

Take-home Exam

Due on Thursday, December 8 (in class)

PY 502, Computational Physics, Fall 2011

Department of Physics, Boston University

Instructor: Anders Sandvik

GENERAL INSTRUCTIONS

This exam counts counts as 25% of your course grade, the homeworks counting for the other 75%.

You must complete this exam independently. You may consult any course material posted on the PY502 web site. You may also use general reference texts on mathematics and Fortran 90 programming, as well as physics texts giving background/theory for the problem. You may not consult any work that discusses details of the solution of the problem of this exam.

Your programs should be written in Fortran 90 and should be e-mailed to py502@buphy.bu.edu. Your written report with graphs should be handed to the instructor on or before the due date. For late submissions, 10% of the total score will be subtracted for each late day. Electronic (pdf) versions of the report are acceptable for late submissions.

Contact the instructor if you need further clarifications.

Good Luck!

SIMULATED ANNEALING OF A SPIN GLASS

Simulated annealing

In materials science, *annealing* refers to a process of heating, e.g., a metal and then cooling it very slowly. At elevated temperatures, crystal defects are removed by thermal motion, and subsequently a more perfect crystal structure can form if the cooling is done sufficiently slowly. If the cooling is done rapidly, which is referred to as *quenching*, there is not sufficient time for a perfect crystal structure to form and hence many defects will be "frozen in" (which some times is desirable, e.g., to harden steel).

Simulated annealing is an optimization method based on the notions of physical annealing. It is particularly useful for problems with a large number of parameters; $P = \{p_1, p_2, \dots, p_n\}$ (with n ranging from just a few to many millions). The problem is to minimize or maximize some function $E(P)$ of these parameters. This function may have a huge number of local minimas (or maximas) and it is in general very difficult to find the global optimum. In simulated annealing the parameters are treated as the degrees of freedom of a statistical-mechanics system, with a partition function

$$Z = \sum_P e^{-E(P)/\Theta}, \quad (1)$$

where $E(P)$, which clearly corresponds to the energy, is the function to be minimized (or, by changing the sign of the exponent, maximized) and Θ is a fictitious temperature. Starting at a high temperature, a Monte Carlo importance sampling of the parameters is carried out. The temperature is gradually decreased, whence the average $\langle E(P) \rangle$ decreases. If the cooling is done sufficiently slowly, $\langle E(P) \rangle$ will reach its global minimum value E_0 in the limit $\Theta \rightarrow 0$. Local minimas are the analogues of crystal defects in physical annealing. If the cooling is done rapidly, or if the simulation is carried out at a very low constant temperature (e.g., $\Theta = 0$, in which case only Monte Carlo updates lowering $E(P)$ will be accepted), the parameters will most likely settle into one of the local energy minimas from which they cannot escape because of the lack of thermal fluctuations to overcome energy barriers separating the minimas. If the cooling is done slowly, the parameters are gradually fluctuating towards lower and lower minimas and will eventually settle into the global minimum (corresponding to the perfect crystal in real annealing, although the parameter values may not correspond to any kind of order).

In practice, it may not be possible to lower Θ sufficiently slowly to guarantee that the global minimum is reached. There may not even be a unique global minimum. However, in many cases one does not really need to find the absolutely optimum solution, since local minimas with $E(P)$ not very much higher than E_0 may in practice be sufficiently good and the strictly optimal solution may be only marginally better. So, in general, simulated annealing is used to obtain almost optimal solutions. It is a very versatile and often used optimization method in science and engineering.

