Title: Quantum thermal state preparation

Speakers: Chi-Fang Chen

Series: Perimeter Institute Quantum Discussions

Date: May 10, 2023 - 3:30 PM

URL: https://pirsa.org/23050090

Abstract: A key subroutine in quantum computing, especially in quantum simulation, is to prepare thermal states or ground states of Hamiltonians. Today, I will talk about a new family of quantum algorithms for this task. Physically, our algorithms distill the essence of system-bath interaction by simulating an effective Lindbladian; computationally, our algorithms are quantum analogs of classical Markov chain Monte Carlo sampling. Given the ubiquity of thermodynamics and the triumph of classical Monte Carlo methods, we anticipate that quantum thermal state preparation will become indispensable in quantum computing.

Joint work with Michael J. Kastoryano, Fernando G.S.L. Brandão, and András Gilyén. https://arxiv.org/abs/2303.18224

Zoom Link: https://pitp.zoom.us/j/91641127738?pwd=cExPM3Bvd3BaYnJYS0U0UlBiVTJ0QT09

Quantum Thermal State Preparation

[CKBA '23], [in progress]

Anthony (Chi-Fang) Chen
with Fernando Brandão, András Gilyén, Michael Kastoryano

May 10th@PI

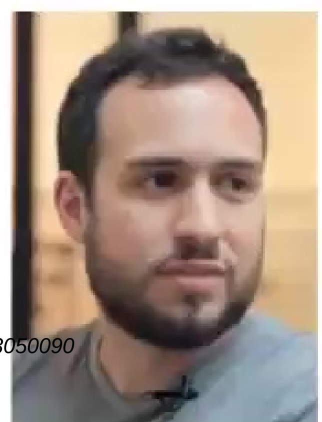

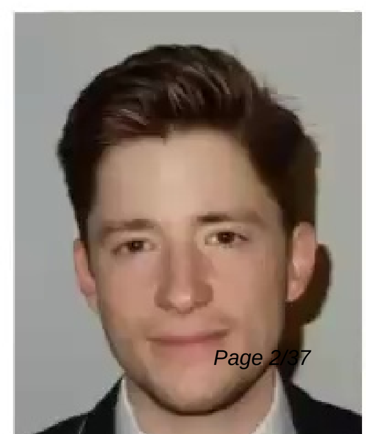

Quantum thermal state preparation (Gibbs sampling):

Given a quantum Hamiltonian and a quantum computer,
prepare the thermal state

$$\boldsymbol{\sigma}_\beta \propto \mathrm{e}^{-\beta \boldsymbol{H}} = \sum_i \mathrm{e}^{-\beta E_i} \left|\psi_i\right\rangle \left\langle\psi_i\right|$$

(in the energy basis)

Killer app: Quantum simulation

"What's the ground energy of CO_2?"

state preparation	$\|\psi\rangle$
time evolution	$e^{\mathrm{i}\boldsymbol{H}t}$
measurement	$\boldsymbol{O}$

Amazon[20' Chamberland,...,B],
Google[21' Lee],
Microsoft[21' von Burg],..

Mysteries of low-energy states

Classically | Quantumly

classical heuristic*

Thermodynamics?
(average-case?)

worst-case QMA-hard*

Easy — Hard

Which assumptions and algorithms?

*tensor network for gapped ground state (area law? [hastings '07])

*

A new angle

"Monte-Carlo style" quantum algorithms

classical Monte-Carlo (Gibbs sampling)

- physics simulation (Ising model)
- Optimization problems/CSPs
- Machine learning, ...

Open system (Thermodynamics)

- quantum chaos
- spin glass
- Thermodynamics

A new angle

“Monte-Carlo style” quantum algorithms

Background: classical Monte-Carlo sampling

Markov Chain Monte Carlo (MCMC)

Implement a Markov chain that converges to the Gibbs state.

