

Worksheet 2: Statistical Mechanics and Molecular Dynamics

Olaf Lenz Florian Dommert Peter Kořovan

November 17, 2012

Institute for Computational Physics, University of Stuttgart

Contents

1	General Remarks	1
2	Statistical Mechanics	2
3	Molecular Dynamics: Lennard-Jones Fluid	3
3.1	Lennard-Jones Potential	3
3.2	Reduced Units	4
3.3	Lennard-Jones Billards	5
3.4	Periodic Boundary Conditions	6
3.5	Lennard-Jones Fluid	9

1 General Remarks

- **This worksheet is still incomplete! The complete worksheet will be there on Friday evening at the latest.**
- Deadline for the report is **Tuesday, 27th November 2012, 10:00**.
- On this worksheet, you can achieve a maximum of 20 points.
- The report should be written as though it would be read by a fellow student who attends to the lecture, but does not do the tutorials.
- To hand in your report, send it to your tutor via email
 - Olaf (olenz@icp.uni-stuttgart.de; Thursday, 14:00 – 15:30)
 - Jens (smiatek@icp.uni-stuttgart.de; Friday, 8:00 – 9:30)

- Please attach the report to the email. For the report itself, please use the PDF format (we will *not* accept MS Word DOC files!). Include graphs and images into the report.
- If the task is to write a program, please attach the source code of the program, so that we can test it ourselves.
- The report should be 5–10 pages long. We recommend to use L^AT_EX. A good template for a report is available.
- The worksheets are to be solved in groups of two or three people. We will not accept hand-in-exercises that only have a single name on it.

2 Statistical Mechanics

Task (2 points)

- Consider a system A that consists of subsystems A_1 and A_2 , for which $\Omega_1 = 10^{20}$ and $\Omega_2 = 10^{22}$. What is the number of configurations available to the combined system? Also, compute the entropies S , S_1 and S_2 .
- By what factor does the number of available configurations increase when 10 m^3 of air at 1 atm and 300 K is allowed to expand by 0.001% at constant temperature?
- By what factor does the number of available configurations increase when 150 kJ is added to a system containing 2.0 mol of particles at constant volume and $T = 300 \text{ K}$?

(Stolen from: Frenkel, Smit: *Understanding Molecular Simulation*)

Task (2 points)

(Thermodynamic Variables in the Canonical Ensemble) Starting with an expression for the Helmholtz free energy F as a function of N , V , T

$$F = -\frac{\ln [Q(N, V, T)]}{\beta} \quad (1)$$

one can derive all thermodynamic properties. Show this by deriving equations for U , p , and S .

(Stolen from: Frenkel, Smit: *Understanding Molecular Simulation*)

Task

(2 points)

(Ideal Gas) The canonical partition function of an ideal gas consisting of monoatomic particles is equal to

$$Q(N, V, T) = \frac{1}{h^{3N} N!} \int d\Gamma \exp[-\beta H] = \frac{V^N}{\lambda^{3N} N!} \quad (2)$$

in which $\lambda = h/\sqrt{2\pi m/\beta}$ and $d\Gamma = dq_1 \dots dq_N dp_1 \dots dp_N$. Derive expressions for the following thermodynamic properties:

- $F(N, V, T)$ (hint: $\ln(N!) \approx N \ln(N) - N$)
- $p(N, V, T)$ (which leads to the ideal gas law)

(Stolen from: Frenkel, Smit: *Understanding Molecular Simulation*)

3 Molecular Dynamics: Lennard-Jones Fluid

On this worksheet, you will do Molecular Dynamics simulations as they are used to simulate molecular systems. The system that is to be simulated is a fluid of particles at different densities that interact via the Lennard-Jones potential.

All files that are required for this tutorial can be found in the archive `lj.tar.gz` that can be downloaded from the lecture's homepage.

After you have implemented the Lennard-Jones potential, you will start with a program very similar to the program that simulated the solar system from the last worksheet and successively improve the program until it is a fully-fledged and fast Molecular Dynamics simulation program. Note that for the people that have visited the lecture "Physik auf dem Computer" this will be the first time where you will actually need the C programming language, as pure Python is too slow for the task. However, note also that there is no good reason to abandon the convenience of Python altogether!

