

Lab 5: Monte Carlo simulations

In this lab, we will be using a Monte Carlo code that we have written ourselves. Monte Carlo codes are usually simple enough so that you can write them yourself.

The problem: Adsorption of H on the (001) Surface of Pd

When hydrogen adsorbs onto a clean (001) surface of Pd, the H atoms sit between the Pd atoms. These possible adsorption sites are marked with an X in Figure 1. The H coverage, θ , is the fraction of possible adsorption sites that are filled. Evidence obtained with LEED [Behm et al, Surface Science **99** (1980) p320] shows that for $\theta = 0.5$ the H atoms form an ordered arrangement. We will model this adsorption problem with a square 2D lattice model that consists of all the sites where H can adsorb.

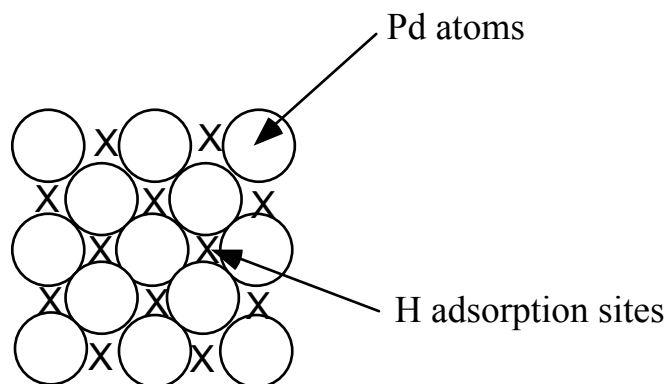



Figure 1

The energy model

As you learned in class, Monte Carlo trajectories are determined by the Metropolis algorithm. If $\Delta E < 0$ for a trial perturbation, the simulation accepts the perturbation; if $\Delta E > 0$ for a trial perturbation, the simulation accepts the perturbation with probability $P \propto \exp(-\Delta E/kT)$. All Monte Carlo simulations need an energy model to calculate ΔE . In principle, if you really wanted to, you could use a first-principles calculation to calculate ΔE for every trial perturbation. In practice, this takes way too long, and simpler energy models are required.

The energy model we will be using is called the "Ising Model". The figure below shows a two-dimensional square lattice on which a nearest-neighbor (NN) and next-nearest neighbor (NNN) interaction are defined. The occupation on each lattice site i is characterized by a spin occupation variable, σ_i , which can take the value -1 or +1. In a model for a binary alloy, +1 and -1 can respectively stand for occupation of that

lattice site by an A or B atom. In this problem, $\sigma_i = +1$ if the site is has an H adsorbed; $\sigma_i = -1$ if the site is vacant.



The Hamiltonian (energy expression) of this lattice model is given below:

$$H = \frac{1}{2} \sum_{i,j}^{NN} V_1 \sigma_i \sigma_j + \frac{1}{2} \sum_{i,j}^{NNN} V_2 \sigma_i \sigma_j - \mu \sum_i \sigma_i$$

The first summation is performed over first-nearest neighbor pairs; the second summation is performed over second-nearest neighbor pairs. The factor of 1/2 accounts for the fact that every pair is counted twice in the sum. V_1 and V_2 are the first and second neighbor interactions. μ is the chemical potential (or magnetic field): it fixes the amount of H adsorbed. At a given temperature it is linearly related to the $\log(P_{H_2})$. Note that $\theta = (1 + \langle \sigma \rangle) / 2$, where $\langle \sigma \rangle$ is the average of the spin occupation variable.

If the parameters V_1 and V_2 are chosen well, and enough neighbors are included in the model, the Ising model can do an excellent job of representing the actual energetics of the system. A first-principles tight binding study found that V_1 is positive and V_2 is negative, with $|V_2| > V_1$. We will therefore use the ratio $V_2/V_1 = -2$ to model this system.

Grand canonical simulations

Monte Carlo simulations on lattice models are typically performed in the grand-canonical ensemble (Fixed volume, temperature, and chemical potential - composition can fluctuate) rather than the canonical ensemble (fixed volume, temperature, and composition - chemical potential can fluctuate). This is largely for computational efficiency. In a canonical simulation, picking a trial perturbation is slow: for every trial spin you flip, you have to flip another one (after making sure it is the correct spin) to

preserve concentration. Both canonical and grand canonical ensembles can be shown to give the same result for thermodynamic averages.

As you may be used to thinking of composition as an independent variable, this may require you to change your thinking of how you plot thermodynamic quantities. In the Monte Carlo simulations we do not pick the concentration directly: we pick μ , and the system will find an equilibrium concentration.

