Phase Estimation and Gibbs Sampling

András Gilyén

Alfréd Rényi Institute of Mathematics
Budapest, Hungary

Summer school on QA and QEC, CWI, Amsterdam
2026 August 24-28

Unbiased Quantum Phase Estimation & Connecting Discrete and Continuous Fourier Transforms

Vanila Quantum Phase Estimation

Given a (black-box) unitary U and one of its eigenvectors $|\psi\rangle$ with unknown eigenvalue $e^{2\pi i\varphi}$ we would like to learn the phase $\varphi \in [0,1)$ by implementing a map $|\psi\rangle|0\rangle \to |\psi\rangle|\varphi\rangle$.

Phase estimation circuit when $\varphi = 0.\varphi_1\varphi_2\ldots\varphi_n$ has (at most) n-bits

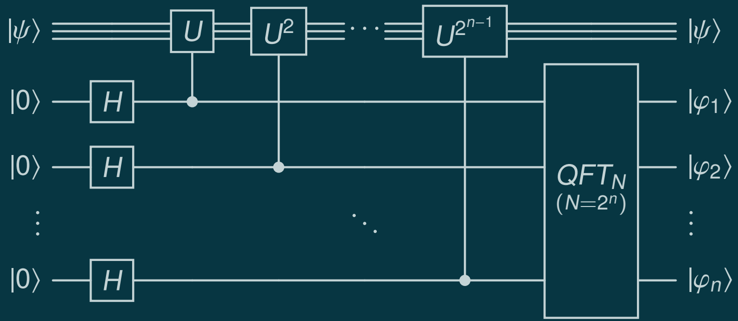

$$|\psi\rangle|0\rangle^{\otimes n} \xrightarrow{H^{\otimes n}} \frac{1}{\sqrt{N}}\sum_{t=0}^{N-1}|\psi\rangle|t\rangle \xrightarrow{\sum_{t=0}^{N-1} U^t \otimes |t\rangle\langle t|} \frac{1}{\sqrt{N}}\sum_{t=0}^{N-1} U^t|\psi\rangle|t\rangle = |\psi\rangle \underbrace{\frac{1}{\sqrt{N}}\sum_{t=0}^{N-1} e^{2\pi i\varphi t}|t\rangle}_{QFT_N^{-1}|N\cdot\varphi\rangle}$$

Quantum Phase Estimation – arbitrary phases

Computing the amplitudes for general φ

$$\frac{1}{\sqrt{N}}\sum_{t=0}^{N-1} e^{2\pi i\varphi t}|t\rangle \quad \overset{QFT_N}{\rightarrow} \quad \frac{1}{N}\sum_{t=0}^{N-1}\sum_{k=0}^{N-1} e^{2\pi i\varphi t} e^{-2\pi i k t/N}|k\rangle$$

$$= \frac{1}{N}\sum_{k=0}^{N-1}\sum_{t=0}^{N-1} e^{2\pi i(\varphi - 0.k_1 k_2 \dots k_n)t}|k\rangle$$

$$= \frac{1}{N}\sum_{k=0}^{N-1} \frac{e^{2\pi i(\varphi - 0.k_1 k_2 \dots k_n)N} - 1}{e^{2\pi i(\varphi - 0.k_1 k_2 \dots k_n)} - 1}|k\rangle \qquad \text{(by geometric summation)}$$

The output distribution in terms of $\Delta := \varphi - 0.k_1 k_2 \dots k_n$

$$\left|\frac{1}{N}\frac{e^{2\pi i(\varphi - 0.k_1 k_2 \dots k_n)N} - 1}{e^{2\pi i(\varphi - 0.k_1 k_2 \dots k_n)} - 1}\right|^2 = \left|\frac{1}{N}\frac{e^{\pi N i\Delta} - e^{-\pi N i\Delta}}{e^{\pi i\Delta} - e^{-\pi i\Delta}}\right|^2 \qquad \left(\text{multiply by } \frac{e^{-\pi N i\Delta}}{e^{-\pi i\Delta}} \text{ under } |\cdot|\right)$$

$$= \left|\frac{1}{N}\frac{\sin(\pi N\Delta)}{\sin(\pi\Delta)}\right|^2 \qquad (e^{ix} - e^{-ix} = 2\sin(x))$$

$$= \left|\frac{\mathrm{sinc}(\pi N\Delta)}{\mathrm{sinc}(\pi\Delta)}\right|^2 \qquad (\mathrm{sinc}(x) = \sin(x)/x)$$

Heavy tail

Although, we get the best two estimates with high probability, the distribution has a heavy tail:

