imperial)

+ 5x Master students
RWTH Physics /
Computer Science

Theory of quantum information science

- Our focus areas:
  1. Mathematical foundations of quantum information
  2. Quantum algorithm development
- Cluster of Excellence: Matter and Light for Quantum Computing (ML4Q)
- Visiting Reader at Department of Computing Imperial College London
- Industry ties with Amazon Web Services Center for Quantum Computing

ML4Q
MATTER AND LIGHT FOR
QUANTUM COMPUTING

erc European Research Council

Imperial College London

EPSRC

aws

Today: Quantum algorithms

Origins of quantum computing

- Understanding physics with computers (1981):

"trying to find a computer simulation of physics seems to me to be an excellent program to follow out (...) nature is not classical, dammit, and if you want to make a simulation of nature, you will better make it quantum mechanical, and by golly it is a wonderful problem, because it does not look so easy"

Richard Feynman

Peter Shor

- First query complexity separation results in 1990s
- Breakthrough prime factorization (1994):

n-bit integer factorization in quantum complexity $O(n^2 \log n)$
versus classical complexity $O\left(\exp\left(1.9 \cdot n^{\frac{1}{3}} (\log n)^{\frac{2}{3}}\right)\right)$

Quantum algorithms research

- Steady progress on quantum algorithm development since 1990s, recent flurry of activities and results

- Breakthrough prize in physics 2023:

  “foundational work in the field of quantum information”

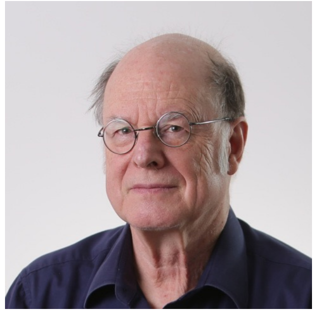

Charles H. Bennett

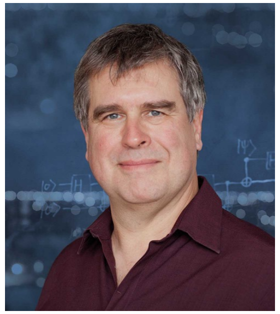

Gilles Brassard

David Deutsch

Peter Shor

- Ultimate goal: Quantify classical-quantum complexity boundary

Classical versus quantum technologies

- *Do algorithms based on quantum components, including*
  - *quantum processing units (QPU)*
  - *quantum random access memory (QRAM)*

provide computational advantages compared to 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 routines whenever possible

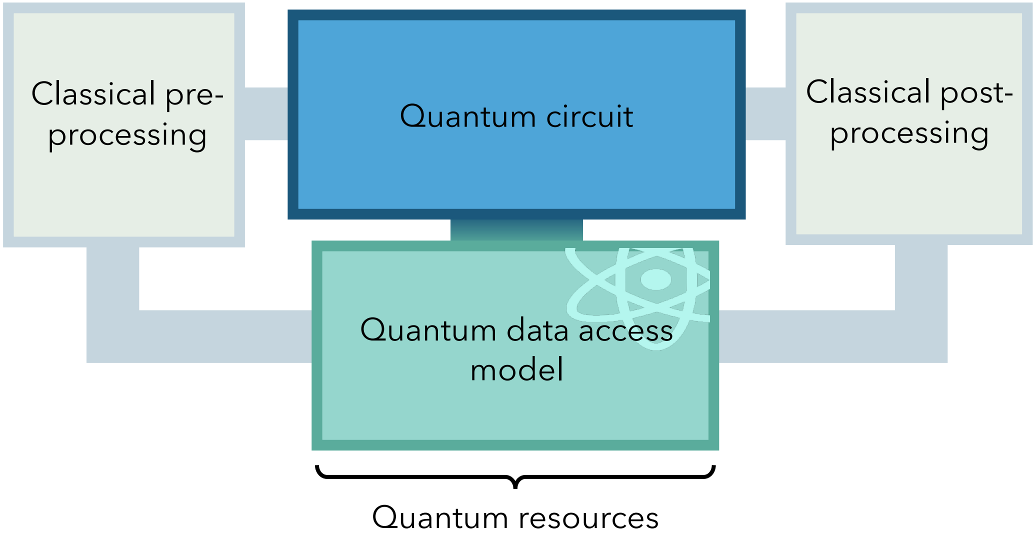

Basics

Classical vs quantum model of computation

- What can be realized within abstract model of quantum mechanics?
- Bloch sphere representation of a two-level system – aka quantum bit – aka qubit:

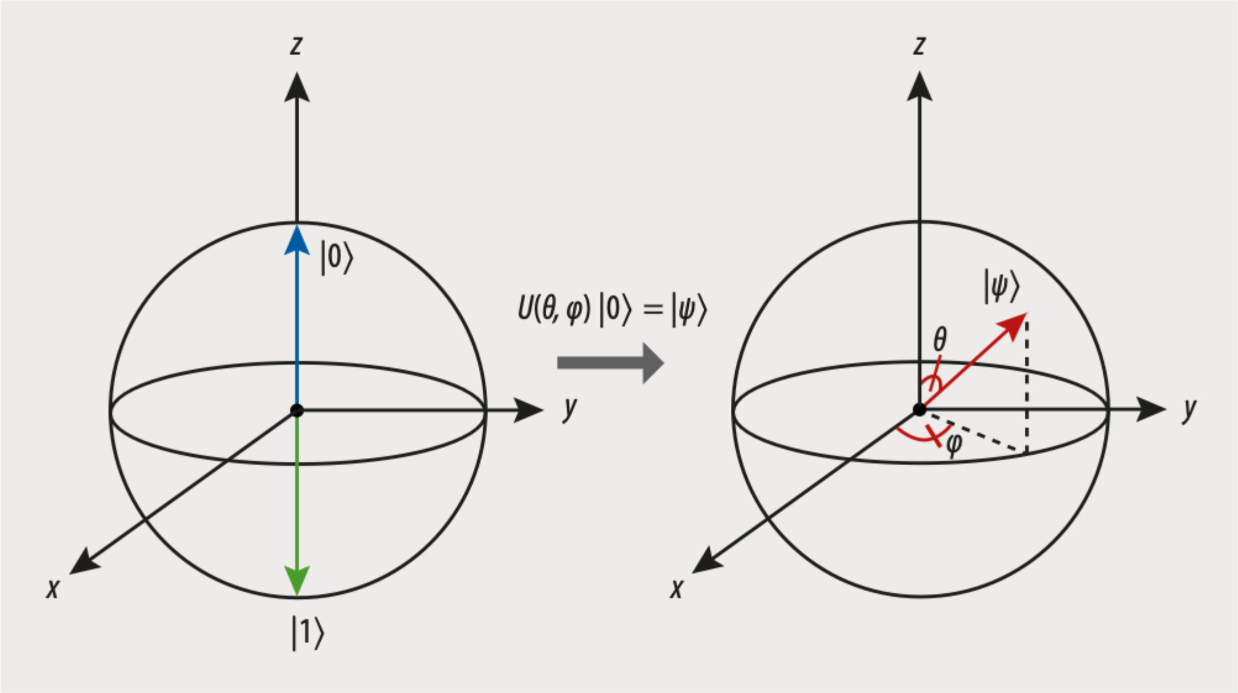

Classical vs quantum model of computation

- What can be realized within abstract model of quantum mechanics?
- Classical versus quantum circuit model:

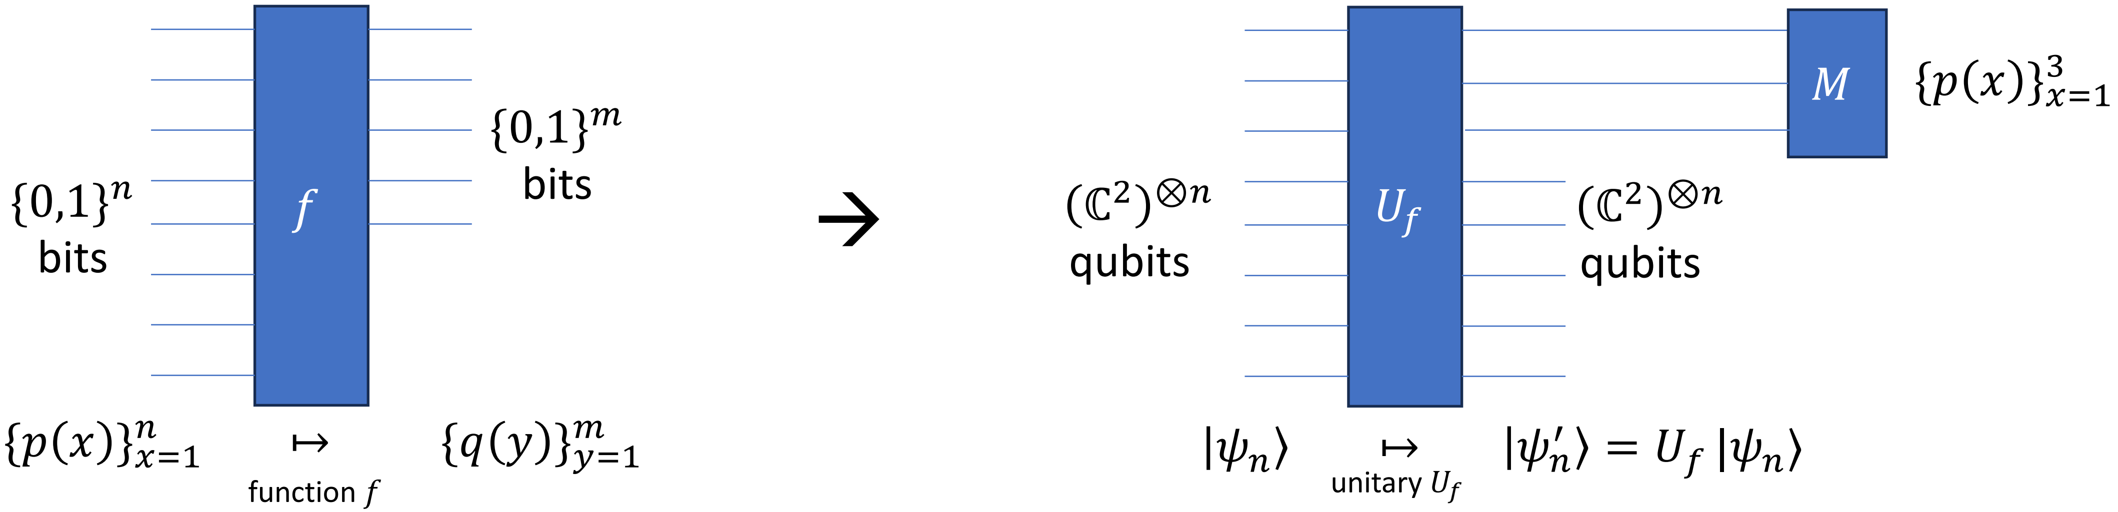

- Schrödinger time evolution operator $U_f = e^{-iH_f t}$ + measurements

Quantum circuit model

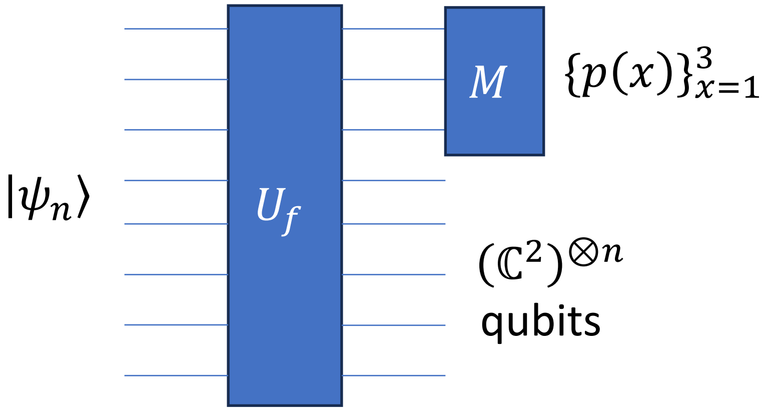

- Where does quantum speed-up come from?

  → precise mathematical reason: complex numbers allow for constructive interference

- Decomposition into set of elementary quantum gates, e.g., single qubit Pauli X, Y, Z and T gates, together with two qubit $CNOT$ gate

- Quantum complexity $=$ minimal number of elementary quantum gates

Quantum hardware

Quantum hardware

- Different technologies: Superconductors, ion traps, neutral atoms, photonics, etc.
- Groundbreaking physics experiments with 6 orders of magnitude improvements in 30 years+

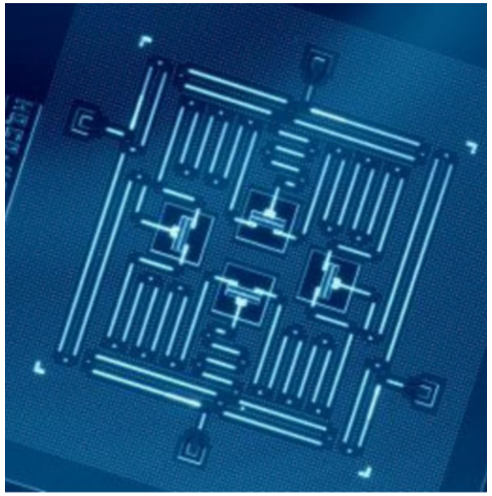

4 IBM superconducting qubits

Alain Aspect

John Clauser

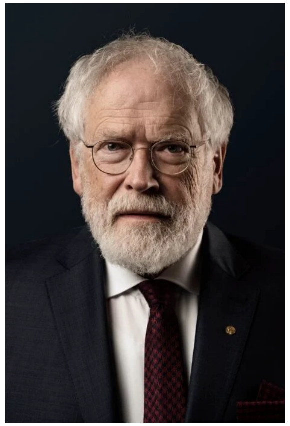

Anton Zeillinger

- Nobel prize in physics 2022:

“experiments with entangled photons, establishing the violation of Bell inequalities and pioneering quantum information science”

Current quantum technologies

- Led by quantum industry, hundreds of billion dollars scale
- IBM Quantum: (up to 433 qubits)