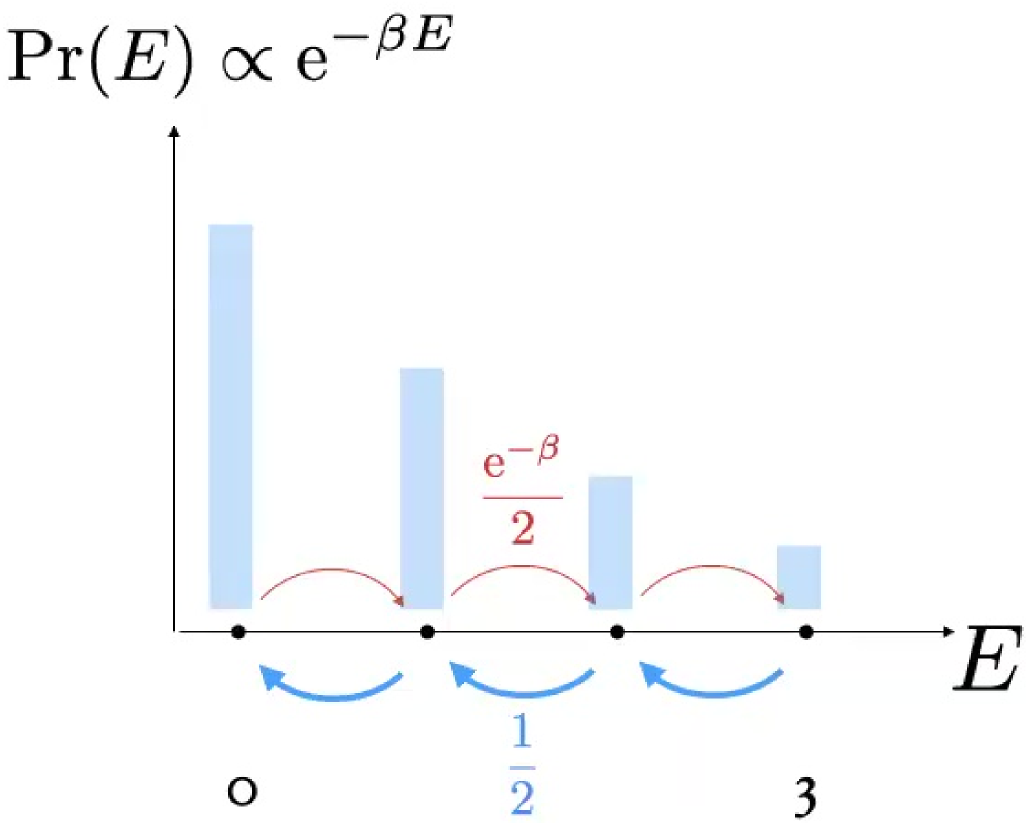

Detailed balance gives Gibbs state

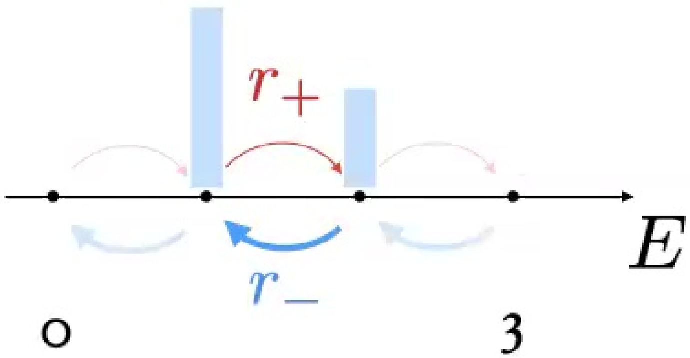

Transition matrix

$\Pr(s_1 \to s_2)$

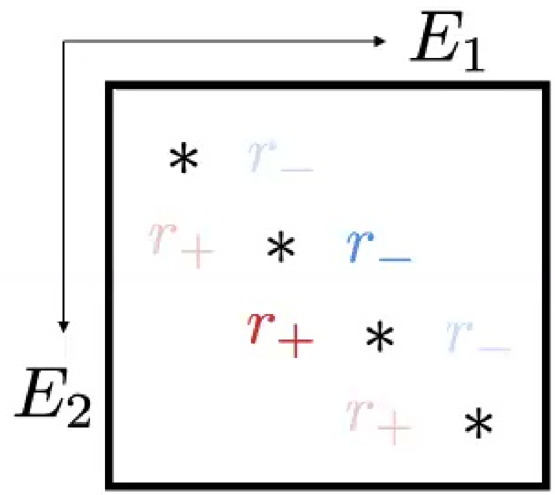

"forward equals to backward"

$$e^{-\beta E} \cdot r_+ = e^{-\beta(E+1)} \cdot r_-$$

Symmetry

$$r_+ / r_- = e^{-\beta}$$

General recipe for detailed balance

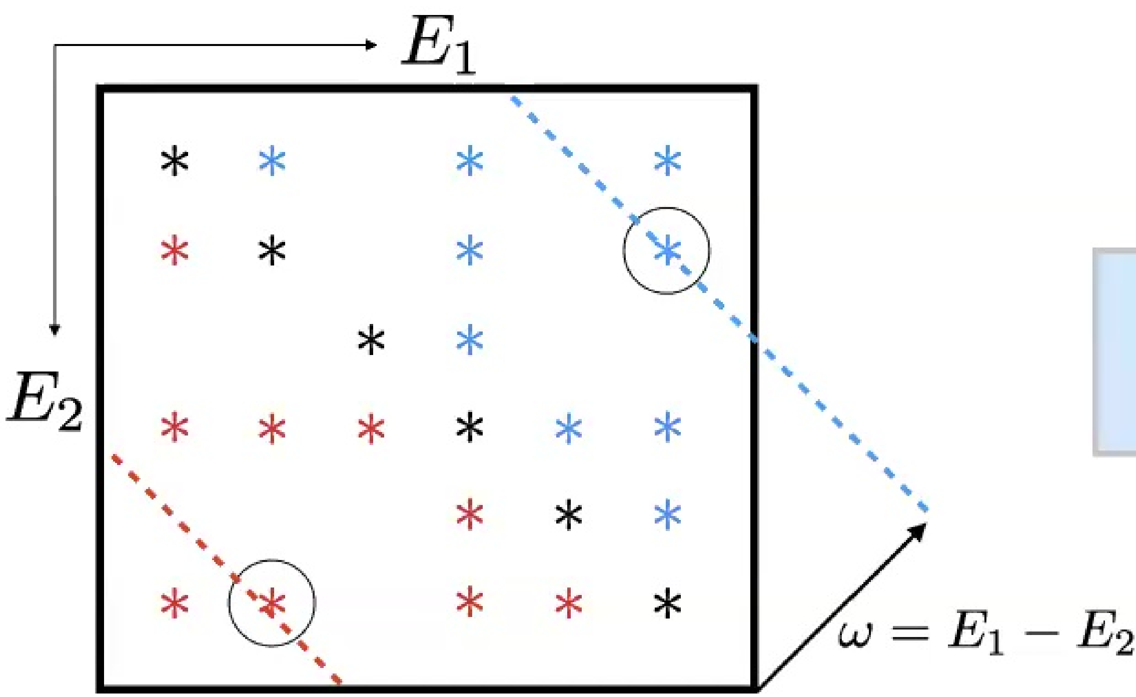

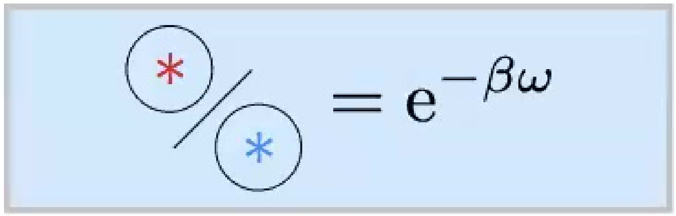

"Rejection" sampling blindly imposes detailed balance
(without knowing the full matrix!)

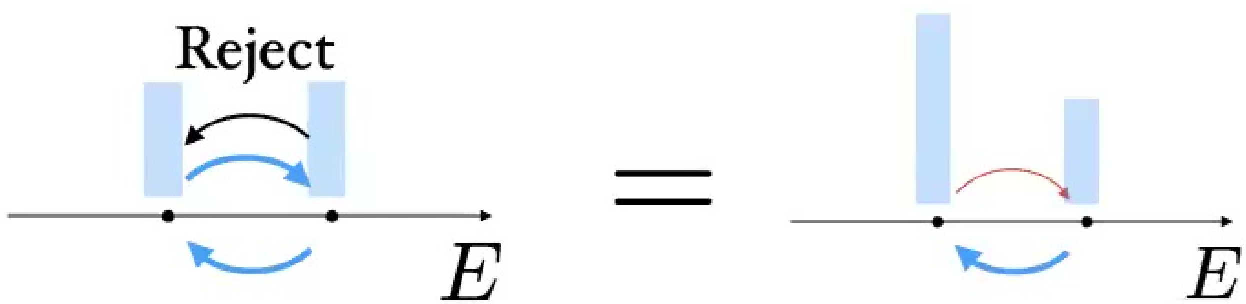

Metropolis-Hastings['53]

parameters:

Hamiltonian	$\boldsymbol{H}$
inv. temperature	β
jumps	$\{\boldsymbol{A}^a\}$

e.g., Ising model

$$H(z) = \sum_{<ij>} J_{ij} z_i z_j \quad \text{where} \quad z_i = \pm 1$$

spin flips $\{\boldsymbol{A}^a\} = \{z_i \to -z_i\}$

Algorithm:

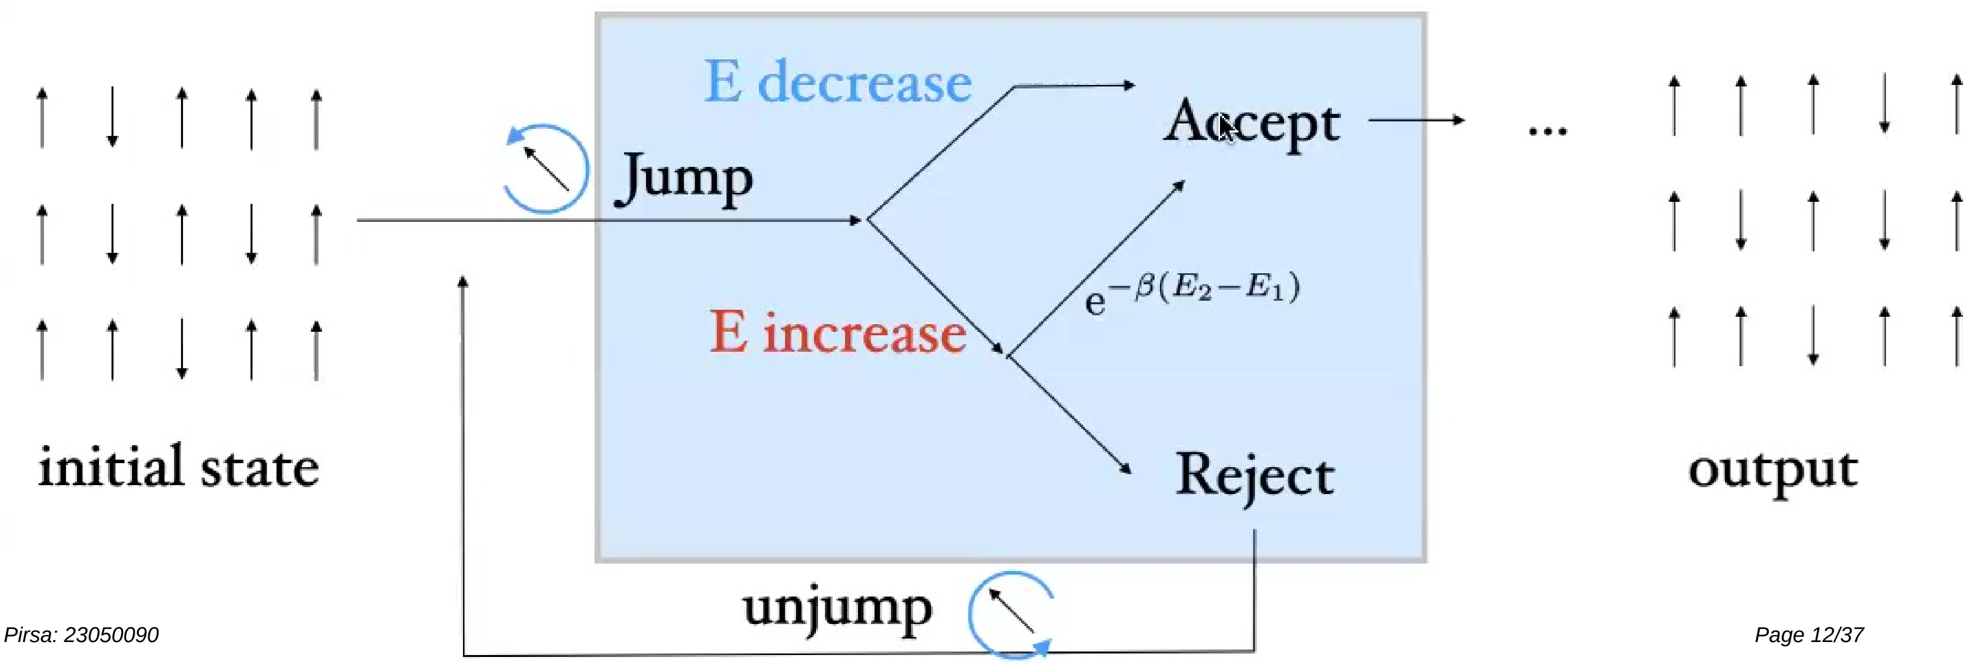

Quantizing Monte-Carlo

	Classical	Quantum
configuration	$z = \uparrow\uparrow\downarrow$	$\|\psi\rangle$
Gibbs dist.	$\Pr(z) \propto \mathrm{e}^{-\beta E(z)}$	$\boldsymbol{\sigma}_\beta \propto \sum_i \mathrm{e}^{-\beta E_i} \|\psi_i\rangle \langle\psi_i\|$
Hamiltonian $\boldsymbol{H}$	$\sum_{<ij>} J_{ij} z_i z_j$	$\sum_{<ij>} J_{ij} \boldsymbol{X}_i \boldsymbol{Y}_j$
Jumps $\{\boldsymbol{A}^a\}$	$\{z_i \to -z_i\}$	$\boldsymbol{X}_i, \boldsymbol{Y}_i, \boldsymbol{Z}_i$
Algorithm	Metropolis-Hastings	? (Today)
Fixed point	detailed balance	
mixing	conditional	

Today: Designing Monte-Carlo style quantum algorithms

Metropolis on the energy basis?

$$\Pr(z) \propto \mathrm{e}^{-\beta E(z)} \longrightarrow \boldsymbol{\sigma}_\beta \propto \sum_i \mathrm{e}^{-\beta E_i} |\psi_i\rangle \langle\psi_i|$$

bit strings $\quad z \longmapsto |\psi_i\rangle \langle\psi_i| \quad$ energy eigenstates

transition probability $\quad \Pr(z_j \to z_i) \to |\langle\psi_i| \boldsymbol{A}^a |\psi_j\rangle|^2 \quad$ transition amplitudes

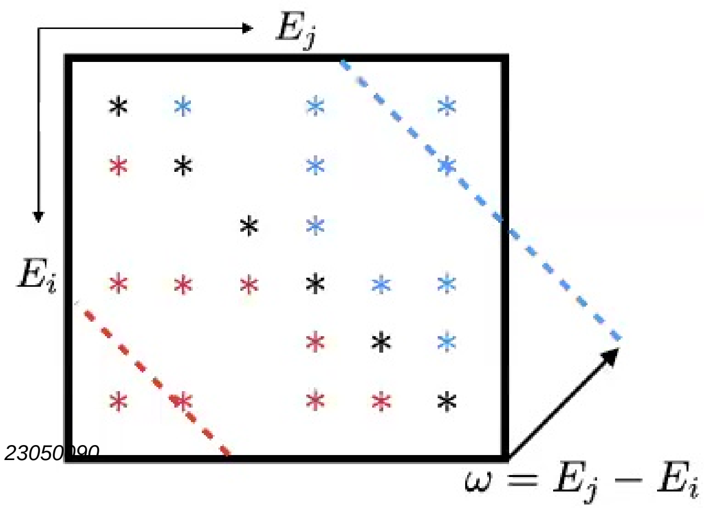

$\longrightarrow$

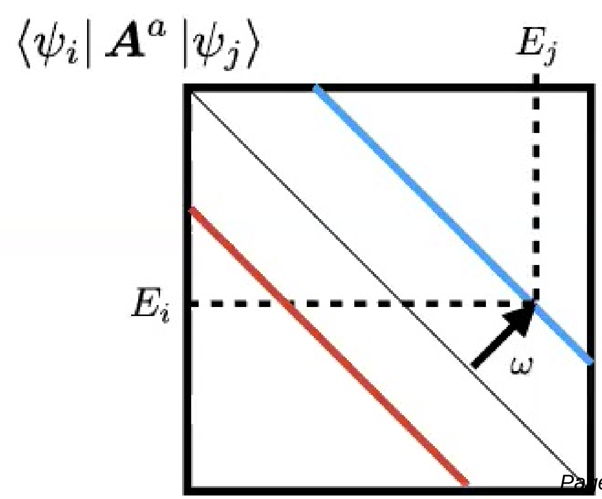

quantum challenges

- (**Algorithm.**) coherently reject a quantum jump
  - without collapsing the state (Mariott-Watrous['05];rewinding)

