Predicting properties of quantum thermal states from a single trajectory

Jiaqing Jiang
Simons Institute, UC Berkeley

with Jiaqi Leng (Simons Institute) and Lin Lin (UC Berkeley)

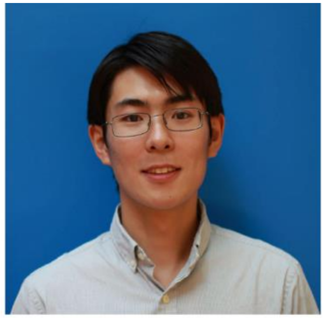

New Frontiers in Quantum Algorithms for Open Quantum Systems,
2026.1.13

Properties of thermal states

Thermal state describes many-body system at finite temperature

$$\rho_{\beta H} := \; e^{-\beta H}/tr(e^{-\beta H})$$

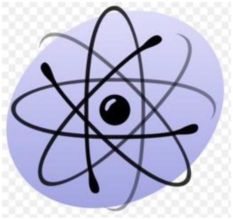

$\boldsymbol{\beta}$ inverse temperature; $\boldsymbol{H}$ local Hamiltonian

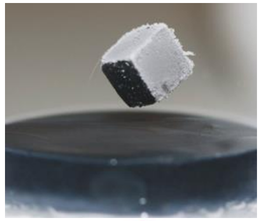

Our task: Properties of thermal state

$\mathrm{tr}(O\rho_{\beta H})$ for observable O

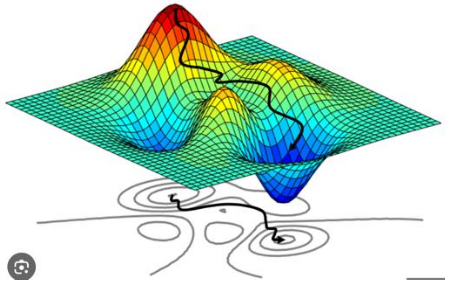

Gradient of Boltzmann machine

Existing methods for estimating $tr(O\rho_{\beta H})$.

- **Classical Markov Chain Monte Carlo** (MCMC)

  (Sign problem; Limited accuracy for general quantum system)

Existing methods for estimating $tr(O\rho_{\beta H})$.

❖ **Quantum Gibbs sampling** (quantum MCMC)

(No sign problem; promising approach for high accuracy)

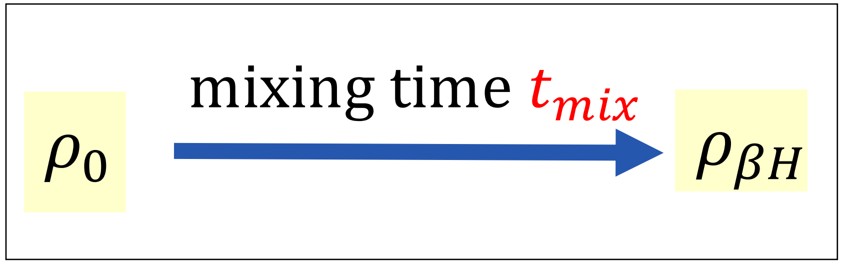

- Davies generator-inspired approach [CKBG23,DLL24,RWW23]
- Quantum Metropolis sampling [JI24,TOV+11]
- System-bath interactions [DZPL25]
- ... (check yesterday's talk!)

Existing methods for estimating $tr(O\rho_{\beta H})$.

❖ **Quantum Gibbs sampling** (quantum MCMC)

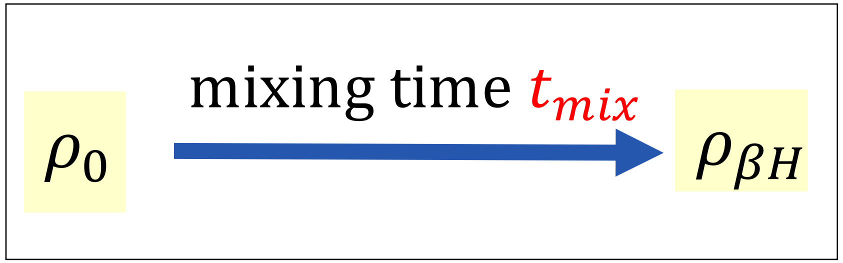

Do better!

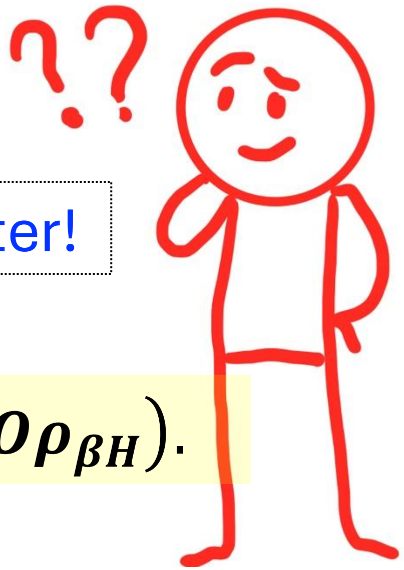

❖ **Our goal:** Use quantum Gibbs sampling to estimate $\boldsymbol{tr(O\rho_{\beta H})}$.

Existing methods for estimating $tr(O\rho_{\beta H})$.

❖ **Quantum Gibbs sampling** (quantum MCMC)

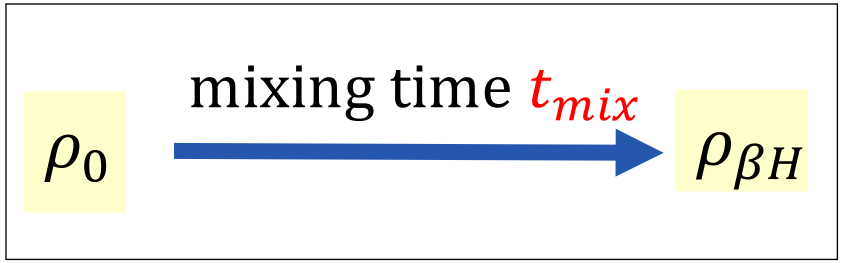

$\boldsymbol{tr(H\rho_{\beta H})}$

Extension discussed later

❖ **Our goal:** Use quantum Gibbs sampling to estimate $\boldsymbol{tr(O\rho_{\beta H})}$.

Outline $tr(\boldsymbol{H}\boldsymbol{\rho}_{\boldsymbol{\beta H}})$

- Multiple-trajectory approach
- Single-trajectory approach (**our approach**)
- Challenge and strategy
- Practical mode (without knowing the spectral gap)
- Extensions

