Frontiers of Quantum Complexity *Instructor: Anand Natarajan, Anurag Anshu*

Lecture 19: Gibbs Sampling

November 6, 2025 *Scribe: Vetri Vel*

1 Background

Recall the Gibbs state at inverse temperature β,

$$\rho_\beta(H) = \frac{e^{-\beta H}}{\mathrm{Tr}\left(e^{-\beta H}\right)} \tag{1}$$

where we usually take H to be a local Hamiltonian. In a previous lecture, we discussed how the Gibbs state shows up in many contexts in physics. For example, the state of maximum possible entropy given expectations of observables has the form of a Gibbs state.

However, just as important as Gibbs states is how to create them. This is where quantum Gibbs samplers come in. They are a family of quantum channels parameterized by time that drive an arbitrary initial state towards a specific Gibbs state.

Quantum Gibbs Sampler: A family of quantum channels parametrized by time $\{\mathcal{E}_t\}_{t\geq 0}$ such that

- A Gibbs state is the fixed point: $\mathcal{E}_t[\rho_\beta(H)] = \rho_\beta(H)$ for some Hamiltonian H
- Any initial state σ is eventually mapped to the Gibbs state, $\lim_{t\to\infty} \mathcal{E}_t[\sigma] = \rho_\beta(H)$

We will focus on Markovian Gibbs samplers where $\mathcal{E}_t = e^{t\mathcal{L}}$ for a superoperator $\mathcal{L}$ called the Lindbladian. In this context, Gibbs sampling is called a quantum thermalization process.

1.1 Recent history

Chen et al. (2023) define a quantum Gibbs sampler with quasi-local Lindbladian to produce a Gibbs state given an arbitrary Hamiltonian [1, 2]. This gives a more physical process to create Gibbs states than previous works where the Lindbladians are non-local. They also give quasi-local Hamiltonians whose unique ground states are canonical purified Gibbs states, also known as thermofield double states. These states are useful in the study black holes or as ground state ansatzes,

$$|\mathrm{TFD}\rangle = (\mathbf{1} \otimes \sqrt{\rho_\beta})\, |\mathrm{EPR}\rangle^{\otimes n} \tag{2}$$

Another recent work shows that a Gibbs sampler-like process called thermal gradient descent can simulate quantum circuits. More precisely, they show that thermal gradient descent finds local minima and that the problem of finding a local minimum of a quantum Hamiltonian is BQP-complete using a modified Feynman-Kitaev construction [3].

Also of interest to complexity theory, Bergamaschi et al. construct a family of commuting local Hamiltonians whose Gibbs states are classically hard to sample from even at a fixed, finite temperature but can be efficiently prepared by a quantum Gibbs sampler [4].

The goal of this lecture is to provide foundations to understand the promising results of the past few years. We will cover:

- Classical Gibbs samplers and the importance of detailed balance
- A mapping from classical Gibbs samplers to quantum Hamiltonians and ground states
- An introduction to quantum Gibbs samplers and Bohr frequencies

1.2 Digression

We will informally discuss how any classical computation can be encoded by a constant temperature Gibbs state on a 2-dimensional geometry [5].

We take the computation and apply Cook-Levin to get a constraint satisfaction problem, represented by a local Hamiltonian H_{CL} on a 2-dimensional geometry. We then consider the Gibbs state at finite temperature, $\rho_\beta(H_{\mathrm{CL}})$. This corresponds to applying each gate with error probability $\sim e^{-\beta}$. Therefore, if we start with a fault-tolerant version of the computation that can handle a constant error rate, the Gibbs state correctly encodes the classical computation. An analog of this construction for quantum computation is unknown.

> **Open question**: Can finite temperature Gibbs states encode quantum computation? How about on a low-dimensional lattice?

It is known that the quantum PCP conjecture implies the hardness of Gibbs states [6]. Therefore if the quantum PCP conjecture holds, it is conceivable that Gibbs states at low temperature could encode quantum computation. It is also believed (assuming QMA $\neq$ NP) that quantum PCP Hamiltonians cannot exist in low-dimensional lattices since their low-energy states must be highly entangled by the NTLS theorem [7]. For example, quantum PCP Hamiltonians cannot exist on 1-dimensional lattices because the area law bounds entanglement [8]. However, this doesn't seem to rule out the possibility of Gibbs states on a low-dimensional lattice that encode quantum computations.