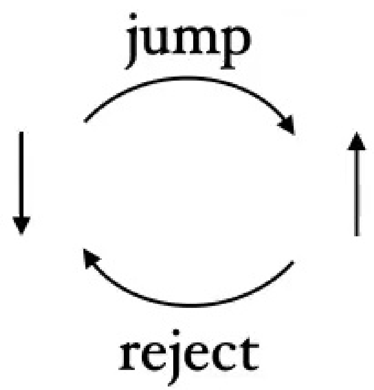

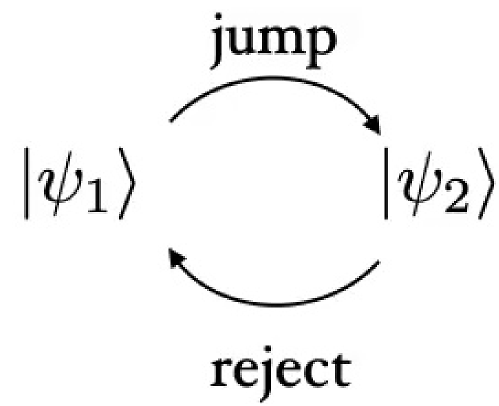

- (**Analysis.**) Energy measurement is uncertain
  - (fixed point?) detailed balance is broken
  - (mixing?) nearby energies can stay entangled

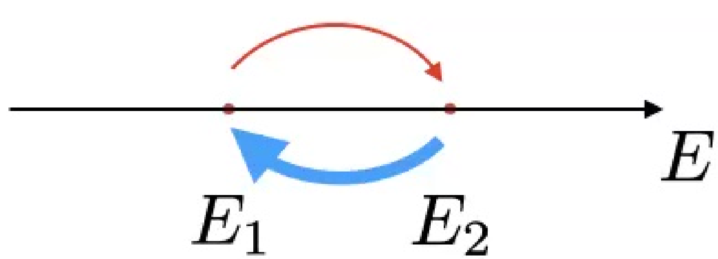

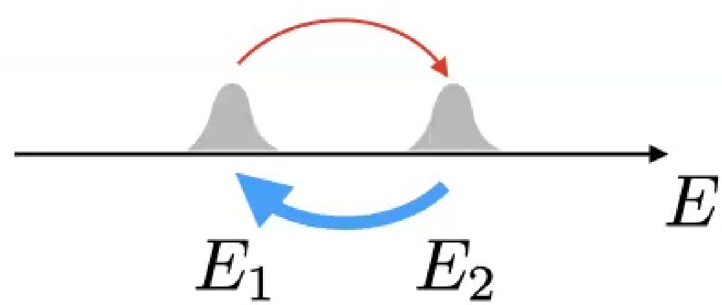

attempts: [Temme et al '09] [Hung et al '19] [Wocian&Temme '21]

Metropolis on the energy basis?

$$\Pr(z) \propto \mathrm{e}^{-\beta E(z)} \longrightarrow \boldsymbol{\sigma}_\beta \propto \sum_i \mathrm{e}^{-\beta E_i} |\psi_i\rangle \langle\psi_i|$$

bit strings $\quad z \to |\psi_i\rangle\langle\psi_i| \quad$ energy eigenstates

transition probability $\quad \Pr(z_j \to z_i) \to |\langle\psi_i| \boldsymbol{A}^a |\psi_j\rangle|^2 \quad$ transition amplitudes

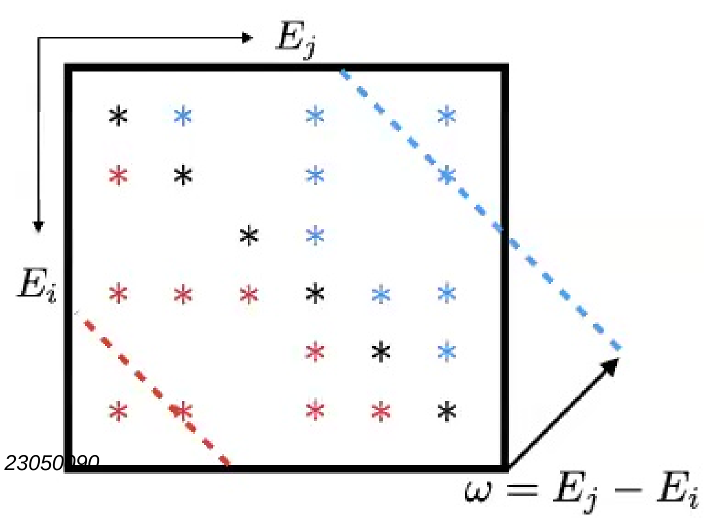

$\longrightarrow$

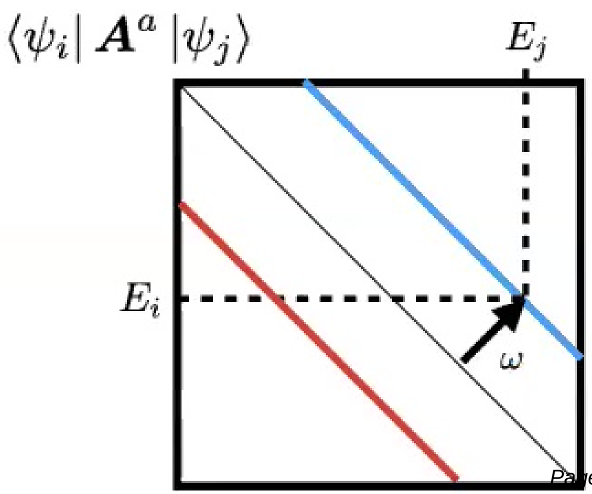

How was it done in Nature ?