Multiple-trajectory approach

$$tr(H\rho_{\beta H})$$

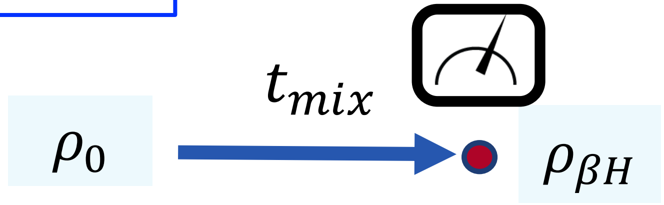

Multiple-trajectory approach

$tr(H\rho_{\beta H})$

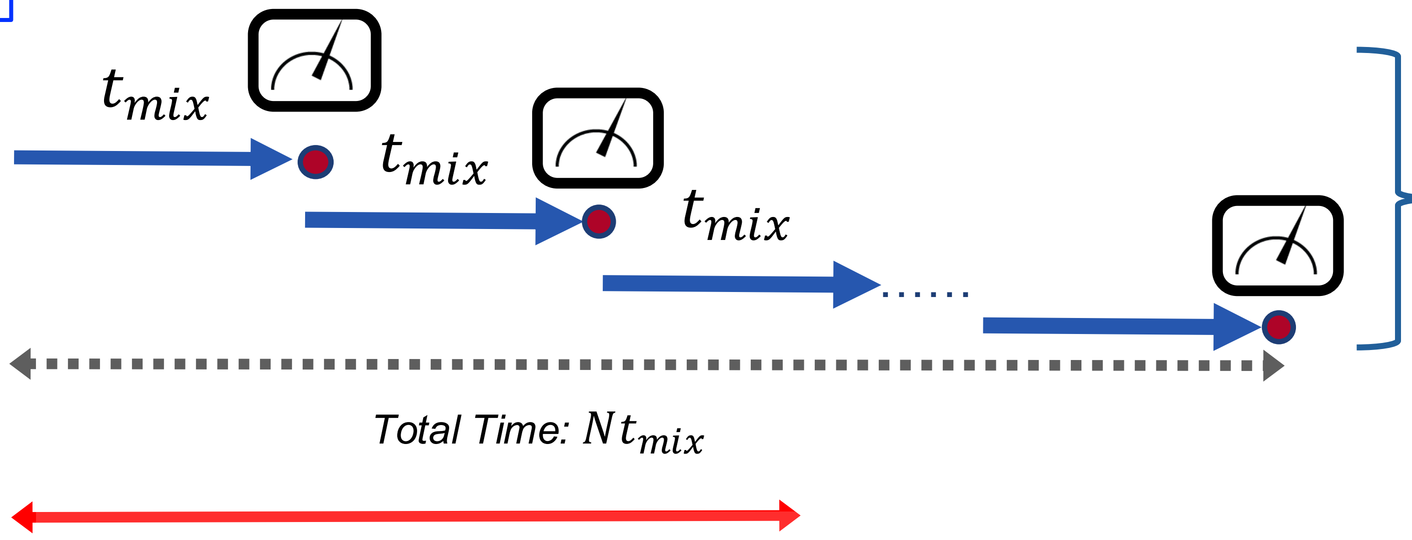

For ϵ precision,
Set $N \sim var_H/\epsilon^2$

Our goal: Reduce the total time

Get (effectively) independent sample with time shorter than t_{mix} !!!

Single-trajectory approach (our method)

$tr(H\rho_{\beta H})$

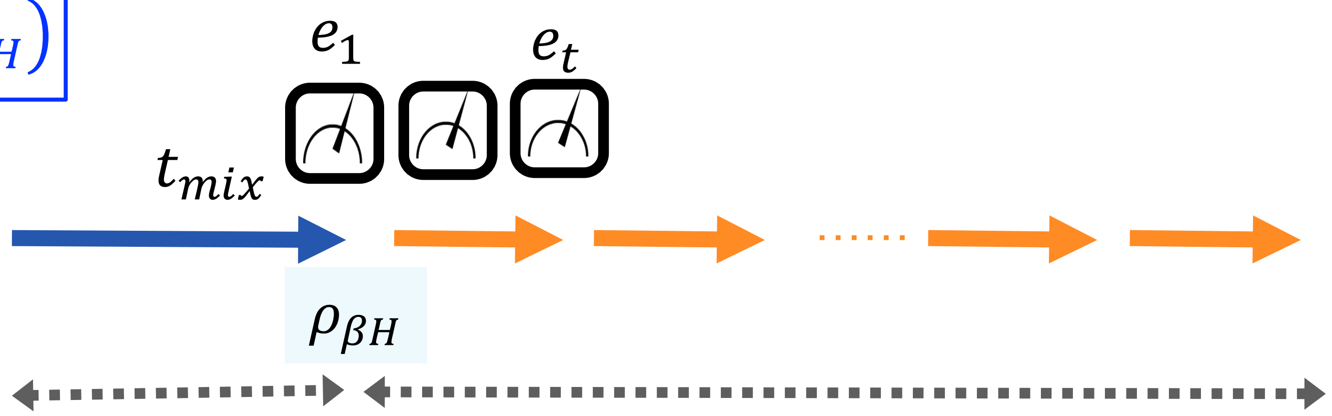

Burn-in Stage

Sampling Stage

- Measure every fixed time, e.g. $e^{1 \cdot L}$, (GQPE)

$tr(H\rho_{\beta \mathrm{H}})$ is estimated by the empirical average

$$\frac{1}{K}(e_1 + \ldots + e_K)$$

How quickly e_t becomes effectively independent from e_1

Autocorrelation time t_{aut}

$tr(H\rho_{\beta H})$

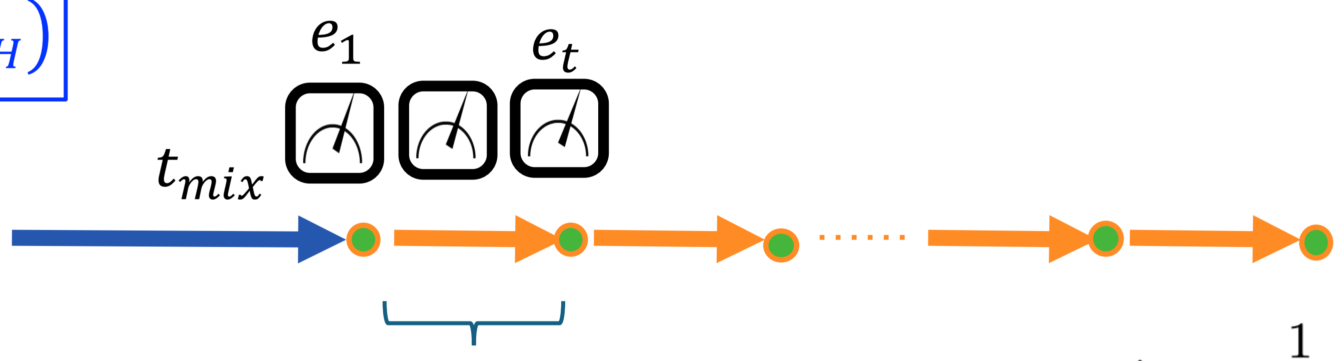

- Effectively independent sample every t_{aut}
- Typically $t_{aut} \ll t_{mix}$, (see next slide)

$$t_{aut} \sim \frac{1}{2} + \sum_{t=1}^{\infty} \frac{Cov(e_1, e_{t+1})}{var_H}$$

