Quantum algorithms for eigenvalue problems

Lin Lin

Department of Mathematics, UC Berkeley
Lawrence Berkeley National Laboratory
Challenge Institute for Quantum Computation

Math Colloquium
University of South Carolina, February, 2022

arXiv:2002.12508 (Quantum 2020), 2102.11340 (PRX Quantum 2022)

Joint work with

(Berkeley->Caltech, IQIM Fellow)

Outline

Solve nature with nature

Figure. A superposition of Feynmans

... if you want to make a simulation of nature (quantum many-body problem), you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy.

– Richard P. Feynman (1981) 1st Conference on Physics and Computation, MIT

Quantum computation meets public attention

Google, Nature 2019
Random circuit sampling
Theory: [Boixo et al, 2018]

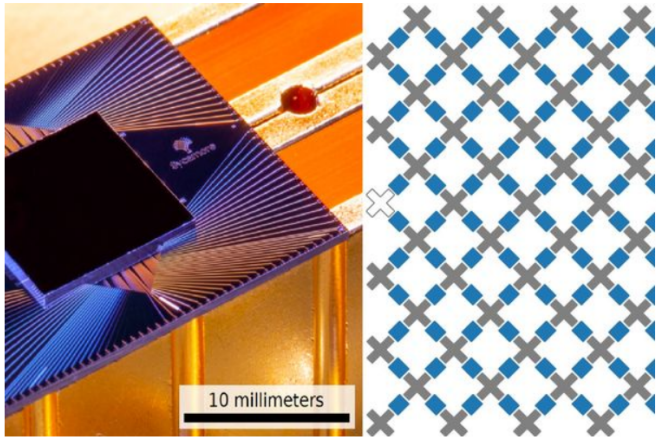

USTC, Science 2020
Boson sampling.
Theory: [Aaronson–Arkhipov, 2011]

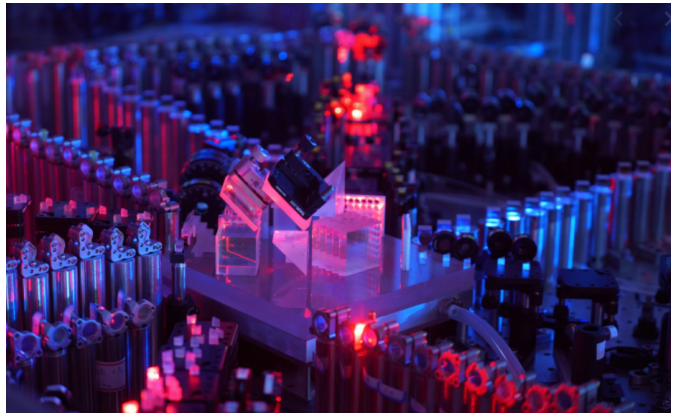

- After about four decades, quantum supremacy has been reached: the point *where quantum computers can do things that classical computers cannot, regardless of whether those tasks are useful.*

- *Is controlling large-scale quantum systems merely really, really hard, or is it ridiculously hard?* – John Preskill (2012)

- Quantum computer does anything useful? called quantum advantage.

A fast growing industry

What is a quantum computer (mathematically)

- $|\psi\rangle \in \mathbb{C}^N \cong (\mathbb{C}^2)^{\otimes n}, \quad N = 2^n$. n : number of qubits.
- Normalization condition $\langle\psi|\psi\rangle = 1$.
- $U \in \mathbb{C}^{N\times N}$ is unitary. $U|\psi\rangle$ is efficient to apply ($\mathrm{polylog}(N)$).
- Quantum computer: $U_K \cdots U_1 |\psi\rangle$, and then classical output by measuring one or a few qubits M times.
- Quantum cost: $MK\,\mathrm{polylog}(N)$.
- Exponential quantum advantage (EQA): if $MK = \mathcal{O}(\mathrm{polylog}(N))$, and classical algorithm scales as $\mathcal{O}(\mathrm{poly}(N))$.

Numerical linear algebra

- Linear systems of equation $Ax = b$
- Least squares problem $\min_x \|Ax - b\|_2$
- Eigenvalue decomposition $Av_i = \lambda_i v_i$
- Singular value decomposition $Av_i = u_i \sigma_i$
- Preconditioner $M^{-1}Ax = M^{-1}b$
- Matrix exponentiation e^{iHt} (Hamiltonian simulation)
- Other matrix functions: $\sqrt{A}, \log A, \ldots$
- Machine learning, e.g. kernel ridge regression $\alpha = (K + \tilde{I})^{-1} y$
- ...

Quantum numerical linear algebra

- Solving numerical linear algebra problems on a quantum computer.
- Many interesting, exciting progresses in the past few years.
- Lecture notes on "Quantum Algorithms for Scientific Computation": arXiv:2201.08309
- Reasonable way towards "quantum advantage".
- Ground state energy: an eigenvalue problem

$$H|\psi_0\rangle = \lambda_0 |\psi_0\rangle$$

$H \in \mathbb{C}^{N\times N}$ Hermitian matrix (Hamiltonian).
Find the algebraically smallest λ_0.

- Under which conditions the cost can be $\mathcal{O}(\mathrm{polylog}(N))$?

Ground state preparation and energy estimation

- Suppose we are given a Hamiltonian that can be succinctly described (local, sparse, etc.):

$$H = \sum_k \lambda_k |\psi_k\rangle \langle\psi_k| , \quad \lambda_k \leq \lambda_{k+1}.$$

- Goal: prepare its ground state $|\psi_0\rangle$ or to estimate its ground state energy λ_0 to some precision.

- QMA-hard of the local Hamiltonian problem without additional assumptions.

- Some physically relevant assumptions:
  1. Good initial guess : $|\langle\phi|\psi_0\rangle| \geq \gamma$.
  2. Spectral gap: $\Delta = \lambda_1 - \lambda_0$. (only needed for preparing the ground state)

Quantum advantage for quantum chemistry

- Quantum many body Hamiltonian (second quantization, $\dim(H) = 2^n$)

$$H = \sum_{ij=1}^{n} h_{ij} c_i^\dagger c_j + \frac{1}{2} \sum_{ijkl=1}^{n} V_{ijkl} c_i^\dagger c_j^\dagger c_k c_l$$

- Ground state energy: useful in predicting material structures, simulating chemical reactions, etc.

- "Poster-child" problem: FeMo cofactor (FeMoco): primary cofactor of nitrogenase for nitrogen fixation.

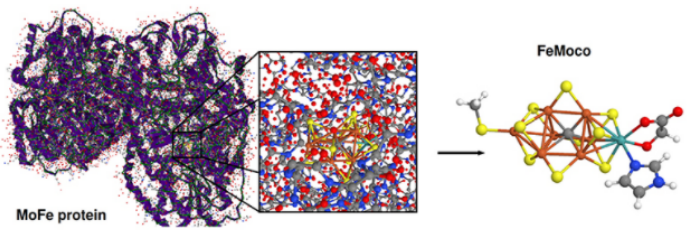

- Strongly correlated quantum chemistry. EQA is debatable.

Input models

- Dimension of Hilbert space is $N = 2^n$.
- Initial vector: U_I to prepare an initial state $|\phi\rangle = U_I |0^n\rangle$.
- Hamiltonian H:
  1. Hamiltonian evolution access: $U_H = e^{-\mathrm{i}\tau H}$ for some τ.
  2. Block encoding access:

$$U_H = \begin{pmatrix} H/\alpha & * \\ * & * \end{pmatrix} \tag{1}$$

- Query complexity: the number of accesses to U_H, U_I.

Outline

Textbook algorithm: quantum phase estimation (QPE)

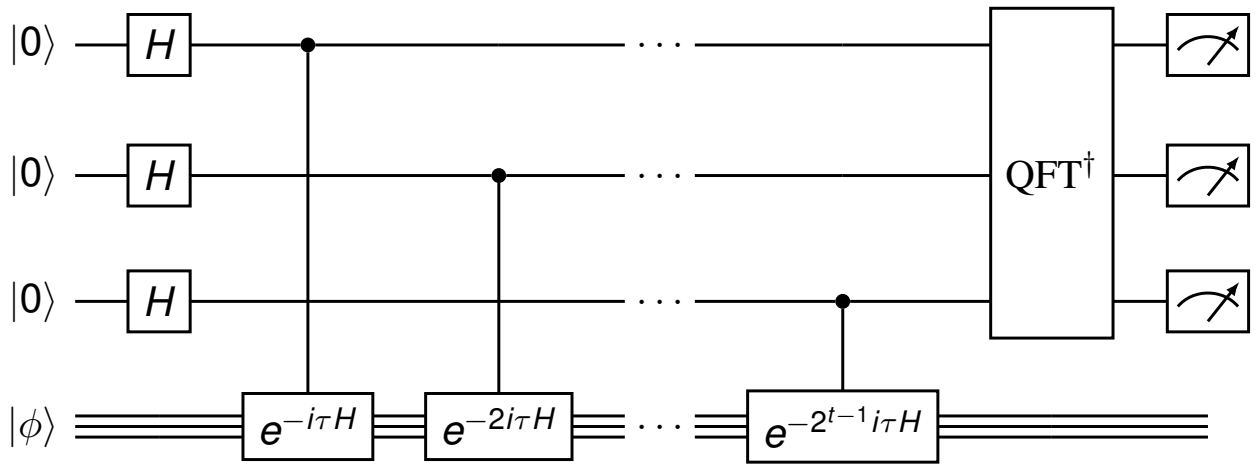

- Requires multiple control qubits (to store eigenvalues in a quantum register)
- Run multiple times and take the lowest value as estimate to λ_0.
- Aliasing problem (textbook analysis: [Nielsen-Chuang 2000])

Query complexities for ground state energy estimation

	H	$\|\phi\rangle$
QPE	$\mathcal{O}(\epsilon^{-1}\gamma^{-2})$	$\mathcal{O}(\gamma^{-2})$
GTC2017	$\mathcal{O}(\epsilon^{-3/2}\gamma^{-1})$	$\mathcal{O}(\epsilon^{-1/2}\gamma^{-1})$
LT2020	$\mathcal{O}(\epsilon^{-1}\gamma^{-1})$	$\mathcal{O}(\gamma^{-1})$

Table: Query complexities for ground state energy estimation.

1. Previous best results: (Ge-Tura-Cirac, 2019)
2. QPE estimates the eigenvalue coherently on a quantum computer (which comes with some additional cost)
3. Lower bound for the overlap: $|\langle\phi|\psi_0\rangle| \geq \gamma$; Target precision: ϵ

(L.-Tong, 2002.12508, Quantum 2020)

Query complexity for ground state preparation and lower bound

	H	$\|\phi\rangle$
QPE	$\mathcal{O}(\Delta^{-1}\gamma^{-2})$	$\mathcal{O}(\gamma^{-2})$
GTC2017	$\mathcal{O}(\Delta^{-3/2}\gamma^{-1})$	$\mathcal{O}(\Delta^{-1/2}\gamma^{-1})$
LT2020	$\mathcal{O}(\Delta^{-1}\gamma^{-1})$	$\mathcal{O}(\gamma^{-1})$

Table: Query complexities for ground state preparation (without knowledge of ground state energy).

Cannot be "substantially" improved!

- $\Delta = \lambda_1 - \lambda_0$: spectral gap
- Idea of the proof: consider **the unstructured search problem** as a ground state preparation problem. Overlap and gap trade-off through an adiabatic path.[1]

[1] Childs et al., 2002, "Quantum search by measurement".

Quantum signal processing

- Quantum signal processing (QSP) / **quantum singular value transformation (QSVT)**.

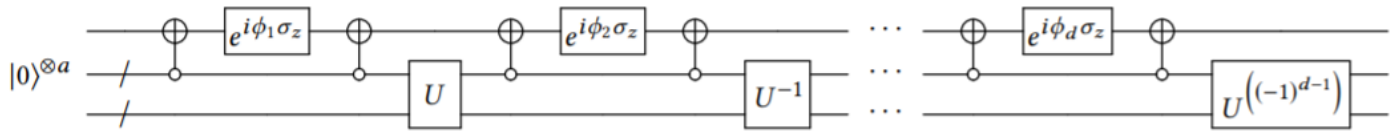

Figure: From [Gilyén-Su-Low-Wiebe, 2019].