Spin glasses

A spin glass is a spin system in which the interactions are random and *frustrated*. Frustration refers to the inability of the interacting spins to minimize simultaneously the energy of all bonds (interacting pairs). An example is shown in Fig. 1. Here three out of the four interactions are ferromagnetic, $J_{12} = J_{23} = J_{41} = -1$, whereas one is antiferromagnetic, $J_{34} = +1$. In such a 4-spin systems with all interactions equal, all $J_{ij} = 1$ or all $+1$, the lowest-energy configurations are the




Figure 1: A frustrated 4-spin Ising system with three ferromagnetic (solid lines) and one antiferromagnetic (dashed line) interaction. Solid and open circles correspond to up and down spins, respectively. The configurations shown, and the ones obtained from them by flipping all spins, are the ones with the lowest energy; $E = -2$. One of the interactions (bonds) is always “unsatisfied”.

ones minimizing all the bond energies, $\sigma_i\sigma_j = -1$, independently of the sign of the interaction. There are two such configurations. In contrast, in the example shown in the figure, the energy is

$$E(\sigma) = -\sigma_1\sigma_2 - \sigma_2\sigma_3 + \sigma_3\sigma_4 - \sigma_4\sigma_1, \quad (2)$$

and in the states with minimum energy, $E = -2$, one of the bonds must be in the high-energy state (an “unsatisfied” interaction). There are four lowest-energy configurations; the ones shown plus the two obtained by flipping all their spins.

In an Ising spin glass with a large number of spins the number of lowest-energy configurations (ground states) grows exponentially with increasing number of spins. It is in general very difficult to find those configurations. At finite temperature, a spin glass model may exhibit a *glass transition*, below which in practice all the configurations cannot be sampled in a Monte Carlo simulation utilizing flips of individual spins. The system “gets stuck” around a local energy minimum from which it cannot escape within reasonable simulation times. There are also spin glasses in nature, and they also exhibit glass transitions. The behavior is similar to that of amorphous materials such as normal glass; thus the name “spin glass”.

Spin glass models are often studied using simulated annealing methods. One interesting aspect is to study the glass transition. Normally the transition temperature T_g is not known, at least not to very high precision, and it is then useful to start the simulation at high temperature and slowly cool it. Above the glass transition such a simulation will be able to explore the full configuration space of the system—it is said to be ergodic. However, as the transition temperature is approached, the cooling rate has to be decreased exponentially fast in order for the simulation to be ergodic. For a very large system one can in practice not achieve ergodic sampling below some temperature close to the glass transition. For ergodic sampling of a finite system at $T < T_g$, exponentially longer simulation times are required with increasing system size N . In practice, the glass transition can be seen in results obtained in several different annealing runs: For $T > T_g$ all simulations will give similar results for measured quantities, whereas for $T < T_g$ different results will be obtained in different runs of practical length. In principle the transition is sharp only in an infinite system, but in practice already modest system sizes will show glassiness in simulations of feasible length.

To identify lowest-energy states of a spin glass model is an optimization problem. Some important optimization problems in science and engineering can be mapped onto Ising and similar models, and often the interactions are frustrated. Spin glasses thus have applications also beyond statistical mechanics.




Figure 2: Cooling scheme showing the inverse temperature $\beta = 1/T$. The simulation consists of m segments of n Monte Carlo steps each (here $n = 5$), so that the total number of Monte Carlo steps is nm . The inverse temperature is increased after each Monte Carlo step; for step $i = 1, 2, \dots, nm$, $\beta = \beta_0 + i\Delta\beta$. Averages are calculated and stored for each n -step segment. Since β changes slightly within such a segment, strictly speaking the average is not quite a correct thermal average. In practice the temperatures β_1, β_2, \dots will be very close to each other ($n\Delta\beta$ is small) and this will not be an important issue.

Summary of the problem

In this assignment you will carry out simulated annealing studies of a three-dimensional Ising spin glass model, with energy function

$$E(\sigma) = \sum_{\langle i,j \rangle} J_{ij} \sigma_i \sigma_j, \quad (3)$$

where the couplings have the same magnitude, 1, but the signs are random; $J_{ij} = \pm 1$. The spins σ_i are located on a simple cubic lattice with $N = L^3$ sites. The particular random coupling sets you should use are given in files on the course web site.

You should write a program which calculates the energy and the magnetization as a function of temperature, using the cooling scheme illustrated in Fig. 2. You can use elements of the Ising Monte Carlo program on the course web site and/or the program you wrote for homework assignment 6.