connected to the performance of $(e_1 + \ldots + e_K)/K$ by chebyshev inequality

Total time $t_{mix} + N t_{aut}$ which is much shorter than $N t_{mix}$

Why typically $t_{aut} \ll t_{mix}$? (a) Intuitive reason

Timescale for effective independence

Why typically $t_{aut} \ll t_{mix}$? (a) Intuitive reason

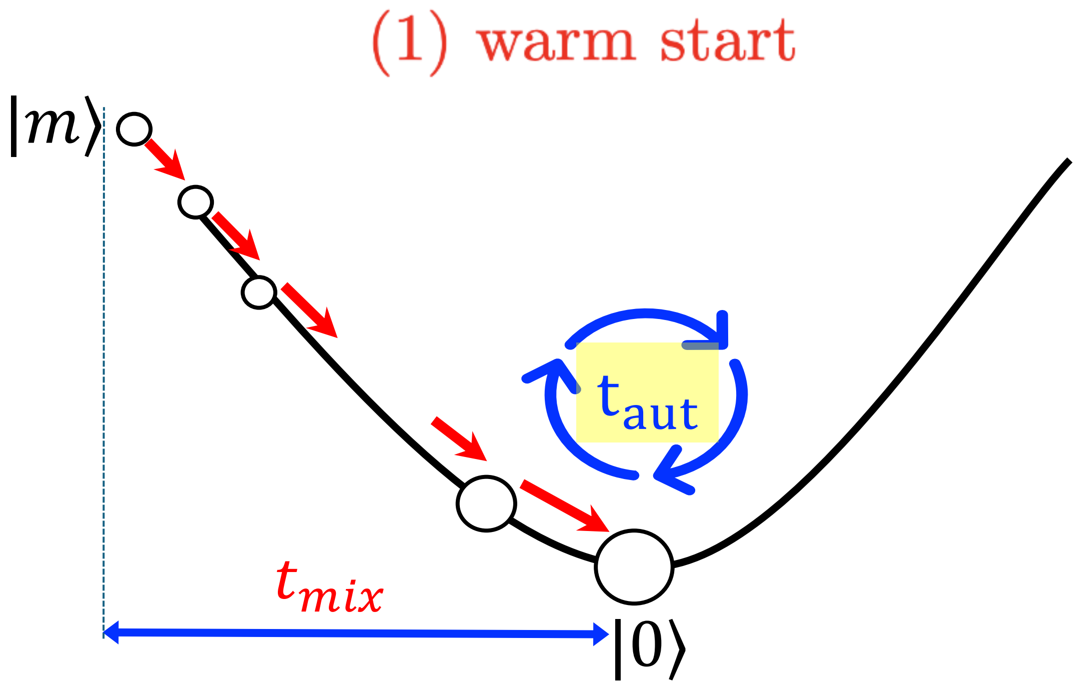

Why typically $t_{aut} \ll t_{mix}$? (a) Intuitive reason

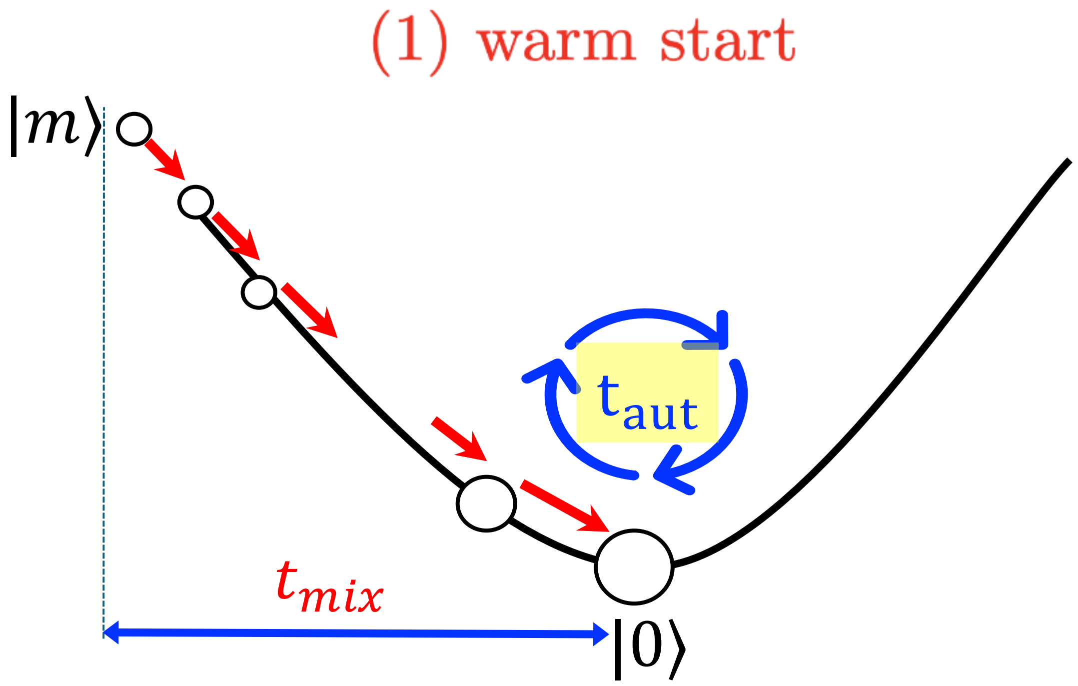

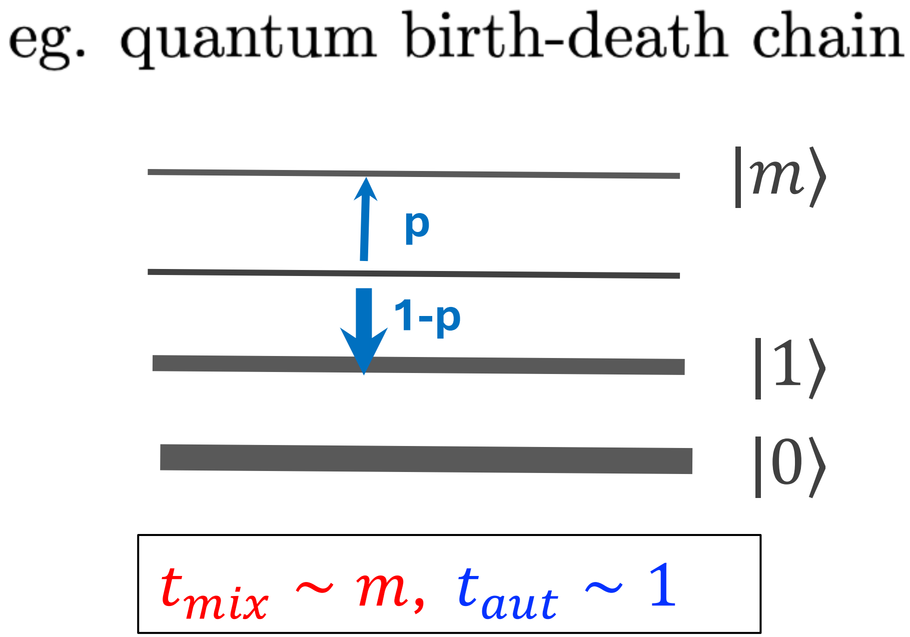

Why typically $t_{aut} \ll t_{mix}$? (a) Intuitive reason

Timescale for effective independence

(1) warm start

(2) t_{aut} is observable-dependent

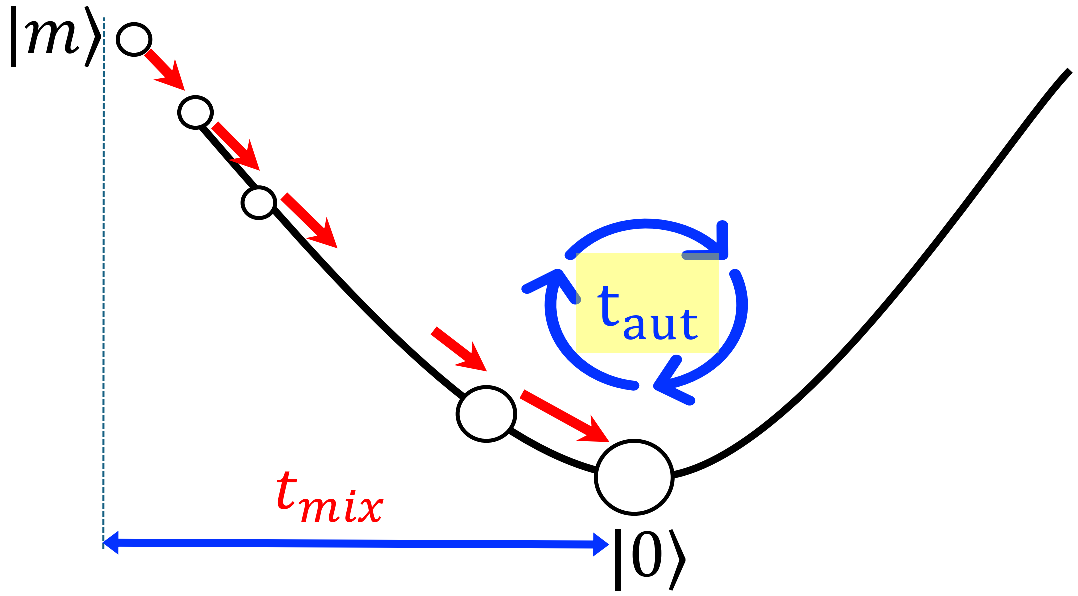

$t_{aut} \ll t_{mix}$ is expected especially if

- The observable is local
- or exhibit certain symmetry

Why typically $t_{aut} \ll t_{mix}$? (a) Intuitive reason

Timescale for effective independence

(1) warm start

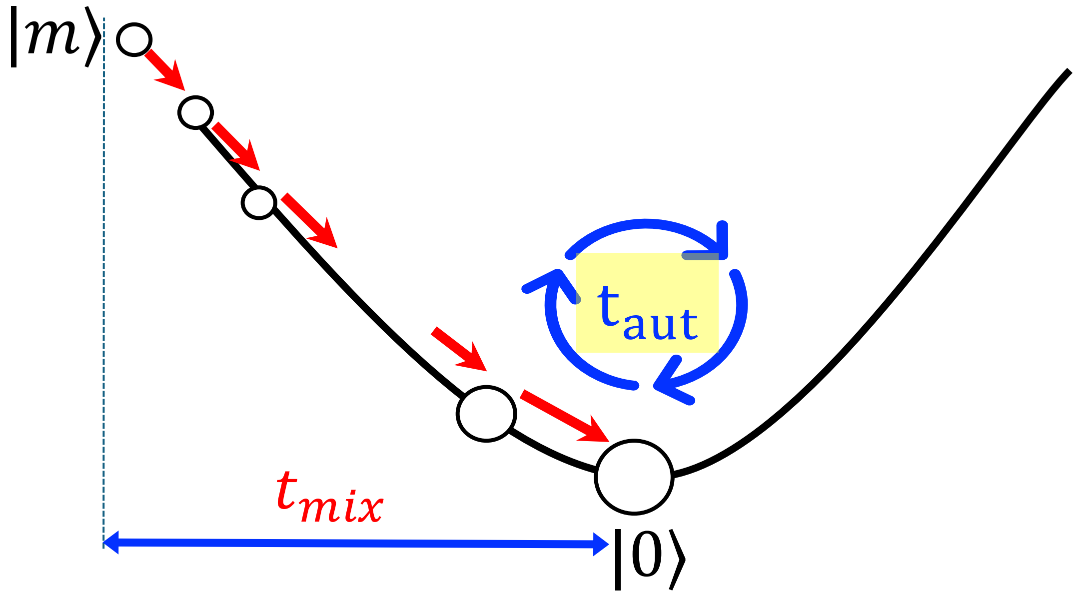