- QSP/QSVT allows us to implement, with a block encoding of H:

$$f(H)\,|\phi\rangle = \sum_{k=0}^{N-1} f(\lambda_k)\,|\psi_k\rangle\,\langle\psi_k|\phi\rangle\,.$$

Low and Chuang, “Optimal Hamiltonian simulation by quantum signal processing”, PRL 2017

Gilyén, Su, Low, Wiebe, “Quantum singular value transformation and beyond”, STOC 2019

Martyn et al, “A Grand Unification of Quantum Algorithms”, PRX Quantum 2021

L., “Lecture notes on Quantum Algorithms for Scientific Computation”, Ch7,Ch8

Eigenstate filtering

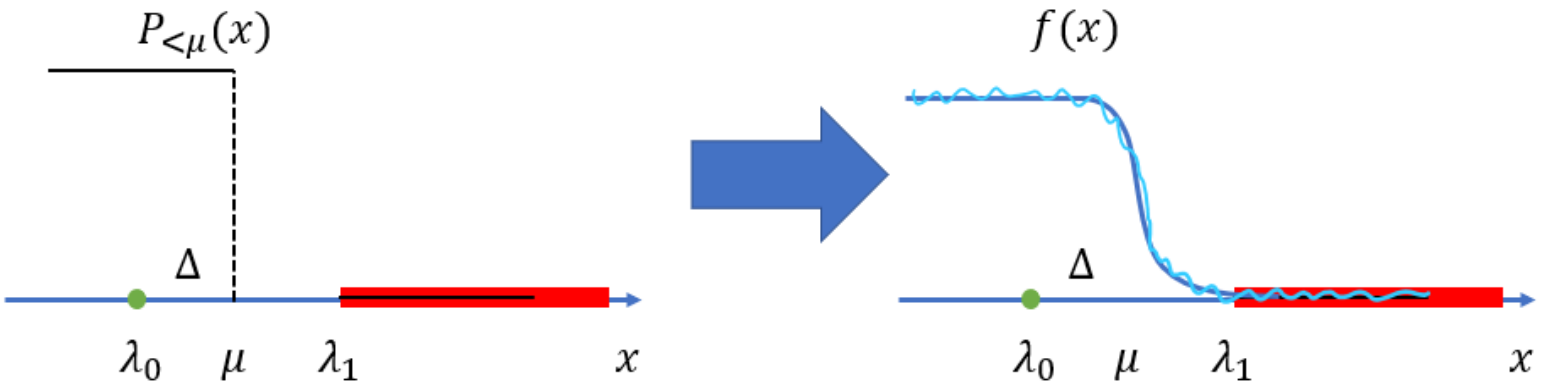

- Polynomial approximation to step functions.
- Ground state energy estimate: find a good μ and implement the filtering with a proper polynomial.
- Convert to quantum circuit via QSVT.

Ground state energy estimation

- Idea: use binary search. Need to solve the following problem: if we know $a \leq \lambda_0 \leq b$, decide $\lambda_0 > (a+b)/2$ or $\lambda_0 < (a+b)/2$.
- **This does not work** because we are essentially asking the quantum circuit to compute a discontinuous function while the output probability distribution is a continuous function of λ_0.
- Need to account for the fuzziness and statistical uncertainty.

The decision problem

Assuming we know $a \leq \lambda_0 \leq b$.

(i) When $a \leq \lambda_0 \leq \frac{2}{3}a + \frac{1}{3}b$, output 0;

(ii) When $\frac{2}{3}a + \frac{1}{3}b \leq \lambda_0 \leq \frac{1}{3}a + \frac{2}{3}b$, output 0 or 1;

(iii) When $\frac{1}{3}a + \frac{2}{3}b \leq \lambda_0 \leq b$, output 1.

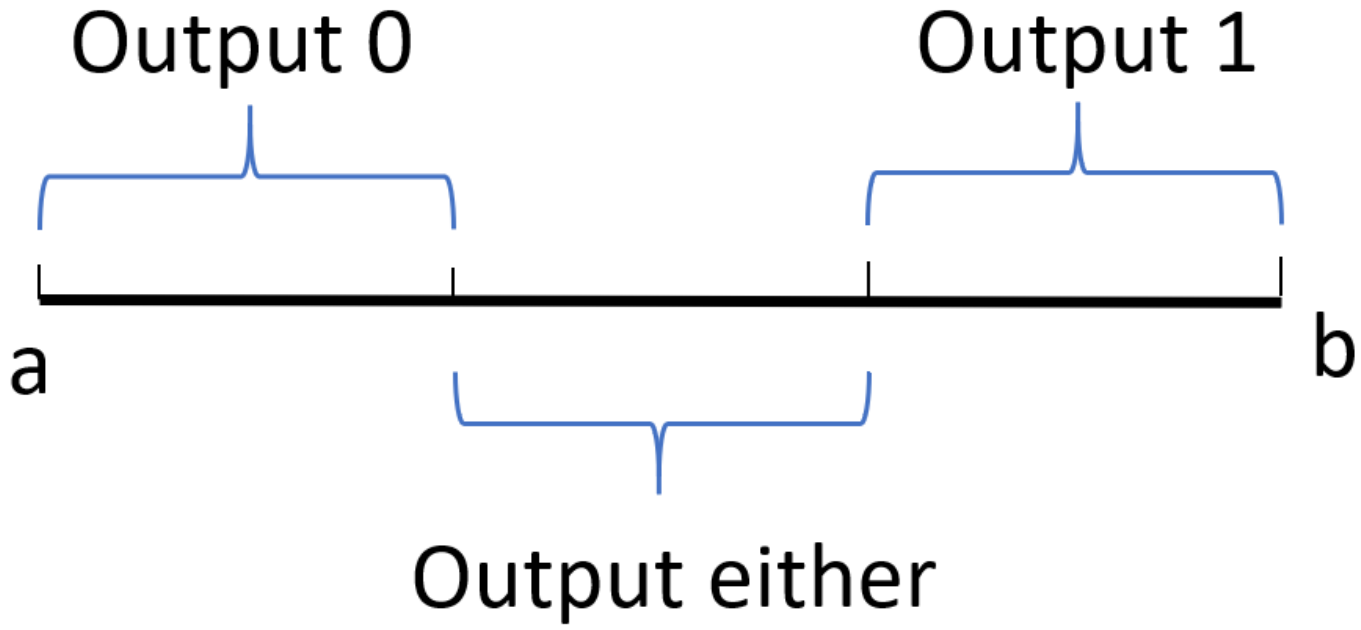

The decision problem

Assuming we know $a \leq \lambda_0 \leq b$.

(i) When $a \leq \lambda_0 \leq \frac{2}{3}a + \frac{1}{3}b$, output 0;

(ii) When $\frac{2}{3}a + \frac{1}{3}b \leq \lambda_0 \leq \frac{1}{3}a + \frac{2}{3}b$, output 0 or 1;

(iii) When $\frac{1}{3}a + \frac{2}{3}b \leq \lambda_0 \leq b$, output 1.

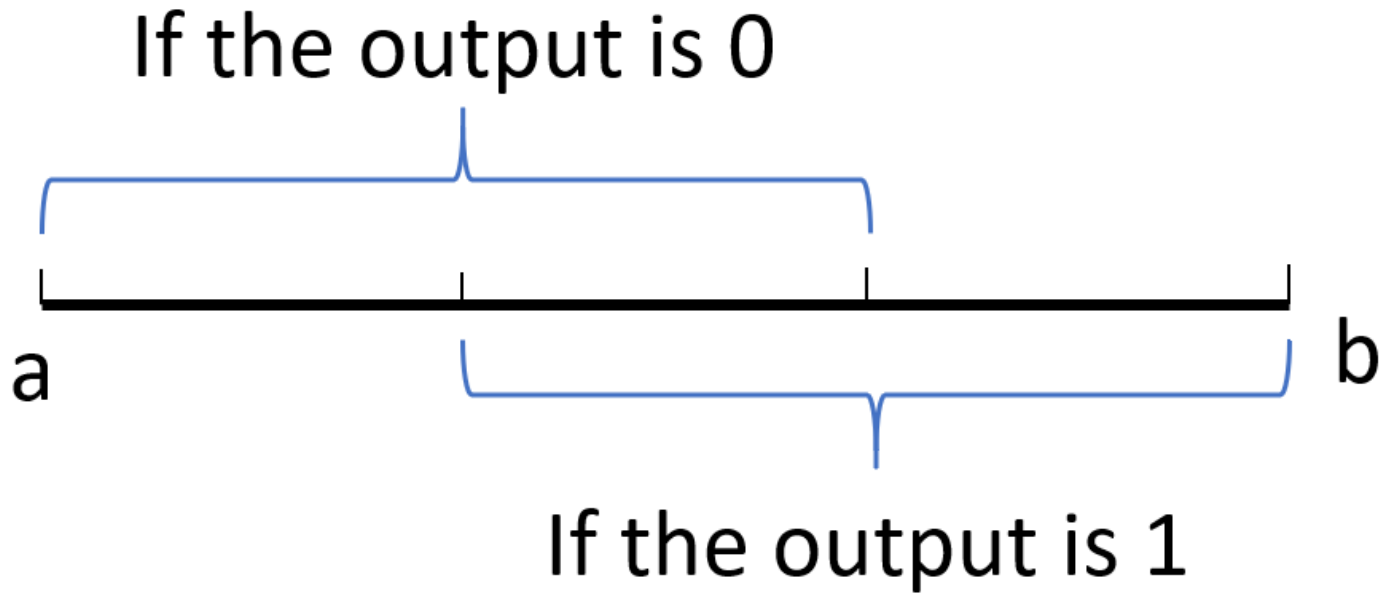

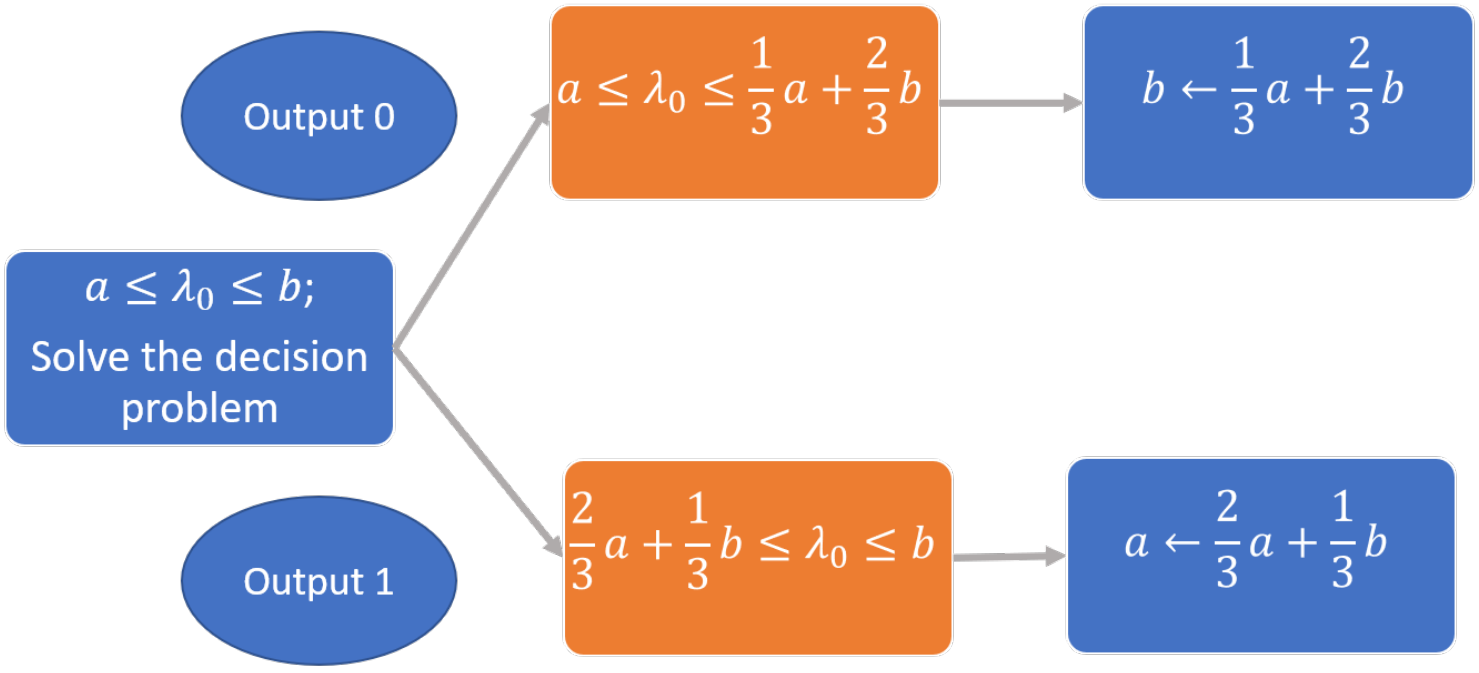

- Always guaranteed: $a \leq \lambda_0 \leq b$, $(b-a)_{\text{new}} = (2/3)(b-a)_{\text{old}}$.
- Solve the decision problem through **eigenstate filtering**: we can estimate $\| f(H) |\phi\rangle \|$ (for a polynomial f) using **amplitude estimation**.

