Quantum Gibbs-state preparation and thermal sampling

Prepare and interpret thermal mixtures while accounting for temperature, normalisation and convergence.

Research 21 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Calculate Gibbs weights for a two-level system.
- Distinguish a purification from a coherent single-system state.
- Separate a correct fixed point from rapid mixing.

Thermal equilibrium is a distribution over energies

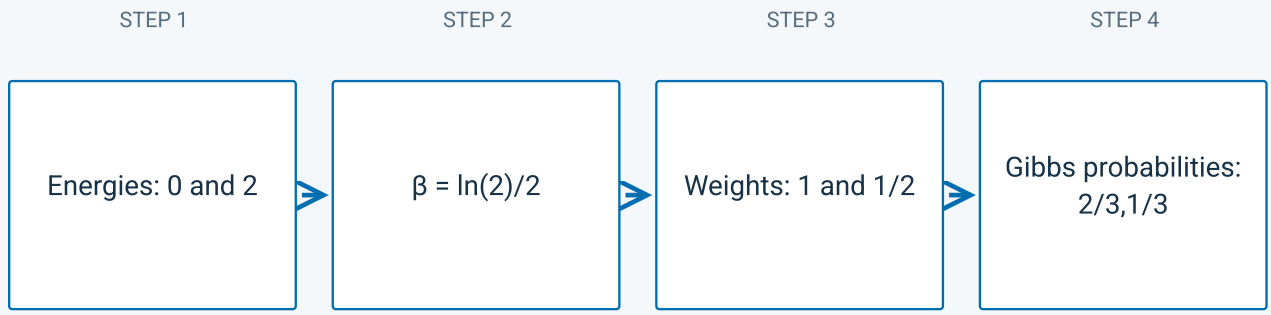
A system at thermal equilibrium generally does not occupy only its lowest-energy state. Higher-energy states also occur, with weights that decrease as energy rises relative to temperature. A Gibbs state captures that distribution. Quantum Gibbs-state preparation asks for a physical quantum state with those thermal weights, including cases where the energy eigenvectors are difficult to describe classically.

The task is broader than preparing a ground state and different from simulating time evolution. A pure energy eigenstate only acquires phase under its own Hamiltonian, while a Gibbs state is a mixture of energy eigenstates. Unitary evolution of an isolated pure state does not magically turn it into the desired mixed thermal state; ancillas, randomness, conditioning or an environment-like process must be modelled.

Review superposition versus mixture, measurement and Hamiltonian-related resource questions. The inverse temperature beta grows as the temperature decreases, in units where the relevant physical constants have been absorbed.

WORKED DIAGRAM · AA13

Normalising two thermal weights



The partition function is 3/2. Doubling β changes the excited weight to 1/4 and probabilities to 4/5 and 1/5. A thermal mixture is not the corresponding coherent one-qubit superposition.

State the target precisely

For Hamiltonian H, the Gibbs state is

$$\rho_\beta = \frac{e^{-\beta H}}{Z}, \qquad Z = \mathrm{Tr}(e^{-\beta H}).$$

Z is the partition function, the normalising sum of thermal weights. The input includes an implementable description or coherent access to H, beta and an error tolerance in a specified state distance. The output may be the mixed state itself, a purification with an auxiliary system, or samples and observables derived from it. These output formats should not be conflated.

An energy shift changes Z but leaves the normalised Gibbs state unchanged. Multiplying H by a scale changes the effective temperature unless beta is adjusted. State energy units and normalisation before comparing resource formulas, because hidden scales can make two expressions look inconsistent.

Work a two-level thermal example

Let the two energies be zero and two. Choose beta equal to the natural logarithm of two divided by two. The excited-state Boltzmann weight is then one half, while the ground-state weight is one. The partition function is one and a half. The thermal probabilities are two thirds for the ground state and one third for the excited state.

A purification is an entangled two-system state with amplitudes square root of two thirds on the joint ground label and square root of one third on the joint excited label. Ignoring the auxiliary system gives the mixed Gibbs state. The purified state is not the same as a single-system coherent superposition with those two amplitudes; the latter would retain off-diagonal coherence.

At beta equal to zero, both energy levels have equal weight and the state is maximally mixed. As beta increases, the lower energy dominates. This limit does not imply low-temperature preparation is easy. Narrow gaps, small target overlap, partition-function factors or slow mixing can make the preparation costly.

