Institute for Theoretical Physics
University of Cologne

Prof. Dr. Simon Trebst
Sagar Ramchandani

Computational Many-Body Physics

Exercise Sheet 3

Winter Term 2025

Due date: Monday, **17 November** 2025

Website: thp.uni-koeln.de/trebst/Lectures/2025-CompManyBody.shtml

The goal of this week's sheet is to improve the Monte Carlo techniques, which you implemented when working on last week's sheet, towards greater numerical efficiency. The principal goal is to reduce the autocorrelation effects in the sampling of the Markov chain. This can be achieved by improving the elementary update steps by either (i) introducing better update techniques (cluster updates) or (ii) alternative ensemble approaches (extended ensembles). We will explore both routes on this exercise sheet.

Exercise 5: Simulated Annealing

One common strategy in performing Monte Carlo simulations is to include a thermalization phase, such that the Markov chain is already fully converged towards its final thermal distribution, before starting physical measurements. The amount of thermalization sweeps that are needed, however, varies between applications.

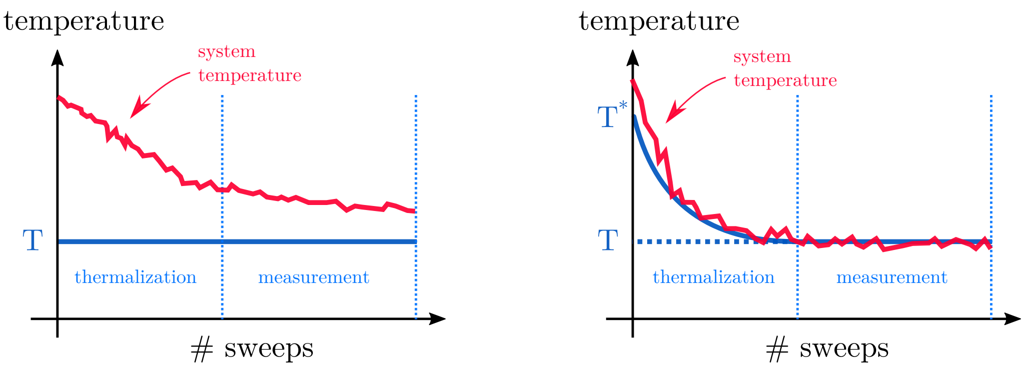

Figure 1: Comparison in temperature curves for normal thermalization (left) and simulated annealing (right). The system's temperature follows the simulated annealing curve much better as the formation of domains is suppressed through slow cooling

One example, which you have encountered on sheet 2, was the Monte Carlo sampling for the Ising model on the two-dimensional honeycomb lattice. When one performs simulations at very low temperatures ($T \lesssim 0.4\, T_c$), one often finds that the data starts to deviate from the expected results. One can track down this discrepancy by looking at the real space configurations and finding a **domain pattern** evolving. The formation of different domains prohibits the formation of a global state and introduces domain walls into the system which in return raise the energy globally. Note that for the honeycomb lattice, local spins at the domain boundaries do need a finite energy to flip, since they have more neighbors which are aligned than those which are not aligned (the local coordination is odd in distinction from e.g. the square lattice). Improving data quality in this case can be achieved by performing a thermalization procedure which suppresses the formation of domains, known as **simulated annealing**.

The new concept that simulated annealing introduces is illustrated in Figure 1. Instead of performing thermalization sweeps at the desired temperature right away, one instead cools down the system slowly from higher temperatures. This process can be compared to cooling down liquids until they start to solidify and crystallize. If cooled down too rapidly, different orientations of crystals are obtained in different places whereas a slow cooling will result in a single crystal orientation.

Your task in this exercise is to revisit your Ising Monte Carlo code for the honeycomb lattice and implement this simulated annealing approach for the thermalization. Play around with different temperature curves (algebraically decaying, exponentially decaying, modulated with sin or cos, etc.) until you are satisfied with your results and document your thermalization by plotting the temperature curve in a plot similar to Figure 1.

Note: Although in this exercise simulated annealing is only used for the Ising model, it can also be applied to other models directly. Eventually, you will find out that simulated annealing can greatly improve your data quality without much effort in many cases!

Exercise 6: Cluster updates for the Ising model

One source of autocorrelation effects in Markov chains arises from the fact that in every Monte Carlo update of a spin system, only a single spin is flipped at a time. Therefore, the vast majority of spins ($N-1$ spins) remain unchanged in a Markov chain step and two subsequent configurations are highly correlated. These autocorrelation effects persist even beyond a single sweep.

To reliably generate a novel and considerably less correlated configuration, one possibility is to flip **clusters of spins** instead of single spins. The goal of this exercise shall be for you to implement such cluster updates for your Ising Monte Carlo simulations (from the previous exercise sheet) in order to sample configurations more efficiently and reduce autocorrelation effects.

In this exercise, you should proceed as follows:

a) First, go back to the Ising Monte Carlo code from sheet 2 for a lattice of your choice (or simply the honeycomb lattice). Perform a binning analysis to determine the autocorrelation time for a fixed temperature.

b) Second, implement a cluster update of your choice, i.e. either the **Wolff** or the **Swendsen-Wang** cluster update to replace your single-spin flip update. Check the energy per spin vs. temperature curve to make sure you implemented the update correctly.

c) Measure again the autocorrelation time for the temperature of part a). Compare the two autocorrelation times and discuss the result. Consider varying the temperature and compare your two update techniques.