Solving the decision problem (i)

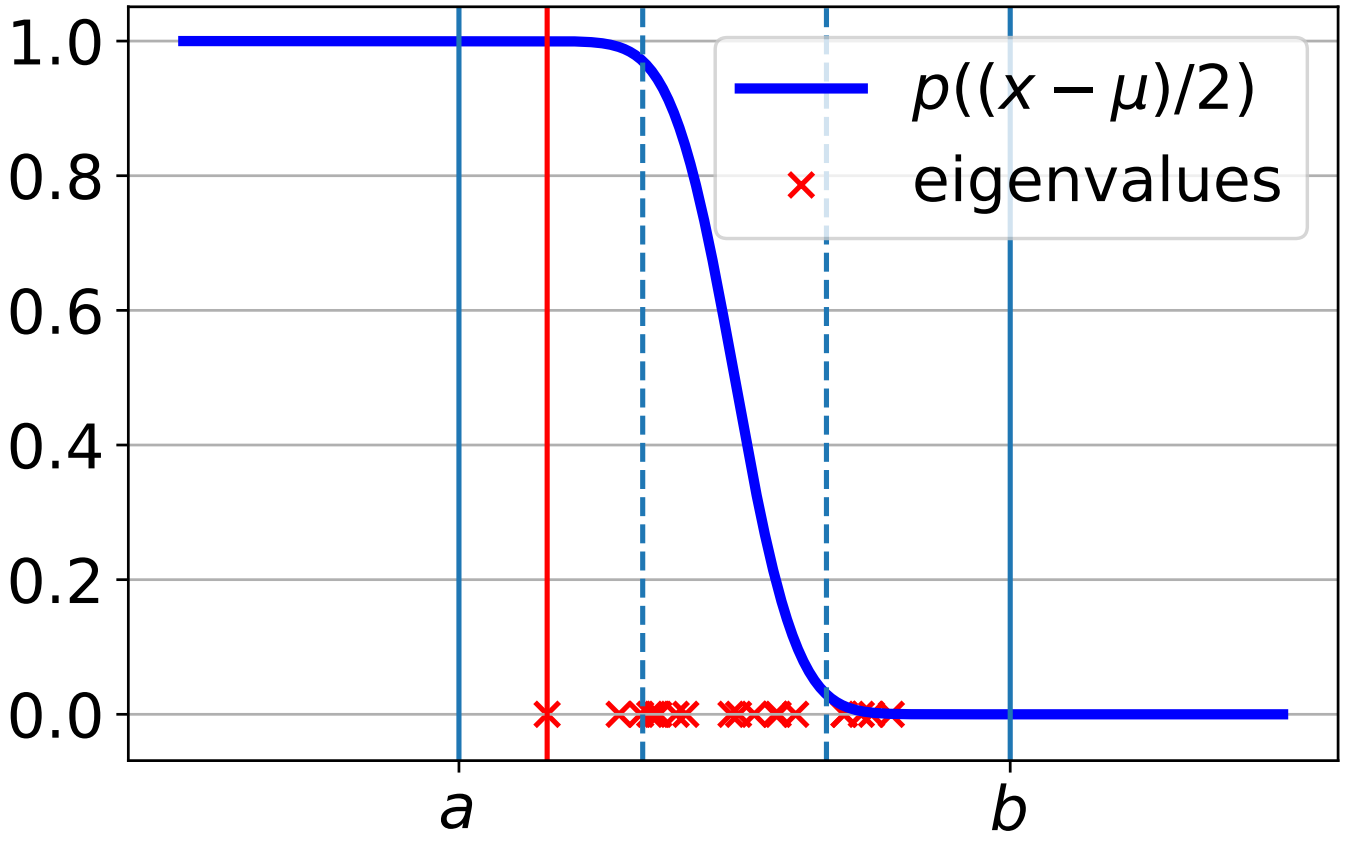

$$\lambda_0 \leq \tfrac{2}{3}a + \tfrac{1}{3}b \implies \|f(H)\,|\phi\rangle\,\| \geq \gamma(1-\epsilon')$$

Solving the decision problem (ii)

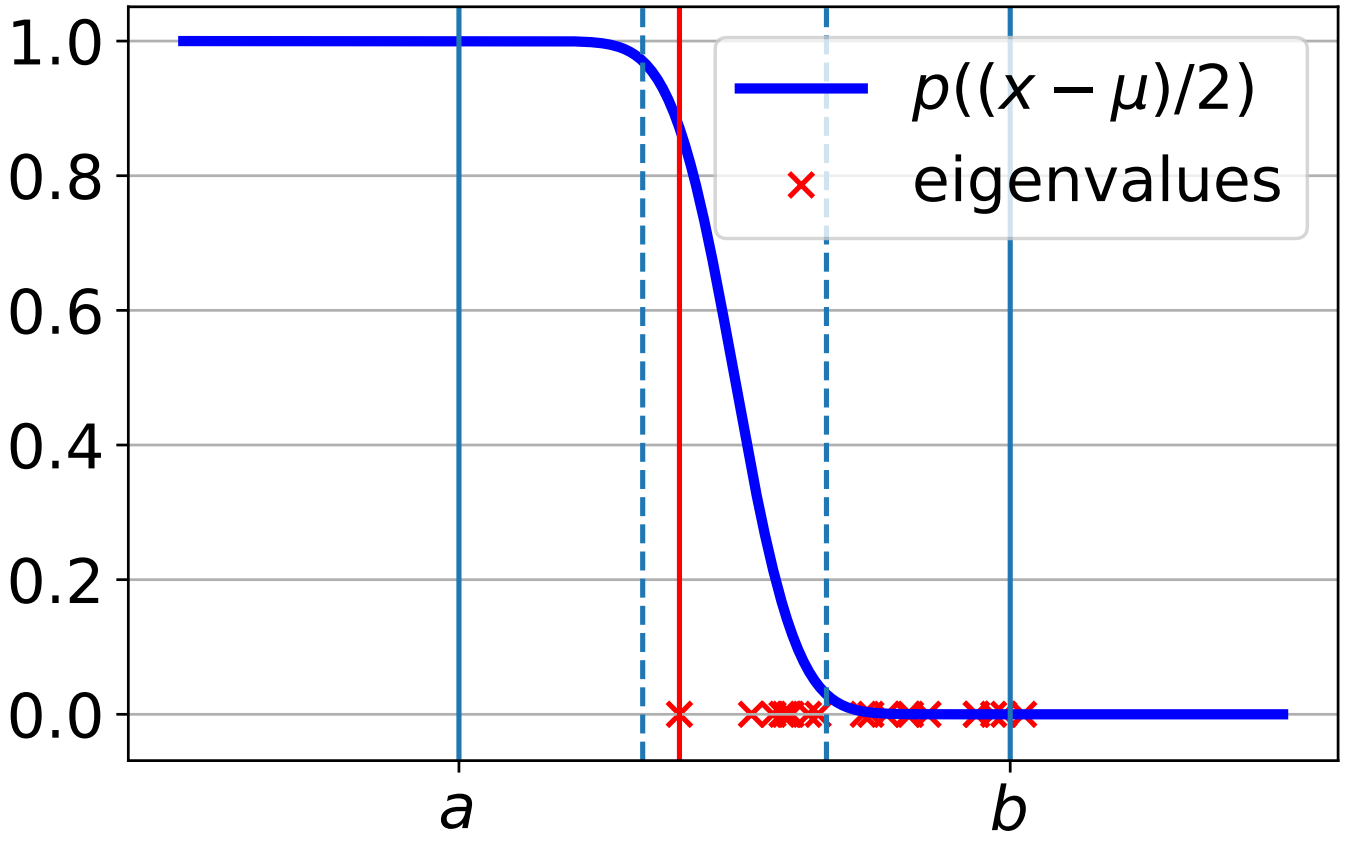

$$\tfrac{2}{3}a + \tfrac{1}{3}b \leq \lambda_0 \leq \tfrac{1}{3}a + \tfrac{2}{3}b$$

Solving the decision problem (iii)

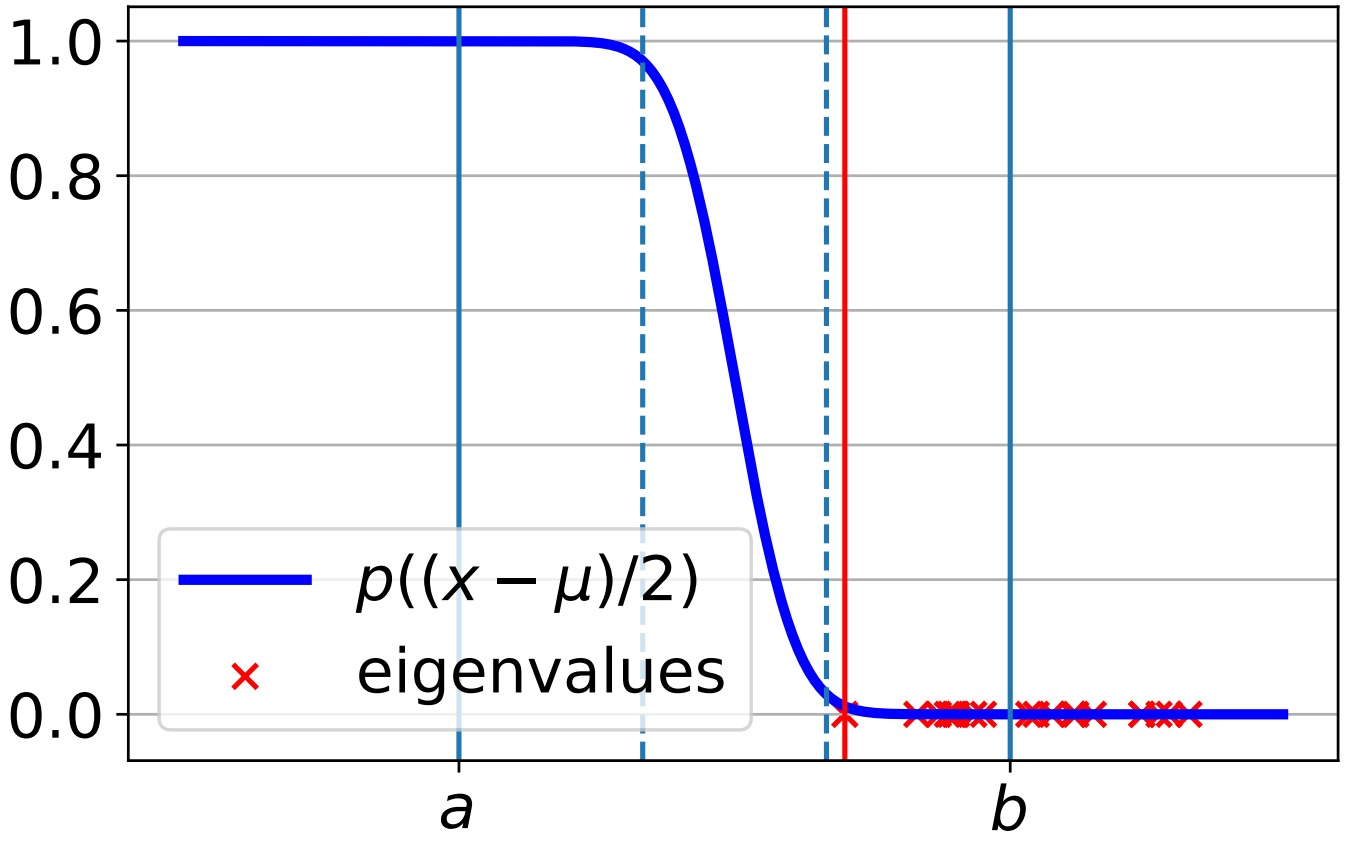

$$\lambda_0 \geq \tfrac{1}{3}a + \tfrac{2}{3}b \implies \|f(H)\,|\phi\rangle\,\| \leq \epsilon'$$

Solving the decision problem

- Only need to distinguish between $\|f(H)\,|\phi\rangle\,\| \geq \gamma(1-\epsilon')$ and $\|f(H)\,|\phi\rangle\,\| \leq \epsilon'$.
- Can use amplitude estimation to do so with overhead $\mathcal{O}(\gamma^{-1})$.
- Error probability can be exponentially suppressed using majority voting (Chernoff bound).