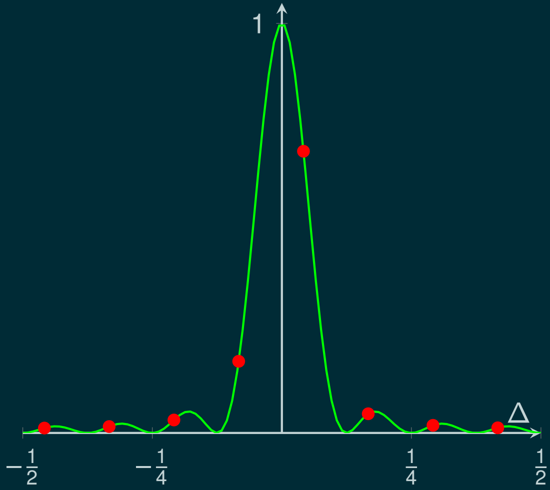

Figure: Plot of $\dfrac{\operatorname{sinc}^2(\pi N \Delta)}{\operatorname{sinc}^2(\pi \Delta)}$ for $N = 8$ and true phase $1/24$.

Quantum Phase Estimation – error probabilities

The probability of obtaining estimate off by Δ

We just computed it for $\Delta \in [-\frac{1}{2}, \frac{1}{2})$:

$$\frac{\operatorname{sinc}^2(\pi N \Delta)}{\operatorname{sinc}^2(\pi \Delta)}$$

- The probability of obtaining the best n-bit estimate is when $|\Delta \mod 1|$ is the smallest. The worst case is when $\Delta_{\min} = \frac{1}{2N}$:

$$\frac{\operatorname{sinc}^2(\pi N \Delta_{\min})}{\operatorname{sinc}^2(\pi \Delta_{\min})} \geq \frac{\operatorname{sinc}^2(\pi N \frac{1}{2N})}{\operatorname{sinc}^2(\pi \frac{1}{2N})} \geq \operatorname{sinc}^2(\pi/2) = (\frac{1}{\pi/2})^2 = \frac{4}{\pi^2} > 40\%$$

- The probability of obtaining one of the two best n-bit estimates corresponds to $\Delta_{\min}$, $\frac{1}{N} - \Delta_{\min}$. The worst case is once again when $\Delta_{\min} = \frac{1}{2N}$:

$$\frac{\operatorname{sinc}^2(\pi N \Delta_{\min})}{\operatorname{sinc}^2(\pi \Delta_{\min})} + \frac{\operatorname{sinc}^2(\pi N (\frac{1}{N} - \Delta_{\min}))}{\operatorname{sinc}^2(\pi (\frac{1}{N} - \Delta_{\min}))} \geq 2\frac{\operatorname{sinc}^2(\pi N \frac{1}{2N})}{\operatorname{sinc}^2(\pi \frac{1}{2N})} \geq 2(\frac{1}{\pi/2})^2 = \frac{8}{\pi^2} > 80\%$$

Boosting

The median trick

Suppose our estimator outputs an ε-precise estimate with probability at least 80%.

- Take s independent estimates, and compute their median.
- The expected number of estimates within ε-precision is at least 80%.
- It is exponentially unlikely in s that at least 50% of estimates are farther than ε. (See the Chernoff bound.)
- When more than 50% of estimates are ε-precise their median is also ε-precise!

Median on the cycle?

- Output the most frequently seen element (in case of a tie, choose one randomly).
- It is exponentially unlikely that the most frequently seen estimate is not one of the two best n-bit estimates (as they have jointly probability $\geq$ 80%).
- The output distribution is exponentially concentrated on the two best estimates!

Unfortunately, we cannot ensure that that we get a unique estimate with high probability!

Unbiased (symmetric) estimator

The random shift trick

Input: $\left|\psi(\phi)\right\rangle = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} e^{i\phi k}|k\rangle$ (for unknown ϕ), and a parameter $n \in \mathbb{N}$

1: Sample a uniformly random n-digit binary number $u \in [0,1)$ and define $\xi := \frac{2\pi u}{N}$
2: Apply multi-phase gate $\sum_{k=0}^{N-1} e^{-i\xi k}|k\rangle\langle k|$ to $\left|\psi(\phi)\right\rangle$
3: Perform inverse Fourier transform over $\mathbb{Z}_N$ and measure the state, yielding outcome j
4: **Return** $\varphi := \frac{2\pi j}{N} + \xi = \frac{2\pi}{N}(j+u)$

Median on the cycle?

Theorem (Unbiased Phase Estimation – Apeldoorn, Cornelissen, G, Nannicini (2022))

If we run the above Algorithm with $n = \infty$ in Line 1, then it returns a random phase $\varphi \in [0, 2\pi)$ with probability density function

$$f(\varphi) := \frac{N}{2\pi} \frac{\mathrm{sinc}^2\left(\frac{N}{2}|\phi - \varphi|_{2\pi}\right)}{\mathrm{sinc}^2\left(\frac{1}{2}|\phi - \varphi|_{2\pi}\right)}.$$

Applications: (almost) optimal coherent tomography, improved estimation algorithms for partition functions, low-depth amplitude estimation, etc.