Also, using the grand canonical ensemble changes the Metropolis algorithm somewhat. Instead of looking at ΔE , we use $\Delta E - \mu \sigma_i$, where σ_i is the site of the trial perturbation.

Running the code

You can input data after being prompted on the screen. It may be easier to create an input file and redirect in into the program when running. The format of the input file should then be:

V_1	V_2			
<i>number of passes for equilibration</i>		<i>number of passes for sampling</i>		
<i>start temperature</i>	<i>end temperature</i>		<i>temperature increment</i>	
<i>start chemical potential</i>	<i>end chemical potential</i>		<i>chemical potential increment</i>	

More Helpful hints:

Before you start calculations it is important to get a sense of how long it will take you to obtain a certain result. This may be part of the input you need to decide whether you actually should *do the calculations* ! Never start up a calculation unless you have some idea whether it will take minutes, hours or days to finish.

μ -range to consider. To get an idea of the chemical potentials in your system, you can either do a few quick short runs, or (better) look at the energy of some possible structures in your system and see where they would cross.

T-range to consider. To get an idea of where the transition(s) is (are), you should do a few quick runs over a coarse temperature range. Once you determine the range of interest, you can do runs with a finer temperature mesh.

How do I look for the phase transitions?

A simple way you can look for first-order transitions is to plot the chemical potential vs. spin. Run the Monte Carlo simulation at a fixed temperature, but vary μ . From thermodynamics, we know that in a 2-phase region, the chemical potential will be constant. Thus, at a fixed temperature, if we plot chemical potential vs. spin, we can find the boundaries of the two-phase region. Be careful! This applies to first-order

transitions, but you should not rely upon this method to find boundaries of second-order transitions.

In lab 4, you learned that one can find a first-order transition by looking for discontinuities in first-order derivatives of the free energy. One can find a second-order transition by looking for discontinuities in second-order derivatives of the free energy. Thus, a first-order transition will show discontinuities in E , a second-order transition will show a discontinuity in C_v .

Interpreting the results: reduced units

Since there is only one energy parameter in the model (V_1) all quantities can be scaled by V_1 . This allows us to work in *reduced units*. If we set $V_1 = 1$, chemical potential will be in units of $\mu/|V_1|$ and temperature will be in units of $k_B T/|V_1|$.

The inner workings of the code

The code sets up a 32x32 square lattice with the initial 2x2 configuration. You can set the size of the lattice with `isize` in the source. Note, phase transitions will change with lattice size!!

Then, basically the code works like this

Start loop over temperatures

Start loop over chemical potentials

A. move to next lattice site, i

B. flip spin of site i .

C. Is $\Delta E - \mu \sigma_i < 0$? If so, keep the spin that you just flipped

If not, is

$$\exp\left(\frac{-(\Delta E - \mu \sigma_i)}{k_B T}\right) < \text{random number between 0 and 1?}$$

If so, keep the spin you just flipped

If not, flip it back.

D. Calculate energy, heat capacity, spin

End loop over chemical potentials

End loop over temperatures

There are some tricks to speed up the code: these are documented in the code itself. If you have more questions about speed-up tricks, ask the instructor.

The assignment

Calculate the complete H on Pd phase diagram as a function of temperature.

1) Determine the configuration with lowest energy at 0 K for $\theta = 0.5$. Can you prove it has the lowest energy ?

2) Determine the phase diagram for this adsorbate system, i.e. determine the stable arrangement (ordered, disordered, phase separation, ...) as a function of temperature and θ . For practical reasons it is best to divide all equations by V_1 so that you get μ/V_1 , $k_B T/V_1$ etc. The temperature on your phase diagram should thus be $k_B T/V_1$.

3) Plot coverage (θ) versus $\mu/k_B T$ for a few interesting temperatures and compare with what you would obtain using Langmuir's isotherm ($\ln[\theta/(1-\theta)] = 2\mu/k_B T$). At what temperature do you expect agreement with Langmuir's isotherm ?

4) The maximum temperature at which an ordered arrangement can exist in your phase diagram will be at $\theta = 0.5$. Experimental evidence indicates that this maximum really occurs for a coverage somewhat below $\theta = 0.5$. How could you modify the Hamiltonian so that this feature is reproduced ? You do not need to perform calculations. Simply give some suggestions.

FAQ for PS 5:**1. What are the units of the code?**

In the code, V_1 is set to 1, and all energy units are normalized by V_1 . The code takes $k_B=1$. Thus, temperature is in units of V_1/k_B . Chemical potential is normalized by V_1 and also divided by spin. If this doesn't make sense, look at the exponential: $\exp\left(\frac{-\Delta E - \mu\sigma}{kT}\right)$, and see how the units cancel out.