(2) t_{aut} is observable-dependent

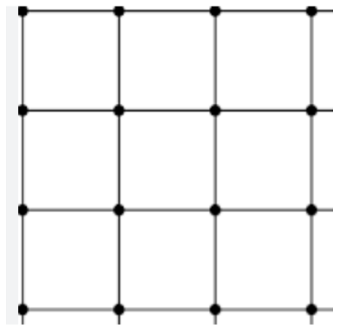

Low−temperature 2D Ising model

Why typically $t_{aut} \ll t_{mix}$? (a) Intuitive reason

(1) warm start

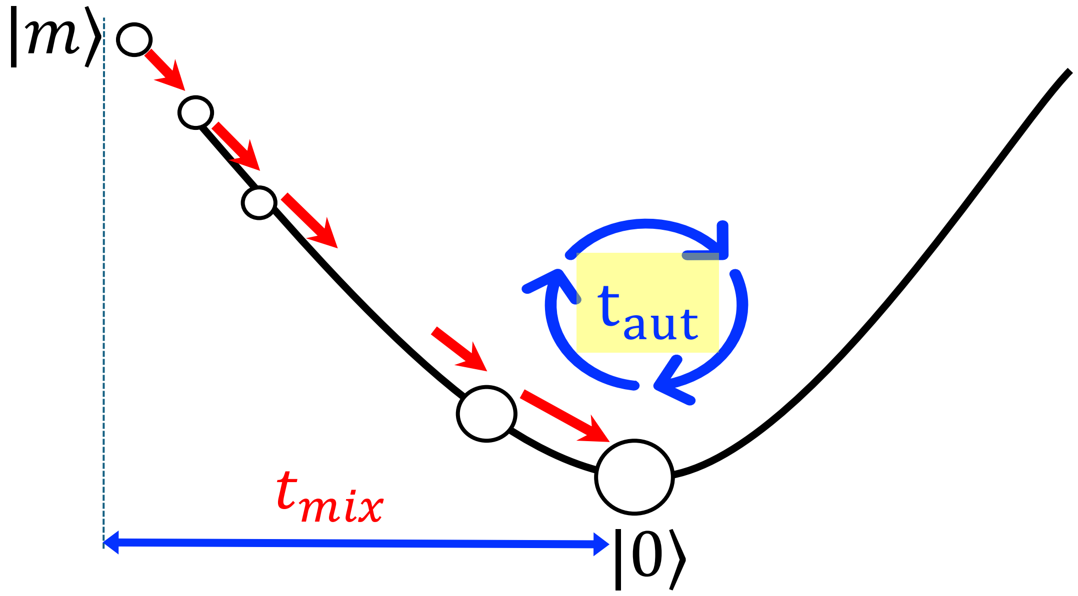

(2) t_{aut} is observable-dependent

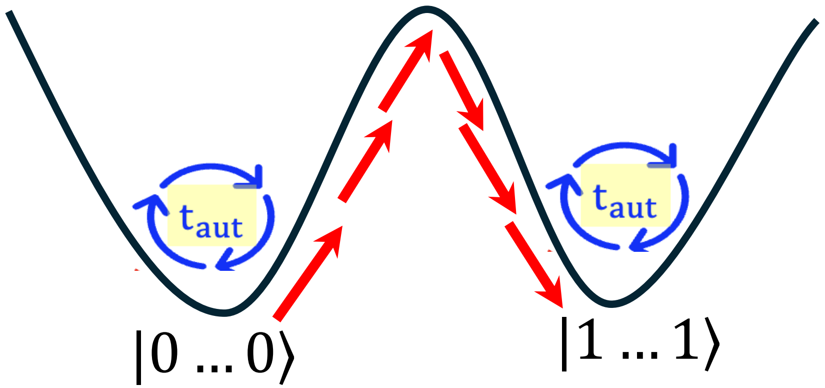

- t_{mix} is exponential
- t_{aut} for energy *might* be polynomial suggested by [GA22]

[GA22] R. Gheissari and A. Sinclair, STOC 2022, "Low-temperature ising dynamics with random initializations,"

Why typically $t_{aut} \ll t_{mix}$? (b) rigorous bound

Our result (general theorem): In quantum Gibbs sampling, for any observable that can be measured by a detailed balanced(DB) channel,

$$t_{aut} \leq 1/gap.$$

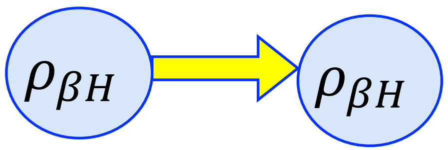

- Includes H; Observables commuting with H;
  (average energy, heat capacity, *partition function…)
  In principle, more general
- For non-DB observable, discussed later