The continuous probability density function of estimates

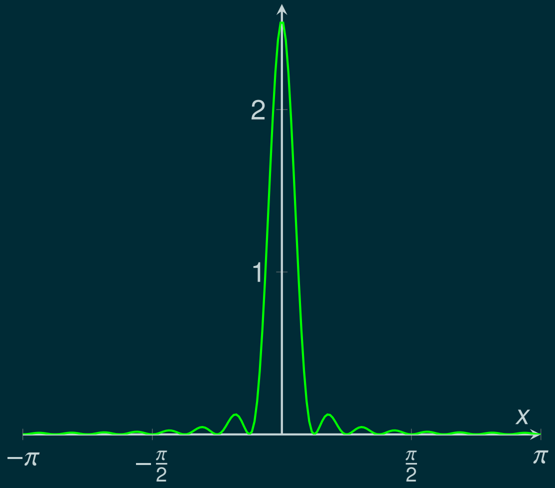

Figure: Plot of $f(\varphi)$ for $x = \phi - \varphi$ and $M = 16$.

Can you boost it while keeping the distribution symmetric?

Connecting Discrete and Continuous Fourier Transforms

Discrete vs. Continuous Fourier Transform

The Continuous Fourier Transform $\mathcal{F}$

$$\hat{f}(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-i\omega t} f(t) dt$$

$\mathcal{F} : \mathcal{L}_2 \to \mathcal{L}_2$ is a unitary transformation (on the Hilbert space of square integrable functions)

Periodic Wrapping of Continuous Functions

Let $f : \mathbb{R} \to \mathbb{C}$, and $r \in \mathbb{R}_+$ be a “period”. We define its wrapping as a function $[0, r] \to \mathbb{C}$

$$f^{(r)}(x) := \lim_{\mathbb{N} \ni K \to \infty} \sum_{k=-K}^{K} f(x + kr).$$

Similarly we define discretized wrapping for $N \in \mathbb{N}$ as a vector in $\mathbb{C}^N$ defined as

$$f_j^{(r,N)} := \lim_{\mathbb{N} \ni K \to \infty} \sum_{k=-K}^{K} f(\frac{j}{N} r + kr).$$

Discrete vs. Continuous Fourier Transform

Connection between Discrete and Continuous Fourier transform

$$\widehat{f(T,N)} = \frac{\sqrt{2\pi N}}{T}\hat{f}\left(\frac{2\pi N}{T},N\right)$$

Terms and conditions apply, but certainly holds for smooth rapidly decaying functions. For more details see Chen, Kastoryano, Brandão, G (2023).

A Commutative Diagram Representation (ignoring the scalar factor)

$$\begin{tikzcd}
f \arrow[r, "\mathcal{F}"] \arrow[d, "\text{discretized periodic}"'] & \hat{f} \arrow[d, "\text{wrapping}"] \\
f(T,N) \arrow[r, "F_N"] & \frac{\sqrt{2\pi N}}{T}\hat{f}\left(\frac{2\pi N}{T},N\right)
\end{tikzcd}$$

Continuous Fourier Transform: shift $\leftrightarrow$ point-wise phase multipl.

Set $T = N = 64$ and $\varphi = 6/37$ – absolute amplitude plot on the circle

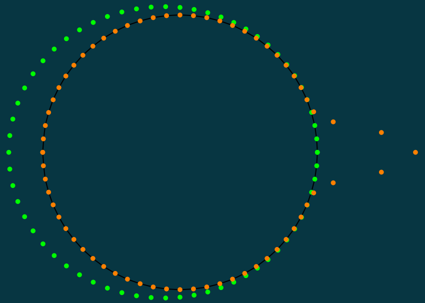

Shifted (discretized) Gaussian

$$f(t) \propto \exp(-(t-32)^2/256),$$

and its Fourier Transform

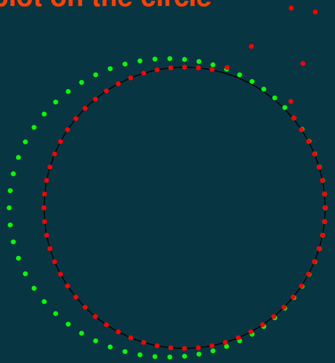

Shifted, phase kicked Gaussian

$$f(t) \propto \exp(\frac{12\pi}{37} it - (t-32)^2/256),$$

and its Fourier Transform

The key observation

In the Gaussian case due to the rapid decay of the tail we have:

$$f^{(T,N)}_{|[0,T)} \approx f^{(T,N)} \qquad \widehat{f^{(T,N)}} = \frac{\sqrt{2\pi N}}{T}\hat{f}^{(\frac{2\pi N}{T},N)} \qquad \hat{f}^{(\frac{2\pi N}{T},N)} \approx \hat{f}^{(\frac{2\pi N}{T},N)}_{|[-\pi,\pi)}$$