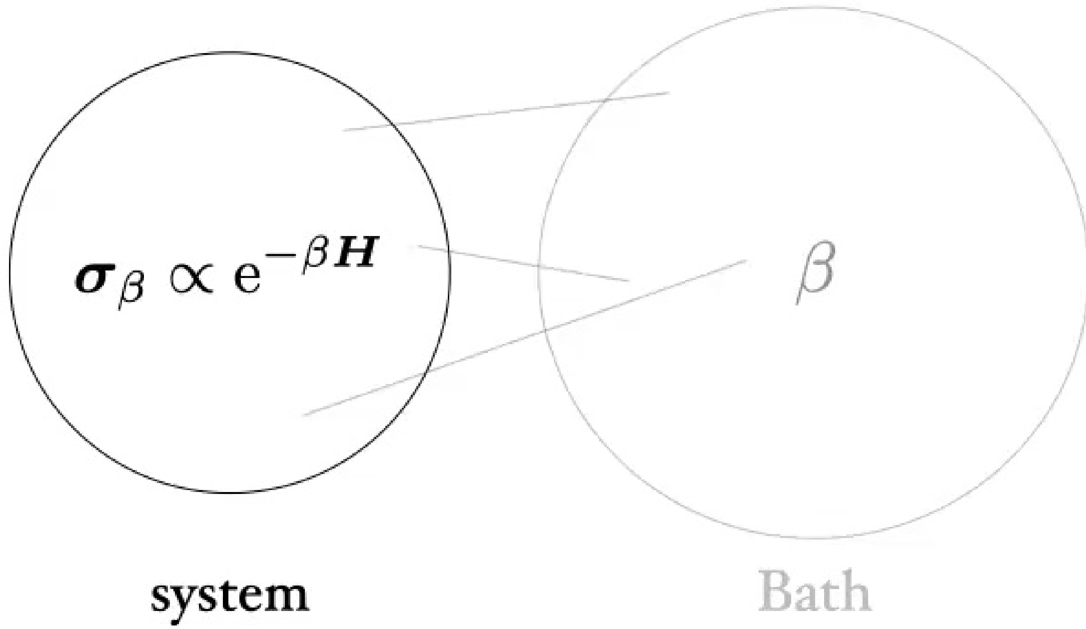

Thermodynamics: coupling to a bath gives the Gibbs state

Continuous walk solves rejection

[Davies '74]

- instead of discrete rejection, lower the rate of unwanted transitions

$$\{\sqrt{\gamma(\omega)}\boldsymbol{A}^a(\omega)\}_{a,\omega}$$

rate x transition ampl.

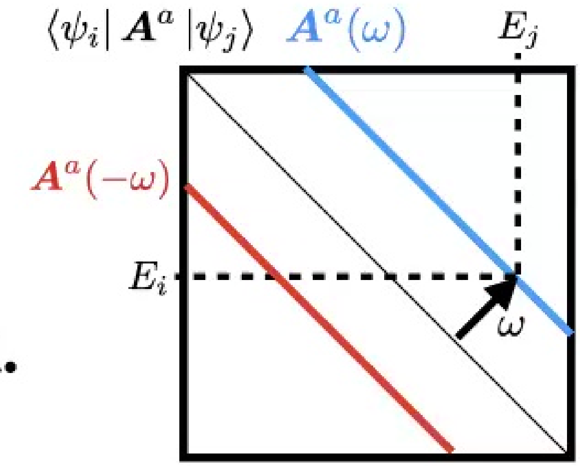

- ensures detailed balance (thus Gibbs fixed point)

$$\gamma(\omega)/\gamma(-\omega) = e^{\beta\omega}$$

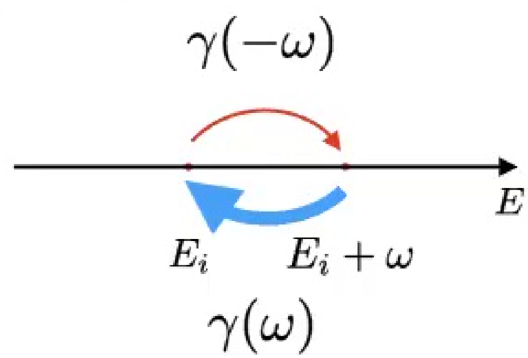

- (still, assuming perfect energy measurements)

Lindbladian formalism

$$\mathcal{L}^{\dagger}[\boldsymbol{\rho}] = \sum_{\boldsymbol{K}} \left(\boldsymbol{K}\boldsymbol{\rho}\boldsymbol{K}^{\dagger} - \frac{1}{2}(\boldsymbol{K}^{\dagger}\boldsymbol{K}\boldsymbol{\rho} + \boldsymbol{\rho}\boldsymbol{K}^{\dagger}\boldsymbol{K}) \right)$$

Lindbladian $\mathcal{L}^{\dagger}$:a continuous-time generator of "quantum" Markov chain (aka Master equation, Liouvillians)

$$\boldsymbol{\rho}(t) = \mathrm{e}^{\mathcal{L}^{\dagger}t}[\boldsymbol{\rho}]$$

non-unitary: have fixed point(s) $\mathrm{e}^{\mathcal{L}^{\dagger}t}[\boldsymbol{\rho}] \rightarrow \boldsymbol{\rho}_{fix}$

Kraus operators $\{K\}$: infinitestimal "quantum" transitions

(aka jump operator, Lindbladian operators)

Main Results: Monte-Carlo style Quantum Algorithms

We efficiently implement a Lindbladian whose fixed point is approximately Gibbs.

- (**universal.**) applies to any Hamiltonian (mixing time varies)
- (**no fine-tuning.**) user only chooses the jumps (e.g., local Pauli XYZ)

Lindbladian Gibbs samplers*

For any n-qubit Hamiltonian $\boldsymbol{H}$, β, and jumps $\{\boldsymbol{A}^a\}$, we define a Lindbladian:
(**Correctness.**) It has approximate Gibbs fixed point

$$\|\boldsymbol{\rho}_{fix}(\mathcal{L}^\dagger) - \boldsymbol{\sigma}_\beta\|_1 \leq \mathcal{O}(\epsilon).$$

(**Efficiency.**) Convergence to the fixed point costs $n + \tilde{\mathcal{O}(1)}$ qubits and Hamiltonian simulation time

$$\tilde{\mathcal{O}}(t_{mix} \cdot \frac{\beta}{\epsilon} t_{mix}),$$

where t_{mix} is the mixing time in trace distance.

Byproducts:

- **(Thermodynamics.)** physically derived Lindbladians
- **(Quantum walk speedup.)** purified Gibbs state

*A recent related result [Rall, Wang, Wocjan '22] handles finite energy resolution differently