You will carry out annealings to find the ground state energy for a system of 4^3 spins, and also investigate the magnetization versus temperature in several different annealing runs. You will then investigate a system with 10^3 spins in a similar way, and in addition study the energy obtained by quenching runs (keeping the temperature fixed at a low value) of different lengths.

Programming task

Write a Metropolis Monte Carlo program to simulate a 3D Ising model with $N = L^3$ spins, with the length L to be given by the user, and periodic boundary conditions. The nearest-neighbor couplings have random signs; $J_{ij} \pm 1$. The couplings are to be read from files supplied on the course web site; physics.bu.edu/py502. The files have N lines in the format

$$\mathbf{s}, J(\mathbf{s}, +\mathbf{x}), J(\mathbf{s}, -\mathbf{x}), J(\mathbf{s}, +\mathbf{y}), J(\mathbf{s}, -\mathbf{y}), J(\mathbf{s}, +\mathbf{z}), J(\mathbf{s}, -\mathbf{z}), \quad (4)$$

where \mathbf{s} is the site number, $\mathbf{s} = 1, \dots, N$, and $J(\mathbf{s}, \mathbf{a})$ is the coupling ± 1 of the spin σ_s to its neighbor in the \mathbf{a} direction. Note that there is redundant information here, as each coupling belongs to two different spins. In the program it is convenient to store the couplings for each neighbor, in the same way as in the files. The lattice is defined as the site label s given by the lattice coordinates $x, y, z = 0, \dots, L - 1$ according to

$$s = 1 + x + yL + zL^2. \quad (5)$$

The cooling scheme illustrated in Fig. 2 should be used, with the number of segments, m , and the number of steps per segment, n , given by the user. The lowest and highest inverse temperatures are β_0 and β_m , respectively. These values should be given by the user. The inverse temperature for a given Monte Carlo step i is

$$\beta = \frac{1}{T} = \beta_0 + i\Delta_\beta, \quad i = 1, \dots, nm, \quad (6)$$

where the increment for each Monte Carlo step is

$$\Delta_\beta = \frac{\beta_m - \beta_0}{nm}. \quad (7)$$

Note that the starting β value is $\beta_0 + \Delta_\beta$ and the change in β over a simulation segments is $\beta_{k+1} - \beta_k = n\Delta_\beta$. Since $\beta_{k+1} - \beta_k$ will be small, it will not matter that the temperature is not constant within a segment. One can think of this as the results for a given segment k corresponding effectively to those of a system at inverse temperature $(\beta_k + \beta_{k-1})/2$, which will not be much different from β_k . We change the temperature continually (instead of doing it in larger steps with a constant β within the segments) in order to realize a uniform slow annealing of the system.

A Monte Carlo step is defined as N attempted flips of randomly selected spins. Expectation values over each n -steps segment, $k = 1, \dots, m$, should be calculated and stored. The quantities to be computed are the energy (3) normalized by the system size; $e = E/N$, and the absolute value $|m|$ of the magnetization,

$$m = \frac{1}{N} \sum_{i=1}^N \sigma_i. \quad (8)$$

The program should carry out a number `nr` of independent annealing runs according to the same cooling scheme, with `nr` given by the user. Each run should be started with all spins $\sigma_i = 1$. All expectation values $\langle e(k, r) \rangle$ and $\langle |m(k, r)| \rangle$, $k = 1, \dots, m$, $r = 1, \dots, \text{nr}$, should be stored and written to files `e.dat` (energy) and `m.dat` (magnetization) at the end of the program. The files should also contain the averages over the independent runs;

$$\langle e(\beta_k) \rangle = \frac{1}{\text{nr}} \sum_{r=1}^{\text{nr}} \langle e(k, r) \rangle, \quad \langle |m(\beta_k)| \rangle = \frac{1}{\text{nr}} \sum_{r=1}^{\text{nr}} \langle |m(k, r)| \rangle. \quad (9)$$

The corresponding error bars $\sigma_e(\beta_k)$ and $\sigma_{|m|}(\beta_k)$, calculated in the standard way, should also be written to the files. Write the files with lines for each β_k , $k = 1, \dots, m$, in the format

$$\beta_k, A(\beta_k), \sigma_A(\beta_k), A(k, 1), A(k, 2), \dots, A(k, \text{nr}), \quad A = e, |m|. \quad (10)$$

This way it will be easy to compare and plot the results obtained in the different runs.