With phase shift

Let us introduce

$$f_\varphi(t) := f(t) \cdot e^{2\pi i \cdot \varphi t}$$

$$f^{(T,N)}_{\varphi\,|[0,T)} \approx f^{(T,N)} \qquad \widehat{f_\varphi^{(T,N)}} = \frac{\sqrt{2\pi N}}{T}\hat{f}_\varphi^{(\frac{2\pi N}{T},N)} \qquad \hat{f}_\varphi^{(\frac{2\pi N}{T},N)} \approx \hat{f}^{(\frac{2\pi N}{T},N)}_{\varphi\,|[2\pi(\varphi-\frac{1}{2}),2\pi(\varphi+\frac{1}{2}))}$$

Note that changing $\varphi \pm 1$ does not change anything due to periodic wrapping!

Compare to vanilla phase estimation

Set $N = 16$ and $\varphi = 6/37$ – absolute amplitude plot on the circle

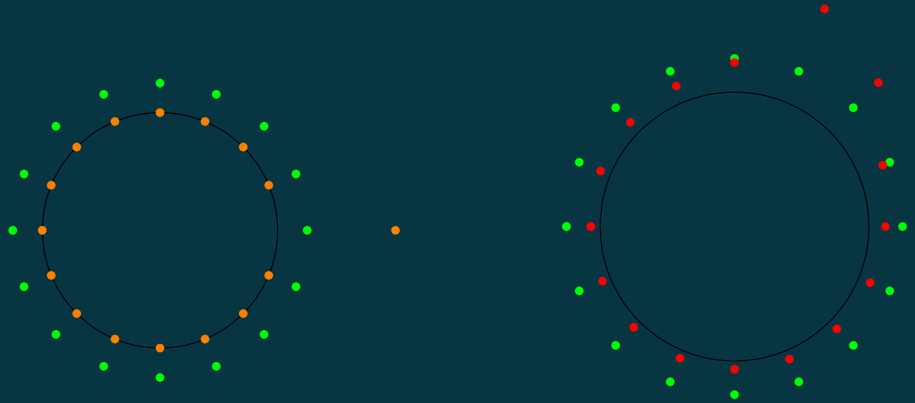

No phase shift, i.e., $\varphi = 0$

Phase shift: $\varphi = 6/37$

Bonus in the Gaussian case: we can increase the resolution cheaply!

Increase N, keep $T = N$ and do not change the Gaussian function just its wrapping.

This does not change the “query” complexity just requires a larger *QFT*.

Gaussian parameters

Let $\sigma \approx \sqrt{\log(1/\delta)}/\varepsilon$ and $f(t) \propto \exp(-t^2/(2\sigma^2))$. Then

$$\hat{f}(\omega) \propto \exp(-\sigma^2\omega^2/2).$$

Implying that the absolute amplitude of $|j\rangle$ in $\widehat{f^{(N,N)}}$ is roughly proportional to ($\omega \leftarrow \frac{2\pi j}{N}$)

$$\exp(-\sigma^2(2\pi j/N)^2/2).$$

Gaussian tail bound

$$\int_x^\infty \frac{1}{\sigma} e^{-t^2/(2\sigma^2)} dt \leq \frac{1}{\sigma}\int_x^\infty \frac{t}{x} e^{-t^2/(2\sigma^2)} dt = \frac{\sigma}{x}\int_x^\infty \frac{t}{\sigma^2} e^{-t^2/(2\sigma^2)} dt = \frac{\sigma}{x} e^{-x^2/(2\sigma^2)}$$

To get δ accuracy we need about $\sqrt{\log(1/\delta)}\sigma \approx \log(1/\delta)/\varepsilon$ uses of U.

If $N = \Omega(\sqrt{\log(1/\delta)}\sigma) = \Omega(\log(1/\delta)/\varepsilon)$, then due to the above tail bound the truncated and the wrapped (discrete) Gaussians are $\mathcal{O}(\delta)$ close to each other.

High accuracy phase estimation in a single run

Idea: use Gaussian amplitudes

- Wrapped Gaussian is almost the same as truncated Gaussian due to rapid decay.
- Fourier transform of a Gaussian is Gaussian, so we get Gaussian noise!
- Choosing parameters appropriately we get an estimator with standard deviation about $1/N$ (up to logarithmic factors) in a single run without garbage ancilla states.
- Initial Gaussian amplitudes can be efficiently prepared – see McArdle, G, Berta (2022)
- Further optimized initial weights can give potential constant improvements: use Kaiser window function from signal processing. See Berry et al. (2022).