3.1 Lennard-Jones Potential

We are going to simulate a system consisting of particles interacting with the Lennard-Jones (LJ) potential, which has been originally derived for a system of noble gas atoms (see Fig. 1 on the following page):

$$V_{\text{LJ}}(r) = 4\epsilon \left(\left(\frac{\sigma}{r} \right)^{12} - \left(\frac{\sigma}{r} \right)^6 \right) \quad (3)$$

On very short distances of the order of the atom radius, the Pauli principle disallows the atoms to overlap, which results in a strongly repulsive potential that should be $\propto e^{-r}$. However, the actual functional form of the core does not really influence the results as long as it is strong enough, therefore the computationally simpler term $\propto r^{-12}$ is used.

On intermediate distances, the interaction is slightly attractive, which corresponds to the van-der-Waals-interaction of induced dipoles. It can be shown that this interaction should be $\propto -r^{-6}$, which is the second term in Equation 3. On long distances, the interaction is negligible.

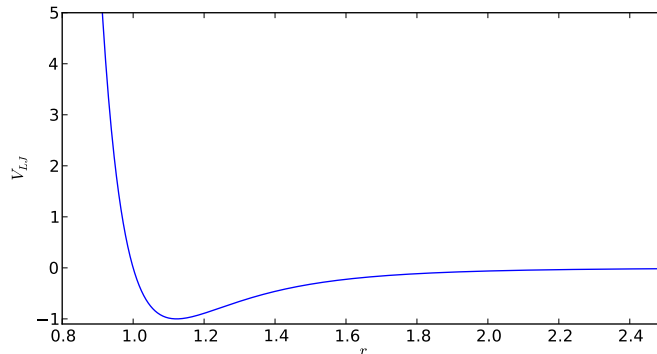


Figure 1: The (12,6)-Lennard-Jones potential ($\sigma = 1$, $\varepsilon = 1$)

3.2 Reduced Units

For simple liquids composed of point-like particles which interact with the same pairwise-additive potential of a general form:

$$U(r) = \varepsilon\phi(\sigma/r), \quad (4)$$

where ε and σ are constants and ϕ is a smooth and differentiable function, it is possible to define a set of dimensionless reduced units such as $x^* = x/\sigma$, $V^* = V/\varepsilon$, $T^* = k_{\text{B}}T/\varepsilon$ and $P^* = P\sigma^3/\varepsilon$. Within this set of reduced units, all systems interacting with potentials of the form given by equation (4) follow one single universal equation of state. This law is called the *theorem of corresponding states*. Therefore, results of simulations of simple fluids are usually given in these reduced units. One particular implication of the theorem of corresponding states is that we can perform all simulations with $\varepsilon = 1.0$ and $\sigma = 1.0$ and if desired, the results can be transferred to other values of interaction parameters. For computational precision, it is desirable to choose the reduced units in the natural units of the problem under consideration, which makes all measured quantities on the order of unity.

Therefore, in all the following tasks, we will assume that $\varepsilon = 1$ and $\sigma = 1$.

Task

(2 points)

- Implement a Python function `compute_lj_potential(rij)` that returns the Lennard-Jones potential (for $\varepsilon = 1$ and $\sigma = 1$) for two particles with the distance vector `rij` (which is a 3d-NumPy-array).
- Implement a Python function `compute_lj_force(rij)` that returns the Lennard-Jones force (as a 3D-Numpy-array) for two particles with the distance vector `rij`.
- Create a distance vector $(d, 0, 0)$ and let d take on 1000 values in the interval $[0.85, 2.5]$. Compute the LJ force and potential for the particles and create a plot of the potential and the first component of the force against d .

3.3 Lennard-Jones Billards

A large part of the Python program `ljbillards.py` from the archive is very similar to the sample solution of `solar_system.py` from the last worksheet. It is almost ready to play 3d-Lennard-Jones billards. It sets up five LJ particles (“billards balls”) at certain positions and with certain velocities and simulates them for 20 time units. The only thing that is missing are the functions from the previous task.

