Quantum Gibbs samplers through the lens of contraction and resources

Quantum Resources Workshop, Tokyo

Daniel Stilck França

QMATH, University of Copenhagen

Joint work with Cambyse Rouzé, Álvaro M. Alhambra, Sam Slezak, and Matteo Scandi
Based on arXiv:2411.04885, 2601.16154

Quantum Gibbs sampling = fixed-point iteration

Gibbs sampler

Given a local Hamiltonian H, construct a quantum channel $\mathcal{T}$ (or semigroup $\mathcal{T}_t = e^{t\mathcal{L}}$) such that

$$\sigma_\beta = \frac{e^{-\beta H}}{\mathrm{tr}(e^{-\beta H})} \qquad \text{is the unique fixed point.}$$

Then iterate until convergence:

$$\rho_0 \overset{\mathcal{T}}{\longmapsto} \rho_1 \overset{\mathcal{T}}{\longmapsto} \rho_2 \overset{\mathcal{T}}{\longmapsto} \cdots \longrightarrow \sigma_\beta .$$

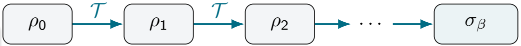

continuous-time version:

$\rho_t \;=\; e^{t\mathcal{L}}(\rho_0)$ with $\mathcal{L}$ reversible w.r.t. σ_β

Trace-distance mixing time:

$$t_{\mathrm{mix}}(\varepsilon) := \inf\{t \geq 0 : \sup_\rho d_{\mathrm{tr}}(e^{t\mathcal{L}}(\rho), \sigma_\beta) \leq \varepsilon\},$$

$$d_{\mathrm{tr}}(\rho, \sigma) := \tfrac{1}{2}\|\rho - \sigma\|_1 .$$

- Classical analogue: Markov chain Monte Carlo.
- Typical question: how many times we need to apply to get convergence?
- Usually done by contraction coefficients.

What changed after Chen et al.?

- Chen–Kastoryano–Brandão–Gilyén / Chen–Kastoryano–Gilyén gave an exact, efficiently implementable, quasi-local noncommutative Gibbs sampler with fixed point σ_β [Chen et al. 23a, Chen et al. 23b].
- This unlocked both:
  - **algorithmics:** efficient preparation in regimes where one expects thermalization to be easy;
  - **analysis:** spectral gaps, rapid mixing, repeated-interaction implementations, and physically motivated system–bath models.
- Since then there has been a lot of activity: high temperature, weak interactions, fermions, 1D chains, repeated interactions, macroscopic baths, and recent fermionic/rapid-mixing advances [Rouzé et al. 24a, Rouzé et al. 24b, Slezak et al. 26, Smid–Berta et al. 25a, Smid–Berta et al. 25b].
- Non-trivial families of quantum channels for which we can bound the contraction coefficients!

A contraction-coefficient dictionary

$$\eta_D(\Phi) := \sup \frac{D(\Phi(\text{input}), \sigma_\beta)}{D(\text{input}, \sigma_\beta)}$$

"how much of a quantity survives one step?"

Quantity	Representative estimate	What it buys
χ^2/variance	$\mathrm{Var}_\sigma(e^{t\mathcal{L}^\dagger} X) \le e^{-2\lambda t}\,\mathrm{Var}_\sigma(X)$	spectral gap / Poincaré; typically polynomial or linear-time preparation
oscillator norm	$\left\lVert\!\left\lvert e^{t\mathcal{L}^\dagger} X \right\rvert\!\right\rVert \le e^{-ct}\, \lVert\!\lvert X \rvert\!\rVert$	rapid mixing $t = O(\log(n/\varepsilon))$ in the high-temperature noncommuting regime
relative entropy	$D(e^{t\mathcal{L}}(\rho)\|\sigma) \le e^{-\alpha t} D(\rho\|\sigma)$	MLSI / entropy decay; Usually the gold standard

- For this talk, gap $\leftrightarrow$ χ^2-contraction and oscillator norm contraction are the main positive results.
- The real missing piece is relative-entropy contraction in genuinely noncommuting settings.

The oscillator norm is a local Lipschitz seminorm

$$\delta_a(X) := X - \frac{I_a}{2} \otimes \mathrm{tr}_a(X),$$

$$|||X||| := \sum_{a \in \Lambda} \|\delta_a(X)\|_\infty .$$

- $\delta_a(X)$ removes the part of X that is invisible on site a.
- $|||X|||$ measures how many sites really matter for the observable.
- Multiples of the identity have oscillator norm 0.
- For a Pauli string of weight w, $|||P||| = w$.