Quantum Gibbs Sampling:
When Coherence Meets Randomness

A random story: the Metropolis-Hastings algorithm

- Markov Chain Monte Carlo methods were invented in the late 1940s by Stanisław Ulam who together with John von Neumann used it to tackle challenging computational problems related to, e.g., the design of the H-bomb.
- They were working at Los Alamos, and the codename “Monte Carlo” of the secret computational project was suggested by their colleague Nicholas Metropolis (who built MANIAC I & II) referring to Ulam’s uncle who often gambled there.
- Nicholas Metropolis, Arianna & Marshall Rosenbluth, Augusta & Edward Teller in 1953 developed an algorithm for studying chemical physics problems through thermodynamical averages (apparently Metropolis “just” built the computer).
- It had also proven useful for other systems with complex degrees of freedom.
- Wilfred K. Hastings gave generalization for non-symmetric transitions. (1970)
- In 2000 it was named as one of the top 10 most important algs. of the 20th cent.

Illustrating the main idea – a simple example

- Suppose we have a regular graph, such as a hypercube.
  A simple random walk will converge to the uniform distribution.
- We want to sample from a target distribution say proportional to $\propto \tau = (1, 2, 2, 4)$

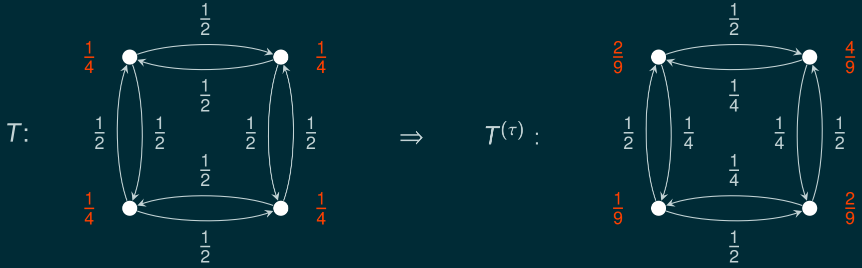

- $T^{(\tau)}$ is *detailed balanced* with respect to $\pi = \frac{\tau}{9} = (\frac{1}{9}, \frac{2}{9}, \frac{2}{9}, \frac{4}{9})$,
  meaning that the total *probability mass exchange* between any two nodes is zero.
- Metropolis algorithm: from x make a transition to y
  - If $\tau_y \geq \tau_x$ accept the move
  - If $\tau_y < \tau_x$ accept with probability $\frac{\tau_y}{\tau_x}$, otherwise reject

detailed balance:
$T^{(\tau)}_{yx} \tau_x = T^{(\tau)}_{xy} \tau_y$

Generic setup of Metropolis in cont. & discrete-time

- We want to sample from a target distribution $\propto \tau \in \mathbb{R}_+^N$ (can be unnormalized)
  Think about Gibbs sampling of an n-spin Ising model $z \in \{-1, +1\}^n$:

$$H(z) = -\sum_{i,j} \alpha_{ij} z_i z_j - \mu \sum_j h_j z_j, \qquad \tau_z = \exp(-\beta H(z)), \qquad N = 2^n$$

- The dynamics is built upon a symmetric "exploratory" Markovian process with transition matrix $T \in \mathbb{R}_+^{N \times N}$ (with non-negative matrix elements)
  For example: pick a random spin and flip it (i.e., random walk on the hypercube)
- Metropolis algorithm: from x make a transition to y according to T
  - If $\tau_y \geq \tau_x$ accept the move
  - If $\tau_y < \tau_x$ accept with probability $\frac{\tau_y}{\tau_x}$, otherwise reject

detailed balance:
$T^{(\tau)}_{yx} \tau_x = T^{(\tau)}_{xy} \tau_y$

- To translate it to a stochastic process take $L := T^{(\tau)} - D$, where $\underbrace{D_{zz} = \sum_{z'} T'_{z'z}}_{decay}$.
  - Continuous-time dynamics: $\exp(tL)$
  - Discrete-time dynamics: $(I + L)^t$

The structure of Metropolis and Glauber dynamics

- Glauber dynamics: accept the move with probability $\frac{1}{1+\frac{\tau_x}{\tau_y}}$
- More generally $T_{yx}^{(\tau)} = \Gamma\left(\frac{\tau_y}{\tau_x}\right) T_{yx}$, where $T \in \mathbb{R}_+^{N\times N}$ is symmetric, $\Gamma : \mathbb{R}_+ \to [0,1]$ and

$$\forall r \in \mathbb{R}_+ \qquad \Gamma(r) = r \cdot \Gamma\left(\frac{1}{r}\right)$$