2 Classical Gibbs Sampling

Classical Gibbs samplers are special cases of Markov chains. Given a probability distribution $p(x_1, \ldots, x_n)$ for individual bits x_i, define the Markov chain M,

$$M = \frac{1}{n}\sum_{i=1}^{n} M_i \quad \text{where} \quad M_i q(x_i, x_{-i}) = q(x_{-i}) p(x_i|x_{-i}) \tag{3}$$

Intuitively, M_i reads all bits other than x_i (denoted as x_{-i}) and updates the conditional distribution on bit i to match $p(x_i|x_{-i})$. Any fixed point q such that $Mq = q$ must satisfy $M_i q = q$ for all i, which means

$$q(x_{-i})p(x_i|x_{-i}) = q(x_i, x_{-i}) = q(x_{-i})q(x_i|x_{-i}) \tag{4}$$

which means that all conditionals must match, $q(x_i|x_{-i}) = p(x_i|x_{-i})$. However, this doesn't imply that $q = p$. Consider the classical analog of a cat state,

$$p = \frac{1}{2}\mathbf{1}(x = 0, \dots, 0) + \frac{1}{2}\mathbf{1}(x = 1, \dots, 1) \tag{5}$$

$$q = a\mathbf{1}(x = 0, \dots, 0) + (1 - a)\mathbf{1}(x = 1, \dots, 1) \tag{6}$$

There is a whole family of q such that the conditionals of q match p. However, we will consider p such that M does have p as a unique fixed point.

Detailed balance: Given a Markov chain $\mathcal{N}$ and fixed point p, we say it satisfies detailed balance if

$$\forall x, x',\ \mathcal{N}_{x' \leftarrow x} p(x) = \mathcal{N}_{x \leftarrow x'} p(x') \tag{7}$$

meaning the net probability flow along each edge in the Markov chain cancels out to zero.

We now show that M as defined in Eq. (3) satisfies detailed balance for fixed point q. Note that there is a transition between x and x' iff they differ at exactly one bit i. Let $x_i = 0$, $x'_i = 1$, $x_{-i} = x'_{-i}$. Then for any fixed point q,

$$M_{x' \leftarrow x} q(x) = \frac{1}{n} p(1|x_{-i}) \cdot q(0, x_{-i}) = \frac{1}{n} p(1|x_{-i}) q(0|x_{-i}) q(x_{-i}) \tag{8}$$

$$M_{x \leftarrow x'} q(x') = \frac{1}{n} p(0|x_{-i}) \cdot q(1, x_{-i}) = \frac{1}{n} p(0|x_{-i}) q(1|x_{-i}) q(x_{-i}) \tag{9}$$

so M satisfies detailed balance because any fixed point q has $q(x_i|x_{-i}) = p(x_i|x_{-i})$.

2.1 Ising model example

We now restrict the p to be the Gibbs distribution of a classical Hamiltonian, $p \propto e^{-\beta H}$ where H is of the form

$$H = \sum_{(i,j) \in E} \alpha_{ij} x_i x_j \tag{10}$$

for a graph like the one in Figure 1. Therefore,

$$p(x_1, \dots, x_n) \propto \exp\left(-\beta \sum_{(i,j) \in E} \alpha_{ij} x_i x_j\right) \tag{11}$$

which gives

$$p(x_i|x_{-i}) = p(x_i|x_{N_i}) = \frac{\exp\left(-\beta x_i \sum_{j \in N_i} \alpha_{ij} x_j\right)}{1 + \exp\left(-\beta \sum_{j \in N_i} \alpha_{ij} x_j\right)} \tag{12}$$

where N_i is the neighborhood of i as illustrated in Figure 1. This works because we can factor out the contribution from all edges in the graph. From the form of Eq. (12), we can argue that Markov chain M defined as in Eq. (3) has unique fixed point p.

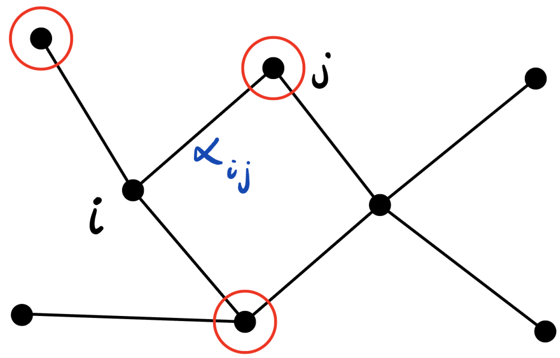

Figure 1: Illustration of a classical Ising model. The red circles indicate the neighborhood N_i of vertex i.