Why this norm?

Trace norm not sensitive enough for local-to-global proofs. The oscillator norm localizes the problem to one-site influences.

$X = Z_2 Z_5$ on a six-site chain

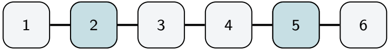

$\delta_2(X) = X$

$\delta_5(X) = X$

$\delta_a(X) = 0$ for $a \notin \{2, 5\}$

Hence $|||X||| = 2$.
It counts how many sites support the observable.

High-temperature rapid mixing in oscillator norm

Contraction statement (local noncommuting Gibbs sampler)

For the quasi-local sampler $\mathcal{L}^{(\beta)}$ and sufficiently high temperature $\beta = O((DJ)^{-1})$, there is a constant $c > 0$ such that

$$\eta_{\mathrm{osc}}(t) := \sup_{X \not\propto I} \frac{\left\| e^{t\mathcal{L}^{(\beta)\dagger}}(X) \right\|}{\|X\|} \le e^{-ct}.$$

This immediately implies

$$\|e^{t\mathcal{L}^{(\beta)}}(\rho) - \sigma_\beta\|_1 \le 4n\, e^{-ct}, \qquad t = O(\log(n/\varepsilon)),$$

so the sampler mixes in the best possible many-body timescale [Rouzé et al. 24b, Rouzé et al. 24a].

- Gap/χ^2 contraction alone only gave polynomial-time mixing.
- Oscillator contraction upgrades this to rapid mixing.
- The same proof strategy also extends to some long-range models.

How the oscillator proof works: anchor + perturbation

Let $X_t = e^{t\mathcal{L}^{(\beta)\dagger}}(X)$. For each site a,

$$\frac{\mathrm{d}}{\mathrm{d}t}\,\delta_a(X_t) = -\lambda\,\delta_a(X_t) + \sum_{b\neq a}[\delta_a, \mathcal{L}_b^{(\beta)\dagger}](X_t) + \delta_a\big((\mathcal{L}_a^{(\beta)\dagger} - \mathcal{L}_a^{(0)\dagger})(X_t)\big).$$

Step 1: depolarizing anchor

At $\beta = 0$, the local generator is just sitewise depolarization:

$$\delta_a \mathcal{L}_a^{(0)\dagger} = -\lambda\,\delta_a, \qquad \lambda = \frac{1}{\sqrt{2}e^{1/4}}.$$

So the infinite-temperature point contracts every local derivative uniformly.

Step 2: control the perturbation

Show that the off-site commutators and on-site thermal correction satisfy

$$\|[\delta_a, \mathcal{L}_b^{\dagger}](X)\|_\infty \le \sum_c \kappa_{a,b}^c \|\delta_c(X)\|_\infty,$$

$$\|\delta_a((\mathcal{L}_a^{\dagger} - \mathcal{L}_a^{(0)\dagger})(X))\|_\infty \le \sum_c \gamma_a^c \|\delta_c(X)\|_\infty.$$

Summing over a turns this into a Grönwall inequality for $\|\!|X_t|\!\|$.

How the oscillator proof closes: locality feeds the bootstrap

Approximate the local term by a spatial truncation $\mathcal{L}_{\beta,r}$ and compare to $\beta = 0$:

$$\sum_{r \geq r_0} \|\mathcal{L}^\dagger_{\beta,r} - \mathcal{L}^\dagger_{\beta,r-1}\|_{\infty\to\infty} \leq \Delta(r_0), \qquad \|\mathcal{L}^\dagger_{\beta} - \mathcal{L}^\dagger_{0}\|_{\infty\to\infty} \leq \eta(\beta).$$

For local Hamiltonians, Lieb–Robinson bounds plus the smooth filter give

$$\Delta(r_0) \text{ fast decaying in } r_0, \qquad \eta(\beta) = O(\beta J).$$

Hence the total influence can be bounded by

$$\kappa \leq 4(2r_0+1)^{2D}\eta(\beta) + f(r_0).$$

Choose first r_0 large and then β small so that $\kappa < \lambda$; then

$$|||X_t||| \leq e^{-(\lambda-\kappa)t}\, |||X||| \, .$$

QIT translation

This is a noncommutative Dobrushin argument: start from a tensor-product contraction, quantify how much locality and temperature spoil it, and keep the spoilage below the depolarizing margin.