- Metropolis algorithm: $\Gamma_M(r) = \max(1, r)$ Glauber dynamics: $\Gamma_G(r) = \frac{1}{1+\frac{1}{r}}$
- Once again detailed balance ensures that the "probability mass" transfer equals in both directions between every pair of states, so fixed point is $\propto \tau$.
- T' is detailed balanced if and only if it is symmetric w.r.t. τ-weighing:

$$T'\mathrm{diag}(\tau) = \mathrm{diag}(\tau)T'^T \quad \Leftrightarrow \quad \mathrm{diag}(\sqrt{\tau^{-1}})T'\mathrm{diag}(\sqrt{\tau}) = \mathrm{diag}(\sqrt{\tau})T'^T\mathrm{diag}(\sqrt{\tau^{-1}})$$

Stochastic quantum transitions ($0 < \rho \in \mathbb{C}^{2^n \times 2^n}$, $\mathrm{Tr}[\rho] = 1$)

1. Quantum Markov chain (q. channel): completely positive and trace preserving

$$\underbrace{\mathcal{T}[\rho] = \sum_j K_j \rho K_j^\dagger}_{\text{completely positive}} \qquad \underbrace{\mathcal{T}^\dagger[I] = \sum_j K_j^\dagger K_j = I}_{\text{trace preserving}}$$

2. Continuous-time quantum Markov chain (q. dynamical semigroup): $e^{t\mathcal{L}}$
   Infinitesimal generator of a continuous-time q. Markov chain, a.k.a., Lindbladian:

$$\mathcal{L}[\cdot] = \underbrace{\sum_{j=0}^{m} K_j[\cdot]K_j^\dagger}_{\substack{\text{transitions} \\ \text{(CP map)}}} \underbrace{-\frac{1}{2}\sum_{j=0}^{m}\{K_j^\dagger K_j, \cdot\}}_{\text{decay term}} \underbrace{- i[C, \cdot]}_{\text{coherent term}} = \underbrace{\mathcal{T}[\cdot]}_{\substack{\text{transitions} \\ \text{(CP map)}}} \underbrace{-\frac{1}{2}\sum_{j=0}^{m}\{\mathcal{T}^\dagger[I], \cdot\}}_{\text{decay term}} \underbrace{- i[C, \cdot]}_{\text{coherent term}}$$

3. (KMS) detailed balanced sup-op: $\rho^{-\frac{1}{4}}\mathcal{T}[\rho^{\frac{1}{4}} \cdot \rho^{\frac{1}{4}}]\rho^{-\frac{1}{4}} = \rho^{\frac{1}{4}}\mathcal{T}^\dagger[\rho^{-\frac{1}{4}} \cdot \rho^{-\frac{1}{4}}]\rho^{\frac{1}{4}}$. $\rho = e^{-\beta H}$
   Lindbladian detailed balanced if $\mathcal{T}$ is and $\langle\psi_i|C|\psi_j\rangle = \frac{i}{2}\tanh(\beta\frac{E_i - E_j}{4})\langle\psi_i|\mathcal{T}^\dagger[I]|\psi_j\rangle$.

Why do we like (KMS) detailed balance?

- Classical detailed balance: probability mass exchange in both direction equals for the distribution π. Equivalently the (probability) transition matrix is similar to a symmetric one

$$\mathrm{diag}(\pi^{-\frac{1}{2}})P'\mathrm{diag}(\pi^{\frac{1}{2}})$$

- The quantum generalization of the detailed balance condition to a CP-map $\mathcal{T}'$ with respect to a full-rank density matrix $0 \prec \rho \in \mathbb{C}^{d\times d}$ is the following equation

$$\rho^{-\frac{1}{4}}\mathcal{T}'[\rho^{\frac{1}{4}}\cdot\rho^{\frac{1}{4}}]\rho^{-\frac{1}{4}} = \rho^{\frac{1}{4}}\mathcal{T}'^{\dagger}[\rho^{-\frac{1}{4}}\cdot\rho^{-\frac{1}{4}}]\rho^{\frac{1}{4}},$$

  i.e., $\mathcal{T}'$ is similar to a self-adjoint map.
- A ρ-detailed-balanced Lindbladian fixes the Gibbs state

$$\begin{aligned}\mathcal{L}[\rho] &= \rho^{\frac{1}{4}}(\rho^{-\frac{1}{4}}\mathcal{L}[\rho^{\frac{1}{4}}\sqrt{\rho}\rho^{\frac{1}{4}}]\rho^{-\frac{1}{4}})\rho^{\frac{1}{4}}\\ &= \rho^{\frac{1}{4}}(\rho^{\frac{1}{4}}\mathcal{L}^{\dagger}[\rho^{-\frac{1}{4}}\sqrt{\rho}\rho^{-\frac{1}{4}}]\rho^{\frac{1}{4}})\rho^{\frac{1}{4}} = \sqrt{\rho}\underbrace{\mathcal{L}^{\dagger}[I]}_{=0}\sqrt{\rho} = 0.\end{aligned}$$

How to handle ambiguity in phase estimation?