Picture: semi-classical Random Walk

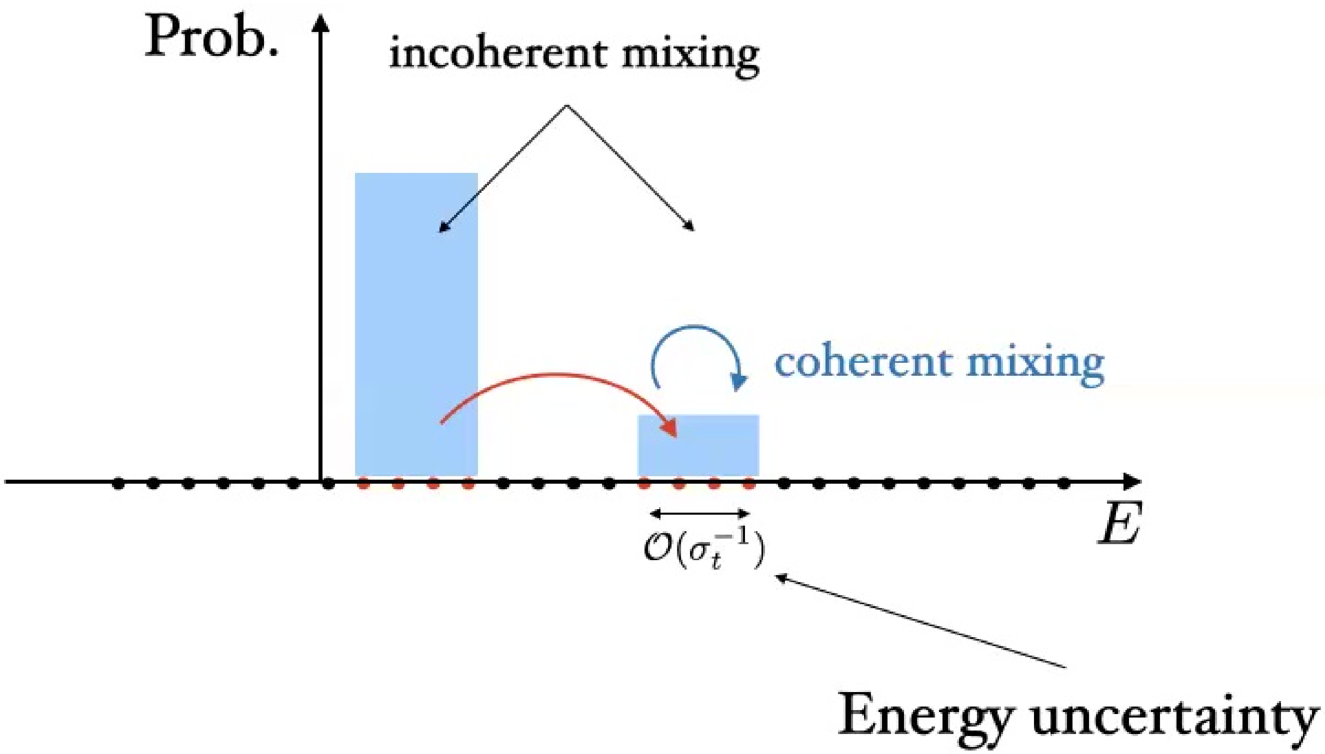

Picture: semi-classical Random Walk

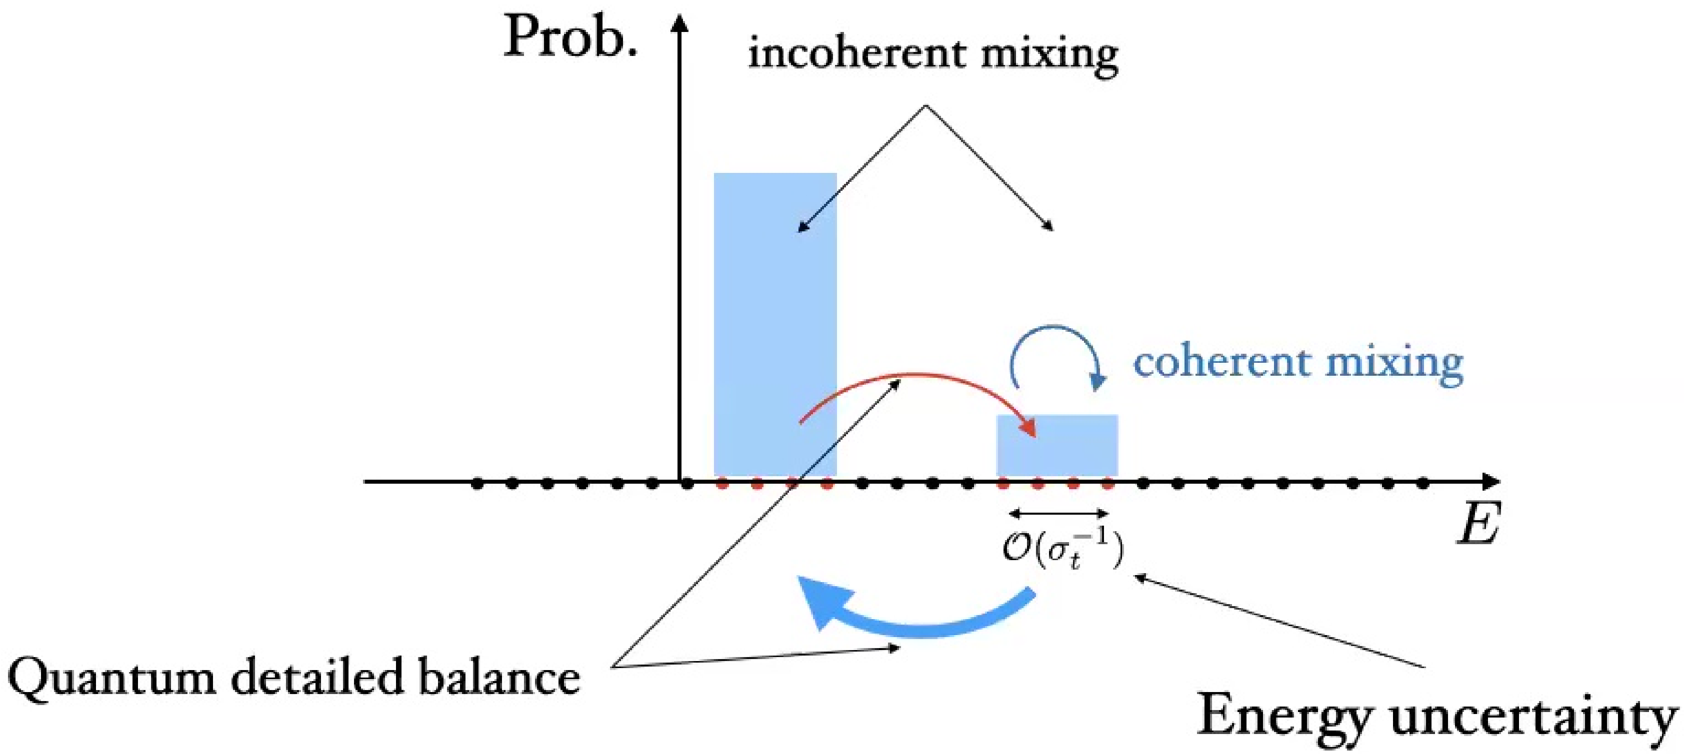

Why is it powerful?

vs. VQE

Suppose $A|0\rangle = \frac{|1\rangle + |-1\rangle}{\sqrt{2}}$, how does the energy gradient look like?