The search process

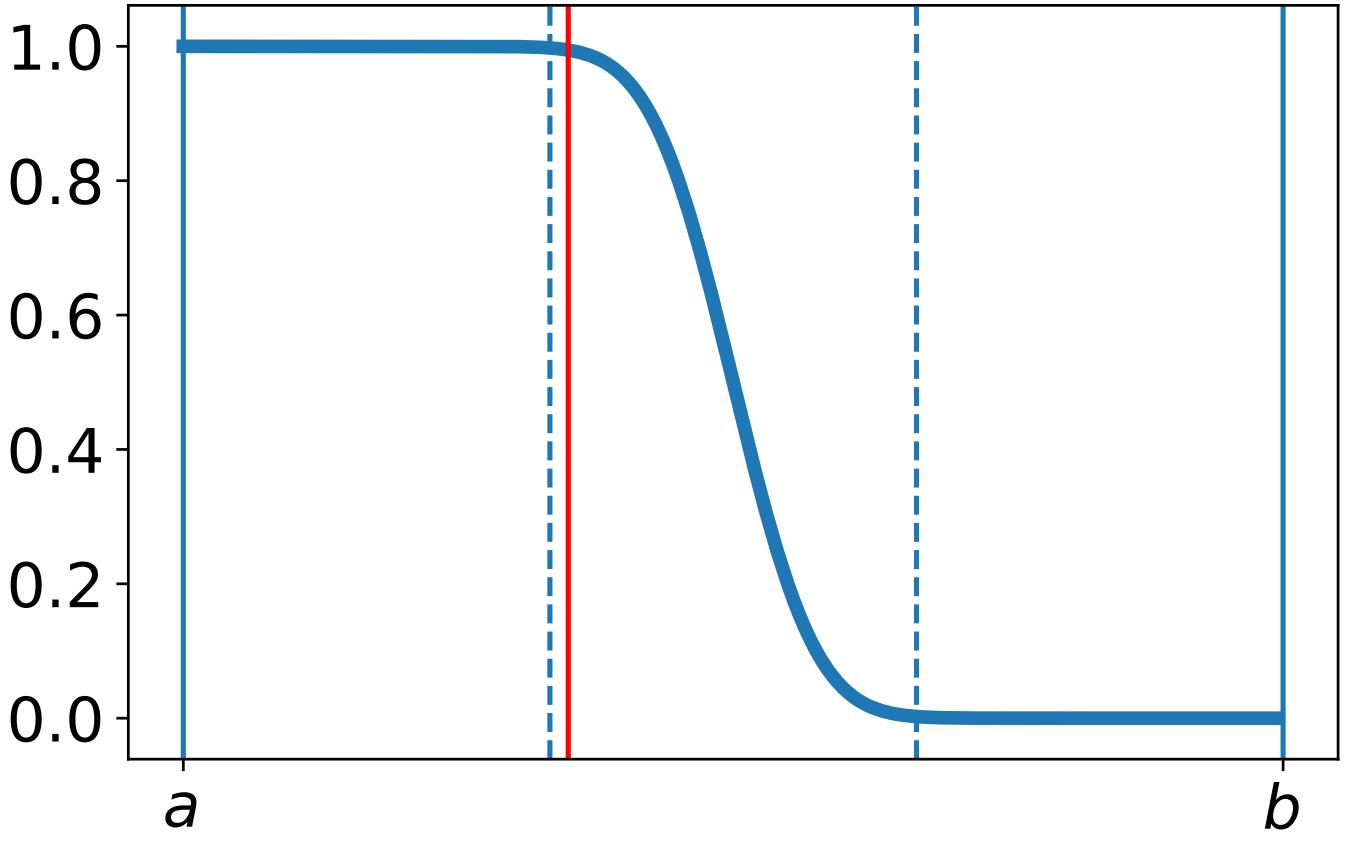

The search process

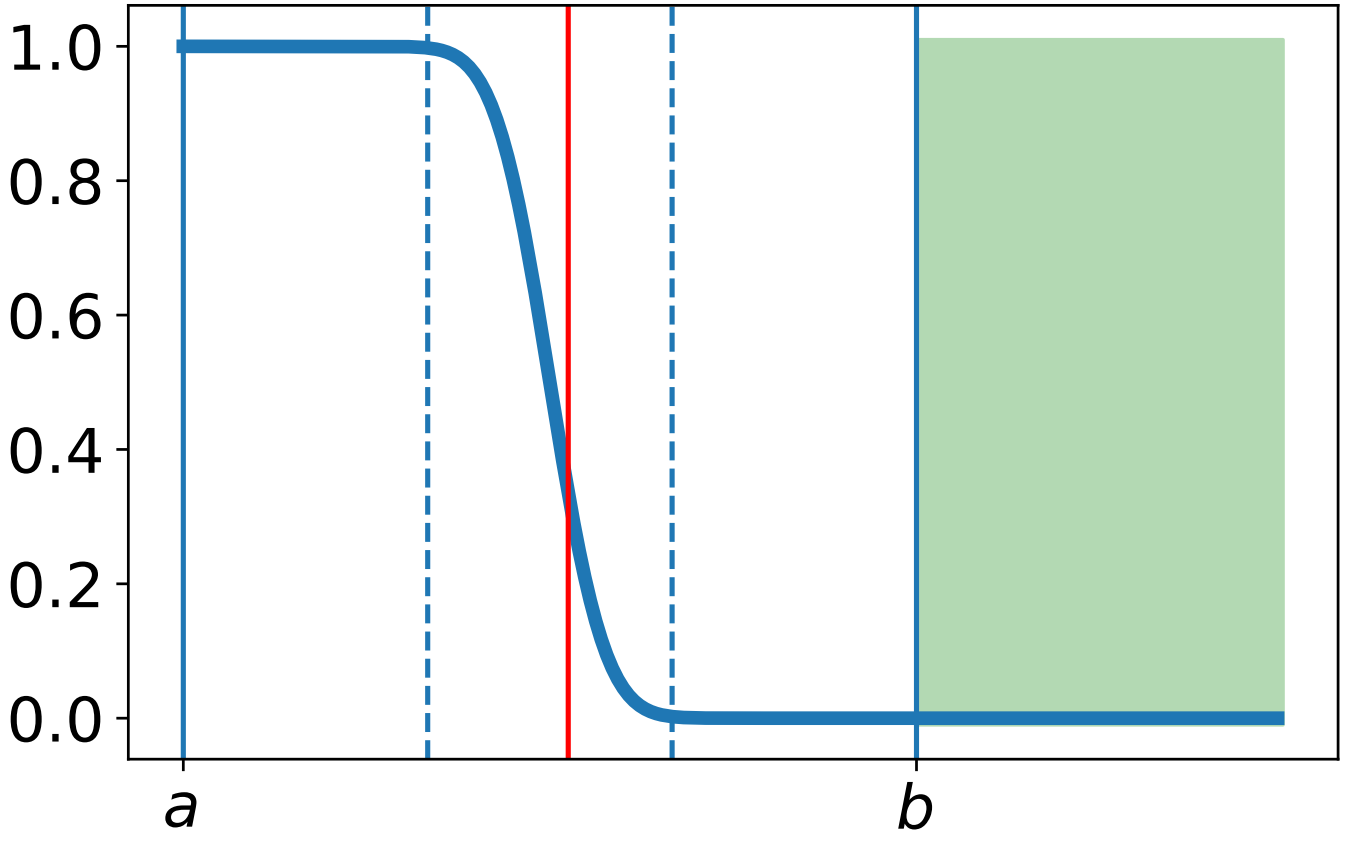

The search process

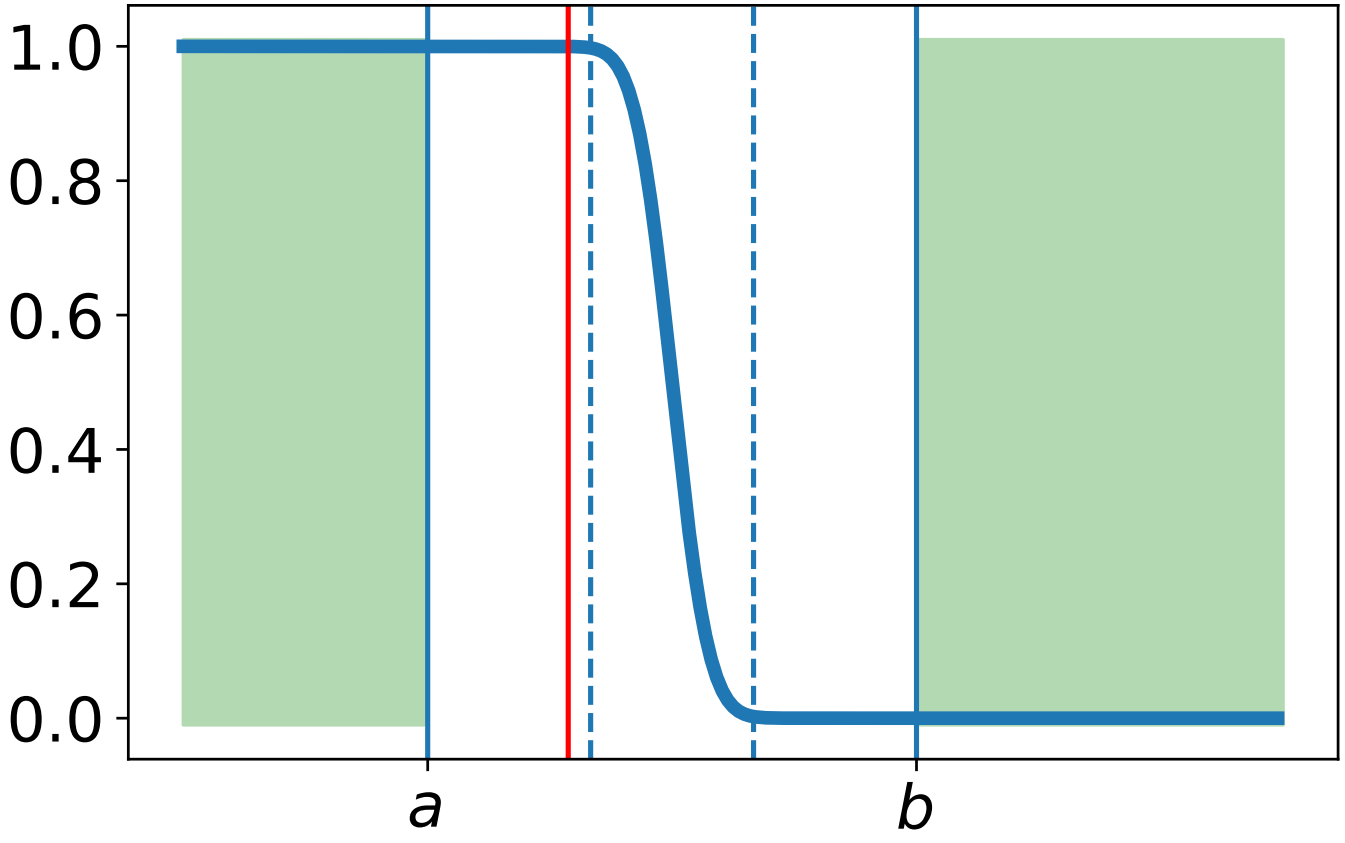

The search process

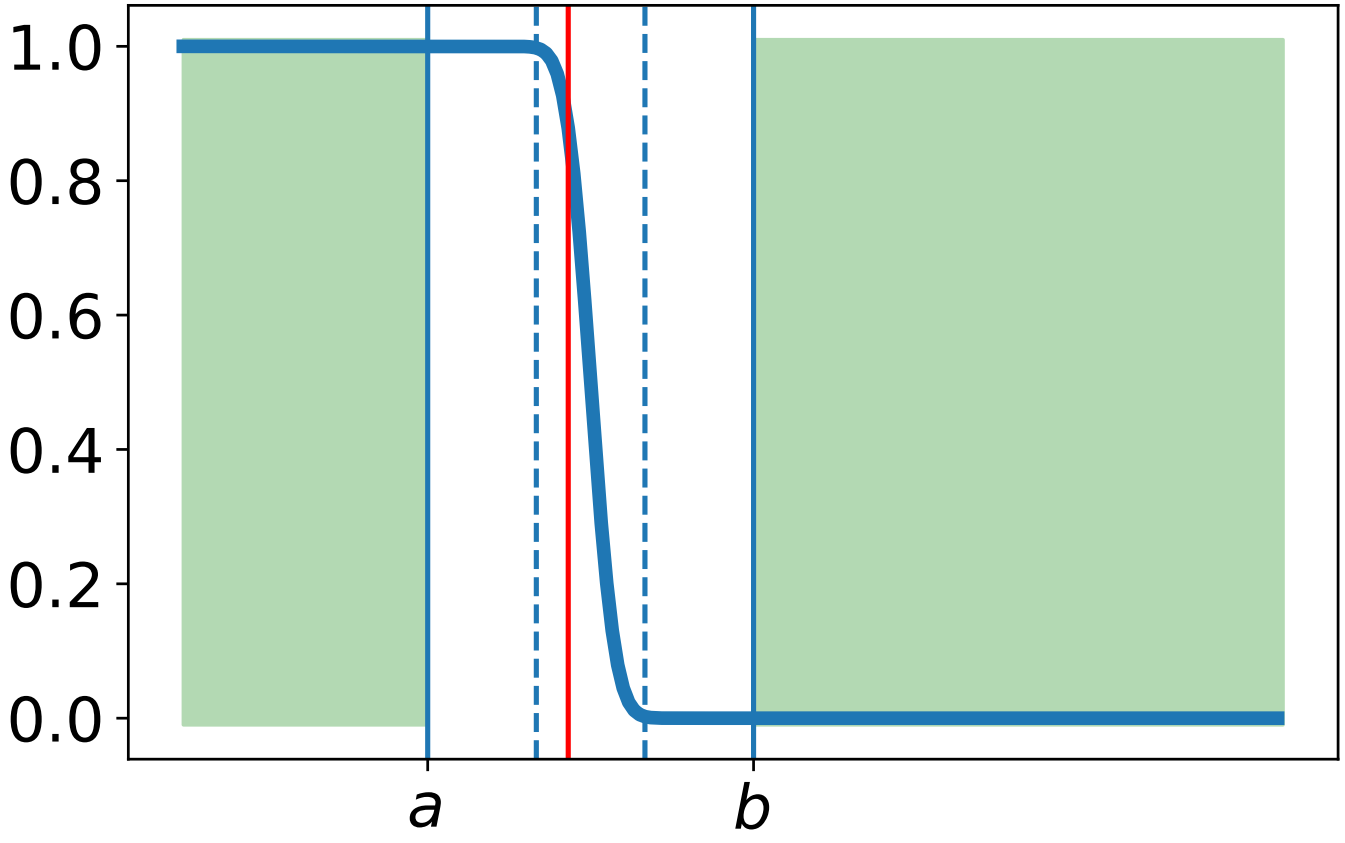

The search process

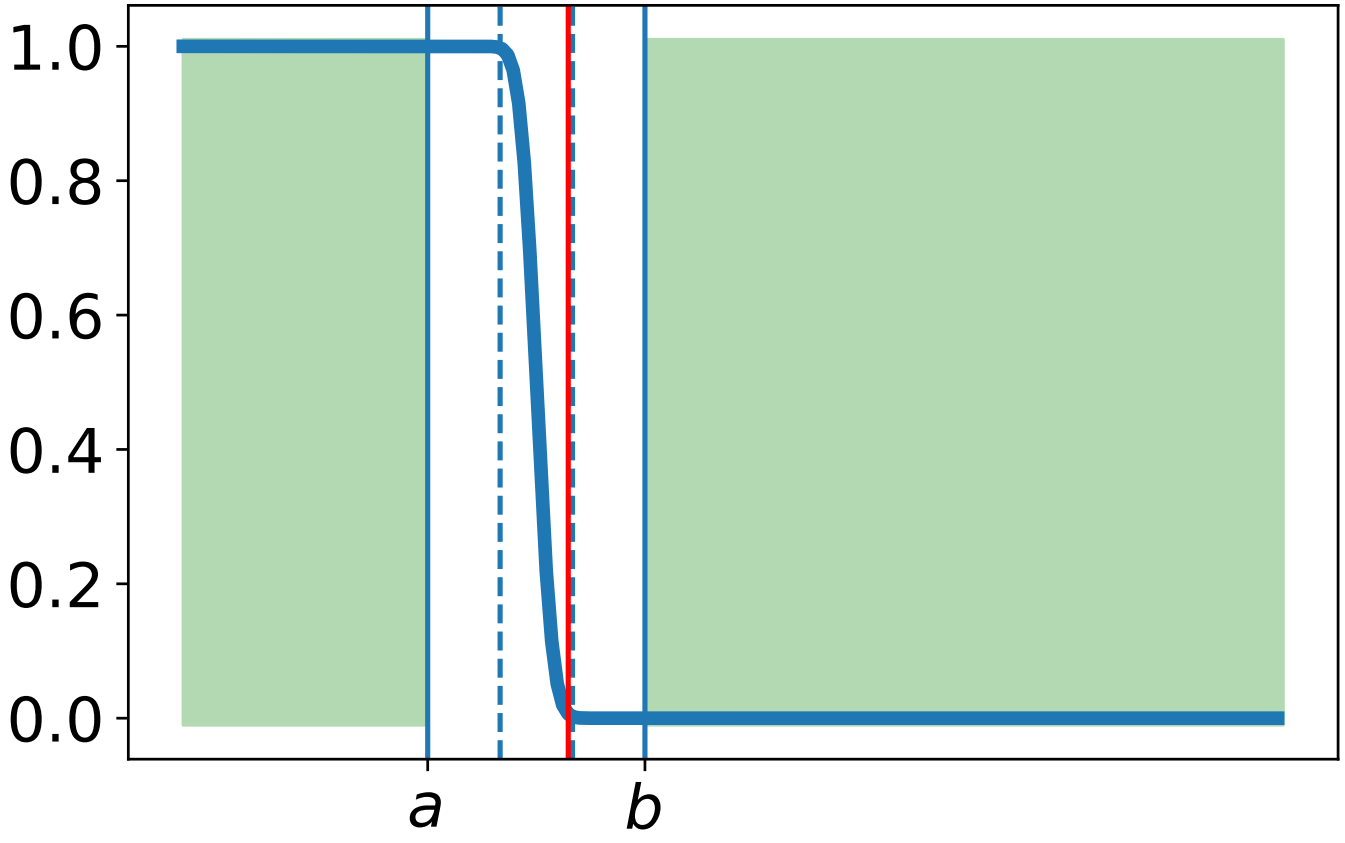

The search process

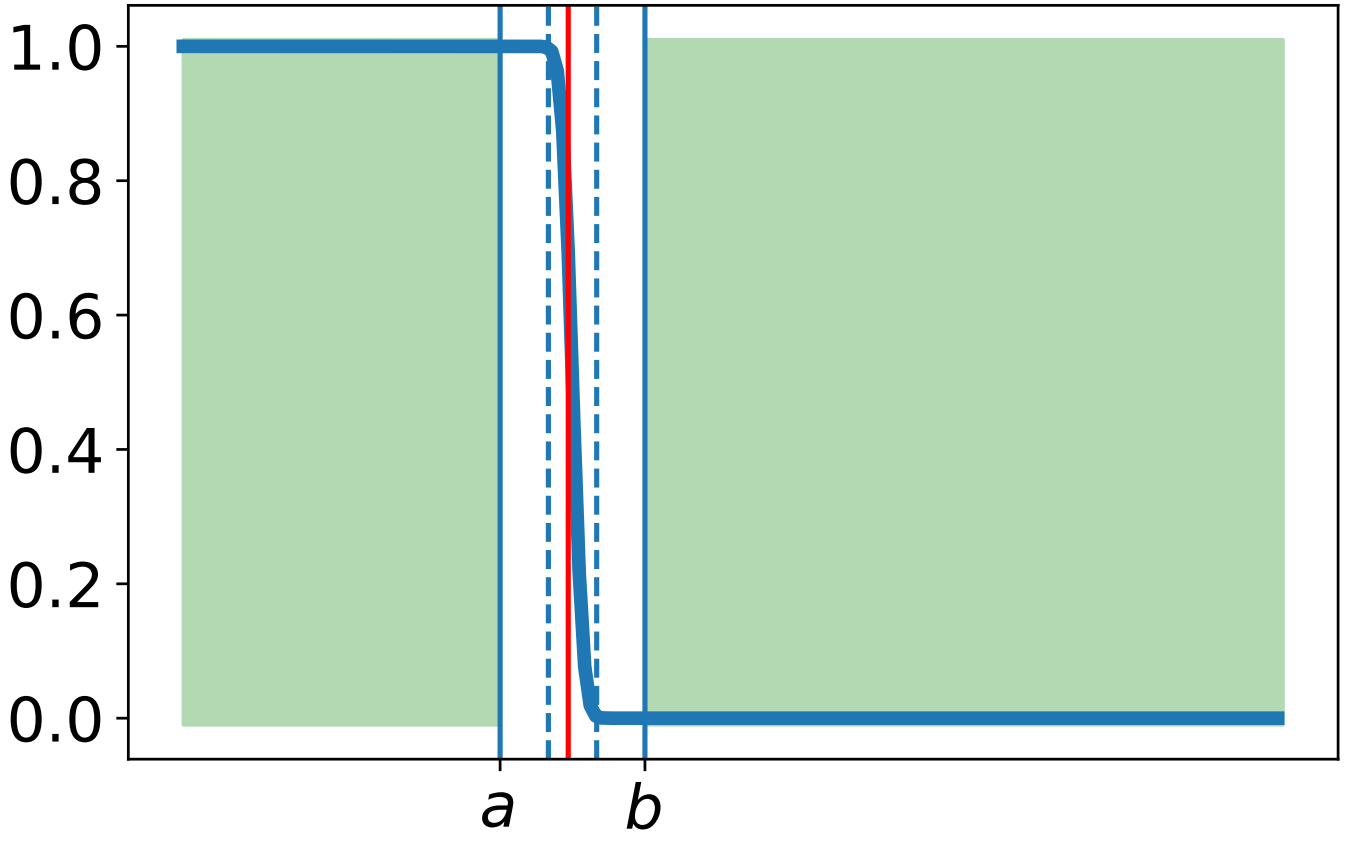

The search process

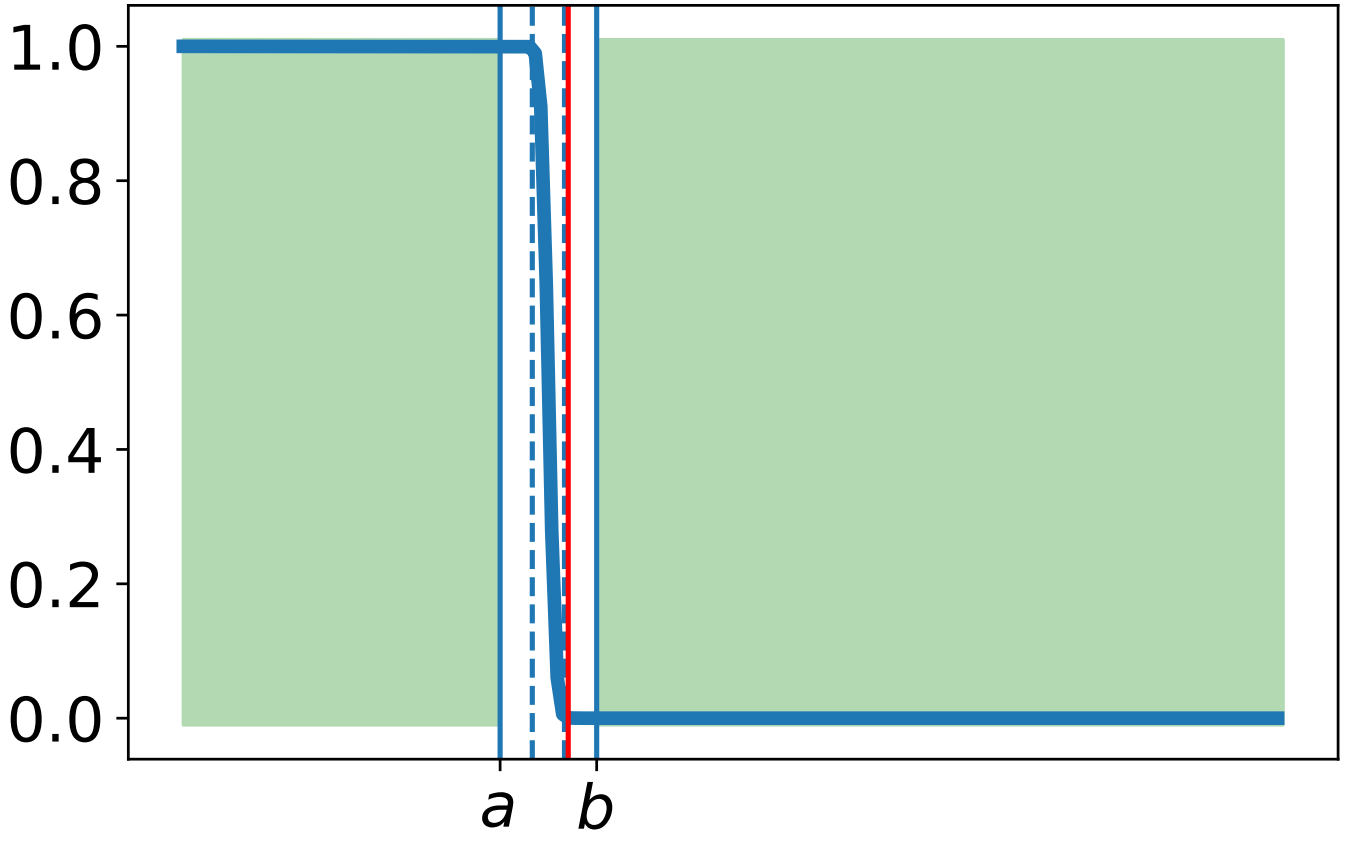

The search process

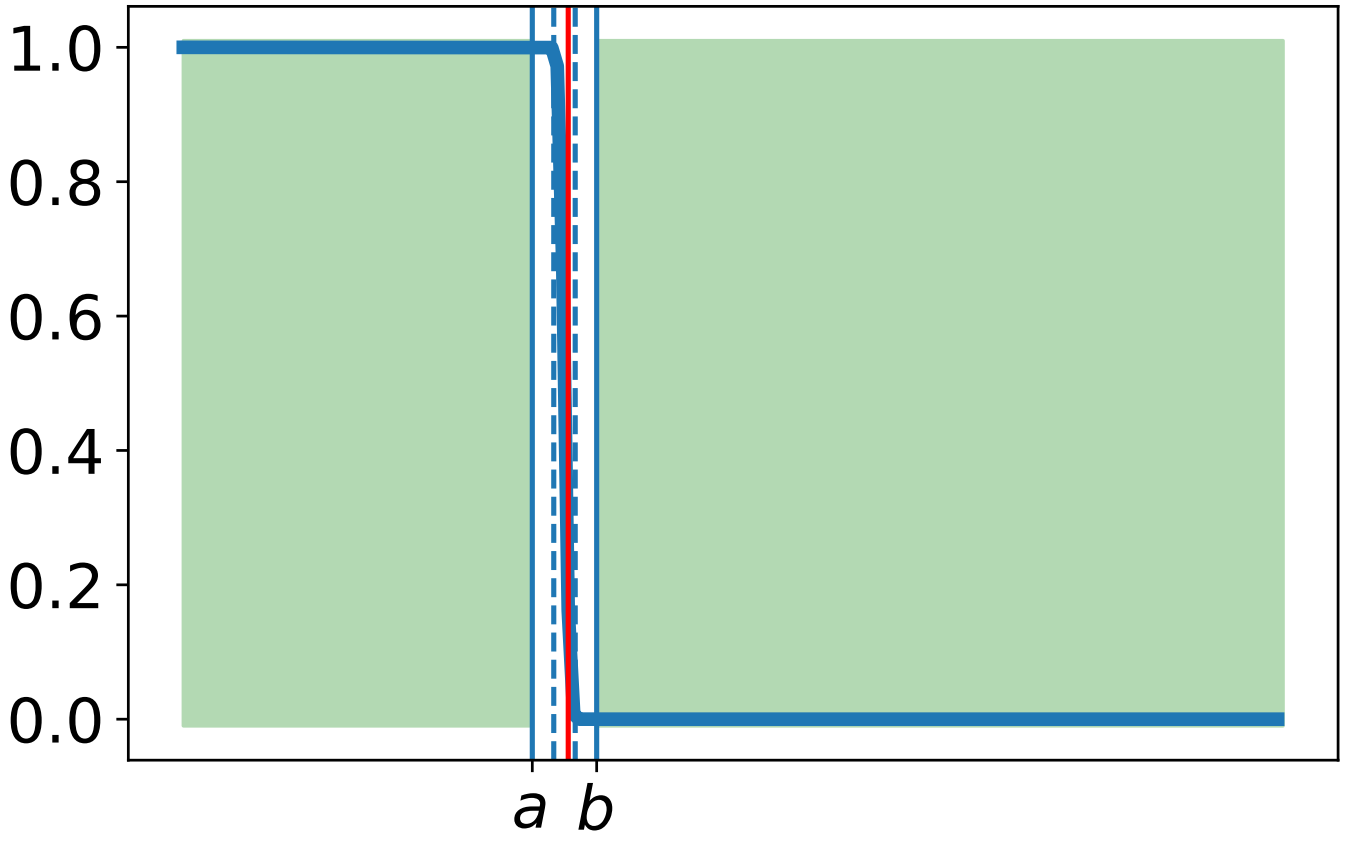

Outline

What is early fault tolerant quantum computer

- The number of logical qubits is limited
- It is undesirable to have a large number of controlled operations
- It is a priority to reduce the circuit depth

Better to run a circuit of depth $\mathcal{O}(D)$ for $\mathcal{O}(M)$ times than to run a circuit of depth $\mathcal{O}(DM)$ for a constant number of times.

(Our current thinking; there may be other definitions)

How to use one ancilla qubits