Why typically $t_{aut} \ll t_{mix}$? (b) rigorous bound

Our result (general theorem): In quantum Gibbs sampling, for any observable that can be measured by a detailed balanced(DB) channel,

observable H

$$t_{aut} \leq 1/gap.$$

Why typically $t_{aut} \ll t_{mix}$? (b) rigorous bound

Our result (general theorem): In quantum Gibbs sampling, for any observable that can be measured by a detailed balanced(DB) channel,

observable H

$$t_{aut} \leq 1/gap.$$

$$\frac{1}{gap} \lesssim t_{mix} \lesssim \frac{1}{gap} \log(\sigma_{min}^{-1})$$

As large as the system-size n

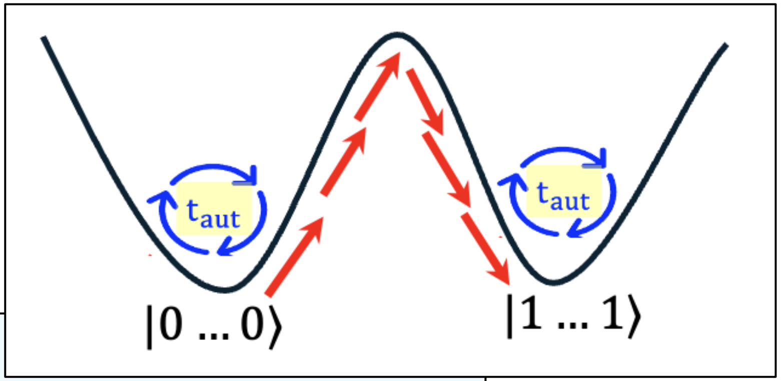

- In theory t_{aut} may differ from t_{mix} by a factor of n;
- In practice the separation can be even larger (e.g. sub-exponential)

Why typically $t_{aut} \ll t_{mix}$? (b) rigorous bound

Our result (general theorem): In quantum Gibbs sampling, for any observable that can be measured by a detailed balanced(DB) channel,

observable H

$$t_{aut} \leq 1/gap.$$

- Our result is **non-trivial** since (**next two slides**)
  1) Measuring H disturbs the Gibbs sampling (controlled)
  2) High cost of QPE (reduced to logarithmic overhead)

[Book, LP17]: D. A. Levin and Y. Peres, Markov chains and mixing times, 2017

Why non-trivial (1):
Measurement disrupts the Gibbs sampling evolution

e_t

classical MCMC: $|x_1\rangle \to |x_2\rangle \dots \to |x_t\rangle \dots\dots$

Quantum MCMC: $|\psi_1\rangle \to \sum_j \alpha_j |\psi_j\rangle \to \dots$

$|\psi_j\rangle$

Our observation:

$(\mathcal{M}\mathcal{N}\mathcal{M})$

The effective quantum channel

New t_{mix} and new t_{aut}?

$\mathcal{M}$ measurement channel; $\mathcal{N}$ Gibbs sampling channel

Why non-trivial (1):
Measurement disrupts the Gibbs sampling evolution

e_t

classical MCMC: $|x_1\rangle \to |x_2\rangle \ldots \to |x_t\rangle \ldots \longrightarrow$

Quantum MCMC: $|\psi_1\rangle \to \sum_j \alpha_j |\psi_j\rangle \to \ldots$

$|\psi_j\rangle$

Our observation:

$$gap(\mathcal{M}\mathcal{N}\mathcal{M}) \geq gap(\mathcal{N})$$

The effective quantum channel

Proof: compare the **Dirichlet form**; use the **contractive property** of DB channel.

$\mathcal{M}$ measurement channel; $\mathcal{N}$ Gibbs sampling channel

Why non-trivial (2): High cost of QPE

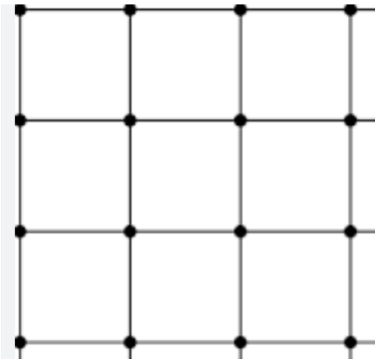

Classical MCMC

- gate cost $\sim n$, depth cost ~ 1

Quantum MCMC

high precision QPE

- gate cost $\sim n\ \epsilon^{-1}$, depth cost $\sim \epsilon^{-1}$

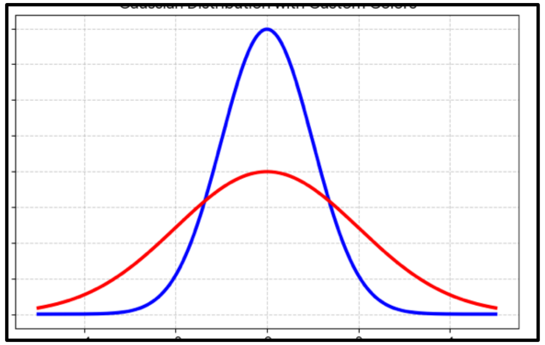

Our solution: unbiased measurement; logarithmic overhead

- **Gaussian filtered QPE** with ~ 1 variance [M19]
- Gate cost $\sim n$ polylog ϵ^{-1}, depth cost $\sim$ polylog ϵ^{-1}
- Logarithmic ancilla qubits

[M19] J. E. Moussa, "Low-depth quantum metropolis algorithm," arXiv:1903.01451,2019

Outline $tr(\boldsymbol{H}\boldsymbol{\rho}_{\boldsymbol{\beta H}})$

- Single-trajectory approach (our approach)
- Challenge and strategy
- Practical setting (without knowing the spectral gap)
- Extensions