- Temme, Osborne, Vollbrecht, Poulin, Verstraete'09; Yung and Aspuru-Guzik'10
  - Require "shift-invariant" or perfect phase estimation $\rightarrow$ unphysical
- Wocjan and Temme'21 (continuous-time quantum Metropolis $\leftrightarrow$ Davies generator)
  - Assume spectrum has periodic gaps ("rounding promise") $\rightarrow$ unphysical
- Chen, Kastoryano, Brandão, **G.**'23 & Jiang, Irani'24. (apx. phase est., ~cont.time)
  - No extra conditions, efficiently implementable $\rightarrow$ Non-Local jumps, apx. Det.-Bal.

  Can be made exactly detailed balanced via Gaussian uncertainty
  Energies are estimated both before and after transitions $\Rightarrow$ non-local jumps
- Chen, Kastoryano, **G.**'23 [CKG23] (builds on WT'21 – continuous-time)
  - Apply damped phase estimation / operator Fourier transform $\rightarrow$ Quasi-Local, D-B ✓
- Ding, Li, Lin'24 [DLL24] and **G.**, Chen, Doriguello, Kastoryano'24 [GCDK24]
  - Filter jumps via a weighted integral over Heisenberg evolution $\rightarrow$ Q-L, D-B ✓

Making a *self-adjoint* CP map exactly detailed-balanced

- Linear constructions: can be analyzed at the level of Kraus-operators
- Davis-Metropolis: do classical Metropolis in the eigenbasis (discontinuous)
- Operator Fourier trans. w. Gaussian damping (Chen, Kastoryano, Brandão, **G.**'24)

$$K(\omega) = \int_{\mathbb{R}} e^{iHt} K e^{-iHt} e^{-i\omega t} f(t) dt, \qquad \rho \to \int_{\mathbb{R}} \gamma_M\left(\omega + \frac{1}{2}\right) K(\omega)\rho K(\omega)^\dagger d\omega,$$

  where f is a Gaussian with variance 1 (or more generally $\sim \beta$).
- Classical Metropolis with uncertain energy estimation (Ceperley, Dewing 1999)
- Coherent reweighing (**G.**, Chen, Doriguello, Kastoryano'24 and Ding, Li, Lin'24):

$$\langle\psi'|K^{(\rho)}|\psi\rangle = \gamma\left(\frac{\beta}{2}(E_{\psi'} - E_{\psi})\right)\langle\psi'|K|\psi\rangle \qquad [\text{note } \gamma(x) = \Gamma(e^x)]$$

- Coherent reweighing can be expressed as integral over Heisenberg evolution

$$K^{(\rho)} = \int_{\mathbb{R}} \underbrace{f(t)}_{\text{exponentially decaying above } |t| \sim \beta} \exp(iHt) K \exp(-iHt) dt. \qquad (\hat{f} = \gamma)$$

What do we / the environment learn when a jump is performed

- Davies: The exact energy difference caused by the jump
- Phase-estimation based Metropolis: approximate energies before and after
- Operator Fourier transform based: approximate energy difference
- Coherent reweighing: Nothing (except which jump was applied)

Parent Hamiltonians from KMS-det.-bal. Lindbladians

- Take a ρ-detailed-balanced Lindbladian

$$\mathcal{L}_i[\cdot] = K_i[\cdot]K_i^\dagger \underbrace{-\frac{1}{2}\{D_i, \cdot\}}_{\text{decay term}} \underbrace{-[C_i, \cdot]}_{\text{coherent term}}$$

- Group together decay and coherent terms: $O_i = \frac{1}{2}D_i + C_i$
- Parent-Hamiltonian:

$$-\widetilde{K}_i \otimes \widetilde{K}_i^* + \widetilde{O}_i \otimes I + I \otimes \widetilde{O}_i^*, \quad \text{where} \quad \widetilde{K}_i = \rho^{-\frac{1}{4}} K_i \rho^{\frac{1}{4}} \text{ and } \widetilde{O}_i = \rho^{-\frac{1}{4}} O_i \rho^{\frac{1}{4}}$$

has the purified Gibbs state as the ground state

$$\sum_i^{-\beta E_i/2} |\psi_i\rangle \otimes \left|\psi_i^*\right\rangle / \sqrt{\mathrm{Tr}[^{-\beta H}]}$$

- We get a **frustration-free quasi-local Hamiltonian**!
- Also works for discrete-time quantum channels with appropriate modifications.