Thermal weights become probabilities only after division by the partition function.

$$\rho_\beta = \frac{e^{-\beta H}}{Z_\beta}, \quad Z_\beta = \mathrm{Tr}(e^{-\beta H}), \qquad H = \mathrm{diag}(0,2),\ \beta = \frac{\ln 2}{2} \quad \Longrightarrow \quad \rho_\beta = \mathrm{diag}(2/3, 1/3).$$

β is inverse temperature, with Boltzmann's constant absorbed into the units. The weights are one and one half; their sum $Z_\beta = 3/2$ supplies the normalisation.

An auxiliary register can purify the mixed thermal state without adding local coherence.

$$|\Psi_\beta\rangle = \sqrt{\frac{2}{3}}|00\rangle + \sqrt{\frac{1}{3}}|11\rangle, \qquad \mathrm{Tr}_{\mathrm{aux}}\,|\Psi_\beta\rangle\langle\Psi_\beta| = \rho_\beta.$$

The partial trace means ignore the auxiliary register. The remaining system is mixed; it differs from a single coherent qubit with amplitudes $\sqrt{2/3}$ and $\sqrt{1/3}$.

Several preparation routes

One approach filters a simple initial state with a response related to the exponential of negative half beta H, then uses an ancilla condition and normalisation. Another constructs a quantum channel designed to have the Gibbs state as a fixed point. Metropolis-inspired methods propose energy-changing moves and acceptance rules. Each route has different access, precision, success and convergence assumptions.

Chowdhury and Somma analyse Gibbs preparation using quantum algorithmic primitives and obtain improved parameter dependence under their stated model. Their bounds explicitly involve quantities such as temperature, normalisation and accuracy, illustrating why "prepare a thermal state" is not a parameter-free primitive. Read the Gibbs-sampling algorithm paper.

Follow a useful workflow

Define H and its energy scale, beta, the required observable and the permitted state error. Decide whether a purification is acceptable or whether only a subsystem should be returned. Choose a preparation method and identify its success or mixing parameter. Implement the coherent access or channel, estimate preparation cost and repeat as needed for the final observable's uncertainty.

If the scientific goal is an expectation such as average energy, state preparation is only one part of the computation. Repeated measurements or another estimation primitive are needed. If the observable has large variance, many fresh thermal preparations may be required. A thermal-state preparation theorem does not automatically give every thermodynamic quantity at the same cost.

Why equilibrium does not imply rapid convergence

A process can have the correct stationary state and still approach it extremely slowly. Classical Markov chains exhibit this through bottlenecks and small spectral gaps; quantum channels have analogous convergence questions. Establishing a fixed point and establishing an efficient mixing time are separate achievements. Do not infer the second from the first.

Low-temperature systems may have multiple nearly degenerate states or phase-transition behaviour. The relevant complexity can depend strongly on those features. General hard instances prevent a universal claim that every Hamiltonian's Gibbs state can be prepared efficiently at every temperature. Useful algorithms exploit promises, structure or favourable parameter regimes.

Estimate a thermal observable

In the two-level example, the average energy is zero times two thirds plus two times one third, or two thirds. This average is neither one of the only possible measured energies nor a claim that every copy has an intermediate energy. It summarises repeated observations from the thermal ensemble. After doubling beta, the mean falls to two fifths because the excited probability becomes one fifth. A computation seeking only this scalar still needs a measurement or estimation procedure after state preparation, with uncertainty determined by the observable and the number of fresh preparations.

Practice with answers

Exercise 1. If beta doubles in the two-level example, what are the new probabilities?

Check the worked answer

The excited weight becomes one quarter and the ground weight remains one. Normalisation gives ground probability four fifths and excited probability one fifth. Lower temperature suppresses the excited level more strongly.

Exercise 2. A channel has the Gibbs state as its unique fixed point. Has efficient thermal preparation been proved?

Check the worked answer

No. We also need a convergence or mixing bound, implementation cost, numerical accuracy and an initial-state analysis. A unique fixed point can be approached too slowly to give an efficient algorithm.

Check your understanding

PAUSE & THINK

What are the Gibbs probabilities after doubling β in the example?

- [ ] 1/2 and 1/2
- [ ] 4/5 and 1/5
- [ ] 2/3 and 1/3

Answer: 4/5 and 1/5. The excited weight becomes 1/4. Dividing weights 1 and 1/4 by 5/4 gives 4/5 and 1/5.

Sources

- Read the Gibbs-sampling algorithm paper