The program creates two plots: A projection of the trajectories onto the x - y -plane (remember: the system is 3d now!) and a plot of the total energy of the system over time. The latter plot can be used to verify that the LJ potential and forces are correct. If they are correct, the energy will be constant, except a few low-amplitude wiggles when the balls hit each other.

Furthermore, the program creates a file `ljbillards.vtf` in the VTF-format. Open the file in a text editor so that you can see what it contains. The file can also be viewed with the 3D-visualization program VMD¹. To load the file into VMD, either execute

```
vmd ljbillards.vtf
```

in the shell, or load it via the in-program dialog **VMD Main** → **File** → **New Molecule...**

When you load the file into VMD, the LJ particles are displayed as points which are almost invisible on the screen. To display the particles as colored spheres, open the dialog **VMD Main** → **Graphics** → **Representations...** and choose **Draw Style** → **Drawing Method** → **VDW** and **Draw Style** → **Coloring Method** → **Index**.

VMD loads the whole trajectory of the system. The slider in the **VMD Main** window can be used to watch the system evolve in time.

¹<http://www.ks.uiuc.edu/Research/vmd/>

To create a high-quality image of the currently visible scene, open the dialog **VMD Main** → **File** → **Render**. Choose **Tachyon** (**internal, in-memory rendering**) and press the button **Start Rendering**. This will create an image `vmscene.png` that contains a raytraced image of the scene.

Task

(2 points)

- Add the Python functions written in the previous task to the program `ljbillards.py`.
- Let the program run and visualize the system with VMD.
- Change the position of the fifth particle (index 4) so that it is hit by the third particle (index 2).
- Although the particles almost behave like billiard balls, some of them do not. Which particles and why? Create a plot of the trajectory or a series of images that supports your hypothesis.

3.4 Periodic Boundary Conditions

Our goal on this worksheet is the simulation of a LJ fluid to determine macroscopic properties of the bulk of the fluid. In practice, however, even the largest modern computers are not able to simulate systems with a macroscopic size with $\approx 10^{23}$ particles, but use systems that are significantly smaller.

One of the problems of fluid simulations with a finite number of particles is how to handle the boundaries of the system. If we would use *open boundary conditions* (PBC) where the available space is infinite, the particles would simply escape into empty space. Therefore, we somehow have to confine the particles in a finite volume. We could create walls on the boundaries that push back the particles when they try to escape. However, this creates boundary effects, and the properties of the fluid close to the walls differ greatly from the bulk properties that we are actually interested in.

The common solution to this problem is to use *periodic boundary conditions*, where the system with a finite size is (virtually) multiplied infinitely into all directions. A particle on the right boundary of the system interacts with the particles close to the left boundary of the system as though there would be a copy of the whole system (a *periodic image* of the system). The same goes for all directions. An alternative view on the situation is that the system lives on a three-dimensional torus. A particle that leaves the system to the right will turn up on the left. This principle is well-known from various computer games.

The advantage of periodic boundary conditions is, that the system is virtually infinite - the environment of the system is the system itself, so there are no boundary effects.

Note that by introducing the periodic boundary conditions, we have also introduced a new parameter of the system, namely the system size L , and we can define the volume of the system $V = L^3$, and the density $\rho = \frac{N}{V}$, where N is the number of particles in the system.

To visualize a system with PBC, VMD has different functions. These function can be accessed via the VMD console that you can open by choosing **VMD Main** \rightarrow **Extensions** \rightarrow **Tk Console**. This console allows you to type the following commands:

pbx box which will draw a box at the boundaries of the system.

pbx wrap -all which will wrap all coordinates of the system into the central image. This is only needed when the coordinates can also be outside the central image in the first place.