Mainly a matter of input model of H:

1. Block encoding access:

$$U_H = \begin{pmatrix} H/\alpha & * \\ * & * \end{pmatrix}$$

   Often require many ancilla qubits, e.g. sparse matrices [Gilyén-Su-Low-Wiebe, 2019]

2. Hamiltonian evolution access: $U_H = e^{-\mathrm{i}\tau H}$ for some τ.
   Can use approximate methods (e.g., Trotter splitting) without ancilla qubits

Three goals of our work

Throughout this part,

$$|\langle\phi|\psi_0\rangle|^2 = \mathrm{Tr}[\rho\,|\psi_0\rangle\langle\psi_0|] = p_0 = \gamma^2.$$

(p_0 can be replaced by its lower bound)

(1) Heisenberg-limited precision scaling: total time is $\widetilde{\mathcal{O}}(\epsilon^{-1}\mathrm{poly}(p_0^{-1}))$;

(2) Using at most one ancilla qubit

(3) Using lower-depth circuit ($\mathcal{O}(\epsilon^{-1}\mathrm{polylog}(\epsilon^{-1}p_0^{-1}))$)

Algorithms	Requirements (1)	(2)	(3)	Other issues
QPE (textbook version) [22, 48]	✓	✗	✗	
QPE (high-confidence) [37, 47, 56]	✓	✗	✓	