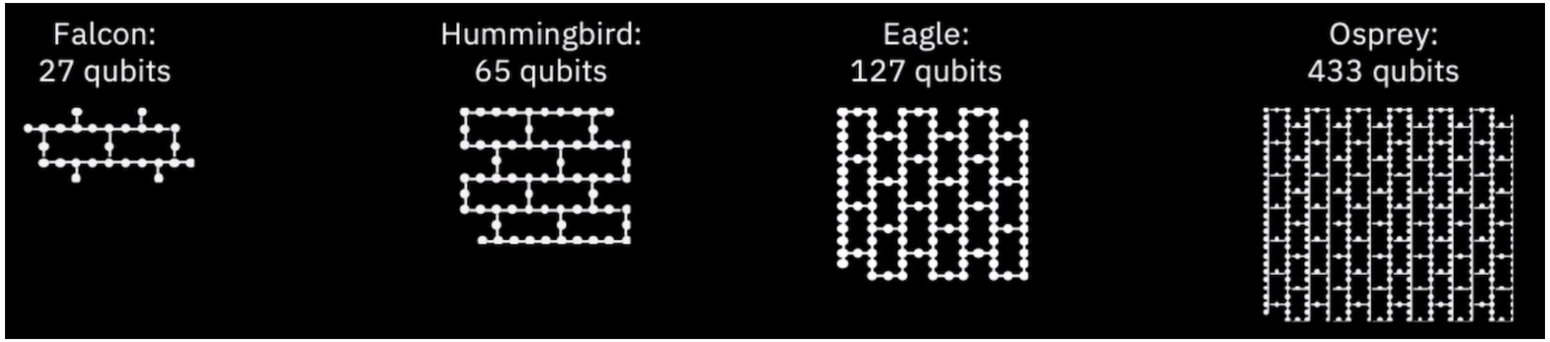

- Google Quantum: (54 qubits)

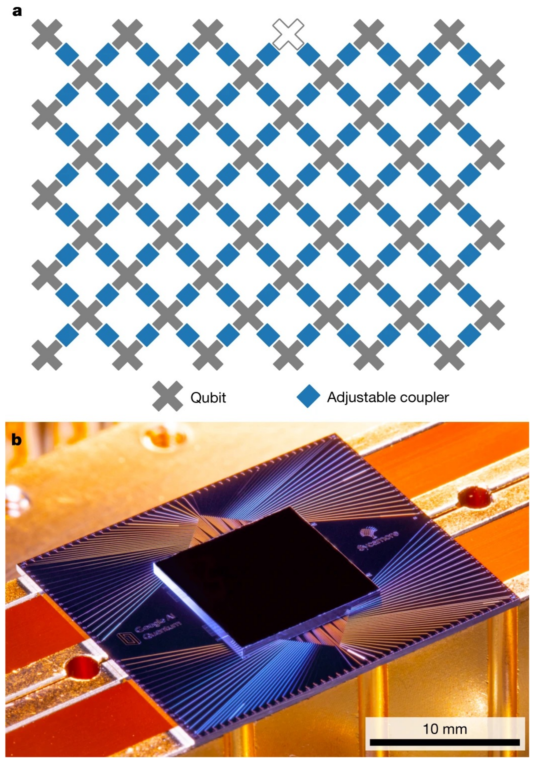

- Severe restrictions: Qubit count, qubit connectivity, one qubit gates (Pauli), two qubit gates (CNOT), read-out errors (measurement), clock speed, and more

Quantum algorithm design

Regimes for quantum algorithm design

- Nascent state of quantum technologies gives noisy and intermediate scale quantum (NISQ) computing, i.e.,
  - Quantum annealers
  - Analogue simulators, not universal, not fully programmable
  - NISQ digital quantum circuits, inbuilt noise resilience, error mitigation, severe scaling limitations, etc.
- Versus what one really wants long-term:

Quantum error-corrected and scaling quantum computer

(roughly two orders of magnitude away)

- Any intermediate regimes of interest?