However, first of all, VMD needs to know about the size of the system. This can be easily done in a VTF file, by adding a line so that it looks like the following:

```
atom 0:1 radius 0.5
pbx 10.0 10.0 10.0
timestep
3.88000009704 2.98000008821 3.97900008821
6.11999990296 5.01999991179 6.01999991179
.
.
```

Long- and Short-Range Interactions

How can we perform simulations with PBC? The trouble is, that there are an infinite number of copies of each particle. If the range of the interactions is not finite, computing the force acting on a single particle becomes very tricky. Such interactions with an infinite range are also called *long-range interactions*, examples for this type of interaction are the Coulomb interaction or gravitation. In fact, there are methods to handle this kind of problem, which you will learn about in the second part of the lecture.

Fortunately, most interactions have a finite range (so-called *short-range interactions*), or at least they converge to zero fast enough so that we can introduce an approximation and truncate the potential at a certain distance (the *maximal interaction range*, or *cutoff*). In that case, interactions with particles that are further away are simply neglected.

In the case of the Lennard-Jones potential, we can define the *truncated LJ potential* as follows:

$$V'_{LJ}(r) = \begin{cases} V_{LJ}(r) - V_{LJ}(r_{\text{cutoff}}) & r \leq r_{\text{cutoff}} \\ 0 & r > r_{\text{cutoff}} \end{cases}. \quad (5)$$

Note that we also shift the potential slightly by $V_{\text{LJ}}(r_{\text{cutoff}})$, so that the potential is continuous at the cutoff. For $r < r_{\text{cutoff}}$ the force due to the truncated LJ potential is the same as due to the full form, otherwise the force is 0. In practice, $r_{\text{cutoff}} = 2.5\sigma$ is often used, which we will also do in the rest of this worksheet.

Minimum Image Convention

Now that we know that we use an interaction with a finite range, we can think of how to compute the forces on the particles. The so-called *minimum image convention* states, that when the maximal interaction range is less than the system size L , to compute the force that a particle B has on particle A, you only have to compute the interaction with a single periodic image of particle B, the *minimum image* that is closest to the particle A. This minimum image of particle B is not necessarily in the same image of the system that particle A is in.

Knowing this, it becomes significantly easier to implement periodic boundary conditions. In particular, the code of the LJ billiards from the previous task has to be changed only at a few places to implement periodic boundary conditions. There are several possibilities to implement PBC.

Note that it is *not* necessary to keep the particles in the central image, by folding them back when they leave the central image. Instead, one can also simply let the particles propagate outside of the central image and take into account the PBC only when computing the interactions. Doing so has a few advantages when it comes to computing observables. For visualization, coordinates that are folded back into the central image are more useful. Therefore, VMD supports the function `pbw wrap` that folds all coordinates into the central image.

Task

(3 points)

- Modify the LJ potential and force such that it implements the truncated LJ potential (with $r_{\text{cutoff}} = 2.5$).
- Extend the program `ljbilliards.py` such that it takes into account periodic boundary conditions.
- What functions need to be modified to implement them?
- To test the PBC, set up two particles in a system with side length $L = 10.0$, with initial positions $\mathbf{x}_0 = (3.9, 3, 4)$ and $\mathbf{x}_1 = (6.1, 5, 6)$ and initial velocities $\mathbf{v}_0 = (-2, -2, -2)$ and $\mathbf{v}_1 = (2, 2, 2)$. Simulate them for 20 time units. With PBC, the particles should collide. Plot the trajectory.

Hints The following Python functions might be useful to implement periodic boundary conditions: `numpy.fmod()`, `numpy rint()` and the modulo operator `%`.

3.5 Lennard-Jones Fluid

Now we are ready to simulate a LJ fluid!

Pure Python

First, let's implement the simulation in pure Python. The Python program `ljfluid.py` from the archive is an almost complete simulation program that should be very similar to your LJ billiards program with PBC.

To simulate a LJ fluid, we just have to increase the number of particles and make the system more dense. When the number of particles becomes large, it is no fun to set up the system manually by defining their coordinates. On the other hand, if we put in the particles randomly, the chances are good that some of the particles overlap (in particular at high densities), which would lead to extremely high forces that would explode the system. Therefore, in our case we simply set up the particles on a cubic lattice with n particles per side.

