Exercises: Simulation Methods in Statistical Physics

Astrid S. de Wijn

February 9, 2015

1 Integrating equations of motion with Runge-Kutta

a) Show that the usual RK2 algorithm is not symplectic, by calculating the jacobian for a general Hamiltonian system and demonstrating that, in general, phase-space volume is not conserved.

b) RK2 and RK4 are described on wikipedia and can be explained using increments. That is to say, the second order version uses two points, k_1 and k_2, and the fourth order one uses k_1, k_2, k_3, and k_4. Derive another second-order RK2 algorithm than the usual one, by choosing different partial steps and Taylor expanding $x(h) = x(0) + b_1 k_1 + b_2 k_2 + O(h^3)$.

c) When you use a time-step size that is much too large, for example when you simulate the Argon fluid in the programming exercise, there will be a drift in the energy. Explain why, on average, the energy will increase. Think about the kinetic energy.

d) Why is it not recommended to integrate Hamiltonian equations of motion with a Runge-Kutta method? If you end up doing this anyway, for whatever reason (probably lazyness), what should you do to make sure that your simulations produce valid results?

e) Describe the idea behind adaptive time-step algorithms and how Runge-Kutta algorithms are particularly suitable for this.

For this second part of this question, you will need to know that the Morse potential is

$$V_{\mathrm{M}}(r) = D(1 - \exp[-a(r - r_0)])^2 \ , \tag{1}$$

where D, a and r_0 are the three parameters. The lennard-jones potential is

$$V_{\mathrm{LJ}}(r) = 4\epsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6} \right] \ . \tag{2}$$

f) Estimate the time scales of these potentials in terms of their parameters and use this to give a rough upper bound for the time step that would be reasonable to use when simulating a system with these potentials.

g) Under what circumstances would you need to use shorter time steps than the ones you have just given? Describe how to test the validity of your choice of time step in a simulation.

2 Hard spheres and resource scaling

Usage of resources is always a balance between the three resources at your disposal: your own time, the cpu time, and the memory. Consider the problem of the dynamics of a large number of hard particles on a line. Whenever two particles encounter one another, they collide instantaneously and elastically.

a) Explain why you could not use a Runge-Kutta or Verlet algorithm for this system. Is it practical to have a constant time step?

b) How many different particles can any particular particle collide with during the entire simulation?

c) Sketch a simple algorithm for integrating the dynamics of this system. Estimate how the memory and cpu loads will scale with the number of particles.

Now consider the problem of hard disk or spheres in 2 or 3 dimensions. This complicates the simulation, because the particles can now move around eachother.

c) How many different particles can any particular particle now collide with? What does this mean for your algorithm? Describe an algorithm that would work in 2 or 3 dimensions and show how the cpu and memory usage scale with the number of particles.

d) If your algorithm was $O(N^2)$ per time step in cpu usage or even slower, find an algorithm that is $O(N)$. If you already had an $O(N)$ algorithm, then describe one that is slower but uses less memory.

e) How much memory is generally available? How many cpu clock cycles? What do you think is the optimal algorithm to use? The fast one with the bigger memory footprint or the slower one with the small footprint?

e) Discuss how to speed this kind of simulation up further, at the cost of more memory and more programming and debugging time.

3 Drawing a variable from a non-uniform distribution

Random number generators usually produce homogeneously distributed random numbers. Often this is not really the distribution you want. Suppose you have at your disposal a random number generator which produces random numbers x distributed uniformly between 0 and 1.

a) Using a method similar to the one described in the lecture for Gaussian random numbers, devise a way to obtain a random number y distributed between 0 and ∞ according to the density

$$\rho(y) = \frac{1}{l} \exp\left(-\frac{y}{l}\right) \ , \tag{3}$$

where l is some positive constant. Prove that your scheme gives the correct distribution. For the Monte-Carlo simulation of the random Lorentz gas, which we have used as an example in the lecture, you would need this distribution for the free-flight times (see exercise 1.1).

b) The other distribution in the Lorentz-gas exercise was

$$\rho(\phi) = \frac{1}{2} \cos\phi, \text{with} -\frac{\pi}{2} < \phi < \frac{\pi}{2} \ . \tag{4}$$

There are other ways to get nonuniform distributions, for instance by first drawing from a uniform distribution, and then rejecting the outcome with some probability that depends on the value of the random number. Find such a way to draw from the distribution in Eq. (4).

4 Lyapunov instability, chaos, and equilibration

Because the equations of motion of physical systems are often much too complicated, we are going to play with Lyapunov instability in a very simple mathematical toy system. The coordinate x is one-dimensional and restricted to the interval $[0, 1\rangle$. The dynamics are discrete and described as a function of the descrete time t by

$$x(t+1) = 2x(t) \mod 1 \ . \tag{5}$$

a) Sketch $x(t+1)$ as a function of $x(t)$.

b) Consider some initial condition $x(0)$ and another very nearby (infinitessimally different) initial condition $\tilde{x}(0) = x(0) + \delta x(0)$. Assume first that both $x(0), \tilde{x}(0) < \frac{1}{2}$. Calculate the difference after one iteration, $\delta x(1) = \tilde{x}(1) - x(1)$. What would have happened if $x(0), \tilde{x}(0) > \frac{1}{2}$?

c) The largest Lyapunov exponent is the exponential rate at which infinitessimally close initial conditions separate, i.e.

$$\lambda = \lim_{t\to\infty} \frac{1}{t} \ln \frac{|\delta x(t)|}{|\delta x(0)|}. \tag{6}$$

What is the Lyapunov exponent of this system?

d) Now consider an ensemble of randomly chosen initial conditions. Do you think these will eventually spread out over the entire interval? Do you think the system will reach some kind of equilibrium?

e) Show that the uniform distribution of initial conditions over the $[0, 1\rangle$ will produce the same distribution in the next time step. This is the ergodic measure of the system.