QPE (semi-classical QFT) [9, 30]	✓	✓	✗	
QPE (iterative) [34]	✓	✓	✓	Needs exact eigenstate ($p_0 = 1$)
The LCU approach [24]	✗	✗	✗	
The binary search approach [39]	✓	✗	✗	
VQE [45, 52, 54]	✗	✓	?	No precision guarantee
QITE [46]	✗	✓	?	Requires state tomography
QEEA [62]	✗	✓	✓	
Krylov subspace methods [31, 53, 63]	✗	✓	?	No precision guarantee
This work	✓	✓	✓	

Figure: Comparison with other ground state energy estimation algorithms. The requirements are (1) achieving the Heisenberg-limited precision scaling, (2) using at most one ancilla qubit, and (3) the maximal evolution time being at most $\mathcal{O}(\epsilon^{-1}\mathrm{polylog}(\epsilon^{-1}p_0^{-1}))$.

(L.-Tong, 2102.11340, PRX Quantum 2022)

Comparison with QPE

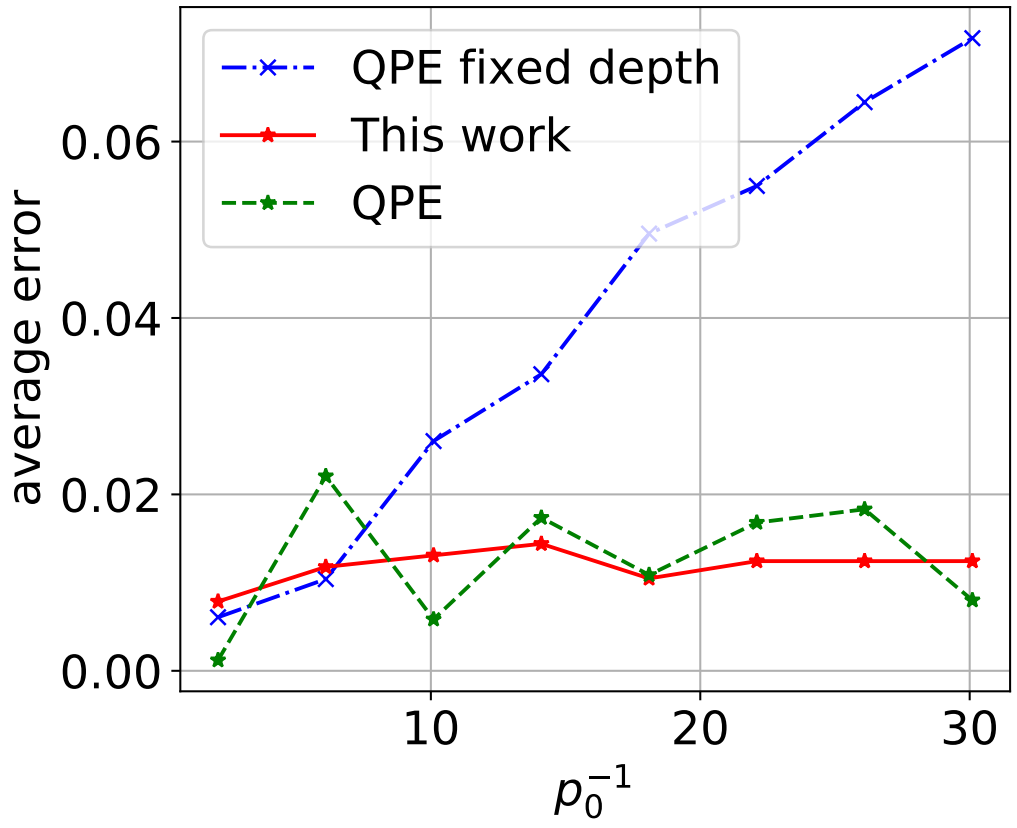

Figure: Average ground state energy estimation error for this method and QPE with fixed max evolution time, benchmarked against QPE with increasing max evolution time. $p_0 = \mathrm{Tr}[\rho\Pi_0]$.