Early fault-tolerant regime

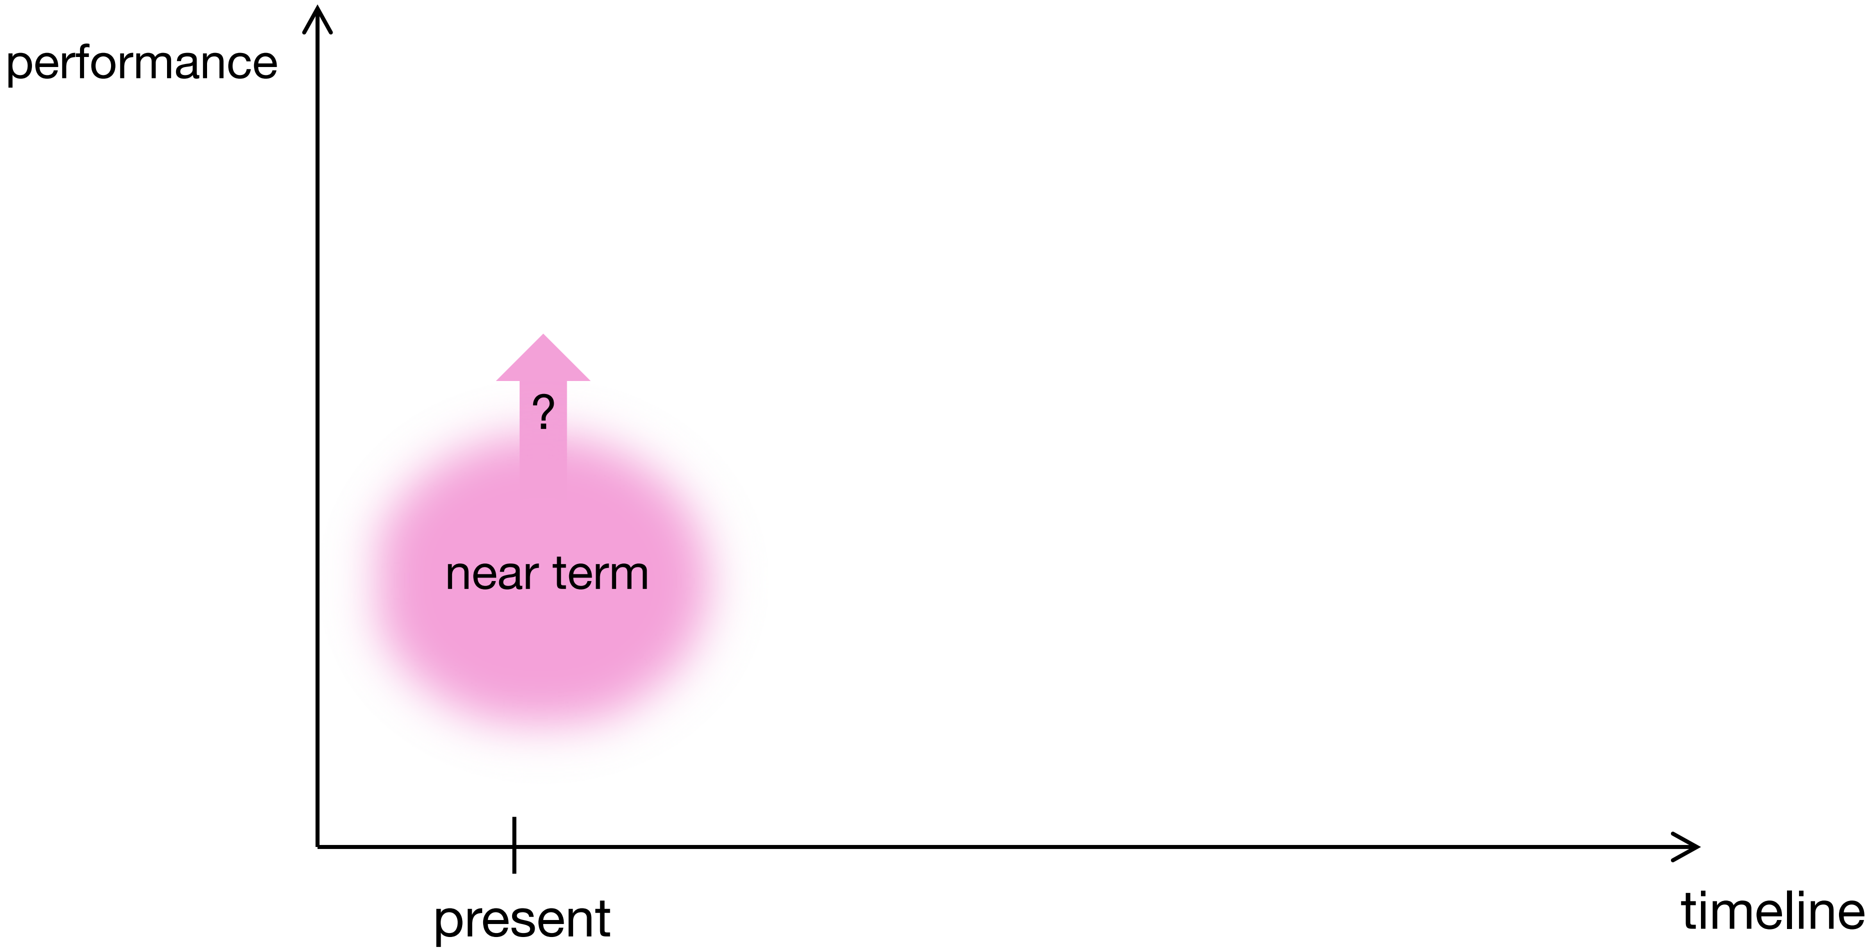

Early fault-tolerant regime

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

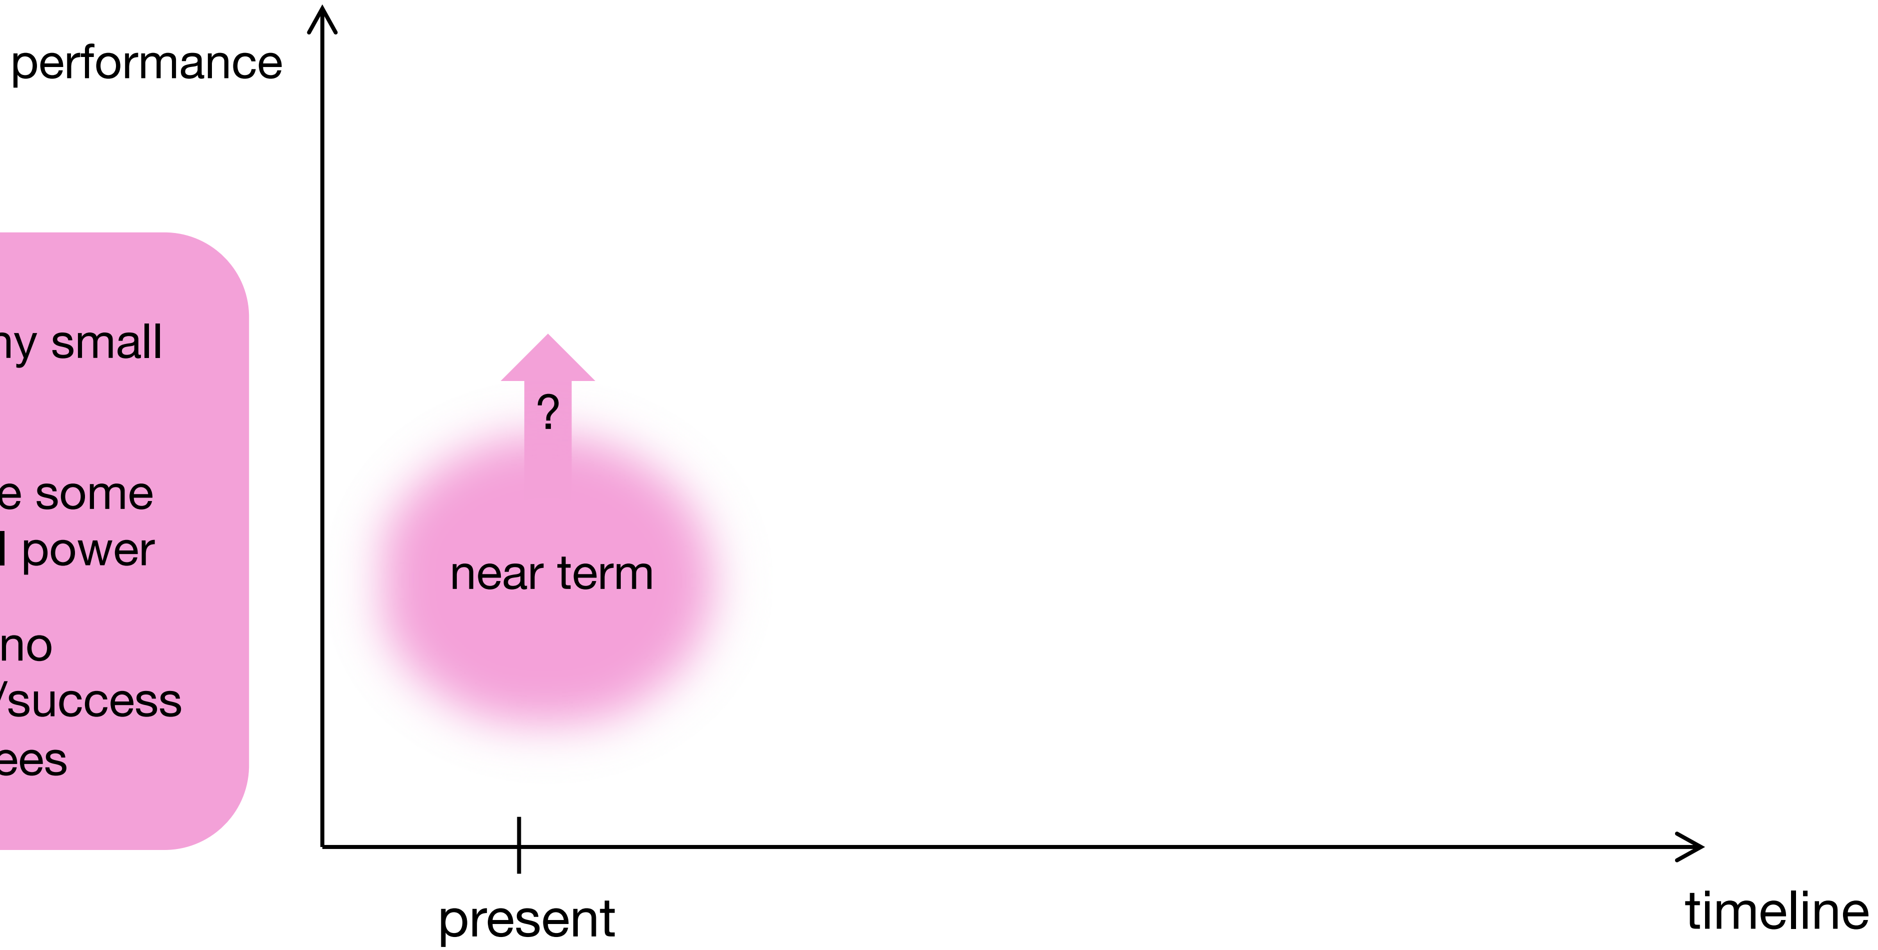

Early fault-tolerant regime

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

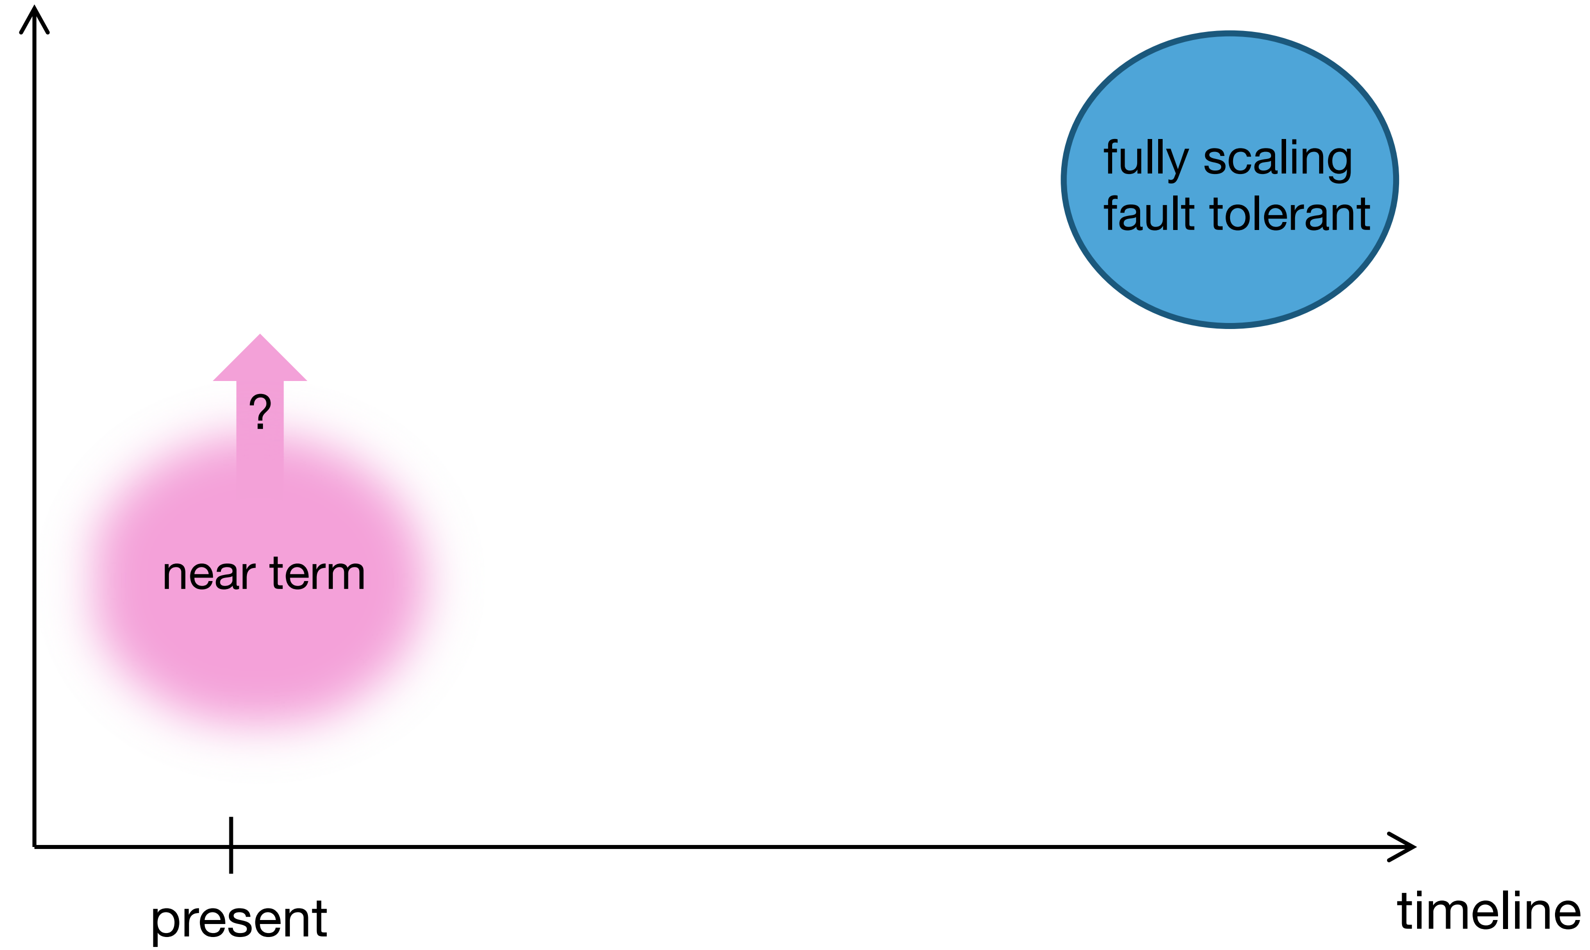

Early fault-tolerant regime

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

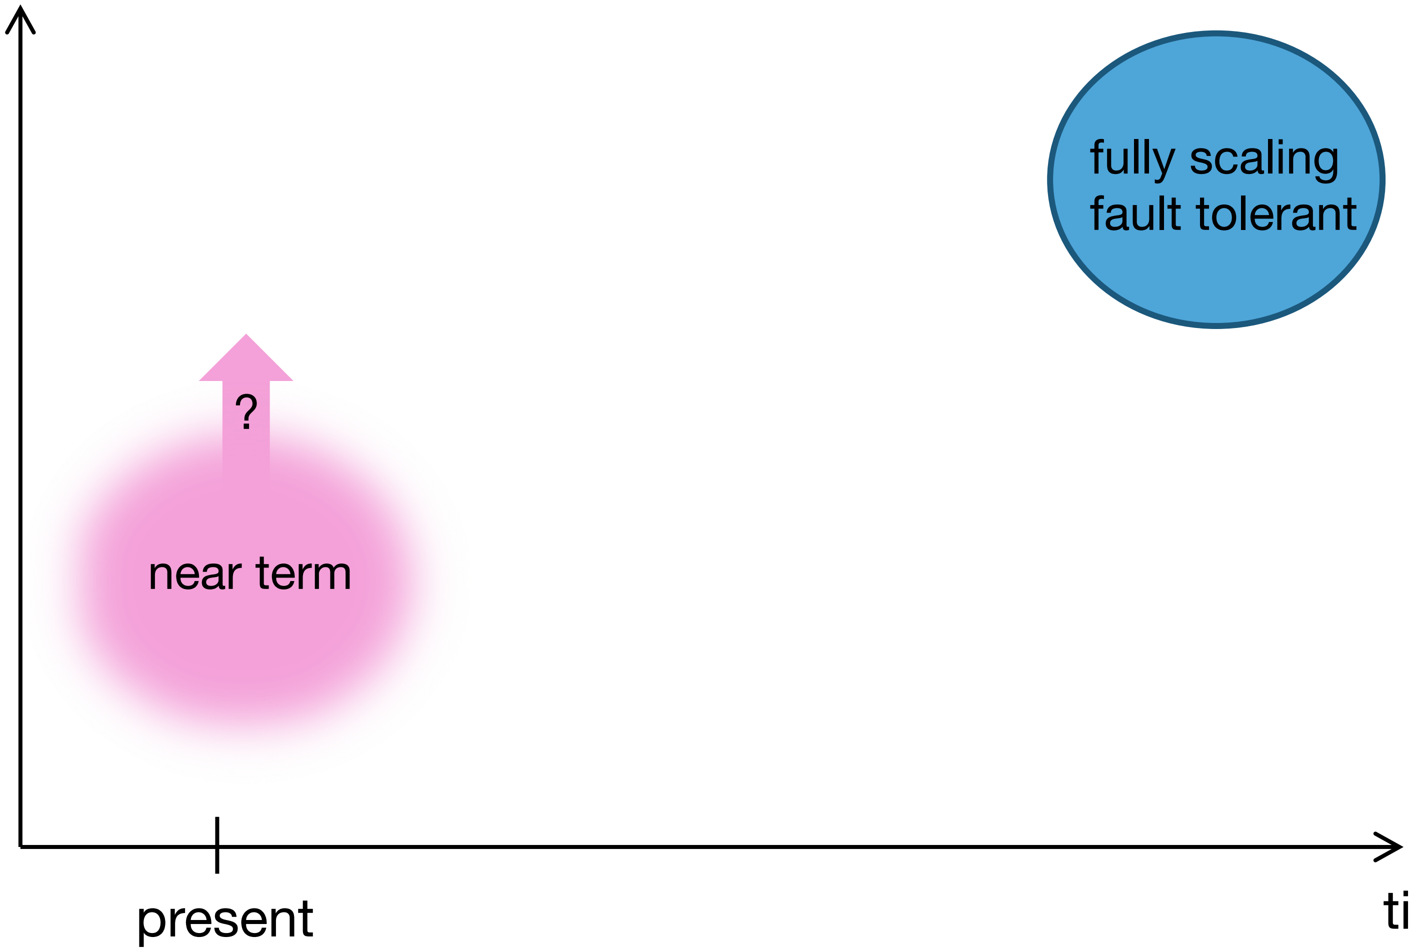

- Large hardware requirements
- Specified in terms of query complexity
- Runtime guarantees
- Caveats can cloud convincing practical advantage

Early fault-tolerant regime

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

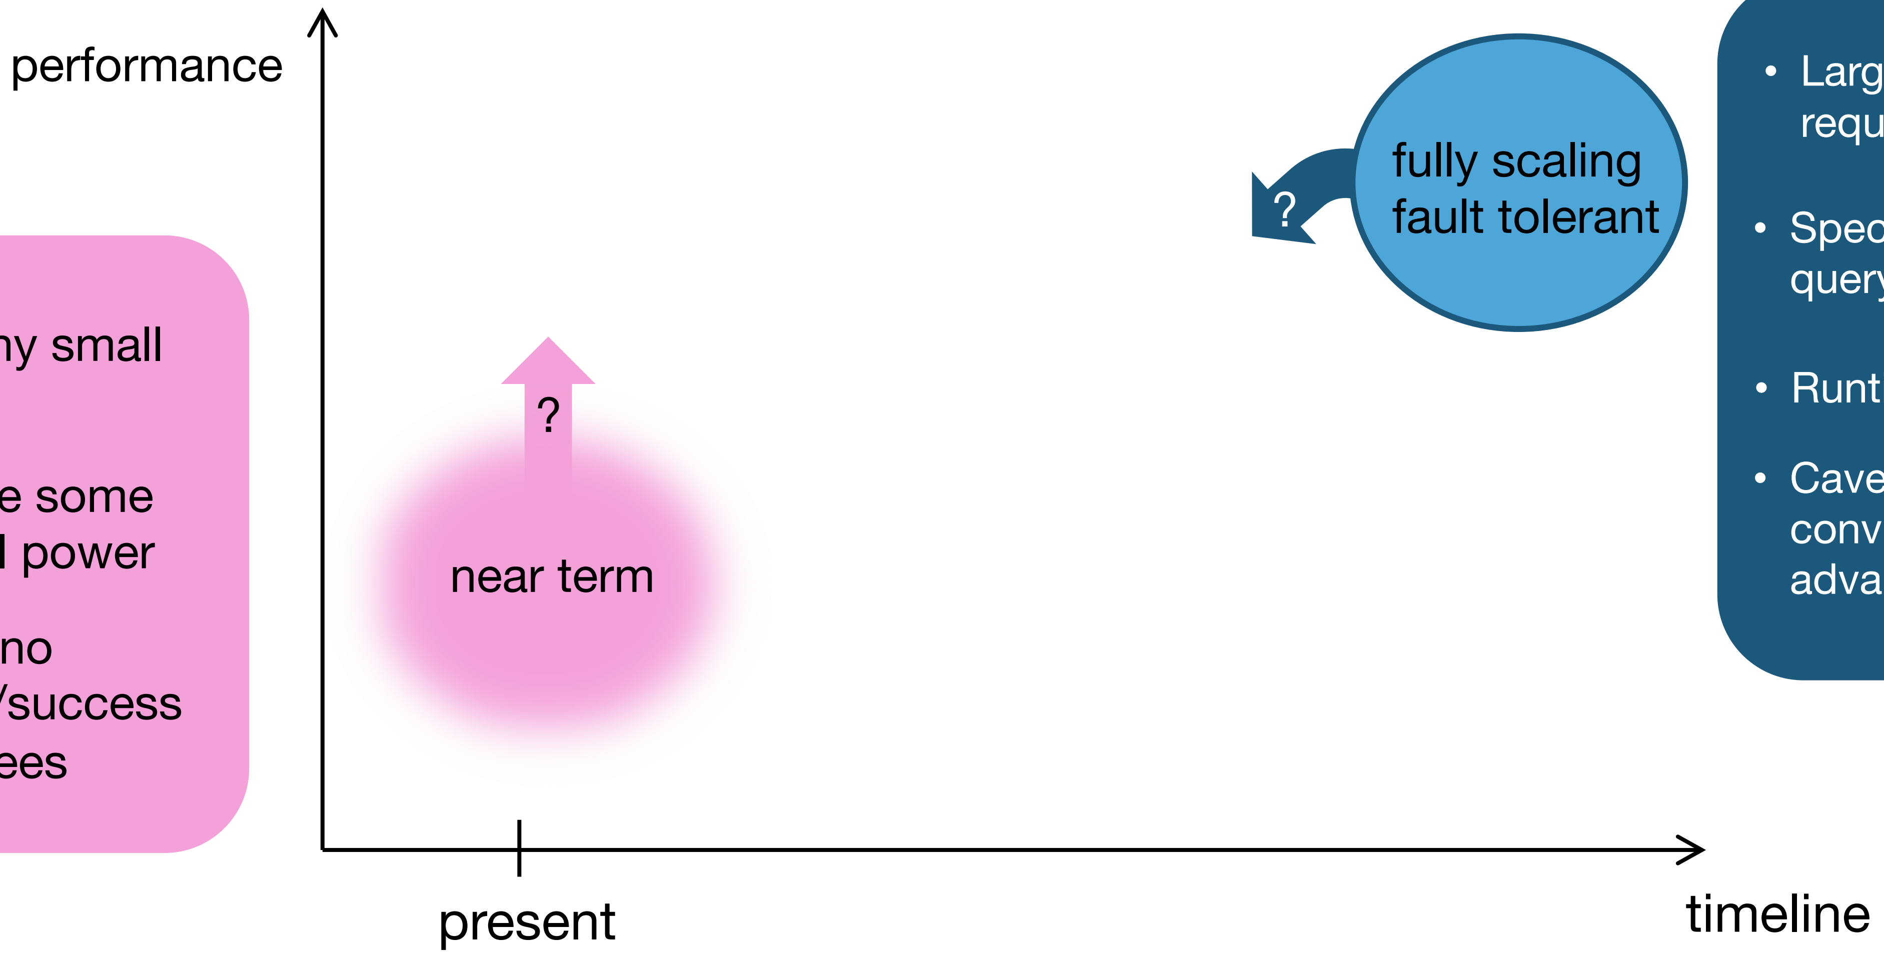

- Large hardware requirements
- Specified in terms of query complexity
- Runtime guarantees
- Caveats can cloud convincing practical advantage