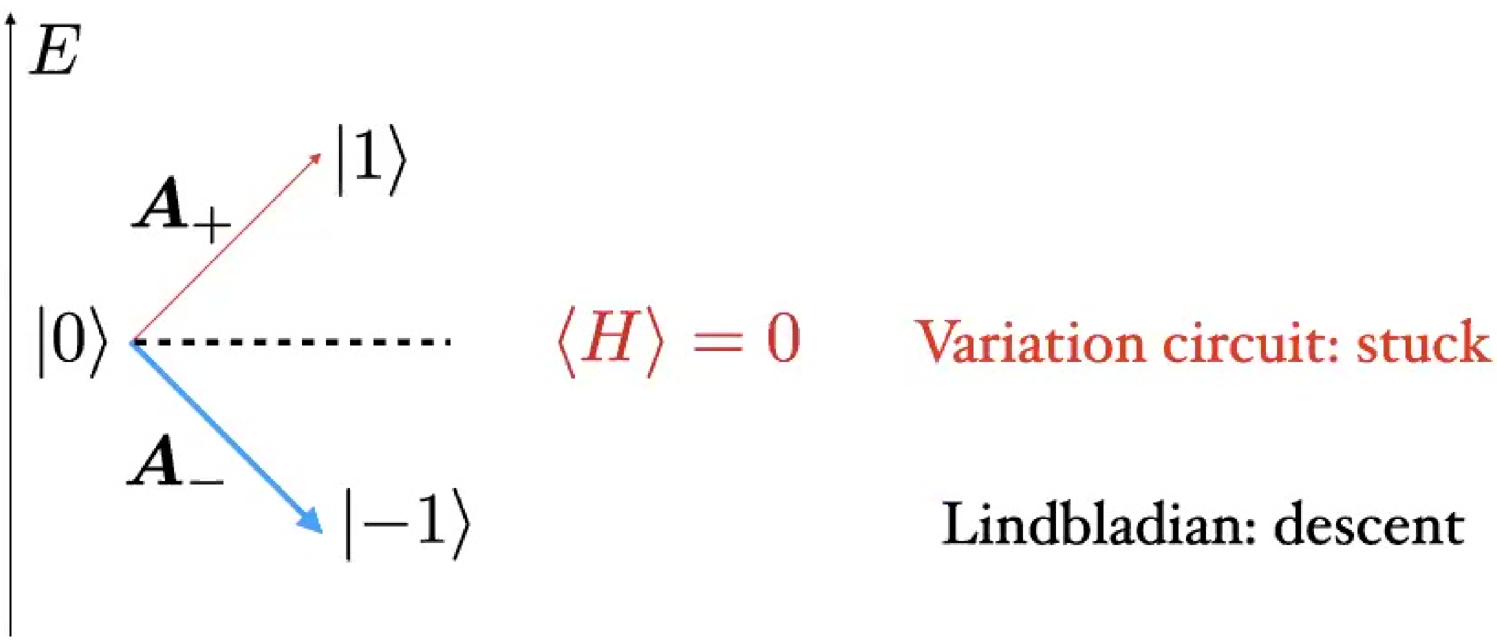

quantum power: phase estimation, "rewinding", "energy filtering",...

VS. imaginary time

$$\boldsymbol{Z} = |1\rangle\langle 1| - |0\rangle\langle 0|$$

$$|+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}$$

$$\boldsymbol{H} = \boldsymbol{Z}_1 + \boldsymbol{Z}_2 + \cdots + \boldsymbol{Z}_n$$

initial state		ground state
$\|+\rangle \|+\rangle \cdots \|+\rangle$	$\longrightarrow$	$\|0\rangle \|0\rangle \cdots \|0\rangle$

Algorithm	imaginary time	Lindbladian
	$e^{-\beta \boldsymbol{H}} \|+\cdots+\rangle \to \|0\cdots 0\rangle$	$e^{\mathcal{L}^\dagger t}[\boldsymbol{\rho}_+] \to \boldsymbol{\rho}_0$
Cost	inverse overlap	mixing time
	$\frac{1}{\|\langle + \| 0\rangle\|^n} = 2^{n/2}$	$t_{mix} \sim n$
feature	"static" (loses prob.)	"dynamic" (Preserves prob.)

Technical ideas.

The secret sauce

"Fourier transform of operator evolution"

$$\boldsymbol{A}^a_{\bar{\omega}} \propto \sum_{\bar{t}} \boldsymbol{A}^a(\bar{t}) \mathrm{e}^{\mathrm{i}\bar{\omega}\bar{t}} f(\bar{t})$$

Naturally, Heisenberg evolution diagnoses energy differences

$$\boldsymbol{A}(t) := \mathrm{e}^{\mathrm{i}\boldsymbol{H}t} \boldsymbol{A} \mathrm{e}^{-\mathrm{i}\boldsymbol{H}t} = \sum_{ij} \mathrm{e}^{\mathrm{i}(E_i - E_j)t} A_{ij} \cdot |\psi_i\rangle \langle\psi_j|$$

A "digital, weighted" version of Fourier integral

$$\boldsymbol{A}(\omega) \propto \int_{-\infty}^{\infty} e^{i\omega t} \boldsymbol{A}(t) dt$$

(More flavors to follow!)

$\gamma(\bar{\omega})/\gamma(-\bar{\omega}) = e^{\beta\bar{\omega}}$

Boosted Fourier transform

Gaussian filter acts as a “boosted phase estimation” with tunable width

$$\boldsymbol{A}^a_{\bar{\omega}} \propto \sum_{\bar{t}} \boldsymbol{A}^a(\bar{t}) \mathrm{e}^{\mathrm{i}\bar{\omega}\bar{t}} f(\bar{t})$$

$$f(t) \propto \exp(-\frac{t^2}{4\sigma_t^2})$$

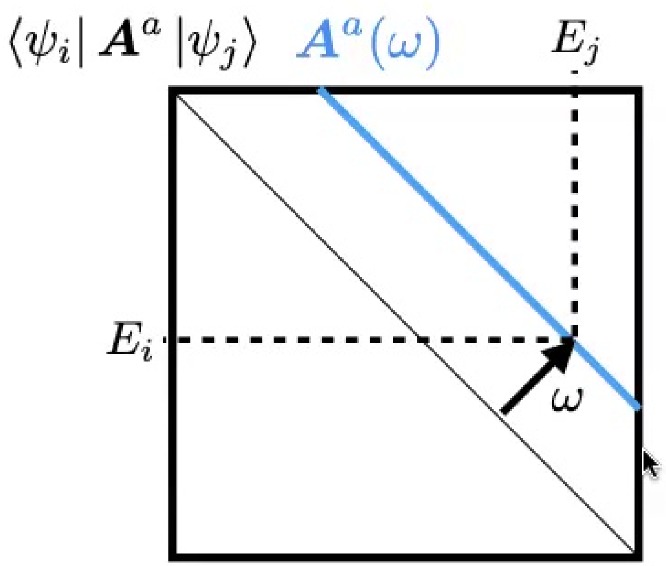

(exact energies)