Practical setting: no prior knowledge of t_{mix}, t_{aut}

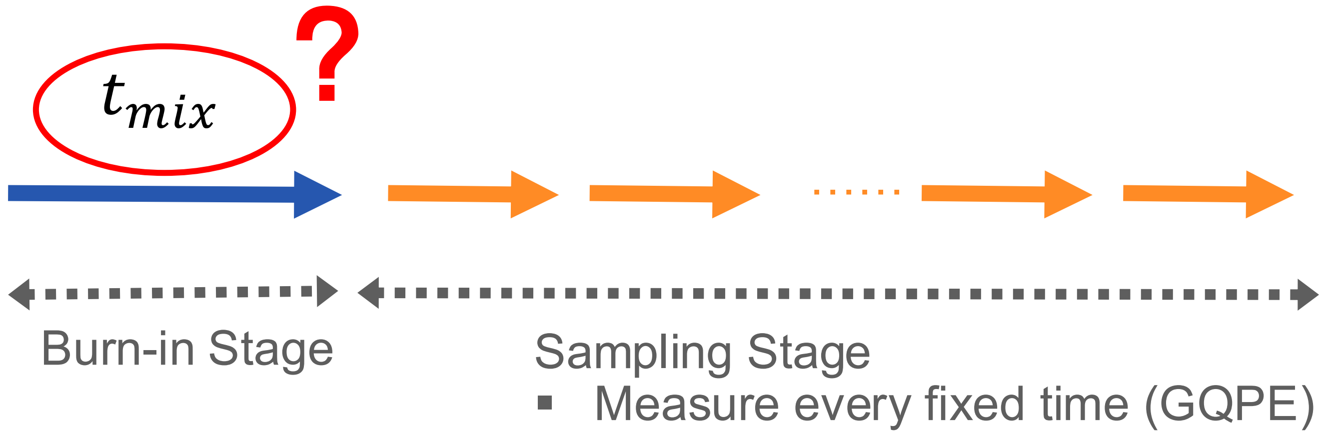

Practical setting: no prior knowledge of t_{mix}, t_{aut}

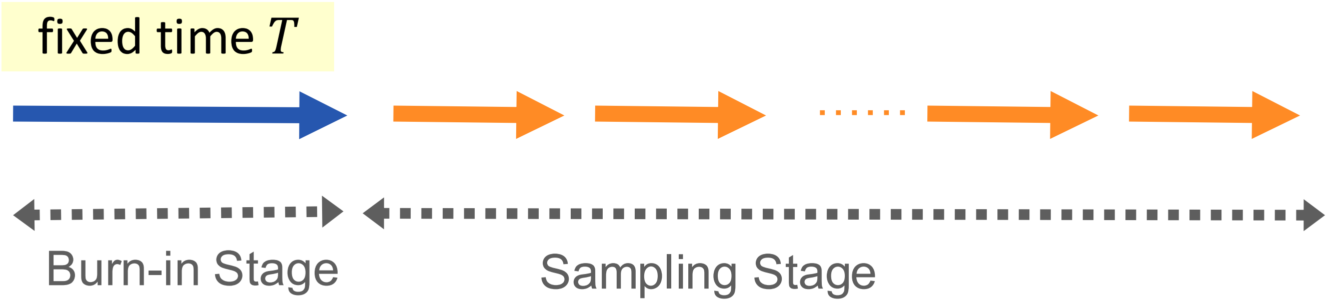

- Measure every fixed time (GQPE)

$$\frac{1}{K}(e_1 + \ldots + e_K) \rightarrow tr(H\rho_{\beta H})$$

slower convergence rate

- Our method can be used as **empirical way** to **verify the convergence** of quantum Gibbs sampling
- (quantum analogy of **Gelman–Rubin diagnostic** in MCMC)

Extension to detailed balanced measurement

Our result can be generalized to any DB observable O.

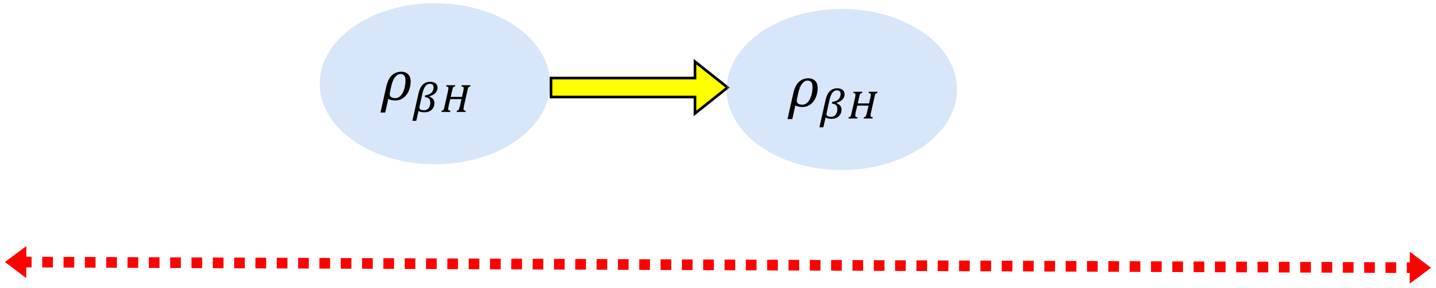

Total time: $t_{mix} + Nt_{aut}$ and $t_{aut} \leq 1/gap$

Potential extension for general observables

Goal: for any observable O, design a measurement s.t.

$\rho_{\beta H} \Rightarrow \rho_{\beta H}$

(1) Fixes Gibbs state (and satisfy DB)
(2) Recover $tr(O\rho_{\beta H})$

Weighted operator Fourier Transform (WOFT) [CKG23]

$O \Rightarrow \widehat{O}(\tau)$

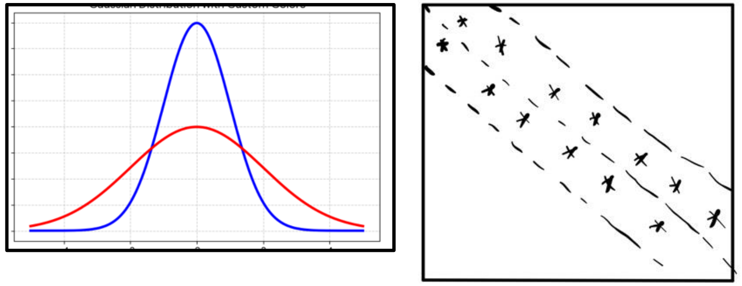

(1) $\left[\widehat{O}(\tau), H\right] \to 0, \quad \text{as } \tau \to 0$

(2) $tr\left(\widehat{O}(\tau)\,\rho_{\beta H}\right) = tr\left(O\,\rho_{\beta H}\right)$

$$\widehat{O}(\tau) := \int_{-\infty}^{+\infty} e^{iHt} O e^{-iHt} f(t) dt$$

[CKG23] C.F.Chen, M. J. Kastoryano, and A. Gilyen, "An efficient and exact noncommutative quantum gibbs sampler,"

Potential extension for general observables

Goal: for any observable O, design a measurement s.t.

$\rho_{\beta H} \Rightarrow \rho_{\beta H}$

(1) Fixes Gibbs state (and satisfy DB)
(2) Recover $tr(O\rho_{\beta H})$

Weighted operator Fourier Transform (WOFT) [CKG23]

$O \Rightarrow \hat{O}(\tau)$

(1) $\left[\hat{O}(\tau), H\right] \to 0, \quad \text{as } \tau \to 0$

(2) $tr\left(\hat{O}(\tau)\, \rho_{\beta H}\right) = tr\left(O\, \rho_{\beta H}\right)$

- Measurement cost for $\hat{O}(\tau)$ is $\sim \boldsymbol{\tau}^{-1}$
- $\tau^{-1} \sim \log \epsilon^{-1}$ if commutator decays exponentially, e.g. gapped system

Summary and open question

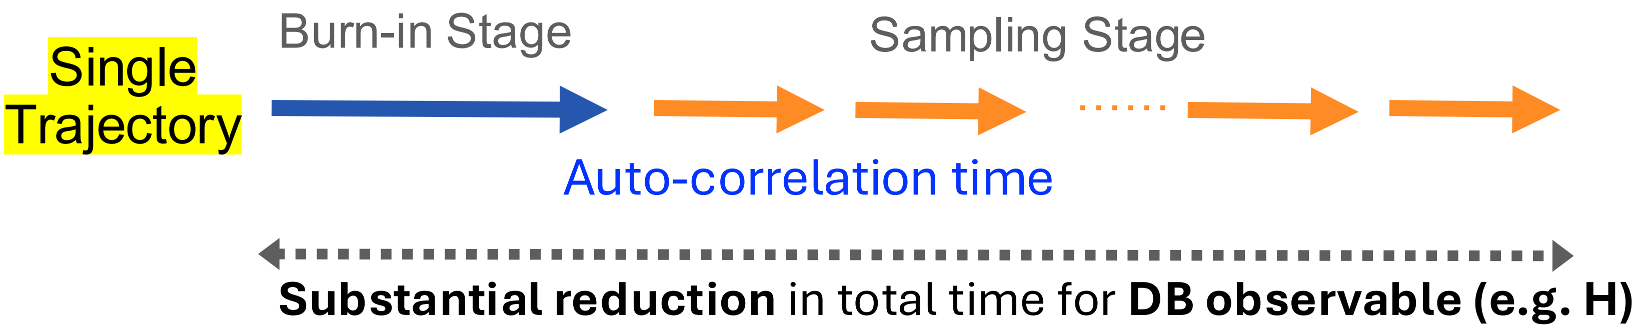

Open Problem: Measure general observable in a DB way?

