Technische Universität München
Lehrstuhl für Informatik V
Univ.-Prof. Dr. M. Bader
M.Sc.(hons.) Kaveh Rahnema

SS 2014
Sheet 9
07.07.2014

Scientific Computing II

Molecular Dynamics

Exercise 13: Multi-Centered Molecules

For single-centered molecules, the force on molecule i equals the sum of all forces between molecule i and all other molecules: $\vec{F}_i = \sum_{j \neq i} \vec{F}_{ij}$.

Using the force, the acceleration of molecule i is given by the following formula:

$$\ddot{\vec{x}}_i = \frac{\vec{F}_i}{m_i} = \frac{\sum_{j \neq i} \vec{F}_{ij}}{m_i}$$

For multi-centered molecules, there are some more values to be considered to be able to represent rotations:

- values considered for single-centered molecules: force $\vec{F}$, mass m, acceleration $\ddot{\vec{x}}$.
- values only to be considered for multi-centered molecules: torque T, moment of inertia I, angular acceleration $\ddot{\omega}$.

Find the formula for the angular acceleration that is analogue to the formular for the acceleration $\ddot{\vec{x}}$.

Solution:

Consider two non-spherical molecules: i with 2 interaction sites and another molecule j with 3 interaction sites. Apart from translation, we also have to consider rotational motion now.

- Compute force between molecule i and molecule j:

  $$F_{ij} = F_{i1,j1} + F_{i1,j2} + F_{i1,j3} + F_{i2,j1} + F_{i2,j2} + F_{i2,j3},$$

  i.e. the sum over all site-site pairs.
- Compute the total force on molecule i: $F_i = \sum_{i \neq j} F_{ij}$
- The force on each site results in a torque τ:

  $$\tau = \sum_{s \in \text{sites } i} d_s \times F_s,$$

  where d_s denotes the distance of site s from the center of mass, and F_s the force on center s.

- The relationship between angular acceleration $\frac{\partial \omega}{\partial t}$ and torque τ is $\tau = I \times \frac{\partial \omega}{\partial t}$).
- Thus, angular acceleration $\frac{\partial \omega}{\partial t} = I^{-1} \times \tau$.

Exercise 14: Linked Cells algorithm

Assume we simulate a molecular dynamics scenario with N molecules. If we explicitely compute the forces between all pairs of molecules, O(N^2) operations are necessary. For short-range potentials, we neglect interactions between particles that have a mutual distance bigger than a certain cut-off radius. This reduces the number of required operations to O(N).

(a) Someone tries to convince you that this is not true:

Assume that a simululation with N molecules requires C operations for the force calculations. Double the number of molecules in the domain. Then we of course have to compute forces for twice as many molecules. But, in addition, the number of molecules within the cut-off radius of a certain molecule doubles, too. Thus, we need 4C operations, which means that the number of operations behaves like O(N^2).

Why is that argumentation wrong?

(b) In 3D, apart from the cell itself, 26 additional cells have to be examined (8 cells in 2D). By reducing the size of the cells, more cells have to be used, but the volume covered by these cells is smaller. As the covered volume corresponds to the number of distance calculations, smaller cell sizes can increase the performance. Calculate the covered volume for $l = rc$, $l = \frac{rc}{2}$, $l = \frac{rc}{4}$ and $l \to 0$ in 2D and 3D.

	$l = rc$	$l = \frac{rc}{2}$	$l = \frac{rc}{4}$	...	$L \to 0$
$2D$					
(% unnecessary)					
$3D$					
(% unnecessary)					

Solution:

(a) The fault is that with doubling the number of molecules in a domain the scenario changes completely, as we now deal with matter of double the density. Thus doubling the problem size doesn't correspond to simply doubling the number of molecules, but rather to doubling the size of the domain, with constant density. In that case the algorithm really needs only twice the number of calculations.

(b) a) $l = rc$: In this case, we have to consider only the directly neighbouring cells:

	2D: 9 cells	3D: 27 cells
A_{cells}	$9rc^2$	$27rc^3$
$A_{\text{in cutoff}}$	πrc^2	$\frac{4}{3}\pi rc^3$
$A_{\text{in cutoff}}$ / A_{cells}	$\frac{\pi}{9} \approx 35\%$	$\frac{4\pi}{81} \approx 16\%$
$\Rightarrow$ unneccessary:	$\approx 65\%$	$\approx 84\%$

b) $l = \frac{rc}{2}$: In this case, we have to consider two neighbouring cells in each direction, thus five cells per dimension:

	2D: 25 cells	3D: 125 cells
A_{cells}	$25(\frac{rc}{2})^2 = \frac{25}{4}rc^2$	$\frac{125}{8}rc^3$
$A_{\text{in cutoff}}$	πrc^2	$\frac{4}{3}\pi rc^3$
$A_{\text{in cutoff}}$ / A_{cells}	$\frac{4\pi}{25} \approx 50\%$	$\frac{32\pi}{375} \approx 27\%$
$\Rightarrow$ unneccessary:	$\approx 50\%$	$\approx 73\%$

c) $l = \frac{rc}{4}$: Here we have to take 9 cells in each dimension into consideration. But now we can get rid of some cells. In 2d, e.g. the smallest distance possible between a particle in the middle cell to a particle in a corner cell is $\sqrt{2(\frac{3}{4}rc)^2} = 1.06rc$. Thus we don't have to search the corner cells for particles within the cutoff-radius.

In 3d, searching all cells along the edges is unneccessary. Moreover we can save 4 more cells per surface of the cube.

	2D: 77 cells	3D: 613 cells
A_{cells}	$77(\frac{rc}{4})^2 = \frac{77}{16}rc^2$	$\frac{613}{64}rc^3$
$A_{\text{in cutoff}}$	πrc^2	$\frac{4}{3}\pi rc^3$
$A_{\text{in cutoff}}$ / A_{cells}	$\frac{16\pi}{77} \approx 65\%$	$\frac{256\pi}{3 \cdot 613} \approx 44\%$
$\Rightarrow$ unneccessary:	$\approx 35\%$	$\approx 56\%$

Here is an overview of all the results:

	$l = rc$	$l = \frac{rc}{2}$	$l = \frac{rc}{4}$	...	$L \to 0$
$2D$	9	6.25	4.81		3.14
(% unneccessary)	65%	50%	35%		0%
$3D$	27	15.63	9.58		4.19
(% unneccessary)	84%	73%	56%		0%

Exercise 15: Parallel Linked Cells algorithm

Now consider the simulation of a nucleation process. The simulation is started with a million molecules in gas phase, so the distribution of the molecules in the domain is homogenous. Over time, small droplets will evolve and the distribution of the molecules will become increasingly inhomogenous.

(a) For the simulation we will use the parallel linked cell algorithm. Describe a scalable implementation of the algorithm.

(b) Initially, every computer is assigned an equally sized part of the domain. Then, during the run of the simulation, a load imbalance will occur. What are suitable criteria for repartitioning?

Solution:

(a)
- The subdomain of each process is surrounded by a layer of halo-cells.

- Communicate boundary particles along spatial dimensions, e.g. first in x-, then in y-, last in z-direction. In 3D, each process has then only 6 instead of 26 communication partners.
- Use Non-blocking, overlapping MPI Send/Receive to allow for maximum parallelism:
  - Start receive operation for both neighbors along spatial dimension
  - Start send operation
  - Wait until all operations finished

(b)
- load imbalance per process (i.e. if load imbalance is small, no repartitioning is needed).
- Load per cell (in order to find partitions with equal load).
- Communication cost required by a given partitioning (e.g. if you cut through a droplet, many particles need to be communicated).