5 Heat capacity (Frenkel and Smit)

(Based on question 8 on page 58 of Frenkel & Smit.)

a) Prove that the heat capacity can be calculated from fluctuations in the total energy in the canonical ensemble through,

$$C_V = \frac{\langle U^2 \rangle - \langle U \rangle^2}{k_\mathrm{B} T^2} \ . \tag{7}$$

b) In a Monte-Carlo simulation, you do not store the total energy, only the potential energy, not the kinetic energy. Show that it is still possible to calculate the heat capacity from energy fluctuations.

c) Heat capacity can also be calculated by differentiating the total energy with respect to temperature. Discuss the pros and cons of this approach over the above expressions.

6 Importance sampling in the Ising model

a) Consider, for example, the three-dimensional Ising model. How many possible states does the system have if it has $N \times N \times N$ sites? Make some simple guesses about computer speeds and the number of operations needed to calculate the terms in the sum. Give an estimate of the order of magnitude of time it would take to numerically calculate the partition function for $N = 5$ with today's computers.

$N = 5$ is nowhere near the thermodynamic limit. Obviously, for large systems, it is not a good idea to consider every state. We can get around this problem by using methods which estimate thermodynamic properties by sampling a small subset of all configurations. Suppose we attempt to estimate the partition function by randomly sampling B states, and calculating

their contribution. Let α_1 through α_k be the results of the sampling. Consider the estimator for a quantity X

$$X_{\text{estimate}} = \frac{\frac{1}{B}\sum_{k=1}^{B} X(\alpha_k)\exp[-\beta E(\alpha_k)]}{\frac{1}{B}\sum_{k=1}^{B}\exp[-\beta E(\alpha_k)]} \ , \tag{8}$$

where $E(\alpha)$ and $X(\alpha)$ are the energy and relevant quantity respectively.

b) Argue that, when B goes to infinity, the expectation value of X_{estimate} goes to the thermal average of X.

c) Recall what happens to the Ising model at low temperatures. What happens to the terms in the sums in equation (8)? How many states (compared to the total number) would you have to sample to get the results to converge to something sensible when the temperature approaches 0?

The reason for the difficulties with the estimator in equation (8) is that all states are equally likely to be sampled, even though some hardly contribute to the partition function. The variance of the terms in the sum is large. Suppose instead that we sample from the range of possible states with some probability $\rho(\alpha)$ that depends on the state α.

d) Write down an estimator for the quantity X, making sure that the expectation value of the estimate is equal to the thermodynamic average. The estimator should be of the form

$$X_{\text{estimate}} = \frac{1}{B}\sum_{k=1}^{B}\xi(\alpha_k) \ . \tag{9}$$

e) The right choice of $\rho(\alpha)$ can greatly improve the convergence of the sum. So far in this problem, $\rho(\alpha)$ has been constant. By calculating the variance of $\xi(\alpha)$, show that $\rho(\alpha) \propto \exp[-\beta E(\alpha_k)]$ is an efficient choice of probability density that works well for physical quantities.

7 Applying thermostats sensibly

When you perform MD simulations in the canonical ensemble, or where for some other reason temperature must be controlled with a thermostat, it is important to consider what sort of dynamics would be realistic.

a) Suppose you would like to know if you can turn your Argon simulation into a crystalline solid. To achieve this, you would have to cool it down very slowly, changing the temperature of the system slowly. This means you would have to apply a thermostat. Would you use the Nosé-Hoover thermostat or a stochastic thermostat like the Langevin thermostat? Why?

b) Consider another simulation of a non-equilibrium system. The setup would be similar to the figure below (taken from my own work). Two parallel rigid surfaces (blue) are made of regularly placed atoms. There are other particles (red) embedded between the two surfaces, keeping them seperated. These could be single atoms, or larger molecules (like in the figure). The interactions between the red particles and atoms in the blue surfaces is governed by some sort of effective interaction, for example Lennard-Jones. The surfaces are kept parallel and orthogonal to the z-axis. One is kept completely fixed. The other

surface is forced to slide with respect to the first. This will inject energy into the system. Without any additional measures, the system would heat up. How would a real system cool down? In the simulation a thermostat is necessary to model this. How would you thermostat this system and why?

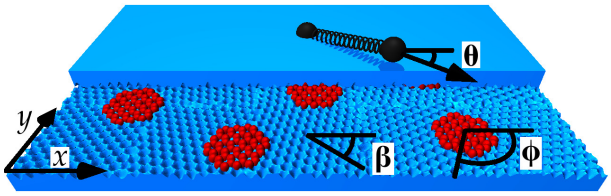

c) Consider the same system as in b), but now with several layers of particles between the two surfaces. Suppose the particles cannot move from one layer to the other. How would you thermostat this system and why?

8 Monte-Carlo trial moves

(Based on question 12 on pages 135 and 136 of Frenkel & Smit.)

a) If one uses a small displacement for displacement trial moves, the fraction of accepted moves goes up. Why is this not efficient?

b) In a simulaiton of a molecule that consists of more than one interaction site, a trial move that rotates the molecule around its center of mass is usually included. Why?

c) Which type of trial move (displacement, insertion/deletion, change of volume) is the most expensive computationally? Why?