Discrete-time construction

- Lindbladian simulation is costly, discrete variant would be attractive
- [GCDK24]: Given ρ-detailed-balanced $\mathcal{T}^{(\rho)}$ CP map, there is a unique natural complementary Kraus-operator such that the CP map

$$\mathcal{T}^{(\rho)}[\cdot] + K[\cdot]K^\dagger$$

  is a ρ-detailed-balanced quantum channel.
-

$$K = \sqrt{\sqrt{\rho}(I - \mathcal{T}^{(\rho)\dagger}[I])\sqrt{\rho}}\rho^{-\frac{1}{2}} = I - \mathcal{S}^-[\mathcal{T}^{(\rho)\dagger}[I]] - \mathcal{S}^-[\mathcal{S}^+[\mathcal{T}^{(\rho)\dagger}[I]]\mathcal{S}^-[\mathcal{T}^{(\rho)\dagger}[I]]] - \dots$$

  where K has a well-behaved Taylor-series. Consequently K is β-quasi-local and under some mild assumptions can be efficiently implemented via QSVT.
- New: Iterate th above construction to first order.
  Slack from trace-preservation gets squared $\Rightarrow$ Double-exponential convergence!
- Use oblivious amplitude estimation to implement quantum channel.

Summary: Quantum Gibbs Sampling

- Typical envisioned usage:
  Prepare several copies of the Gibbs state and measure some observables
- Dissipative Gibbs state preparation – properties of recent constructions (Chen et al., Ding et al., Gilyén et al.) that satisfy exact (KMS) detailed balance:

	Method	Efficient implement.	Local	Energy uncertainty	Ergodicity preserving	Kraus rank
Cl.	Glauber & Metropolis	$\checkmark$	$\checkmark$	0	$\checkmark$	$=$ orig.
Q.	Davies generator	-	-	0	$\checkmark$	$\leq d^2$
	Phase estimation	$\checkmark$	-	$\mathcal{O}(\frac{1}{\beta})$	$\checkmark$	$\leq d^2$
	Operator Fourier transf.	$\checkmark$	quasi	$\mathcal{O}(\frac{1}{\beta})$	$\checkmark$	$\leq d^2$
	Interpolated	$\checkmark$	quasi	σ	$\checkmark$	$\leq d^2$
	Coherent reweighing	$\checkmark$	quasi	∞	$\checkmark$	$=$ orig.

(Also there is a large and rapidly growing number of approximate versions!)

(Quantum) Shadow Tomography

Basic goal: Measure many observables using a few copies of a mixed state ρ
Problem: Quantum measurements disturb the quantum state

- Aaronson'17: General Quantum Shadow Tomography
  Idea: large joint measurements reduce disturbance.

  $$\widetilde{\mathcal{O}}\left(\frac{\log^4(M)\log(D)}{\varepsilon^4}\right)$$

  copies suffice (see van Apeldoorn & Gilyén'18) for simultaneously estimating the expectation value of M bounded observables to precision ε.
  However, requires joint measurements and gate complexity is $\mathrm{poly}(D+M)$!
- Huang, Kueng & Preskill'20: Classical shadows
  Idea: use randomized single-copy measurements and classical post-processing. Local Pauli measurements require $\mathcal{O}(4^k\log(M/\delta)/\varepsilon^2)$ samples for general k-local bounded $\|O_i\| \leq 1$ observables, while global Clifford measurements require $\mathcal{O}(r\log(M/\delta)/\varepsilon^2)$ samples for general rank-r bounded $\|O_i\| \leq 1$ observables.

Shadow Tomography of Thermal States

- Applies for general observable, but needs the extra assumption that we have access to the Hamiltonian H such that
$$\rho \propto \exp(-\beta H).$$
- Idea: a detailed-balanced channel preserves ρ statistically.
- Consider a quantum channel Q as a measurement instrument with Kraus operators K_0 and
$$K_{\pm} := \frac{I \pm A}{2},$$
corresponding to a bounded Hermitian $\|A\| \leq 1$. Then, applying Q on ρ we have
$$\mathbb{E}[\mathrm{sign}(\text{label of Kraus operator})] = [\rho A]\,.$$
- Use ρ-detailed-balanced discrete-time quantum Gibbs sampler with Kraus ops.
$$K'_{\pm} := \int_{\mathbb{R}} f(t) \exp(itH) K_{\pm} \exp(-itH).$$
Expectation is preserved as $[\rho \exp(itH) K_{\pm} \exp(-itH)] = [\rho K_{\pm}]$.
If $\hat{f}(\nu)^* = \exp(-\beta\nu/2)\hat{f}(-\nu)$), there is a K_0 that makes Q detailed-balanced. (Gilyén, Chen, Doriguello, Kastoryano'24)

Sample-Optimal Shadow Tomography of Thermal States

The implementation cost of each ρ-detailed-balanced measurement channel is dominated by its required total $\widetilde{\mathcal{O}}(\beta)$ (controlled) Hamiltonian simulation time.

Chen-G.'26 `arXiv:2603.16845`

$$r = \mathcal{O}\left(\frac{\log(M/\delta)}{\epsilon^2}\right)$$

independent copies

(optimal when assuming subexponential time black-box Ham. sim.)