Note that here the different runs play the role of the bins into which we normally subdivide simulations. Since the temperature is not fixed within the consecutive segments of each run, the program will not strictly calculate thermal expectation values. However, the difference $\beta_{k+1} - \beta_k$ will normally be small, so that in practice the results will be very close to thermal expectation values. The reason for doing several independent runs is that we want to check for non-ergodicity (“glassiness”) by looking for inconsistencies (different results) between the runs.

In addition to the energy and magnetization expectation values, also keep track of the lowest energy $E_0(\beta_k)$ (total energy, not divided by N) that the system has reached up to segment k [so that $E_0(\beta_k) \leq E_0(\beta_{k-1})$]. Check for a minimum in the energy after each measurement. The minimum energies, along with their averages over the `nr` runs, should be written to a file `e0.dat` with lines in the format (10). Remember to reset your stored E_0 (e.g., to 0, which is always higher than the negative minimum energy) after each run, so that the last line of the file `e0.dat` will contain the lowest energies obtained for each independent run.

Results to be produced

Two different lattice sizes will be studied; $L = 4$ and 10, with the couplings given on the course web site; `physics.bu.edu/py502/homework5/1x.in`, with `x=4,10`. For the different lattice sizes you have the following tasks:

L=4 lattice—With $\beta_0 = 0, \beta_m = 4$ (i.e., cooling from a very high temperature $T = 1/\Delta_\beta$) and $m = 200$, carry out annealing runs for different segment lengths n , going at least to $n = 10^4$ but larger if you can (the larger your n the easier it will be to see the features asked about below). Do `nr=10` independent runs.

Check of the program: After segment $k = 5$ ($\beta_5 = 0.1$) you should have $\langle E \rangle / N \approx -0.27$.

- What is the lowest energy (total energy E , not divided by N) of this spin glass? For your largest n , plot the average energy per site, $\langle e(\beta_k) \rangle$, versus β_k and also a few of the results $\langle e(\beta_k, r) \rangle$ for the individual runs $r = 1, \dots, nr$.
- For your largest n , plot the average magnetization $\langle |m(\beta_k)| \rangle$ versus β_k and also some representative results for the individual runs. What are the signs of glassiness in the magnetization? Approximately what is the glass transition temperature T_g ? Do you see any changes in T_g for increasing segment length n (slower cooling)?
- Can you draw any conclusions regarding the uniqueness of the ground state from your results?

L=10 lattice—Again use $\beta_0 = 0, \beta_m = 4$ and $m = 200$. Carry out `nr=10` annealing runs for several different segment lengths n , starting with $n = 1$ and going at least to $n = 10^4$ and larger if you can. Also carry out quenched runs, which you can do with your program if you set $\beta_0 = \beta_m$ to a relatively high value (low temperature). Use $\beta_0 = \beta_m = 4$ and choose the same segment lengths n as in the annealing runs. You can then compare the lowest energies reached in annealing and quenching runs with the same number of Monte Carlo steps.

- Graph the average of the lowest energy reached, $\langle E_0(\beta_m) \rangle$ (with error bars), as well as the lowest of the minimum energies of the individual runs $\langle E_0(\beta_m, r) \rangle$ (i.e., the very lowest energy reached) versus the segment length n . What is the lowest energy of this spin glass, based on your results for large n ? Comment on why or why not you are confident that this really is the lowest energy.
- Make a plot similar to the one described above for the lowest energies obtained in your quenching runs. Based on the results, what can you say about the ability of a quenched simulation to reach the lowest energy?
- For your largest n , plot the average magnetization $\langle |m(\beta_k)| \rangle$ versus β_k and also a few representative

results for the individual runs. Approximately what is the glass transition temperature T_g ? Do you see any substantial changes in T_g as the cooling rate is decreased?