The parameters of this program therefore are not the number of particles N and the system size L , but the number of particles per side of the cubic lattice n and the density ρ .

To measure the timing of a program, use the Unix command `time` as follows:

```
> time python ljfluid.py
.
.
.
real    0m2.866s
user    0m2.778s
sys     0m0.057s
```

The interesting time in our case is the user time.

Task

(2 points)

- Extend the program `ljfluid.py` to set up the particles on a cubic lattice with a given number of particles per side n at a given density ρ .
- Perform simulations of the LJ fluid with $n = \{3, 4, 5\}$ particles per side of the cubic lattice (time step $\Delta t = 0.01$, end time $t_{\max} = 1.0$, density $\rho = 0.7$) and measure their timings.

Hints To measure the timings, it is a good idea to comment out the code that plots the energy of the system!

Pure C/C++

The performance of the simulations in the previous example is not very satisfactory. With these timings, we will not be able to do simulations of several thousands of particles.

Therefore, let's try to implement the simulation in the programming language in C. The program `ljfluid.cpp` is a C program (that uses a few C++ things to simplify it) that does pretty much the same as the Python code in the previous task. However, it is missing the function to measure the total energy of the system.

Study the code! Now that we have already seen the Python code, implementing the program in C is not very hard anymore. Note, however, that writing and debugging the code has actually taken the author several hours, even with the Python program as a prototype.

To compile it, execute

```
g++ -O3 -o ljfluid ljfluid.cpp
```

Once it has successfully compiled, you can execute the program as follows

```
./ljfluid
```

Task

(2 points)

- Given the Python code from the previous task, implement the function `compute_energy()` in the C-program `ljfluid.cpp` and let the program output the total energy of the system.
- Run the program and measure the timing for $n = 5$. Compare the timing with the pure Python program from the previous task.
- Run the program and measure timings for $n = \{5, 6, 7, 8, 9, 10, 11, 12\}$. Plot the timing t against the number of particles N in the system.
- What is the scaling behaviour of the program?

Mixing C/C++ and Python: Cython

Now, we will mix the C-code and the Python-code to see that it is possible to gain the benefits of both worlds.

The subdirectory `cython1` in the archive contains some files that are necessary for this.

`c_lj.cpp` contains some C-code from the previous task, namely the time-critical functions of the LJ simulation. They are copied from `ljfluid.cpp` almost unmodified. The

file `lj.pyx` is written in the language CYTHON² and provides the “glue” between C and Python. `setup.py` contains a (very short) python script that helps to compile `lj.pyx` and `c_lj.cpp` into the Python module `lj`. To compile the files, execute the following command

```
> python setup.py build_ext --inplace
```

When you have done this, you have generated a Python module `lj.so` that can now be loaded into a Python program. `ljfluid.py` loads this module and executes the simulation, calling the C-functions from Python. Compare this program to the pure Python program and look for the differences!

Task

(2 points)

- Copy the function `compute_energy()` from the previous task into `c_lj.cpp`, uncomment the corresponding lines in `lj.pyx` and let `ljfluid.py` measure the energy.
- Run the program and measure the timing for $n = 5$ and $n = 10$. Compare the timings with the pure C program.

Neighborhood Lists

In the task before the previous task, we have noticed that even though we have gained a lot of speed by implementing the core functions in C, the scaling behavior is very unfavourable and would not allow for large systems to be simulated. To do that, it is necessary to modify the actual algorithm. How the algorithm can be modified is out of the scope for this worksheet (we will discuss it in the tutorials), still we provide a program (in the subdirectory `cython2`) that contains the modifications and introduces so-called cell lists and Verlet lists.

The program files are the same as in the previous task, the program can be compiled and run as before. Only a single new function has been added that is called in the main loop of the program `ljfluid.py`, and the function to compute the forces has been slightly modified. Still, the program shows a significantly better performance.

Task

(1 point)

Run the program and measure the timings for $n = \{5, 10, 15, 20\}$.

²<http://cython.org>