Repeated interactions as a Gibbs sampler

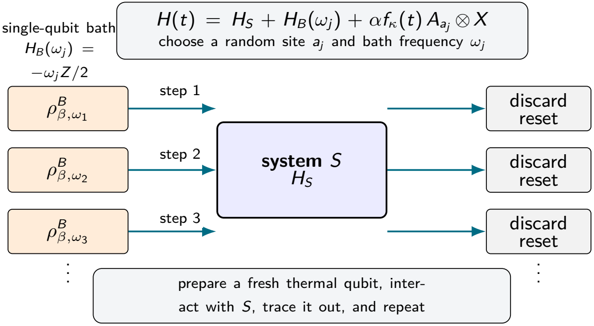

$f_\kappa(t) \propto e^{-t^2/(2\kappa^2)}$ is a Gaussian switching pulse of width κ;
$A_a = A_a^\dagger$ is the local system coupling operator chosen at site a.

One collision step

$$\Phi_{\alpha,\kappa}(\rho) = \mathbb{E}_{a,\omega}\ \mathrm{Tr}_B\left[U_{\alpha,\kappa}(\rho \otimes \rho^B_{\beta,\omega})U^\dagger_{\alpha,\kappa}\right].$$

Effective Gibbs sampler

$$\Phi_{\alpha,\kappa} \approx e^{\alpha^2 \mathcal{L}^{\mathrm{RI}}_\kappa}, \qquad \mathcal{L}^{\mathrm{RI}}_\kappa = -i[H_{\mathrm{LS}},\cdot] + c_\kappa\, \mathcal{L}^G_\kappa.$$

$$\sigma_\beta \propto e^{-\beta H_S}.$$

- Collision dynamics define the one-step CPTP map.
- Iterating that map gives the repeated Gibbs sampler.
- The weak-coupling limit recovers a KMS Lindbladian with fixed point σ_β.
- κ sets the tradeoff: sharper filtering versus less locality.

Why monotonicity is needed beyond quasi-locality

Physical samplers are harder to analyze

- Repeated interactions and weak system–bath couplings are simple/physical.
- But the regime where they accurately approximate thermalization is typically less local.
- Existing gap/MLSI proofs live at the quasi-local end of the family.

Accuracy-locality tradeoff

- Repeated interactions: larger Gaussian width κ improves accuracy.
- Macroscopic-bath KMS generators: smaller coupling α improves accuracy.

Better accuracy means worse locality.

quasi-local point
good for gap / MLSI proofs

same Gibbs fixed point | vary κ, α | different locality scale

accurate physical point
good approximation to the real dynamics

Question: can contraction coefficients be *transported* along the family?

The monotonicity mechanism: convolution in the Dirichlet form

Mechanism

For a one-parameter family $\{\mathcal{L}_\delta\}_{\delta \geq 0}$ with the same fixed point,

$$\mathcal{E}^{(\delta)}_{X,Y}(s) := -\langle Y(s\beta), \mathcal{L}^\dagger_\delta(X(s\beta))\rangle_{\sigma_\beta},$$

$$\mathcal{E}^{(\delta)}_{X,Y} = f_\delta * \mathcal{E}^{(\infty)}_{X,Y}, \qquad f_\delta = g_{\delta,\delta'} * f_{\delta'}.$$

Variational formulas for gap and MLSI then give

$$\lambda_{\mathrm{gap}}(\mathcal{L}_\delta) \geq \|g_{\delta,\delta'}\|_1 \, \lambda_{\mathrm{gap}}(\mathcal{L}_{\delta'}),$$

$$\alpha_{\mathrm{MLSI}}(\mathcal{L}_\delta) \geq \|g_{\delta,\delta'}\|_1 \, \alpha_{\mathrm{MLSI}}(\mathcal{L}_{\delta'}).$$

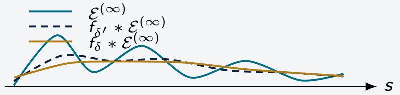

more smoothing = convex averaging of the shifted Dirichlet form

[Slezak et al. 26]

Takeaway

For the Gaussian-filtered families we care about, $\|g_{\delta,\delta'}\|_1 = 1$, so passing to the more accurate / less local generator does not decrease the gap or MLSI.

What monotonicity buys us

- **Repeated-interaction Gibbs sampling:** inherit spectral gaps/χ^2 contraction from the quasi-local sampler, even though the accurate effective Lindbladian is non-local [Slezak et al. 26].
- **Physical KMS Lindbladians from system–bath couplings:** first polynomial-time thermalization results beyond the commuting setting in several regimes [Slezak et al. 26].
- **If MLSI is available quasi-locally,** the same argument transports it as well.