2.2 General Gibbs sampler

By examining Eq. (12), we can interpret the Markov chain M in a more general and physical way. It's similar to a Markov chain $M' = \sum_{i=1} M'_i$ where M'_i flips bit i with probability 1 if it decreases energy ($\Delta E < 0$) and with probability $e^{-\beta \Delta E}$ if it increases energy ($\Delta E > 0$). We have a good understanding of classical Gibbs samplers. In particular,

$$\text{Static properties of the Gibbs state} \equiv \text{Dynamic properties of } M \tag{13}$$

For example, the Gibbs sampler converges rapidly to the Gibbs state if and only if the Gibbs state is local and doesn't have long-range correlations. However, not much is known for quantum Gibbs sampling. A lot of effort has gone into defining quantum Gibbs samplers such that the analog of statement (13) holds.

3 Quantum-Classical Correspondence

We can define the pure state

$$|p\rangle = \sum_x \sqrt{p(x_1, \ldots, x_n)}\, |x_1, \ldots, x_n\rangle \tag{14}$$

Theorem 3.1. ***QC correspondence:*** *If $p(x)$ is the unique fixed point of M, then $|p\rangle$ is the unique ground state of H defined as*

$$H = -P^{-1/2} M P^{1/2} \tag{15}$$

where P is the diagonal matrix whose entries are $p(x)$.

We note that the locality and spectral gap of H match those of M because these properties are invariant under similarity transformations. However, we have to check that H is actually Hermitian. This is where detailed balance is necessary,

$$H_{x,x'} = -\sqrt{\frac{p(x')}{p(x)}} M_{x \leftarrow x'} \tag{16}$$

and we note that by detailed balance, $M_{x\leftarrow x'}p(x') = M_{x'\leftarrow x}p(x)$, which yields

$$H_{x,x'} = -\sqrt{\frac{p(x')}{p(x)}}\frac{p(x)}{p(x')}M_{x'\leftarrow x} = -\sqrt{\frac{p(x)}{p(x')}}M_{x'\leftarrow x} = H_{x',x} \tag{17}$$

so H is indeed Hermitian. It is then straightforward to see that $|p\rangle$ is the unique ground state of H,

$$\begin{aligned} H\sum_x \sqrt{p(x)}\,|x\rangle &= -P^{-1/2}MP^{1/2}\sum_x \sqrt{p(x)}\,|x\rangle \\ &= -P^{-1/2}M\sum_x p(x)\,|x\rangle \\ &= -P^{-1/2}\sum_x p(x)\,|x\rangle \\ &= -\sum_x \sqrt{p(x)}\,|x\rangle \end{aligned}$$

Since the max eigenvalue of M is 1 with unique distribution $p(x)$, the minimum eigenvalue of H is -1 with unique state $|p\rangle$. This correspondence means that understanding ground states of Hamiltonians of the form in Eq. (15) reduces to understanding the distribution $p(x)$. Often, we consider Gibbs states of classical Hamiltonians, $p = e^{-\beta H_{\text{cl}}}$. For example, if we plug in the classical 2-dimensional Ising model, we get a Hamiltonian whose ground state has long-range entanglement and the classical phase transition (parametrized by β) corresponds to a quantum phase transition.

4 Quantum Gibbs Sampling

Now we consider quantum Hamiltonians. In this case, the notion of flipping a bit and accepting with some probability becomes hard to formalize. The starting point for quantum Gibbs sampling is the Davies generator. It arises from a system S weakly interacting with a bath where the interaction can be written as

$$H_{SB} = \sum_a A_a \otimes B_a \tag{18}$$

where each A_a acts on the system and B_a acts on the bath. We interact the system with the bath and then trace out the bath assuming that it doesn't change. The state of the system ρ_S evolves according to a quantum channel,

$$\rho_S \to e^{t\mathcal{L}}[\rho_S] = \rho_S + t\mathcal{L}[\rho_S] + O(t^2) \tag{19}$$