Exercise 7: Extended Ensembles Simulation

In this third exercise we will employ extended ensemble simulation techniques to (re)investigate the two-dimensional Ising model, now performing an "all-temperature" calculation allowing us to directly estimate the density of states $g(E)$. This technique will allow us to calculate all observables at all temperatures from a single simulation.

The exercise is structured into two parts. At first, you will apply the extended ensembles approach to your already existing Ising model Monte Carlo simulation by implementing the **Wang-Landau algorithm.** In a second part, you will then use the Wang-Landau algorithm to investigate the first-order phase transition in the **Q-state Potts model**.

a) Start the exercise by implementing the **Wang-Landau algorithm** for the 2D Ising model on the square or honeycomb lattice with periodic boundary conditions. In principle, you should be able to use and expand your single spin-flip Metropolis code that you used on the last sheet. The square lattice can either be implemented in terms of the lattice library on sheet 2, or you can use a matrix of size $L \times L$ to store the Ising variables.

Check your results by calculating estimates for the **density of states** $g(E)$ from Wang-Landau sampling for systems of linear size $L = 8, 12, 16$. Normalize the calculated density of states such that $g(E_{\min}) = 2$ and plot $\ln g(E)$ for the different system sizes to check if it agrees with your expectation.

Note: A difficulty of this part is to think about which energies can be reached and consequently construct suitable bins for the energy. Furthermore, allow to constrain your calculations to a certain energy interval $[E_{\text{min}}, E_{\text{max}}]$ to investigate only part of the configuration space (you can e.g. only use $E \in [E_{\text{min}}, 0]$ as the distribution is symmetrically around $E = 0$).

b) From the calculated estimate of the density of states calculate the following thermodynamic observables and plot them in the temperature range $T \in [0, 4]$ (showing data for all three system sizes in one plot):

- the **energy** $U(T) = \frac{1}{Z} \sum_E g(E) E \exp(-\beta E) = \langle E \rangle_T$
- the **specific heat** $C_v(T) = dU/dT = \left(\langle E^2 \rangle_T - \langle E \rangle_T^2 \right) / T^2$
- the **free energy** $F(T) = -T \log Z$
- the **entropy** $S(T) = \left(U(T) - F(T) \right) / T$

where Z is the partition function $Z = \sum_E g(E) \cdot \exp(-\beta E)$. Discuss the results in terms of numerical accuracy of your code as well as physical features such as phase transitions.

Up to this point, you have implemented the Wang-Landau algorithm for simulations of the Ising model on the square lattice. We now want to switch models to the **Q-state Potts model** which can be defined for Q-state Potts spins $\sigma_i \in \{1, 2, \ldots, Q\}$ on sites of the square lattice as

$$\mathcal{H} = -\sum_{\langle ij \rangle} \delta\left(\sigma_i, \sigma_j\right) \tag{1}$$

where the sum again runs over all nearest-neighbor bonds as in the Ising model.

c) Implement a Wang-Landau Monte Carlo simulation for the Q-state Potts model on the $L \times L$ square lattice and calculate the density of states $g(E)$ for the 10-state Potts model ($Q = 10$) in the energy range $E_{\min} = -2L^2$ to $E_{\max} = 0$ which should be normalized such that $g(E_{\min}) = Q = 10$. Check your results and discuss features of the density of states in comparison to density of states of the Ising model.

d) Plot **thermodynamic averages** of the specific heat, free energy and entropy for the Potts model as above for the Ising model. Discuss similarities and differences between the two models.

e) To investigate the phase transition in the Potts model, plot the ***canonical* distribution function** $P(E) = g(E)\exp(-\beta E)$ in proximity of the thermal phase transition, which occurs at temperature

$$T^* = \frac{1}{\ln(1+\sqrt{Q})} \tag{2}$$

for the *infinite* system ($L = \infty$). Can you observe a double-peak structure indicative of a first-order phase transition?

f) *(Optional exercise I)* From the precise location of the double-peak structure in the canonical distribution function $P(E)$ one can determine estimates of the **finite-size transition temperature** $T^*(L)$ for a given linear system size L and in principle extrapolate to the infinite system to recover $T^* = T^*(L = \infty)$.

Your task is to calculate the canonical distribution function for more large system sizes (such as $L = 8, 16, 24, 32, 48, 64, \ldots$) and plot your estimates of $T^*(L)$ versus the inverse system size $1/L^2$. Extrapolate your data to the infinite system size limit – can you recover the analytical estimate above?

g) *(Optional exercise II)* Calculate the **local diffusivity** $D(E)$ of the random walk in energy space for the Ising and Potts model simulations above by running simulations with *fixed* weights $w(E) \approx 1/g(E)$, which you have obtained from Wang-Landau sampling. To obtain an estimate for the local diffusivity

- Record two histograms during the sampling process – the energy histogram $h(E)$, which is incremented for every step, and the histogram $h^+(E)$, which is incremented only if the last extremal energy you have visited is $E_{\min}$ (and *not* $E_{\max}$).
- From these two histograms calculate the fraction $f(E) = h^+(E)/h(E)$, which estimates how much time on average the energy random walker spends at a given energy E diffusing towards higher energies. Plot this fraction $f(E)$.
- Calculate the derivative df/dE of this fraction (and plot it).
- Plot an estimate of the local diffusivity via

$$D(E) \propto \left(h(E) \cdot \frac{df}{dE} \right)^{-1} .$$

h) *(Optional exercise III)* Apply the Wang-Landau and/or optimized ensemble approach to some other classical spin models that we have discussed in the lecture such as the Ising antiferromagnet on the square or triangular lattices.