2. What is the initial configuration of the spins at the beginning of each simulation? What are the configurations as the simulations are running?

The spins are initialized to the ordered 2x2 configuration. You can see this by looking at the file `final_config`, which outputs the final configuration – just set the steps for equilibration and sampling to 0. The file `config` will output the config at the end of each step (at the end of each temperature or chemical potential step).

3. How are chemical potential (μ), spin (σ) and coverage (θ) related.

Be careful! They are not all the same thing.

Changing chemical potential will fix the spin (but chemical potential is not the spin itself).

Coverage is related to spin by

$$\theta = (1 + \langle \sigma \rangle) / 2.$$

In our energy model, we use "spins" because of convenience. Note however, you should use coverage when presenting your data.

4. I still don't understand the chemical potential? Shouldn't there be two chemical potentials instead of one? Can I have negative chemical potentials?

The chemical potential for species α is defined as $\mu^\alpha = \frac{dG}{dn^\alpha}$, at constant T, V,

and number of atoms of the other species β, γ, δ , and so on. That is, it is the change in free energy when you add an atom of a particular type (or subtract and

atom) to the system. You should be able to see why chemical potential will control the composition of the system.

You may wonder why there are not two chemical potentials in this system, one for vacancy, and one for H. The reason for this is that in the grand canonical ensemble for a two component system (sometimes called the semigrand canonical potential), only one of these is independent. Thus, we can define the chemical potential as the difference between the two chemical potentials:

$\mu = \mu^{vac} - \mu^H$. A positive or negative chemical potential will control whether the system is thermodynamically favorable with respect to adding either adding H atoms or removing H atoms.

5. Does the Hamiltonian give the “free energy”? What kind of energy is this?

The Hamiltonian only gives the energy. It is not “free” energy without entropy terms. The Hamiltonian gives the potential energy.

Note, there are no kinetic energy terms.

6. How is heat capacity (Cv) calculated?

Heat capacity is defined as the derivative of energy with respect to temperature $\frac{dE}{dT}$. However, in statistical mechanics, one can obtain the heat capacity through

fluctuations in the energy: $C_v = \frac{\overline{E^2} - (\overline{E})^2}{kT^2}$.

7. My results are all messed up. I see no transitions whatsoever.

The most common error is to reverse the interactions. Make sure $V_1=1$ and $V_2 = -2$

8. Can I see what the ground state is by setting T=0 and $\mu=0$?

Remember you are running a Monte Carlo code. It will calculate

$\exp\left(\frac{-\Delta E - \mu\sigma}{kT}\right)$ at every trial perturbation. If T and μ are 0, this quantity is undefined, and your simulation goes beserk.

9. I don't understand phase transitions. What is a first and second order phase transition? What are examples of first and second order phase transitions? How do I mark this on my phase diagram?

Phase transitions can be observed by observing discontinuities in a function $Q(C)$, where Q is an observable and C is a control variable. If the function Q is a first-derivative of the free energy, you have a first-order phase transition. If the function Q is a second-derivative of the free energy, you have a second-order phase transition. This code outputs a few observables, Q , for you to examine. These include Energy, C_v , concentration, as well as slope of energy, C_v , and concentration curves. If you do not know which is a first-derivative or second derivative of the free energy, you should review your classical and statistical thermodynamics. Phase diagrams do not usually say whether a phase transition is first or second order, so you should just manually mark them on the phase diagram you turn in.

Examples of both:

Melting (where the underlying topology is changes a lot) is almost always a first order transition. There is a latent heat of melting, which shows up as a discontinuity of an E vs. T plot. You saw this in lab 4. Order-disorder transitions (where the underlying topology of the lattice is unchanged) are often second order phase transitions, but not always.

10. How do I know if I have enough Monte Carlo steps?

This is one of the things you can not know before testing. You can look at the final pictures or look at acceptance rates as a function of time.

In practice, for this lab, 1000-3000 equilibration steps are fine, and 2000-5000 sampling steps are fine.

11. How do I change my cell size? How do I recompile?

Find the line where it says:

```
parameter(isize=24)           !size of the lattice
```

This is line 38. Edit the number after `isize`. Here, we have chosen a 24x24 lattice. To create the new code, type

```
Armageddon$ make
```

Make sure that the "Makefile" file is in the directory that you have the source code. This creates a new file MC. You can test to make sure you have increased the cell size by checking the "out" file after you run the code.