Comparison with QPE

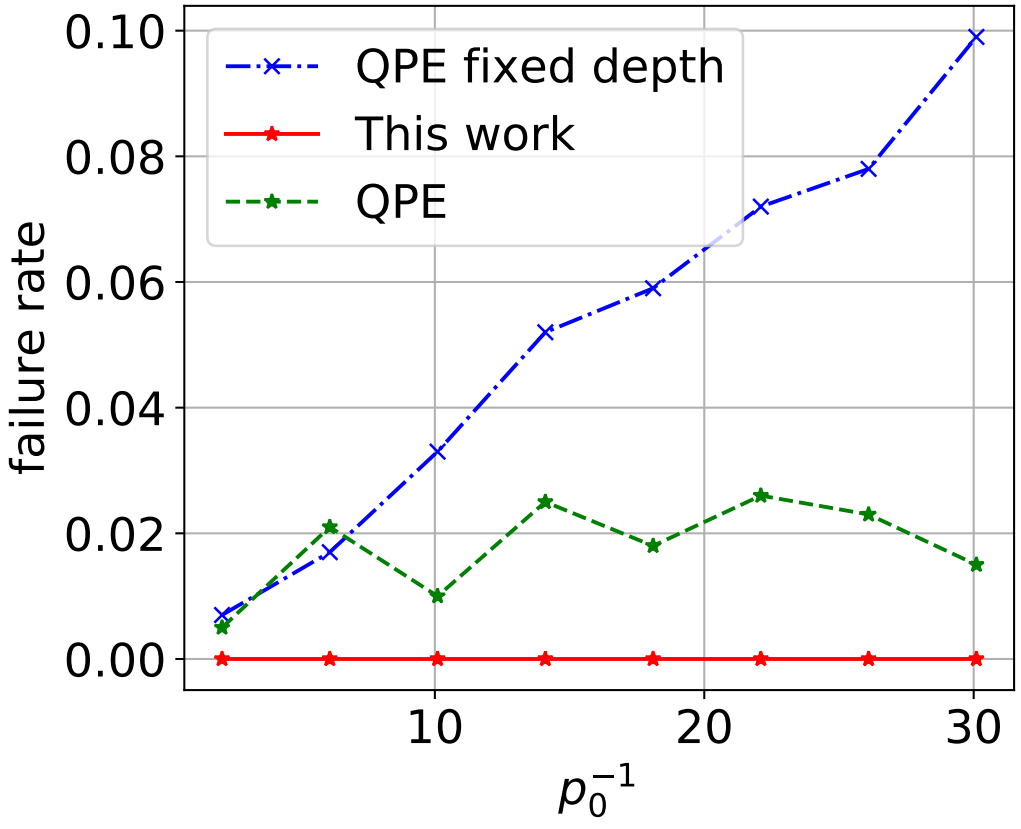

Figure: Failure (error beyond a certain threshold) rates

Key quantity: cumulative distribution function (CDF)

- Eigendecomposition of the Hamiltonian

$$H = \sum_k \lambda_k \Pi_k,$$

$p_k = \mathrm{Tr}[\Pi_k \rho]$ is the population of the initial state ρ in the k-th eigensubspace.

- **Cumulative distribution function (CDF)** of the spectral measure:

$$C(x) := \mathrm{Tr}[\rho f(x - \tau H)] = \sum_{k:\tau\lambda_k \leq x} p_k, \quad f(w) = \begin{cases} 1, & w \geq 0, \\ 0, & w < 0. \end{cases}$$

- Discontinuous function. Each jump corresponds to an eigenvalue

Statistically estimate the cumulative distribution function (CDF) of the spectral measure

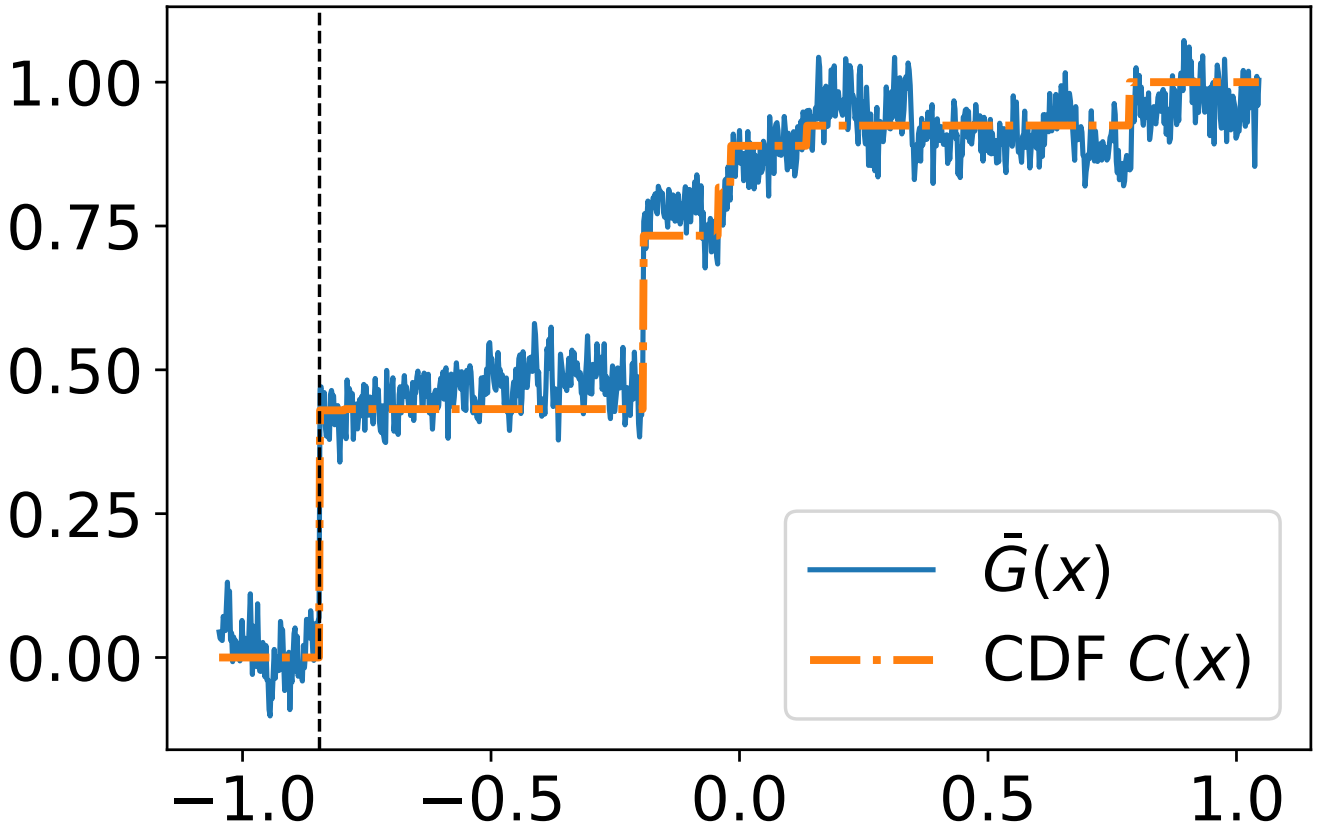

Figure: Evaluating the CDF by sampling from the quantum circuit. Note that we do not need to re-sample for each point.

The CDF: the numerical result

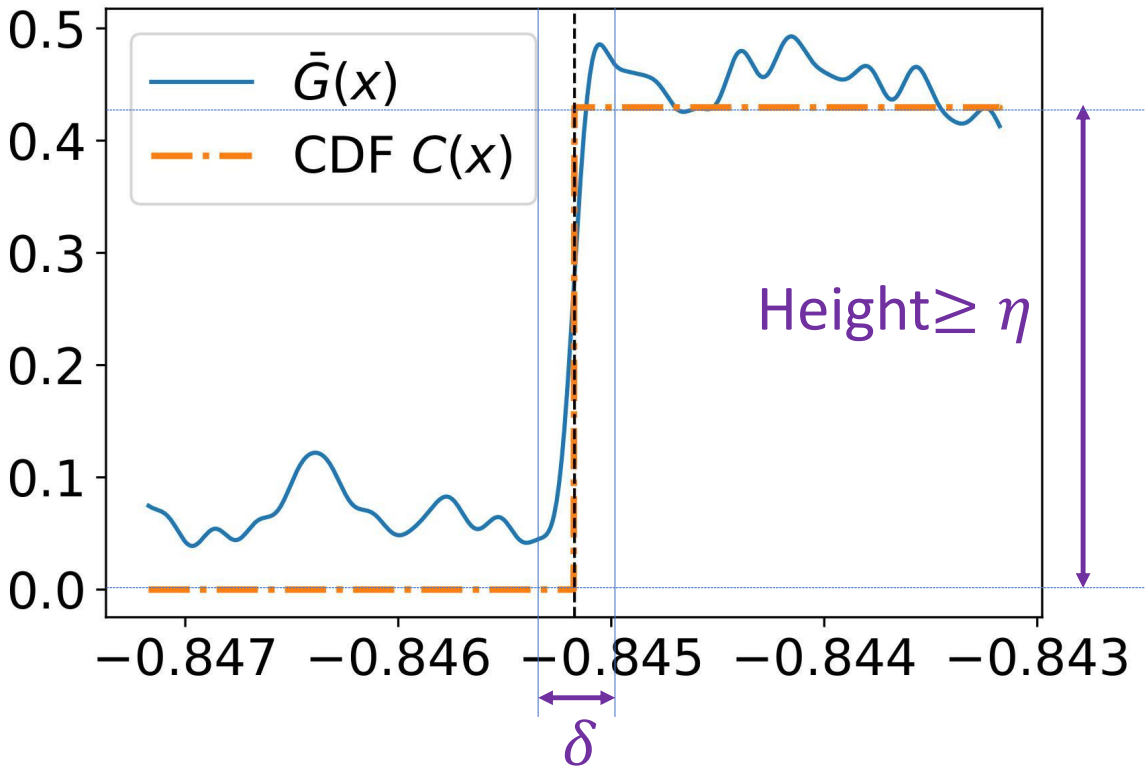

Figure: Zoom-in around the ground state energy

A simplified circuit

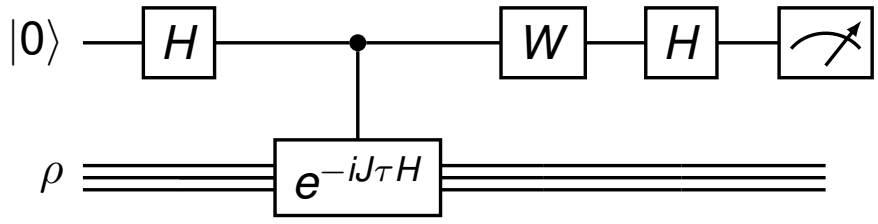

- Used in Kitaev's algorithm (iterative QPE).
- Denote the measurement outcome (± 1) by X (for $W = I$) and Y (for $W = S^\dagger$)

$$\mathbb{E}[X|J] = \mathrm{Re}\,\mathrm{Tr}[\rho e^{-iJ\tau H}]$$
$$\mathbb{E}[Y|J] = \mathrm{Im}\,\mathrm{Tr}[\rho e^{-iJ\tau H}]$$

- We could estimate $\mathrm{Tr}[\rho e^{-iJ\tau H}]$ for $J = 1, 2, \ldots, d$ and do statistical inference [Somma2019, arXiv:1907.11748] (we are not doing this).

Introducing additional randomness

- If we estimate all $\mathrm{Tr}[\rho e^{-iJ\tau H}]$ for $J = 1, 2, \ldots, d$ then the total evolution time is

$$\tau + 2\tau + \cdots + d\tau = \frac{(d+1)d}{2}\tau.$$

- We need to choose $\tau d = \epsilon^{-1}$ and this results in $\mathcal{O}(\tau^{-1}\epsilon^{-2})$ total evolution time. Cannot saturate the Heisenberg limit.

- **Idea:** sample J from a distribution and look at the output probability distribution. Sample large τ less frequently

Random evolution time

- From random variables X and Y define $Z = X + iY \in \{\pm 1 \pm i\}$,

$$\mathbb{E}[Z|J] = \mathrm{Tr}[\rho e^{-iJ\tau H}]$$

- With random evolution time J

$$\mathbb{E}[Ze^{i(\theta_J + Jx)}] = \sum_{j=-d}^{d} \mathrm{Tr}[\rho e^{-ij\tau H}] \Pr[J = j] e^{i(\theta_j + jx)},$$

viewed as a function of $x \in \mathbb{R}$.

- For a function $f(w) = \sum_{j=-d}^{d} A_j e^{ijw}$, with $\sum |A_j| = 1$, we can choose $\Pr[J = j] = |A_j|$ and $\theta_j = \arg A_j$, then

$$\mathbb{E}[Ze^{i(\theta_J + Jx)}] = \mathrm{Tr}[\rho f(x - \tau H)].$$

- The expectation value can be evaluated using Monte Carlo sampling.

- Do not need to re-sample for each x (can reuse Z)!

- If a function f can be approximated by a finite Fourier sum then $\mathrm{Tr}[\rho f(x - \tau H)]$ can be approximately computed using this method, up to an approximation error.