Early fault-tolerant regime

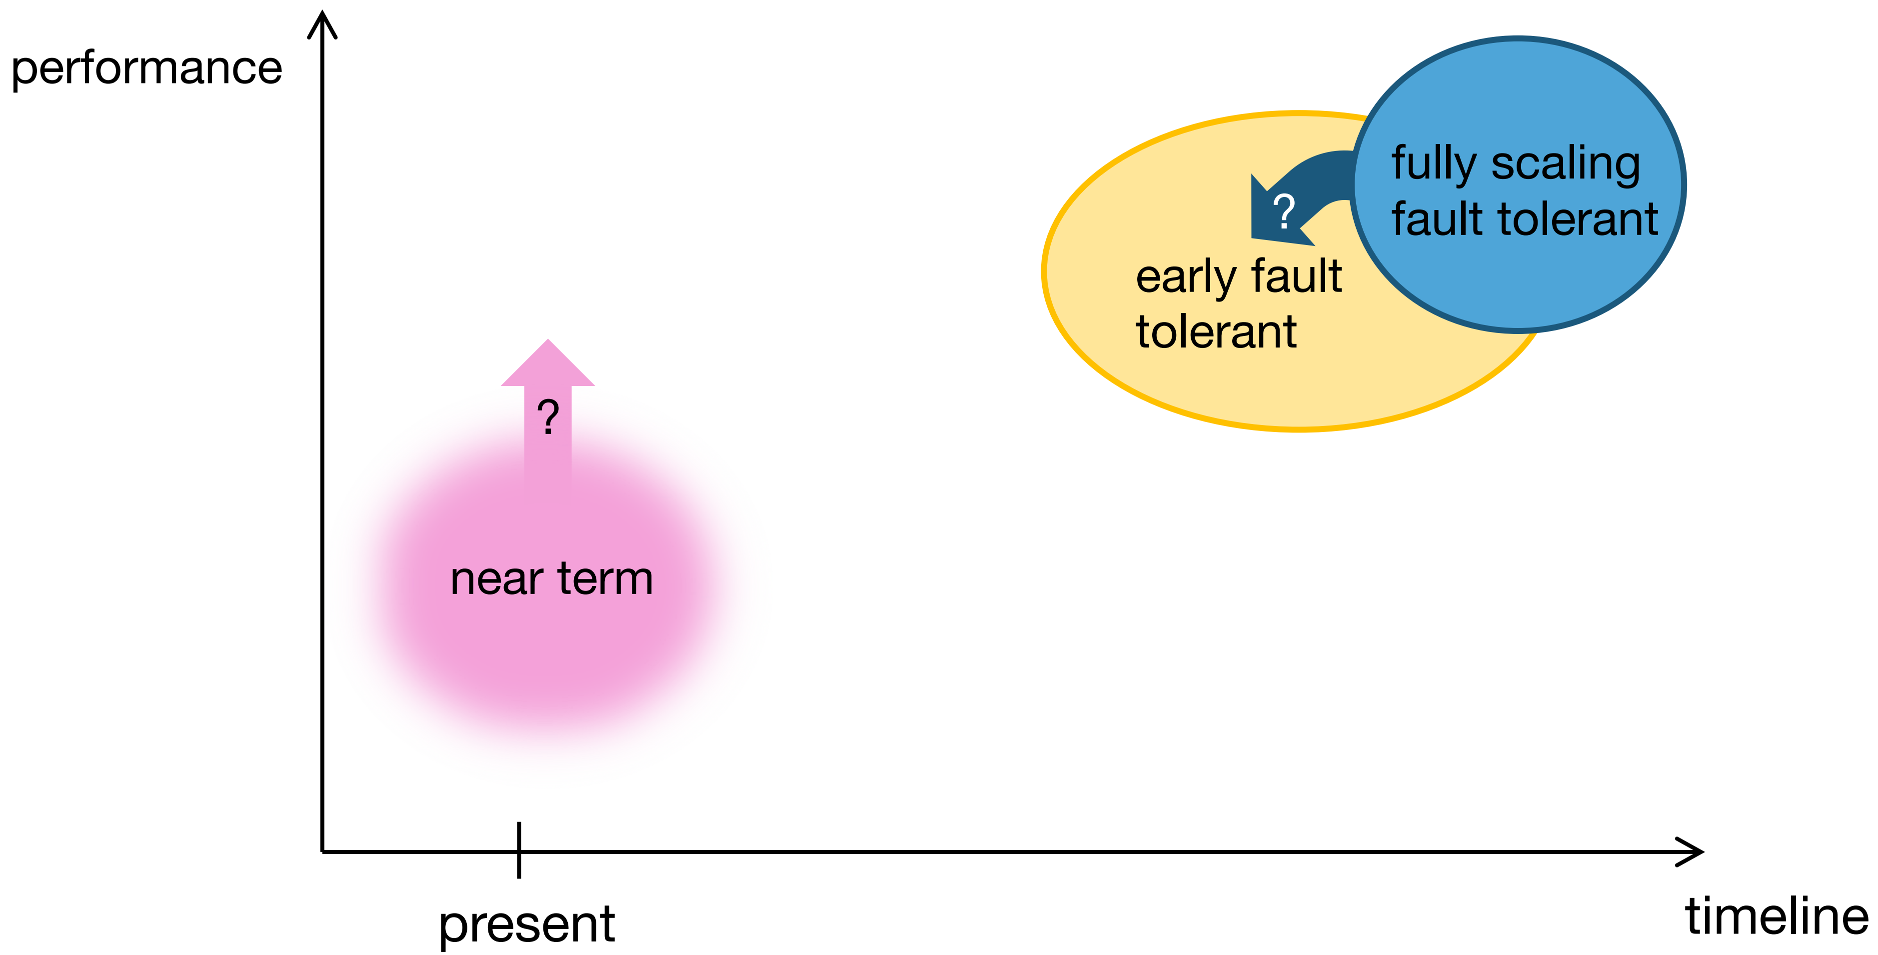

Early fault-tolerant regime

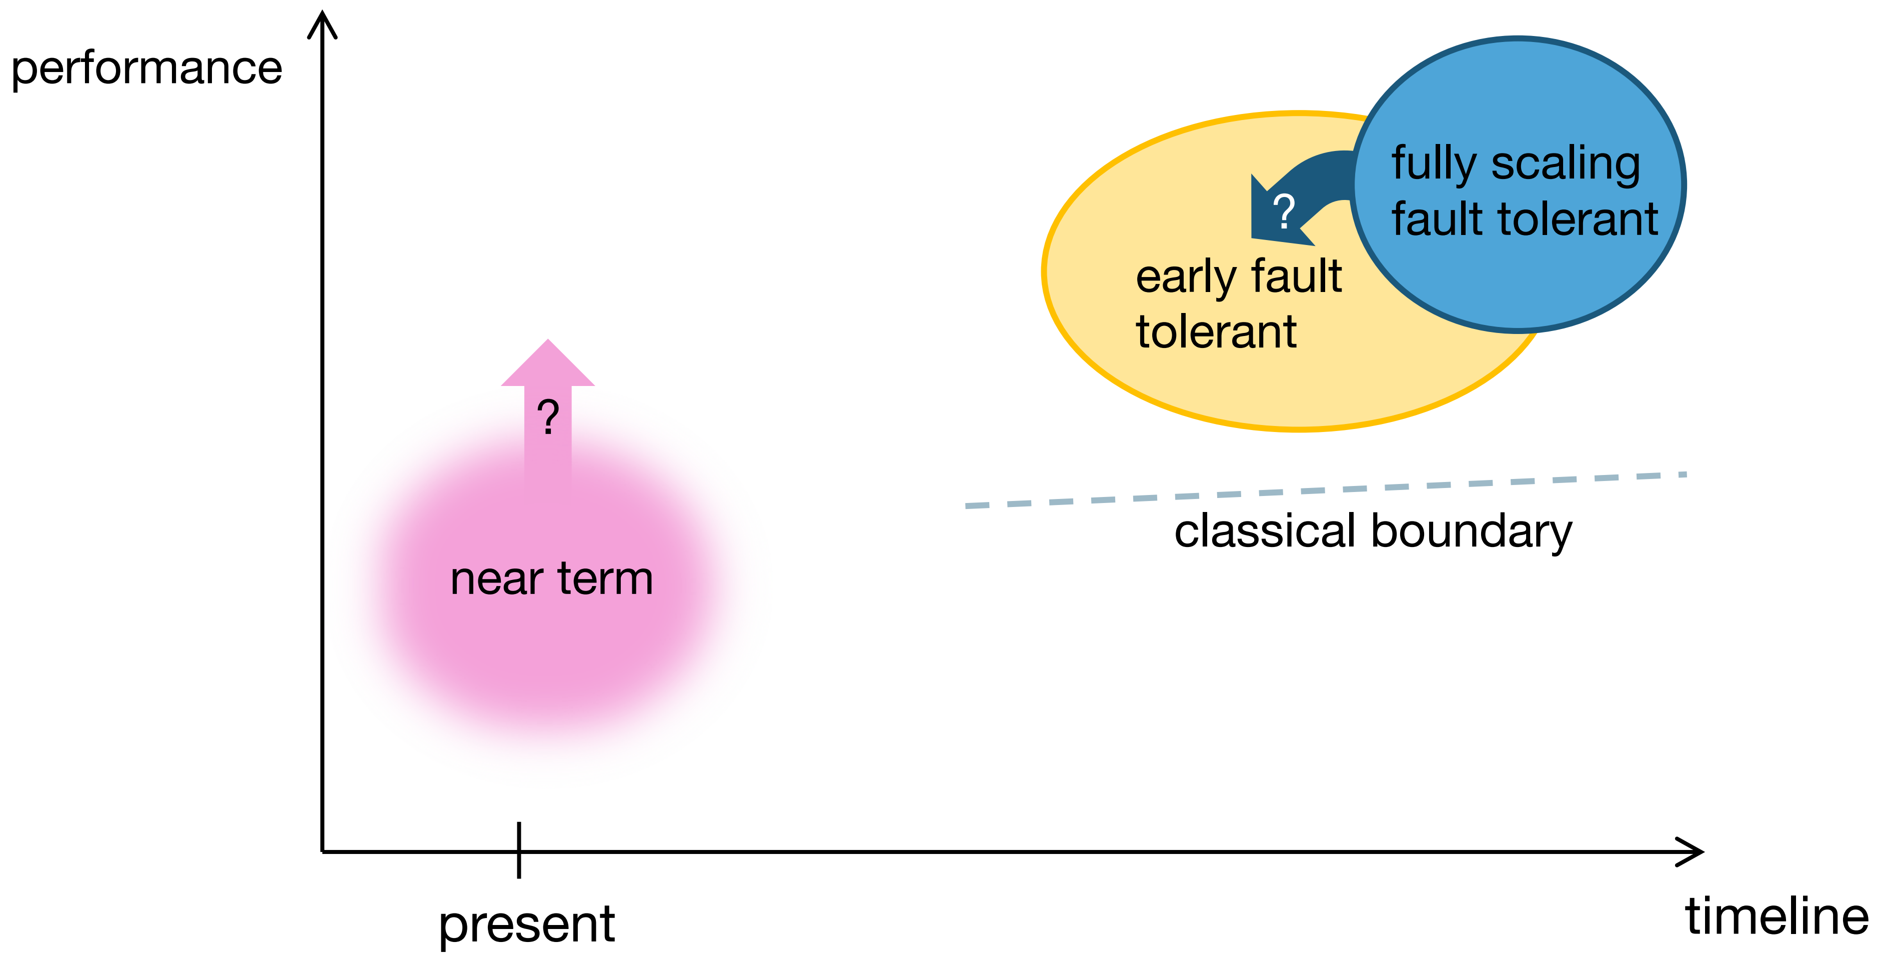

Early fault-tolerant regime

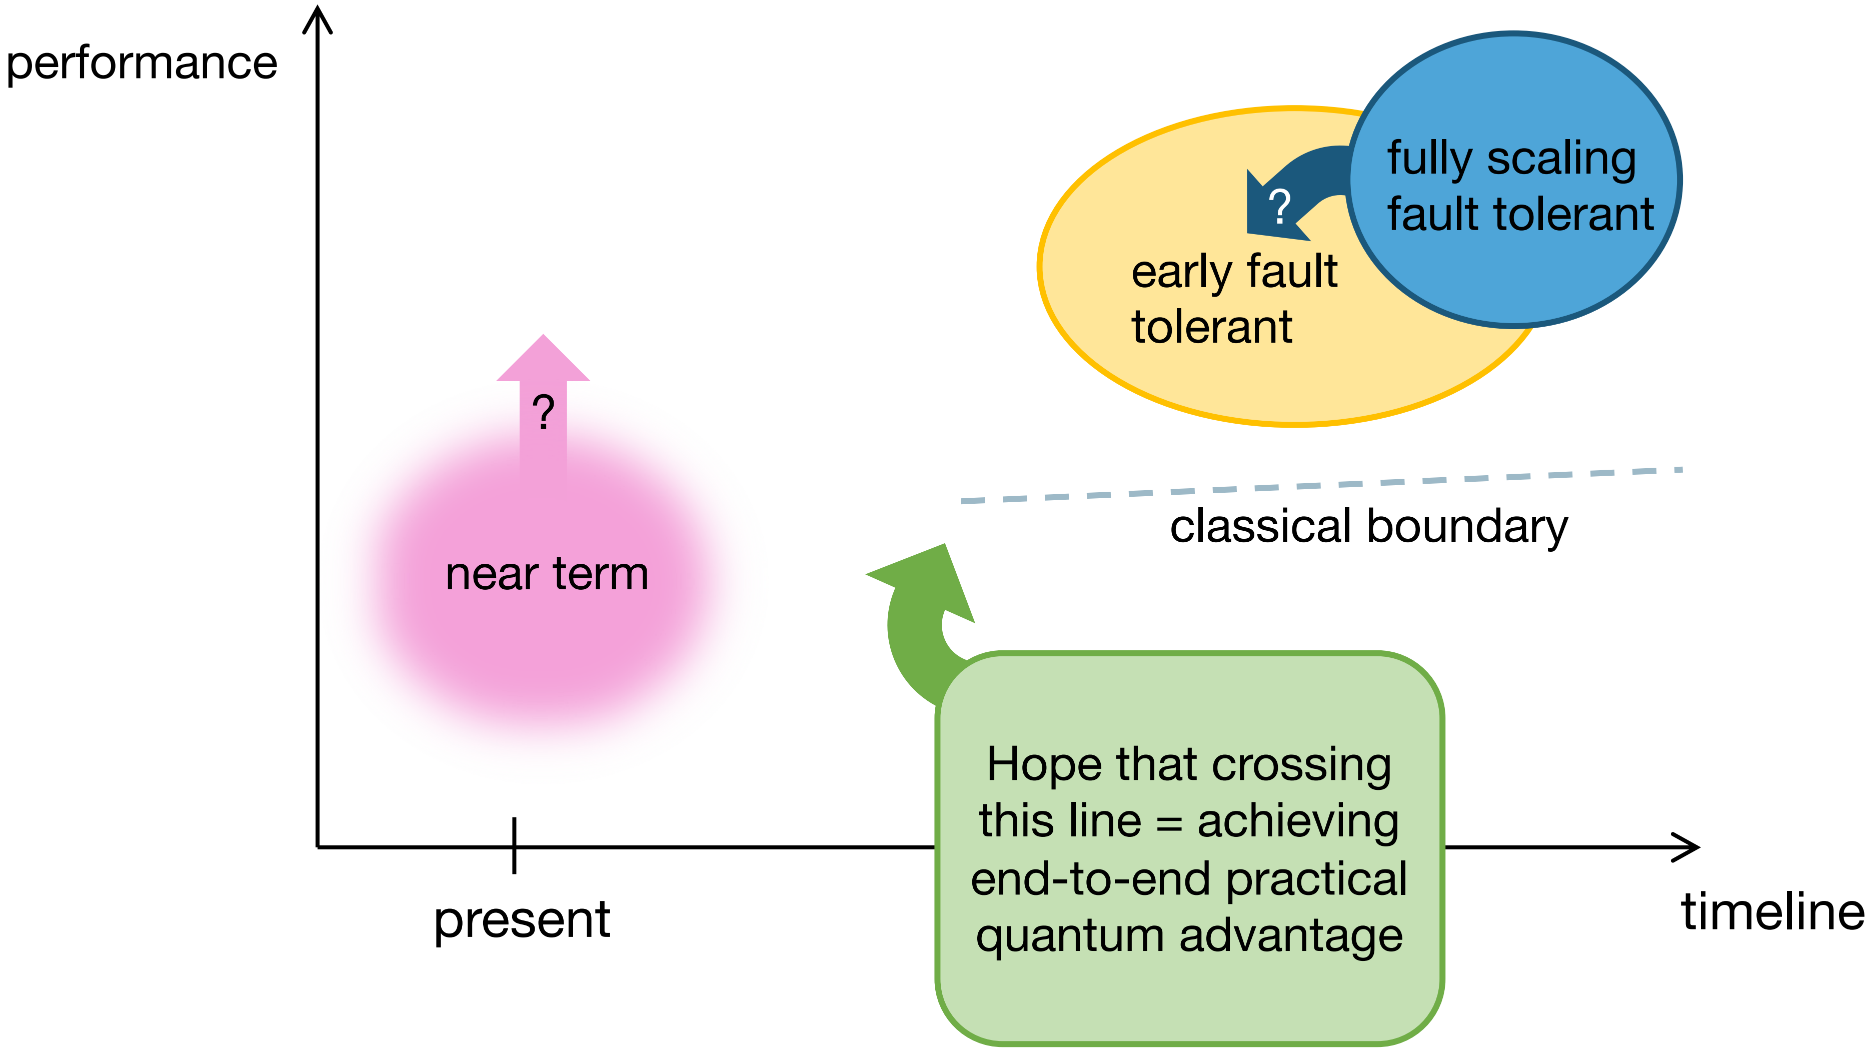

Early fault-tolerance characteristics

- Limited number of logical qubits, with limited 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

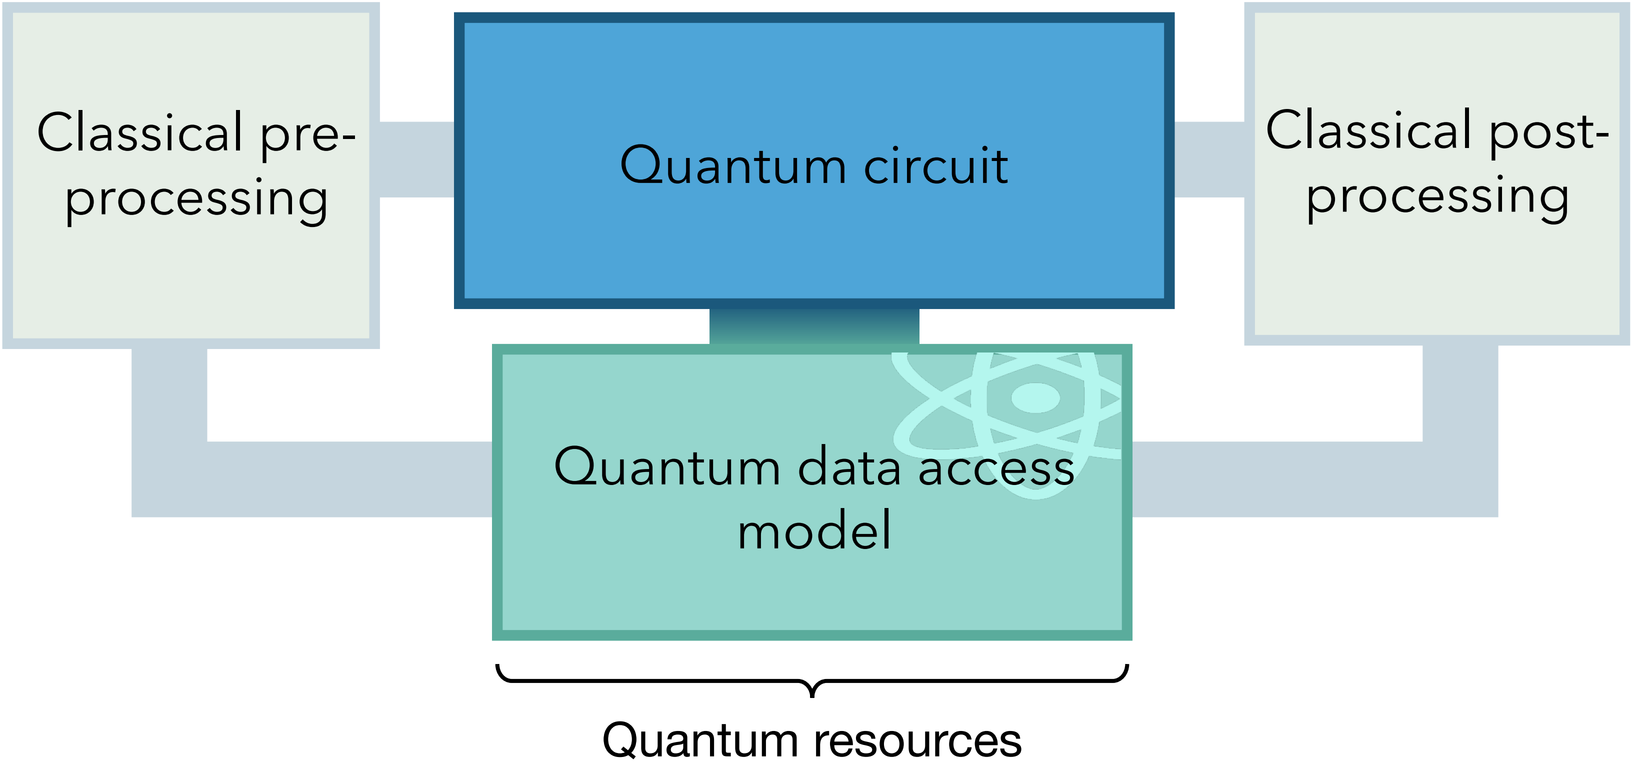

- Resource estimates for comparison with state-of-the-art classical methods

Quantum Algorithms Wiki

Quantum algorithms:

A survey of applications and end-to-end complexities

Alexander M. Dalzell*[1], Sam McArdle*[1], Mario Berta[1,2,3], Przemyslaw Bienias[1], Chi-Fang Chen[1,4], András Gilyén[5], Connor T. Hann[1], Michael J. Kastoryano[1,6], Emil T. Khabiboulline[1,7], Aleksander Kubica[1], Grant Salton[1,4,8], Samson Wang[1,3], and Fernando G. S. L. Brandão[1,4]

[1] *AWS Center for Quantum Computing, Pasadena, CA, USA*
[2] *Institute for Quantum Information, RWTH Aachen University, Aachen, Germany*
[3] *Imperial College London, London, UK*
[4] *Institute for Quantum Information and Matter, Caltech, Pasadena, CA, USA*
[5] *Alfréd Rényi Institute of Mathematics, Budapest, Hungary*
[6] *IT University of Copenhagen, Copenhagen, Denmark*
[7] *Department of Physics, Harvard University, Cambridge, MA, USA*
[8] *Amazon Quantum Solutions Lab, Seattle, WA, USA*

- Brand new (05/10/2023) – available at arXiv:2310.03011

Example applications

Quantum simulation for scientific computing

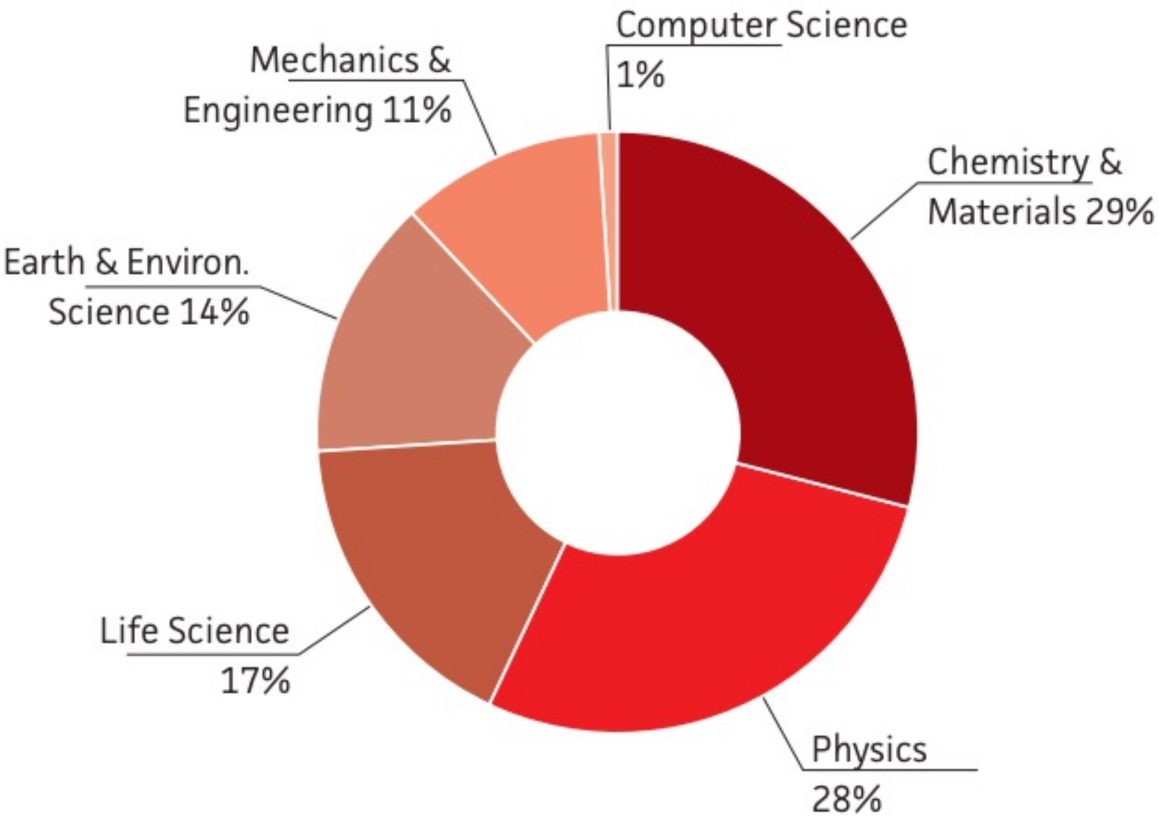

Swiss National Supercomputing Centre Annual Report 2022

Example: Ground state energy estimation

Randomized quantum algorithm for statistical phase estimation