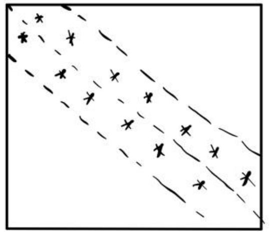

1) WOFT, more numerical experiments?

2) Techniques from quantum Gibbs sampling

- [CKBG23,DLL24,RWW23]
- [JI24, DZPL25, TOV+11]

3) ϵ^{-1}-lower bound for the overhead?

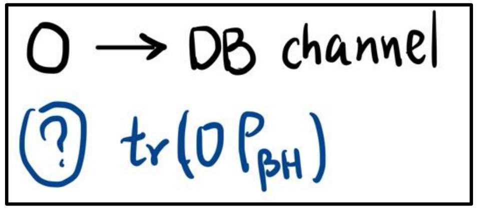

Thanks for listening. Questions ?

Appendix Q&A

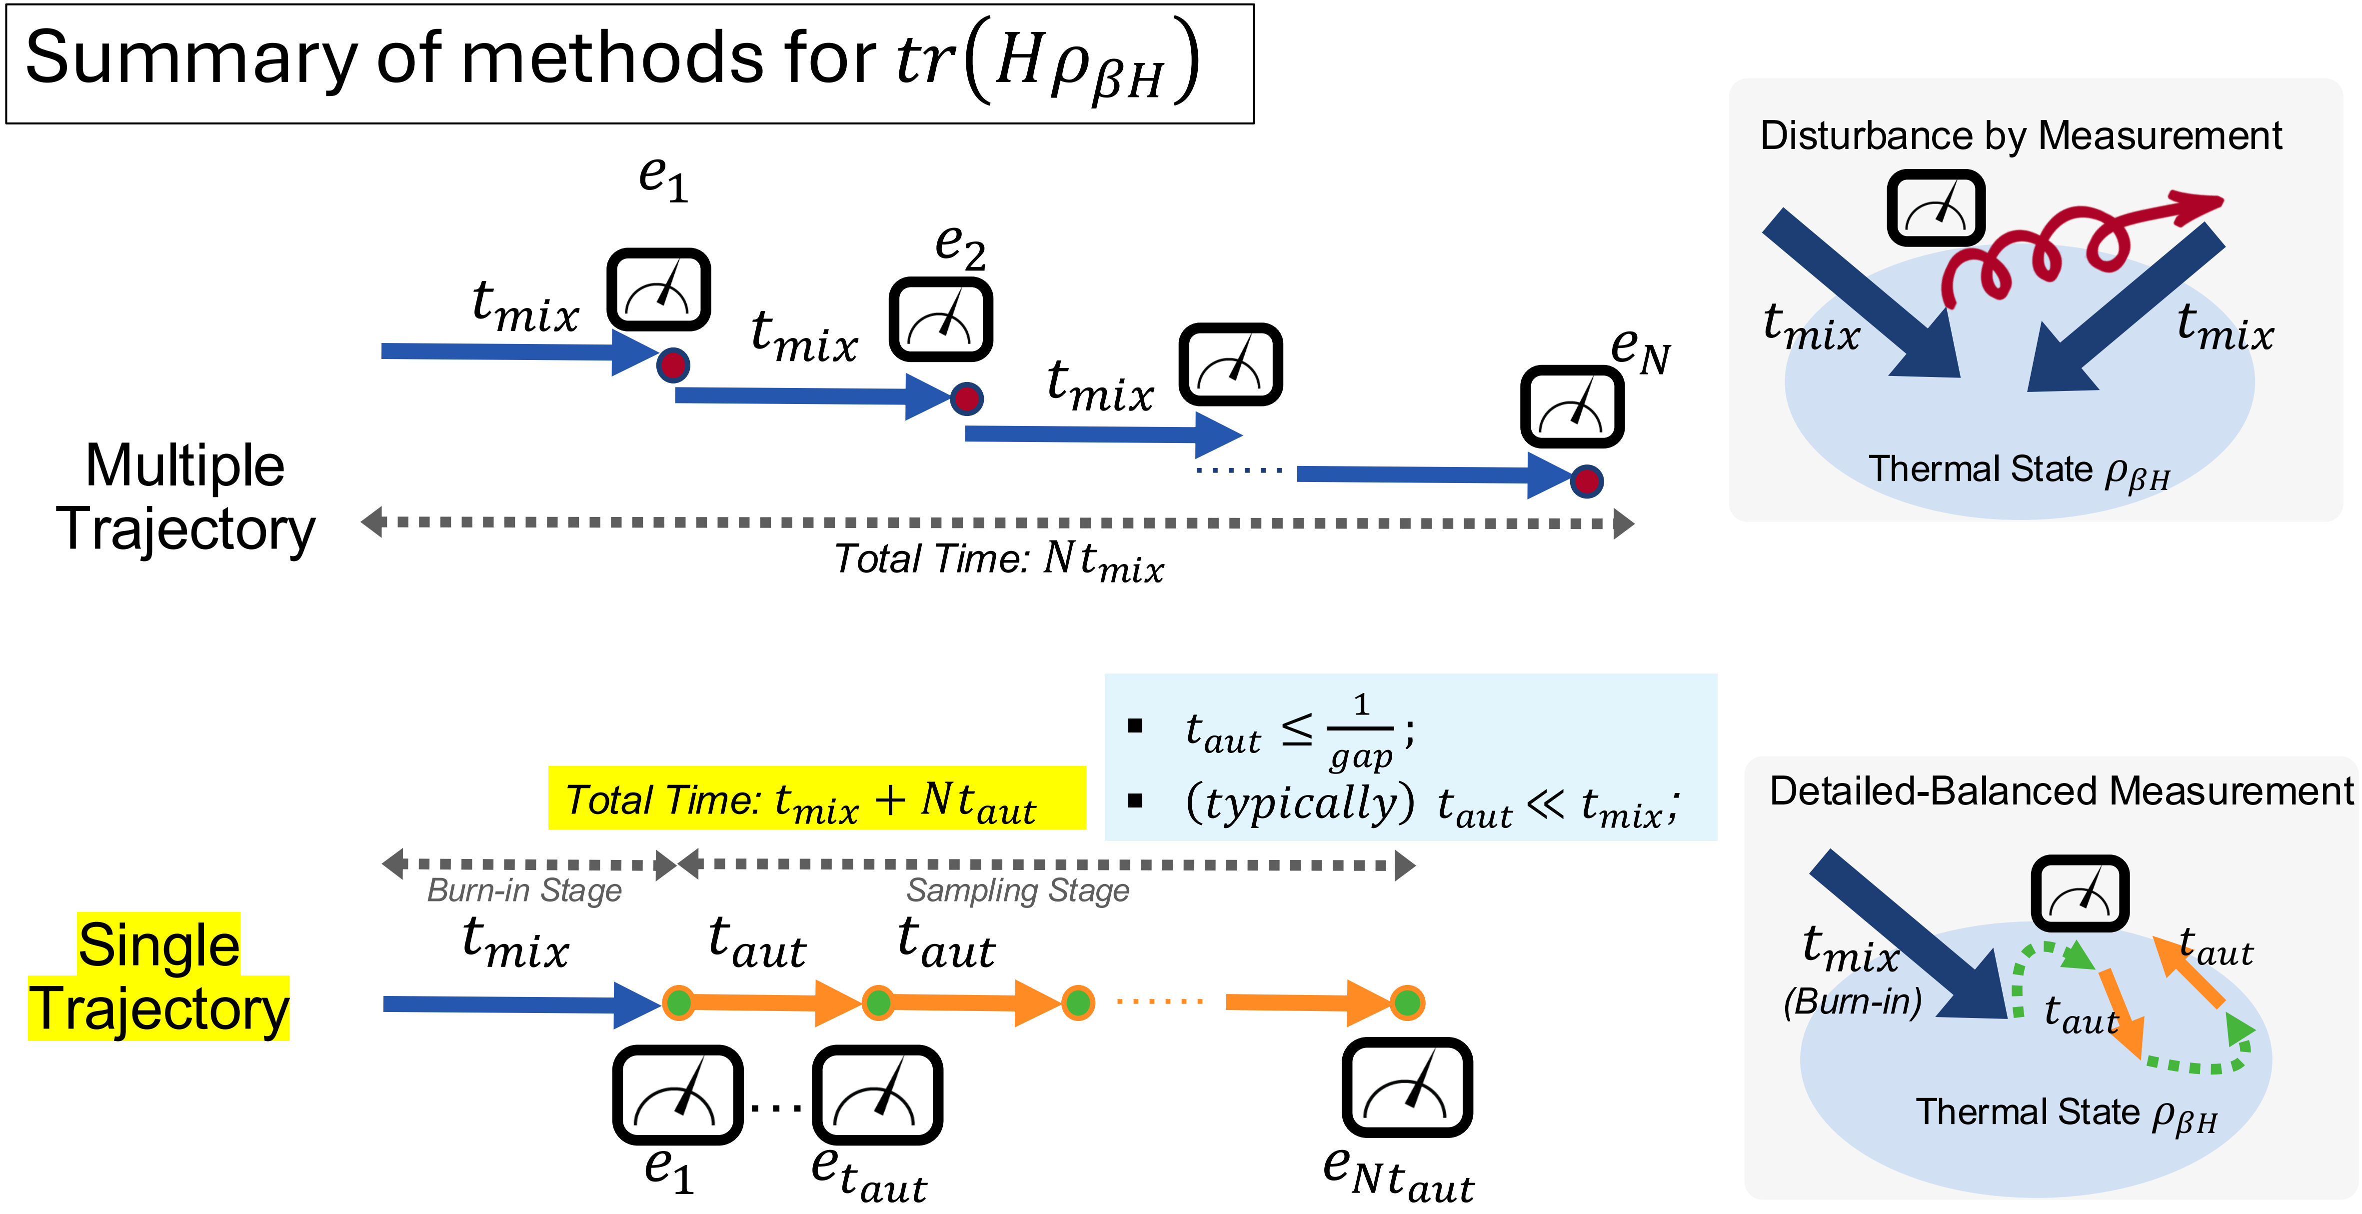

Summary of methods for $tr(H\rho_{\beta H})$
e_1
e_2
e_N
t_{mix}
t_{mix}
t_{mix}
Multiple Trajectory
Total Time: Nt_{mix}
Disturbance by Measurement
t_{mix}
t_{mix}
Thermal State $\rho_{\beta H}$
$t_{aut} \leq \frac{1}{gap}$;
$(typically)$ $t_{aut} \ll t_{mix}$;
Total Time: $t_{mix} + Nt_{aut}$
Burn-in Stage
Sampling Stage
Single Trajectory
t_{mix}
t_{aut}
t_{aut}
e_1
$e_{t_{aut}}$
$e_{Nt_{aut}}$
Detailed-Balanced Measurement
t_{mix}