for small t. Where $\mathcal{L}$, called the Lindbladian, can be written as

$$\mathcal{L}[X] = -i[H_S, X] - \sum_\omega \gamma(\omega)\left(\sum_a A_a(\omega)XA_a^\dagger(\omega) - \frac{1}{2}\{A_a^\dagger(\omega)A_a(\omega), X\}\right) \tag{20}$$

for

$$A_a(\omega) = \sum_{E-E'=\omega} \Pi_E A_a \Pi_{E'} \tag{21}$$

where E, E' are energy eigenvalues of H_S and the Bohr frequencies ω are the set of all energy differences in the spectrum. A quantum channel must be trace-preserving and we can verify that indeed $\mathcal{L}^\dagger[\mathbf{1}] = 0$ so $e^{i\mathcal{L}}$ is trace-preserving. It turns out that the fixed point of this channel is the Gibbs state, $\mathcal{L}[\rho_\beta] = 0$. Instead of measuring energy which would collapse the state, $A_a(\omega)$ play the role of shifting the Bohr frequencies of operators. For example, consider restricting an operator to a well defined Bohr frequency,

$$X(\omega) = \sum_E \Pi_{E+\omega} X \, \Pi_E \tag{22}$$

then

$$X(\omega)A_a(\omega') = \left(\sum_E \Pi_{E+\omega} X \, \Pi_E\right)\left(\sum_{E'} \Pi_{E'+\omega'} A_a \Pi_{E'}\right) \tag{23}$$

$$= \sum_E \Pi_{E+\omega}(XA_a)\Pi_{E-\omega'} \tag{24}$$

$$= (XA_a)(\omega + \omega'). \tag{25}$$

The jump operators at each Bohr frequency can be written as

$$A_a(\omega) \propto \int_{-\infty}^{\infty} e^{i\omega t} e^{iHt} A_a e^{-iHt} \, dt \tag{26}$$

which is nonphysical and makes it practically impossible to implement. The contribution of recent works was to make the operators quasi-local and more implementable by introducing a filter $f(t)$ that suppresses the contribution of large times t,

$$\hat{A}_a(\omega) \propto \int_{-\infty}^{\infty} f(t) e^{i\omega t} e^{iHt} A_a e^{-iHt} \, dt \tag{27}$$

We would expect that this makes the Lindbladian only approximate, but they replace the $[H_S, X]$ term with $[B, X]$ for a carefully chosen B to cancel out certain errors and make the Lindbladian exact [1].

4.1 Importance of detailed balance

In this context, the detailed balance condition is

$$\mathcal{L}^\dagger[X] = \rho_\beta^{-1/2} L[\rho_\beta^{1/2} X \rho_\beta^{1/2}] \rho_\beta^{-1/2} \tag{28}$$

and we note that taking X to be $\mathbf{1}$, detailed balance implies that the fixed point is ρ_β. For detailed balance to hold for $\mathcal{L}$ of the form in Eq. (20), detailed balance must hold at each Bohr frequency, which requires

$$\gamma(-\omega) = \gamma(\omega) e^{\beta\omega} \quad \text{for } \omega > 0 \tag{29}$$

which physically means that the channel tends to decrease energy more than it increases energy.

References

[1] Chi-Fang Chen, Michael J. Kastoryano, and András Gilyén. An efficient and exact noncommutative quantum gibbs sampler, 2025.

[2] Chi-Fang Chen, Michael J. Kastoryano, Fernando G. S. L. Brandão, and András Gilyén. Quantum thermal state preparation, 2023.

[3] Chi-Fang Chen, Hsin-Yuan Huang, John Preskill, and Leo Zhou. Local minima in quantum systems, 2023.

[4] Thiago Bergamaschi, Chi-Fang Chen, and Yunchao Liu. Quantum computational advantage with constant-temperature gibbs sampling. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, page 1063–1085. IEEE, October 2024.

[5] I. J. Crosson, D. Bacon, and K. R. Brown. Making classical ground-state spin computing fault-tolerant. *Physical Review E*, 82(3), September 2010.

[6] Matthew B. Hastings. Obstructions to nonlocal strategy and the quantum pcp conjecture. *Physical Review A*, 88(3):032331, 2013.

[7] Anurag Anshu, Nikolas P. Breuckmann, and Chinmay Nirkhe. Nlts hamiltonians from good quantum codes. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, pages 1090–1096. ACM, 2023.

[8] Matthew B. Hastings. An area law for one-dimensional quantum systems. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(08):P08024, 2007.