Physical Review Letters (2022) with Campbell and Wan
Quantum Information Processing (QIP) 2022

Quantum many body systems

- Consider n-qubit Hamiltonian

$$\mathrm{H} = \sum_{l=1}^{L} \alpha_l P_l^n \text{ with } P_l^n \ n\text{-qubit Pauli operator,}$$

i.e., $\mathrm{P}_l^n = \mathrm{P}_1 \otimes \cdots \otimes P_n$ with $P_i \in \{X, Y, Z, 1\}$

- Native example: Ising model on two-dimensional square lattice

$$H_{Ising} = C \cdot \sum_{j \in J} Z_{i,j} \otimes Z_{i+1,j} + Z_{i,j} \otimes Z_{i,j+1}$$

- General fermionic or bosonic systems from condensed matter physics and computational chemistry can be mapped *efficiently* to qubits

Problem: Ground state energy estimation

- Given n-qubit Hamiltonian

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

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

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

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

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

Early fault-tolerance approach

1. Minimize number of qubits needed – only one ancilla

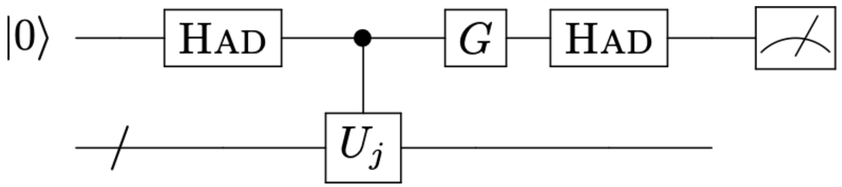

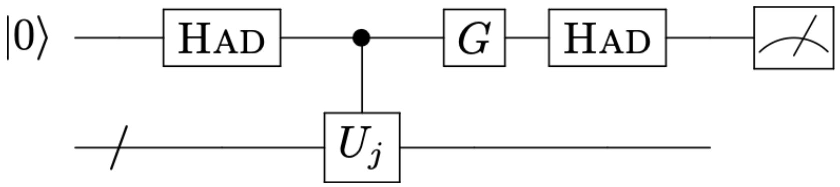

2. Trade-off gate versus sample complexity
3. Decrease error by solely taking more samples
4. Independent of the number L of Pauli terms in H

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 $|\psi\rangle$ 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 $\exp(i\theta P_l)$

- Note: Ansatz state η-overlap bottleneck vs classical methods

Basic idea

- Cumulative distribution function (CDF) relative to $|\psi\rangle$ is

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

- Evaluate $C(x)$ via quantum?