Summary of the algorithm

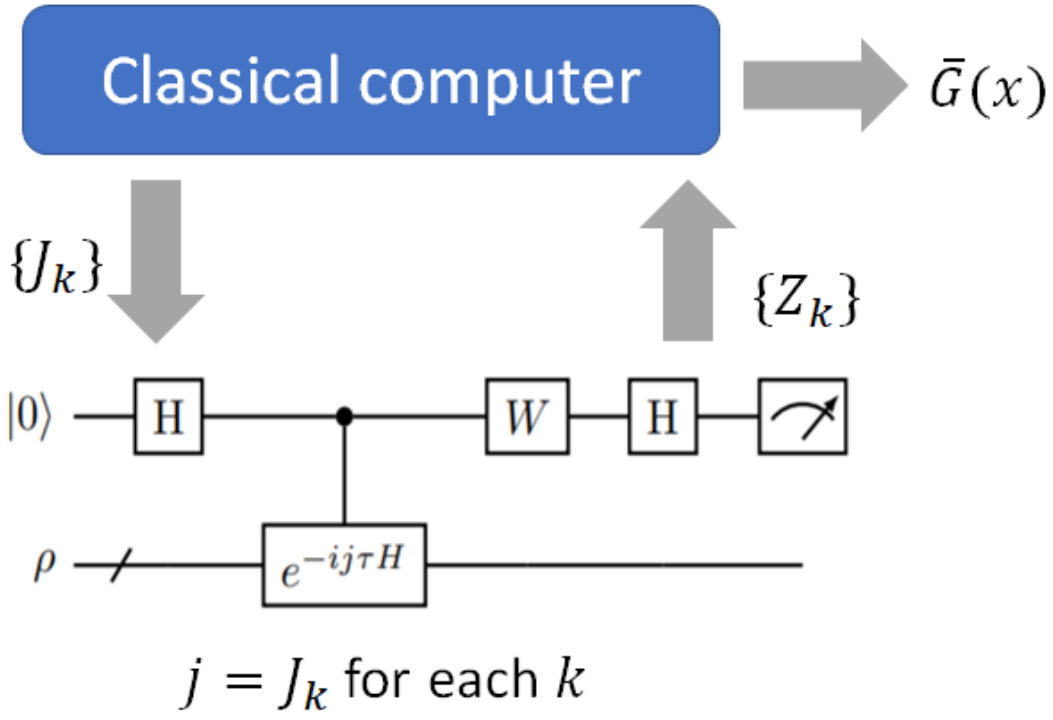

Figure: Schematic representation of the algorithm to compute the CDF. $\bar{G}(x)$ is the approximate CDF.

Ground state energy estimation

- Only need to find the first jump in the CDF; can be done using a binary search;
- The transition length is $\delta = \tau\epsilon$ to ensure precision ϵ;
- The max (coherent) evolution time is $\tau/\delta = \mathcal{O}(\epsilon^{-1})$ (compare with $\mathcal{O}(\epsilon^{-1}p_0^{-1})$ for QPE, lower circuit depth);
- Monte Carlo accuracy $\mathcal{O}(p_0)$, therefore $N_s = \mathcal{O}(p_0^{-2})$;
- Total evolution time is $\mathcal{O}(N_s\tau d) = \mathcal{O}(\epsilon^{-1}p_0^{-2})$ **(Heisenberg scaling)**;

Outline

Conclusion

- Fault-tolerant quantum algorithms with performance guarantees for ground state energy estimation.
- Assumption of a good initial guess is crucial.
- Access model: block encoding and Hamiltonian evolution. Combine the best of both worlds?
- The CDF could be of use in itself.
- Quantum algorithms: (somewhat) traditional numerical analysis in a new context

Acknowledgment

Thank you for your attention!

Lin Lin
https://math.berkeley.edu/~linlin/

U.S. DEPARTMENT OF ENERGY | Office of Science
NSF
AFOSR AIR FORCE OFFICE OF SCIENTIFIC RESEARCH
Google AI Quantum
SIMONS FOUNDATION

The control-free version

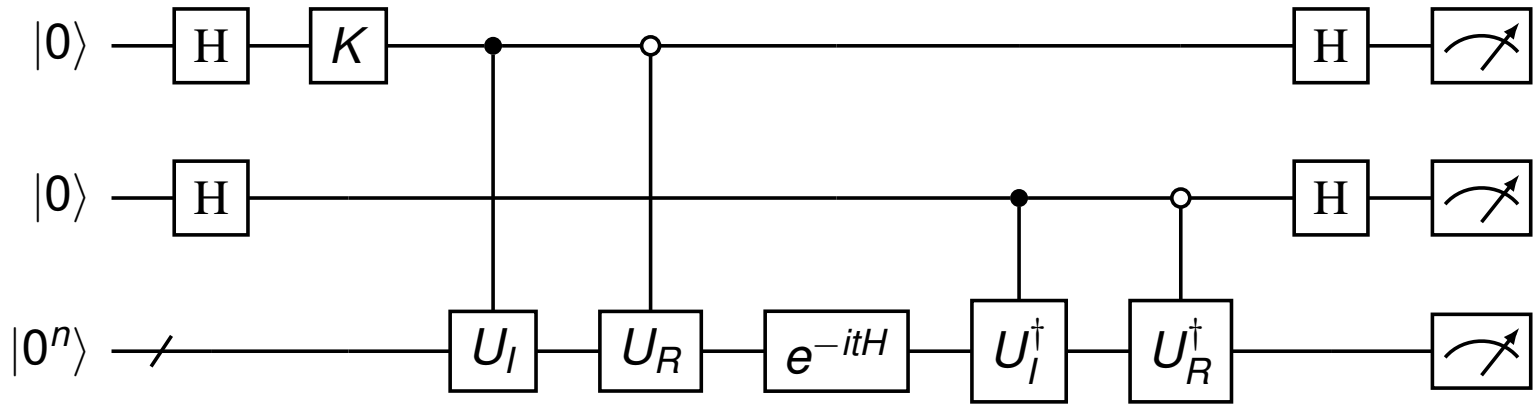

- We assume we have a reference eigenstate $|\psi_R\rangle$ that is easy to prepare and corresponds to a known eigenvalue.[1,2]
- Example: for system with particle number conservation we can use the vacuum state.

[1] Huggins, Lee, Baek, O'Gorman, Whaley, 2019, arXiv:1909.09114

[2] Russo, Rudinger, Morrison, Baczewski, 2020, arXiv:2007.08697

Fourier approximation

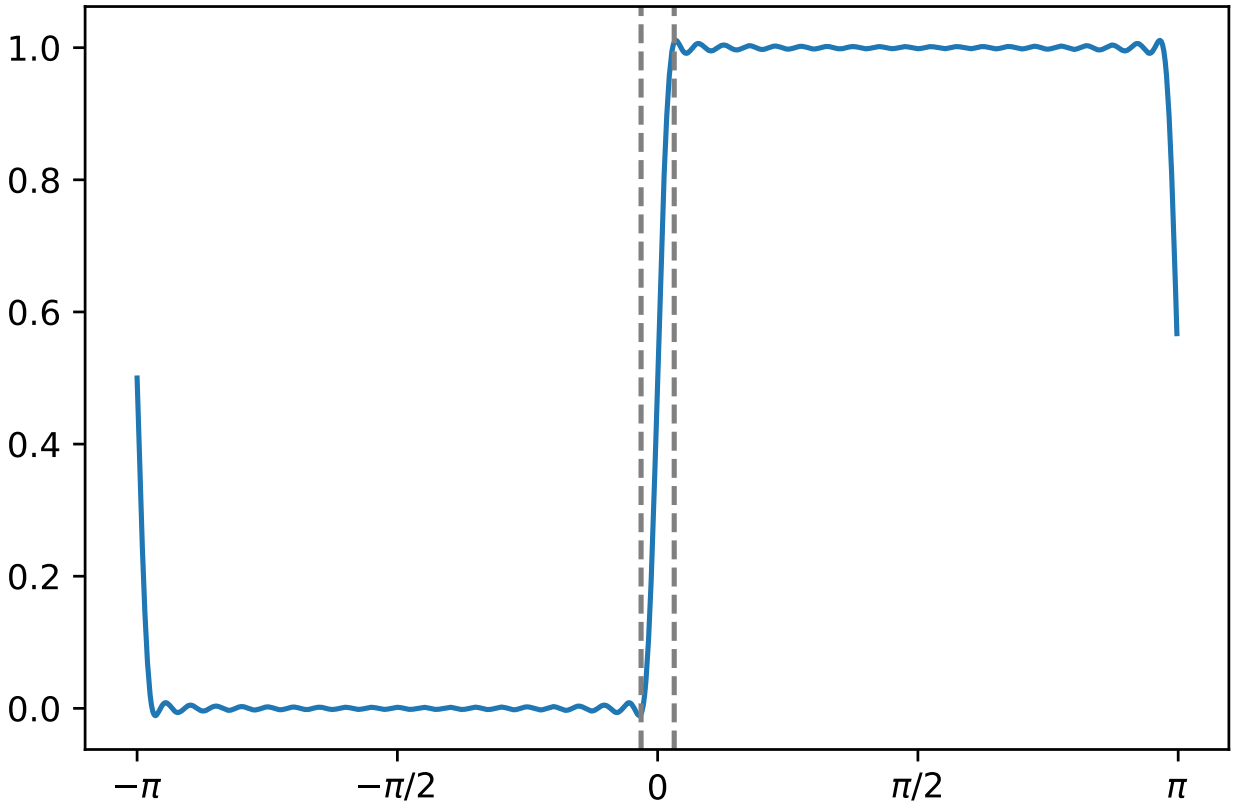

Figure: Our construction of the approx Heaviside function. The number of terms is $2d + 1$ where $d = \mathcal{O}(\delta^{-1} \log(\epsilon'^{-1}))$.

Fourier approximation

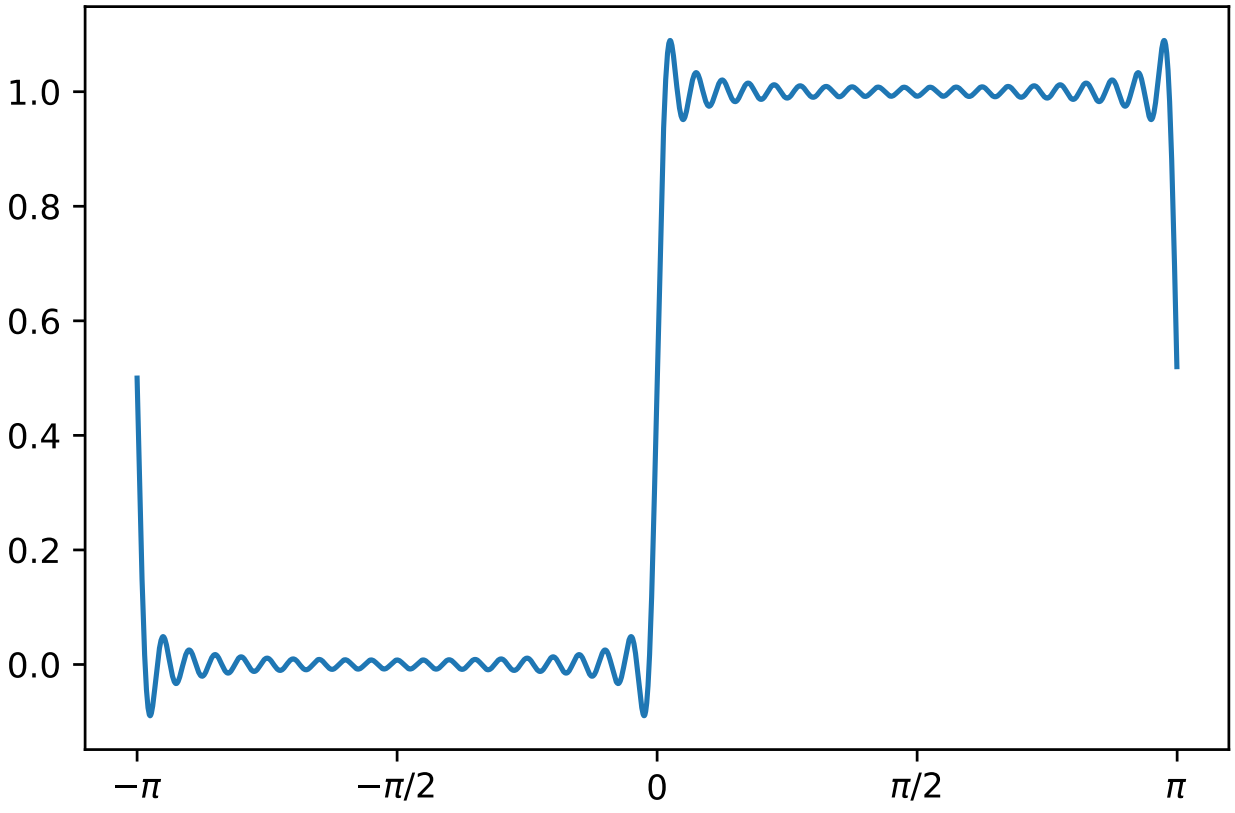

Figure: Directly truncating the Fourier expansion of the Heaviside function. Note the Gibbs phenomenon.