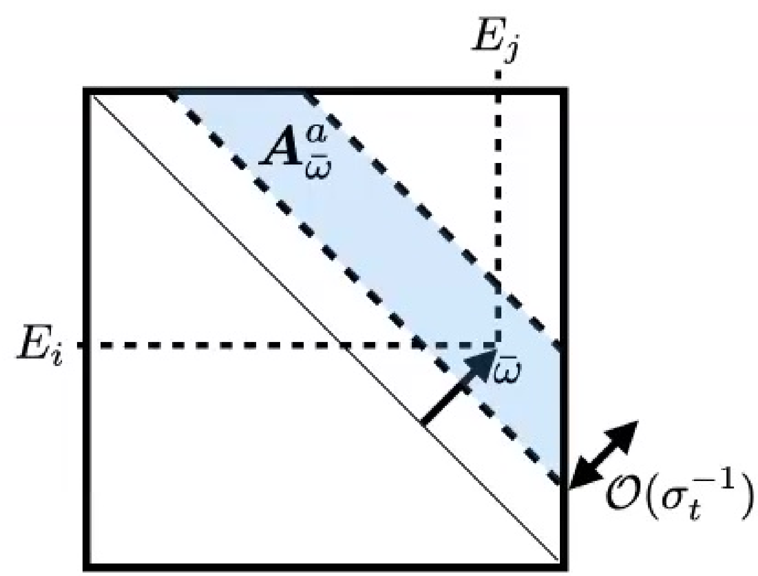

(uncertain energies)

$$\bar{\omega} \pm \mathcal{O}(\sigma_t^{-1})$$

Simple circuit

"Quantum Fourier Transform for operators"

$$\boldsymbol{A} \to \boldsymbol{A}_{\bar{\omega}} \otimes |\bar{\omega}\rangle \quad \text{where} \quad \boldsymbol{A}^a_{\bar{\omega}} \propto \sum_{\bar{t}} \boldsymbol{A}^a(\bar{t}) e^{i\bar{\omega}\bar{t}} f(\bar{t})$$

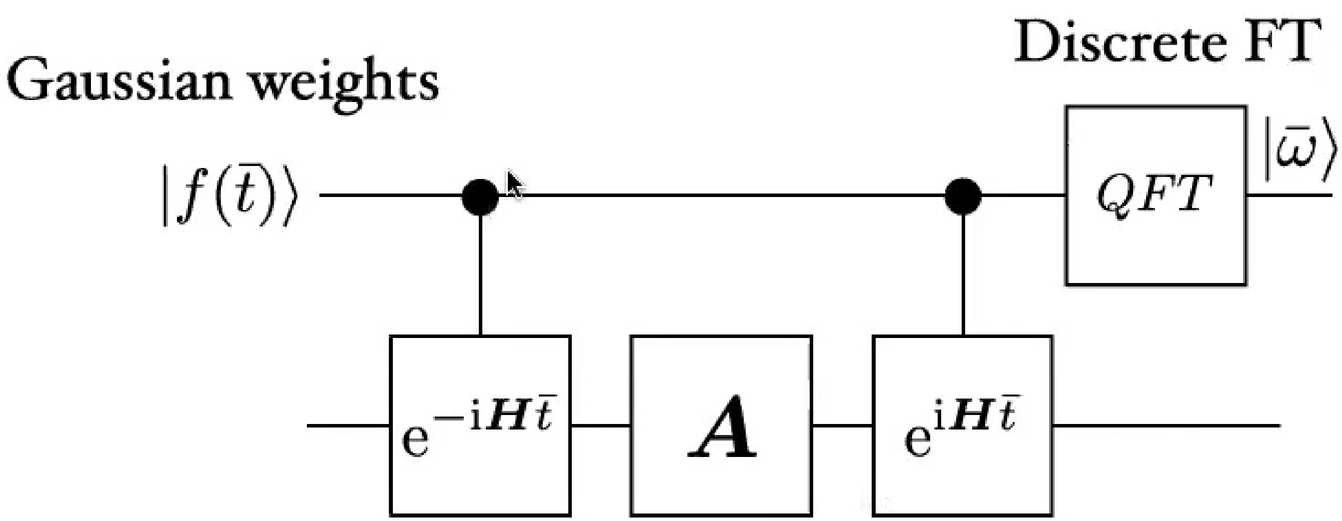

(Feeds into Lindbladian simulation)

Symmetry

(Discrete) Fourier transform satisfies certain "exact symmetry"*

$$A^a_{\bar{\omega}} = (A^a_{-\bar{\omega}})^\dagger$$

Key to quantum approximate detailed balance!

(Like in Metropolis sampling, we start from a reversible/symmetric Markov chain)

* if the filter function is real and the jump is Hermitian

Quantum Approx. detailed balance

$$\left\| \boldsymbol{\sigma}^{1/4}\mathcal{L}[\boldsymbol{\sigma}^{-1/4}\cdot\boldsymbol{\sigma}^{-1/4}]\boldsymbol{\sigma}^{1/4} - \boldsymbol{\sigma}^{-1/4}\mathcal{L}^{\dagger}[\boldsymbol{\sigma}^{1/4}\cdot\boldsymbol{\sigma}^{1/4}]\boldsymbol{\sigma}^{-1/4} \right\|_{2-2} \leq \delta$$

(super-operator norm)

- Analog of classical detailed balance

$$\sigma_i P_{ij} = \sigma_j P_{ji}$$

- δ captures the size of "anti-Hermitian" parts; $\delta = 0$ recovers detailed balance
- approx. detailed balance gives approx. Gibbs fixed point

$$\|\boldsymbol{\rho}_{fix}(\mathcal{L}^{\dagger}) - \boldsymbol{\sigma}\|_1 \lesssim t_{mix}\cdot\delta$$

Finale.

Today: Monte-Carlo style Quantum Algorithms

Task: given a Hamiltonian, prepare the ground state or thermal state

We efficiently implement a Lindbladian whose fixed point is approximately Gibbs.

- (**universal.**) applies to any Hamiltonian (mixing time varies)
- (**no fine-tuning.**) user only chooses the jumps
- (**General theory.**) extends to purified Gibbs state and other Natural Lindbladians.

Followup: exact detailed balance

- Claim [in progress]

For any $\boldsymbol{H}$, β, and a set of jumps $\{\boldsymbol{A}^a\}$, we define a Lindbladian such that (**Correctness.**) The Lindbladian has exact Gibbs fixed point

$$\boldsymbol{\rho}_{fix}(\mathcal{L}^\dagger) = \boldsymbol{\sigma}_\beta$$

(**Efficiency.**) Convergence to the fixed point costs $n + \tilde{\mathcal{O}}(1)$ qubits and Hamiltonian simulation time

$$\tilde{\mathcal{O}}(t_{mix} \cdot \beta),$$

where t_{mix} is the mixing time in trace distance.

- Energy-time uncertainty may not obstruct exact Gibbs sampling!

$$\Delta E \sim \beta^{-1}$$

- Much cleaner theory!
- In contrast, Nature seems erroneous...

Followup: Conditions for mixing

Only known results for non-commuting Hamiltonians

- Quantum chaos (Eigenstate Thermalization Hypothesis) [CB '21, SM '21]

New front: Generalizations of classical results

- Non-commuting Hamiltonians on lattices with decay of Gibbs correlation?
- All-to-all interacting Hamiltonians (e.g., SYK, random Paulis)?
- quantum spin glass: Prove slow mixing?

Future works

- Theory: new algorithms to play with!
  - optimizing Hamiltonians via Gibbs samplers
  - dynamic angle for area law and gapped ground state
- Numerics: new data to take
  - mixing time in lattice models
  - new quantum-inspired classical algorithms?
- Experiment: long way to implementation
  - Analog-friendly adaptations? (Fewer ancillas, hardware restrictions…)