- Two algorithmic ingredients:

  (A) Hadamard test
  (B) Importance sampling

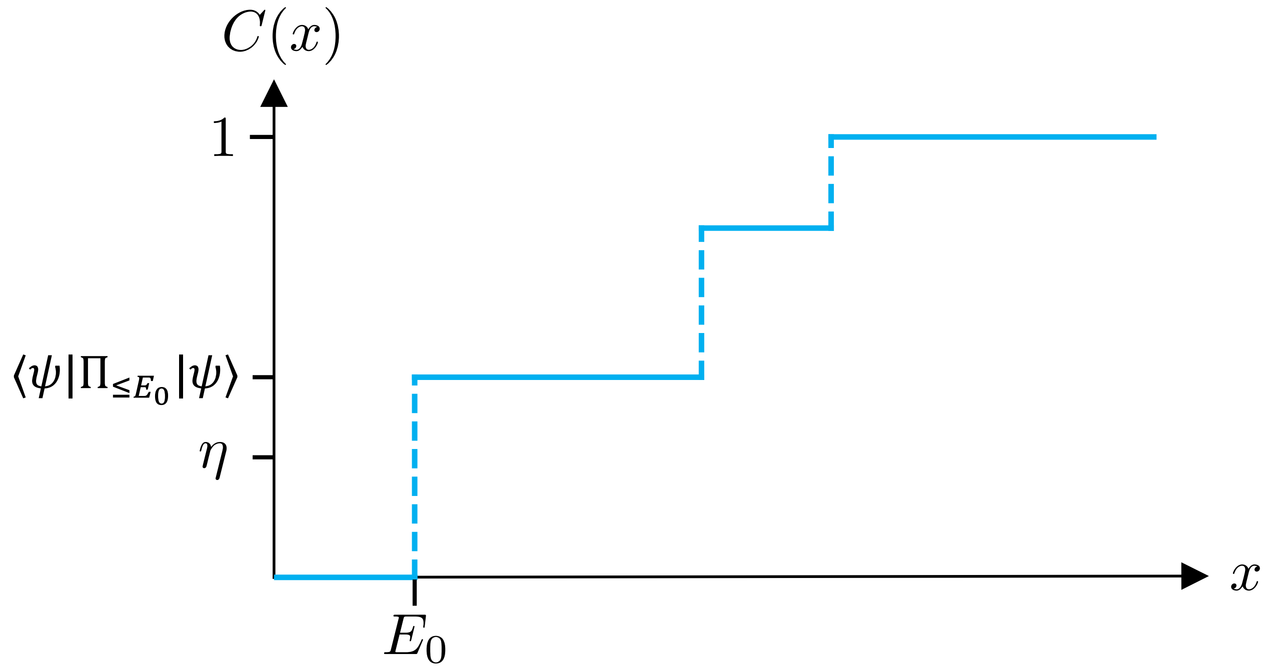

Workhorse A: Hadamard test

- Input: n-qubit state $|\psi\rangle$ together with n-qubit unitary U
- Quantum circuit:

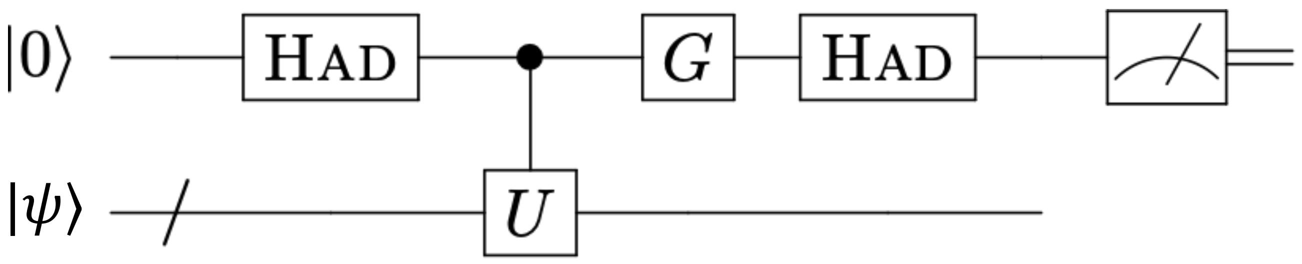

$$X = \begin{cases} +1 & \text{if } |0\rangle \\ -1 & \text{if } |1\rangle \end{cases}$$

- Output is unbiased estimate of $\langle\psi|U|\psi\rangle$ from

$$G = 1 \;\Rightarrow\; \mathbb{E}[X] = Re(\langle\psi|U|\psi\rangle)$$
$$G = S^* \;\Rightarrow\; \mathbb{E}[X] = Im(\langle\psi|U|\psi\rangle)$$

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 $(|\psi\rangle, U_j)$:

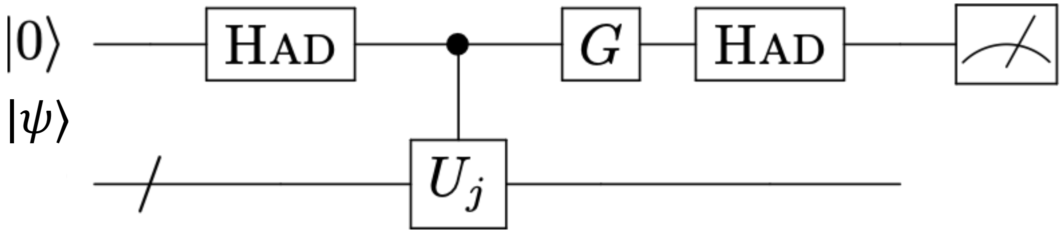

$\Rightarrow$ estimate $\langle\psi|U_j|\psi\rangle$

- Take average of samples, number required is $[A^2\sigma^{-2}]$ for variance $\sigma > 0$

CDF via Fourier series

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

$$C(x) \approx (p * F)(x) = \sum_{j\in S} \hat{F}_j e^{ijx} \cdot \langle\psi|e^{it_jH}|\psi\rangle$$

with runtimes $t_j = j \times$ normalization

- Hadamard test + importance sampling + Hamiltonian simulation:

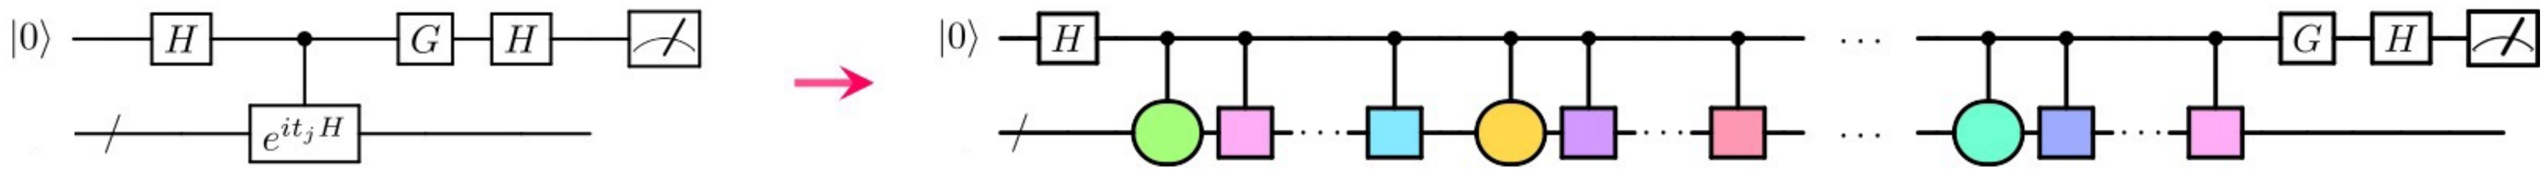

FeMoco benchmark – resource trade-offs

- Li et al. FeMoco Hamiltonian with 152 spin orbitals: 152+1=153 qubits
- Chemical accuracy $\Delta = 0.0016$ Hartree, one-norm $\lambda = 1511$
- State-of-the-art *qubitization*

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

- Ansatz state η-overlap bottleneck + classical methods scale polynomial!

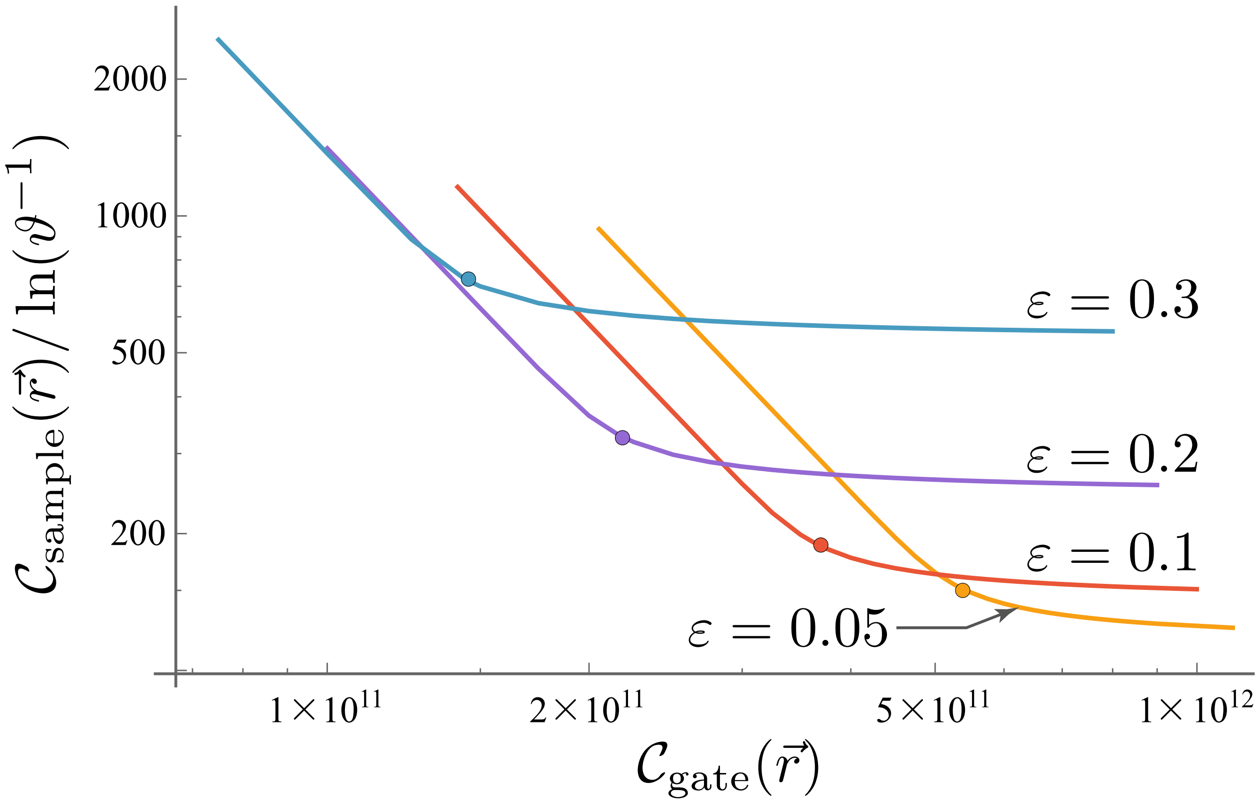

Example: Linear algebra on classical data

Qubit-efficient randomized quantum algorithms for linear algebra

arXiv:2302.01873 (2023) with McArdle and Wang
Quantum Computing Theory in Practice (QCTIP) 2023
Theory of Quantum Computation, Communication and Cryptography (TQC) 2023

Data comes via
classical description

"Early fault-tolerant algorithms for classical data"