Current reach

High-temperature local lattices, weakly interacting fermions, and 1D chains all fall under this template once a quasi-local gap bound is known [Rouzé et al. 24a, Rouzé et al. 24b, Slezak et al. 26, Smid–Berta et al. 25a].

But note the limitation

This is currently a χ^2/gap story. We still do *not* know how to transport rapid mixing in oscillator norm this way, and we certainly do not yet have a general noncommuting relative-entropy contraction theorem.

The efficient regimes we can prove look resource-poor

Spin/qudit high temperature

- Bakshi–Liu–Moitra–Tang: sufficiently high-temperature Gibbs states are separable [Bakshi et al. 24].
- In the qubit/Pauli setting, the decomposition is into stabilizer product states.
- **Consequence:** in those regimes, no entanglement and no magic survive.

Bounded-degree fermions

- Ramkumar–Cai–Tong–Jiang: high-temperature Gibbs states are mixtures of fermionic Gaussian states [Ramkumar et al. 25].
- Smid et al, Tong et al: weakly interacting Fermi–Hubbard is efficiently samplable at any temperature [Smid–Berta et al. 25a].
- **Consequence:** current fermionic positive results stay close to Gaussian / weak-coupling structure.

Interpretation: the regimes where we can currently prove efficient thermalization also look structurally simple. Cluster expansions drive the structural proofs.

A call to arms: can resources themselves have contraction coefficients?

The question I would like this community to attack

Can we prove contraction inequalities for genuine resource monotones, not only for χ^2 or oscillator norms?

$$M(e^{t\mathcal{L}}(\rho)) \leq e^{-\alpha_M t}\, M(\rho) \qquad \text{or at least} \qquad \eta_M(t) := \sup_\rho \frac{M(e^{t\mathcal{L}}(\rho))}{M(\rho)} < 1$$

for monotones such as

$$M \in \{E_R \ \text{(relative entropy of entanglement)},\ \text{relative entropy of magic / mana}\}$$

$$\text{or fermionic non-Gaussianity.}$$

- In the easy high-temperature regimes, this would explain why resources disappear.
- More interestingly, it could reveal a connection between:
  - the threshold for **efficient Gibbs sampling**, and
  - the threshold for **resource sudden death**.

Takeaways

1. Quantum Gibbs sampling is naturally a contraction problem for a fixed-point channel.
2. Today we control two useful notions: χ^2/gap contraction and oscillator-norm contraction.
3. The efficient regimes we can prove are mostly resource-poor.
4. Can we repeat this story for contraction of resource monotones?

Thank you!

References I

C.-F. Chen, M. J. Kastoryano, F. G. S. L. Brandão, and A. Gilyén.
Quantum Thermal State Preparation.
arXiv:2303.18224, 2023.

C.-F. Chen, M. J. Kastoryano, and A. Gilyén.
An efficient and exact noncommutative quantum Gibbs sampler.
arXiv:2311.09207, 2023.

C. Rouzé, D. S. França, and Á. M. Alhambra.
Efficient Thermalization and Universal Quantum Computing with Quantum Gibbs Samplers.
In *STOC 2025*; arXiv:2403.12691.

C. Rouzé, D. S. França, and Á. M. Alhambra.
Optimal quantum algorithm for Gibbs state preparation.
arXiv:2411.04885, 2024.

S. Slezak, M. Scandi, Á. M. Alhambra, D. S. França, and C. Rouzé.
Polynomial-time thermalization and Gibbs sampling from system-bath couplings.
arXiv:2601.16154, 2026.

References II

A. Bakshi, A. Liu, A. Moitra, and E. Tang.
High-Temperature Gibbs States are Unentangled and Efficiently Preparable.
In *FOCS 2024*; arXiv:2403.16850.

A. Ramkumar, Y. Cai, Y. Tong, and J. Jiang.
High-Temperature Fermionic Gibbs States are Mixtures of Gaussian States.
arXiv:2505.09730, 2025.

Š. Šmíd, R. Meister, M. Berta, and R. Bondesan.
Polynomial Time Quantum Gibbs Sampling for Fermi-Hubbard Model at any Temperature.
arXiv:2501.01412, 2025.

Š. Šmíd, R. Meister, M. Berta, and R. Bondesan.
Rapid Mixing of Quantum Gibbs Samplers for Weakly-Interacting Quantum Systems.
arXiv:2510.04954, 2025.

Z. Ding, Y. Zhan, J. Preskill, and L. Lin.
End-to-End Efficient Quantum Thermal and Ground State Preparation Made Simple.
arXiv:2508.05703, 2025.