(Burn-in)
t_{aut}
t_{aut}
Thermal State $\rho_{\beta H}$

How does this work compared to “recover map”

Incomparable, no correlation analysis

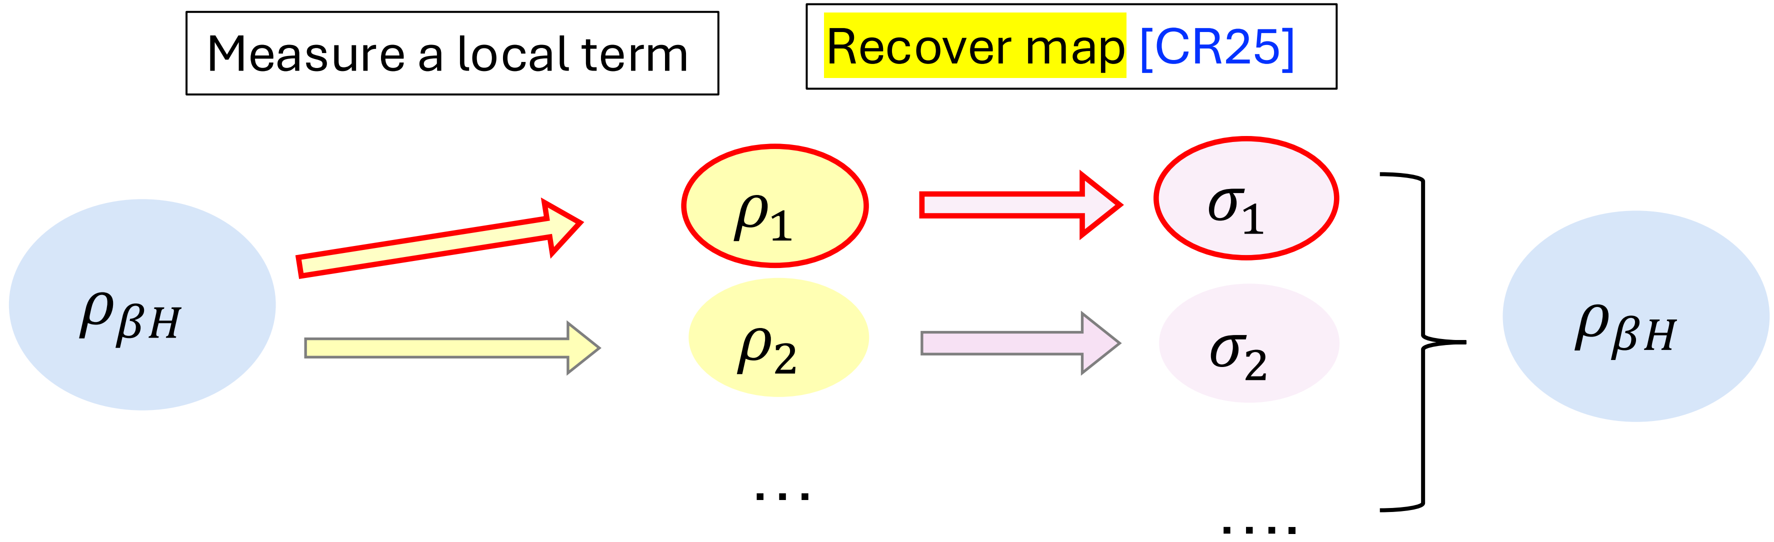

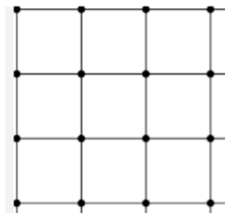

[CR25] Chen, Chi-Fang, and Cambyse Rouzé. "Quantum Gibbs states are locally Markovian."

Purifying the Gibbs state? Pros & Cons

Assumptions

$$gap(\mathcal{L}) = \Delta, \qquad e^{t\mathcal{L}}(\rho_0) \to \rho_{\beta H}$$

$$\Updownarrow$$

$$gap(H_{\mathcal{L}}) = \Delta, \qquad H|\rho_{\beta H}\rangle = |\rho_{\beta H}\rangle$$

} prepare $|\rho_{\beta H}\rangle$

Cons Assume gapped path from I to $H_{\mathcal{L}}$ **gap** Δ^*

For $tr(H\rho_{\beta H})$, need to know Δ

Pros The cost of measurement $\sim \Delta\epsilon^{-1}$ instead of $var_H \epsilon^{-2}$