Hardware efficient
&
Provable guarantees

Quantum algorithms for classical data

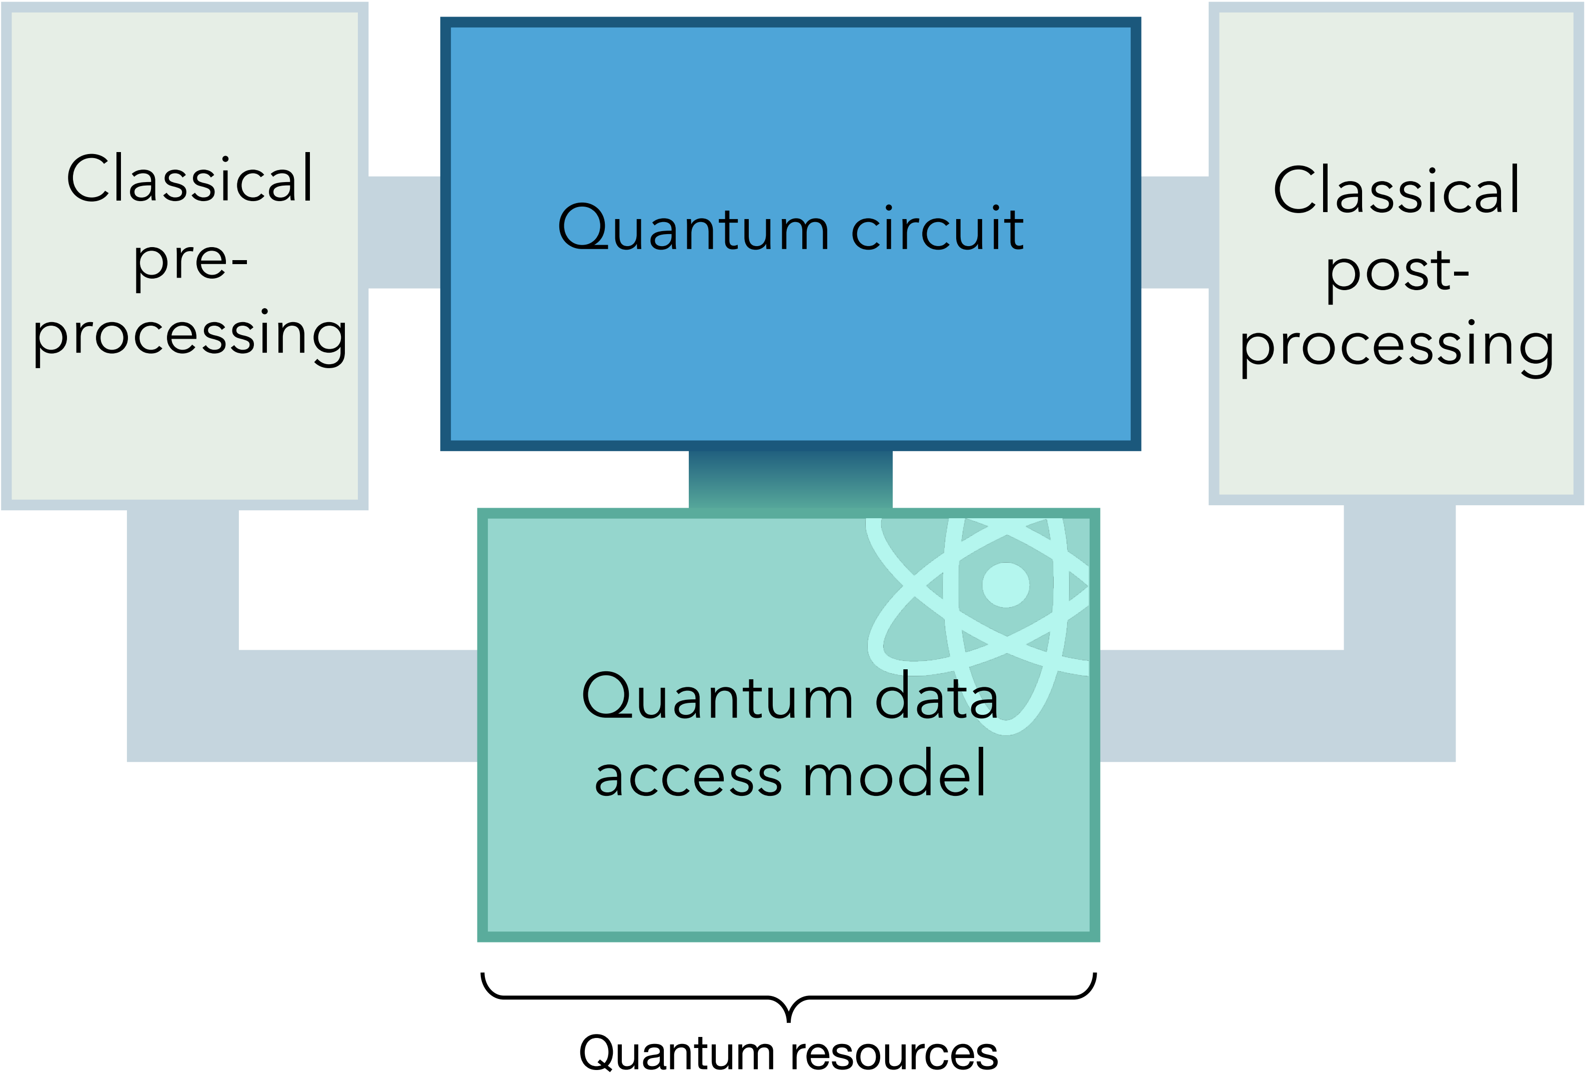

Idea I: Parallelize quantum sub-routines

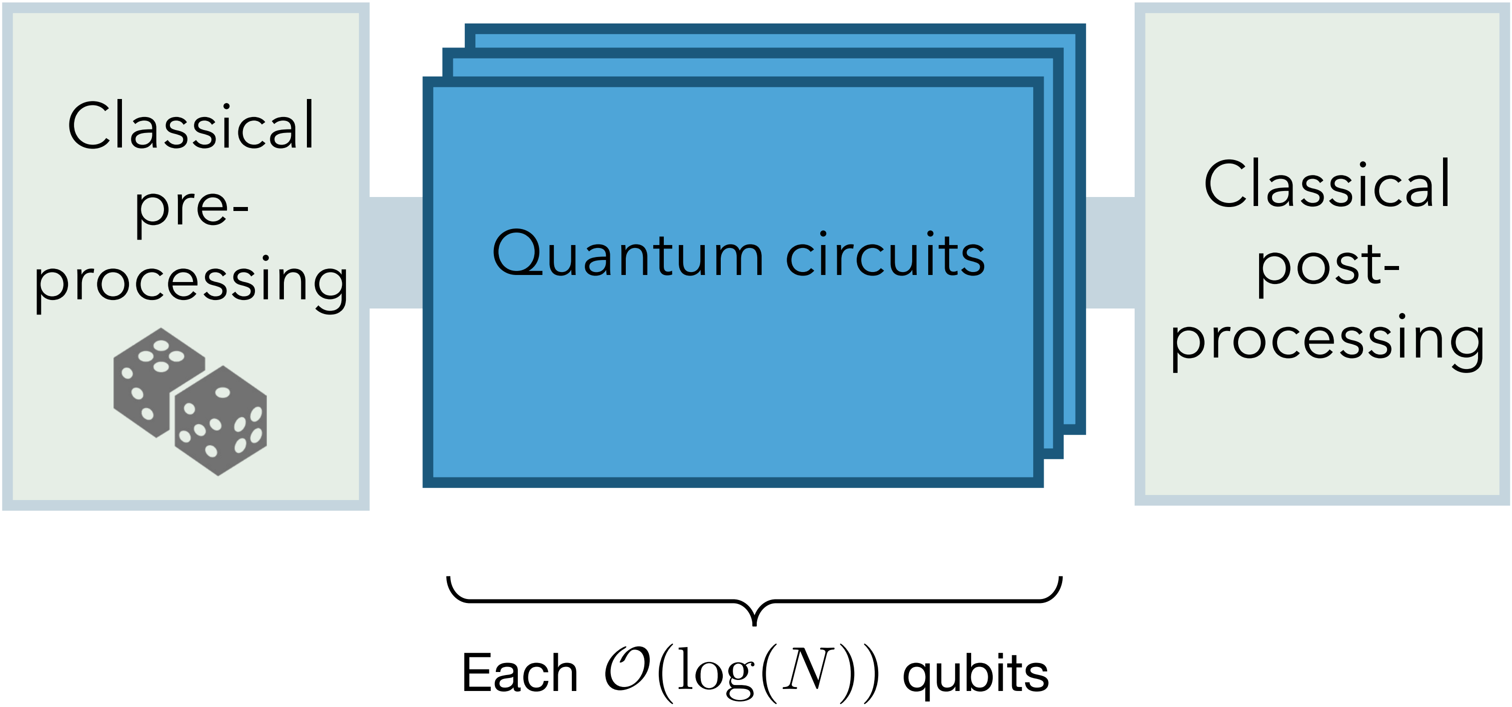

Idea II: Classical instead of quantum access

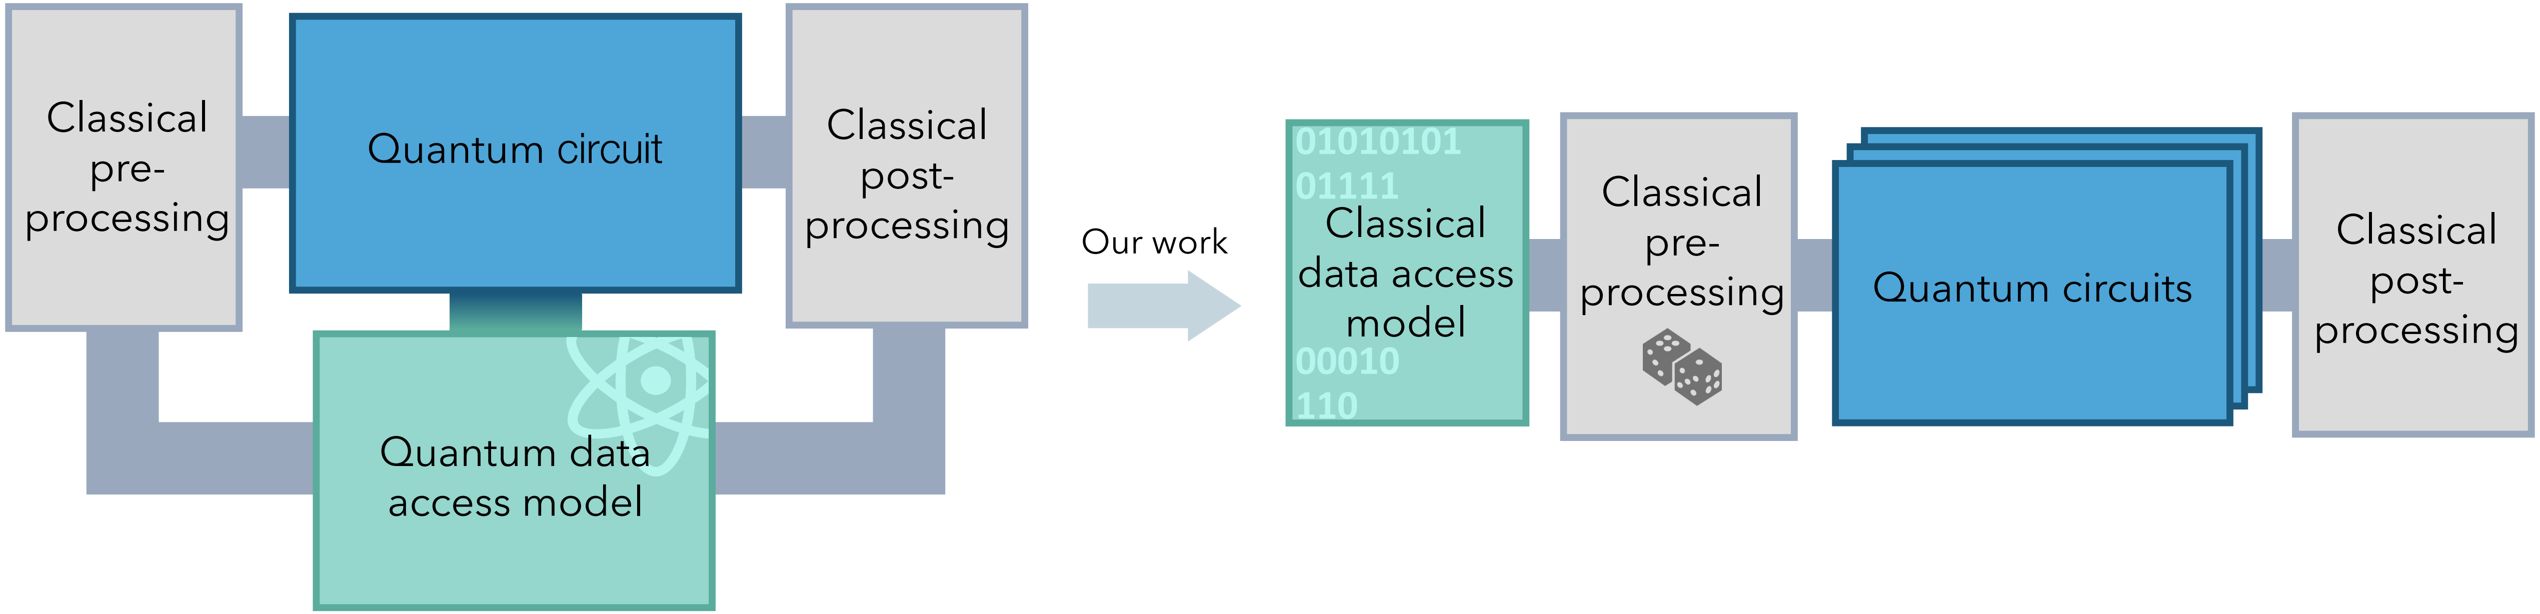

Conclusion

Quantum algorithms for early fault-tolerance

- Use as few qubits and quantum routines as possible, use classical methods whenever sufficient
- Early fault-tolerant methods can even be competitive with state-of-the-art (non-qubit aware) schemes in terms of asymptotic complexities

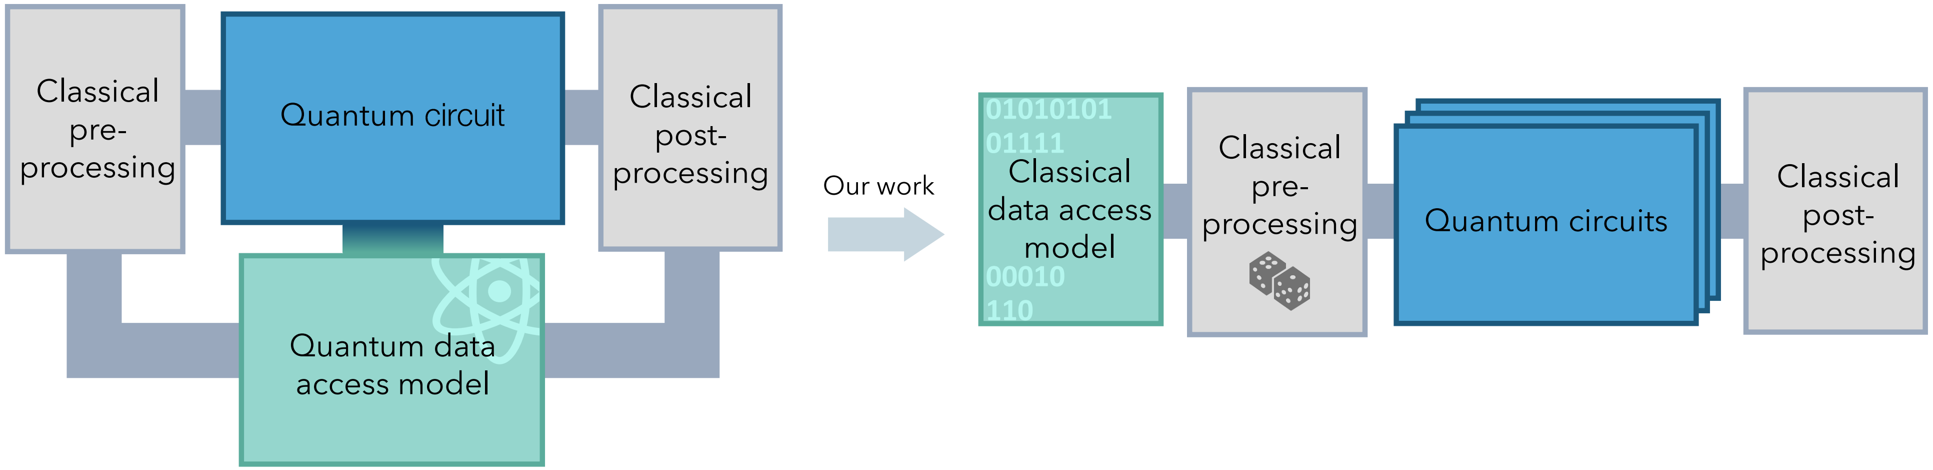

Outlook

- Quantum resource counts for applications featuring end-to-end complexity analyses, quantum speed-up?

  → upcoming popular article *DPG Physik Journal* (November issue)
- 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?
- 50-100 error corrected qubits could allow for truly insightful experiments

Thank you!

RWTH AACHEN UNIVERSITY

Some references

Paper references to some of our work

- A randomized quantum algorithm for statistical phase estimation
  QIP22, 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
  QIP23, arXiv:2209.12887 (2022) with McArdle, Gilyen
- Sparse random Hamiltonians are quantumly easy
  QIP23, arXiv:2302.03394 (2023) with Chen, Dalzell, Brandão, Tropp

Extra content ground state energy

Hydrogen chain benchmark – scaling

- For length N chain, one-norm estimate $\lambda \approx O(N^{1.34})$
- Our work $C_{gate} = \tilde{O}(N^{2.68}\Delta^{-2})$
- Comparison to state-of-the-art *qubitization*:
  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: $C_{gate} = \tilde{O}(N^{0.68})$

Fourier series lemma (Heaviside function)

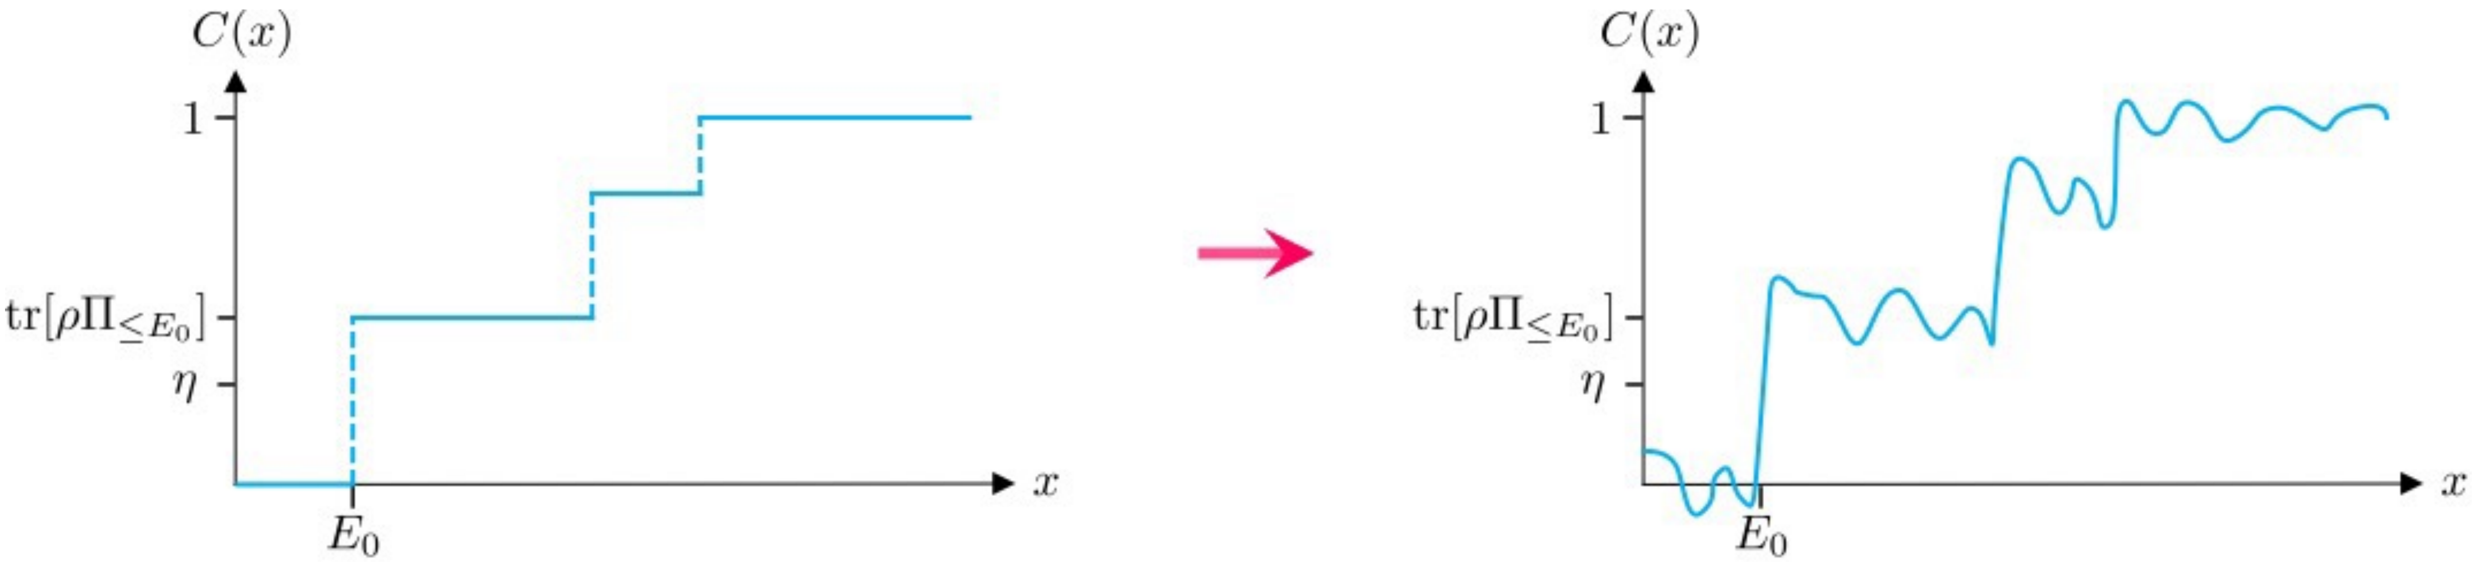

- Improved Fourier series approximation of Heaviside function
- Technical contribution:

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:

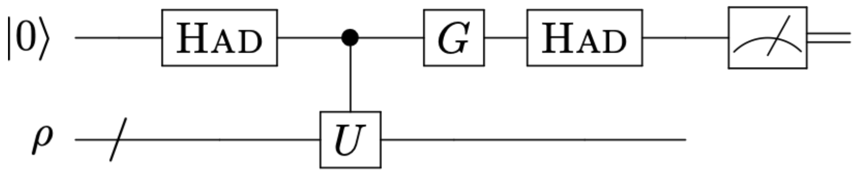

  I. $\mu(r) \coloneqq \sum_k b_k \leq \exp(t^2 r^{-1})$
  II. $COST(C - U_k) = r$ controlled single qubit Pauli rotations $\forall k$

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

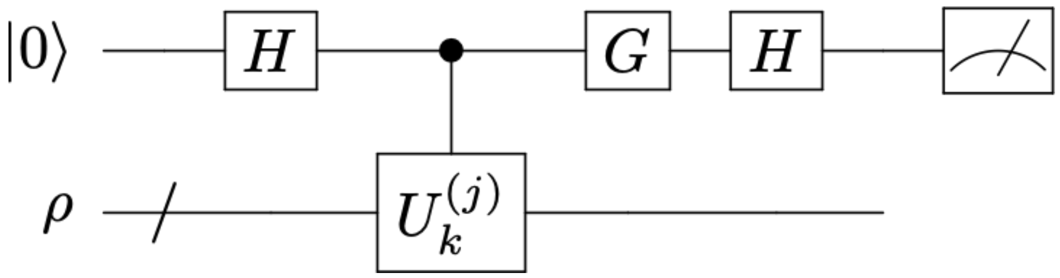

- $e^{it_jH} = \sum_k b_k^{(j)} U_k^{(j)}$ decomposition for runtime vector $\vec{r} = (r_j)_{j} \in \mathbb{N}^{|S|}$ as:

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

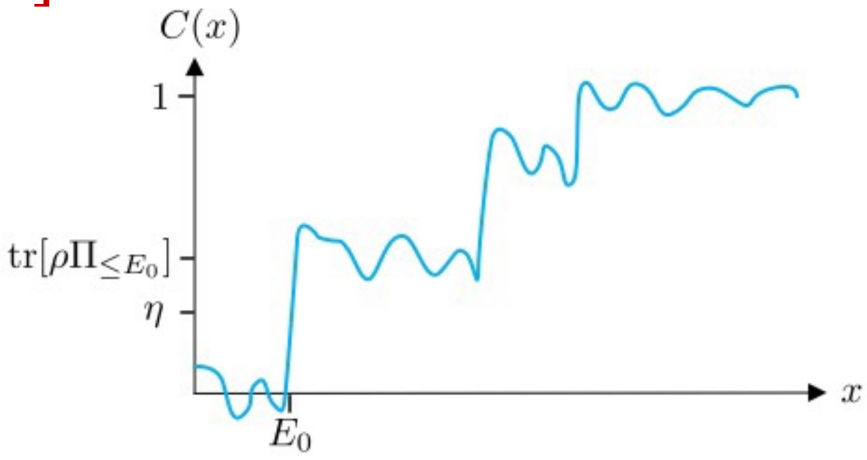

- As $\mu_j \leq e^{t_j^2 r_j^{-1}}$ choosing $r_j = 2t_j^2 \; \forall j$ gives $\mu_j \leq \sqrt{e}$:

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

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

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:

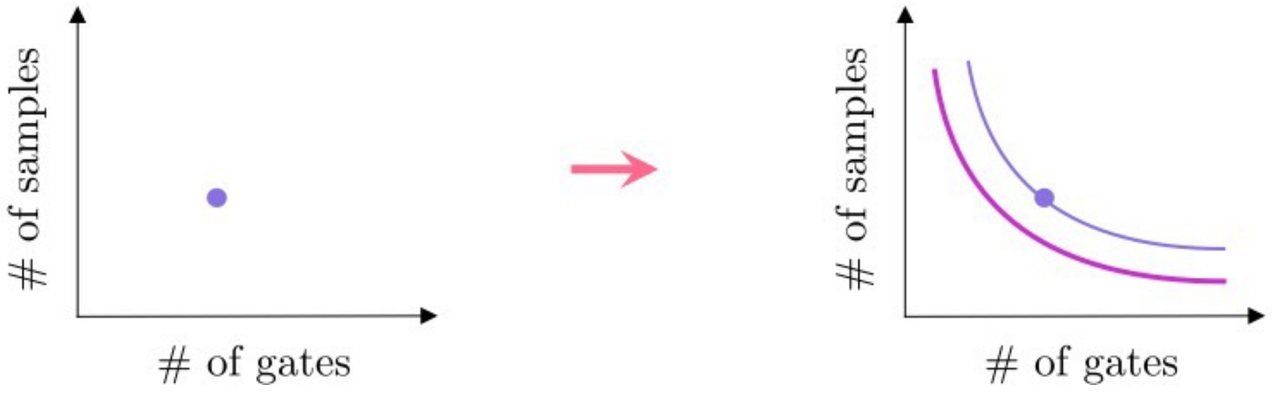

Extra: Proof Fourier series lemma

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

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

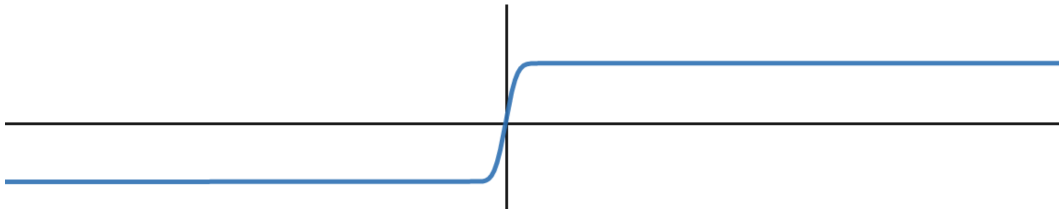

- Fourier series: $\Theta(x) \approx \mathrm{erf}\big(\beta \sin(x)\big) \approx \sum_k c_k T_k \left(\cos\left(\frac{\pi}{2} - x\right)\right)$

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

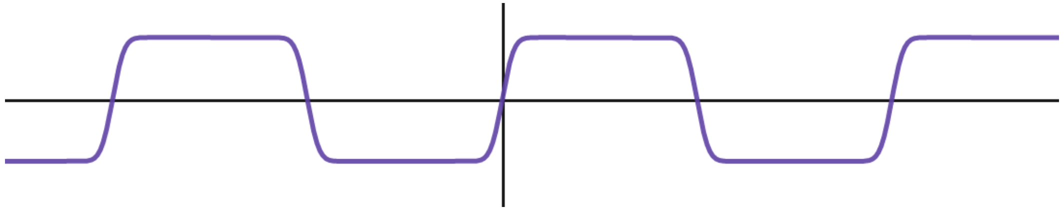

Extra: Proof random compiler lemma

- 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

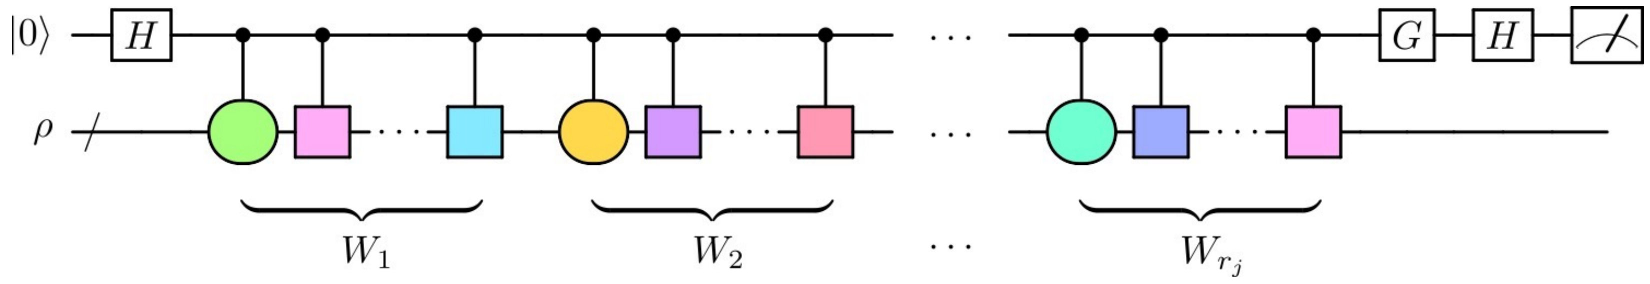

Extra: qDRIFT comparison

[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