

GROMACS - Bug #362

MD problem when use "nolist=-1" option with tabulated potential

10/23/2009 07:53 AM - Lanyuan Lu

Status: Closed	
Priority: Normal	
Assignee: Erik Lindahl	
Category: mdrun	
Target version: 4.0	
Affected version - extra info:	Difficulty: uncategorized
Affected version:	
Description	
Created an attachment (id=398) Input files	
<p>To whom it may concern,</p> <p>I reported a strange MD problem some days ago in the gmx-user list, which happens when I tried to use "nolist=-1" with tabulated potentials. Sometimes, a simulation crashes or has "inf" energy in log file or gives unusual RDFs on certain machines, but I can run the same input files on another machine with reasonable output. I attached my original email at the end. However, after the email I investigated the problem a little bit more and found it was more complicated than I had thought. I did tests on three machines with a number of different MD systems. The machines are:</p> <p>A: Intel i7 gcc 4.1.2 (a local desktop). B: Intel Xeon icc 10.1 (a local cluster) C: AMD Barcelona icc 10.1 (Texas Advanced Computing Center, "Ranger" cluster).</p> <p>I compiled the Gromacs on A/B and the package on C was compiled by TACC. The Gromacs version is 4.0.5.</p> <p>In my last email, I said I got trouble with A/B, but fine with C. However, after testing more systems, I found sometimes I got trouble on C, and fine with A/B. It seems to me on which machine to get crash/wrong trajectory is totally random. Typically, if I use "nolist=1", I can always get the correct trajectory. Here "trouble" means immediate crash/inf energy in log file/obviously wrong structure, and the choice of the three seems to be also random. Interestingly, sometimes I resubmit my job on machine C and previous dead job can pass. Another important factor is the buffer size. Some buffer size works for some machines, but not others.</p> <p>I attached a simple example to demonstrate the problem. I have a box with 1000 particles and in top file I use the epsilon and sigma of Oxygen. In mdp file the LJ cutoff is 1.0 and rlist is 1.3. The table is just a copy of the table12-6.xvg in Gromacs. I tried to run 1000 steps NVT with dt 0.002 and T 298. On A/B the simulation can finish and RDF is correct. But I tried C and it gave inf energy. Because of the random nature of C mentioned above, I submitted the same job twice, both with inf log as attached. It's also strange that I could not use rlist=1.2, which gives a 0.2 buffer. In this case it crashes on all of A/B/C.</p> <p>Thanks very much for your help.</p> <p>Regards, Lanyuan Lu</p> <p>Hi,</p> <p>Some months ago I reported a strange conflict between the GMX 4 new option nolist=-1 and tabulated potentials. At that time I contacted Berk and he could run my test system without any problem. So I thought I had something wrong for my system configurations. However, recently I re-checked the problem and found I got similar issues from two different software/hardware configurations. It seems the problem is sort of general. I'm wondering if any developer or people who are familiar with the code can help me to figure out the problem.</p> <p>The problem is like the follows. When I use nolist=-1 option with tabulated nonbond interaction, sometimes I got a segmentation fault. I used nonebond tables with space 0.004 nm and I smoothed the curves so no complain of wrong numerical derivatives from GMX were found. I tested the simulation on three machines:</p> <p>A: Intel i7 gcc 4.1.2 B: Intel Xeon icc 10.1 C: AMD Barcelona icc 10.1</p> <p>All tests were done using one CPU core. It seems system A/B always gave me troubles and C was fine.</p> <p>I first used rlist =1.3 and rvdw=1.2 to give a 0.1 buffer. And I found the potential energies from A/B were weird with values sometimes as 1.0e29. And the potential energy with C is always around 1.0e3. When I set rlist=1.28, I can optimize the "Average neighborlist lifetime" about 11, as Berk once suggested. But in this buffer size, simulations with A/B got segmentation fault immediately. Interestingly, I tried to try many conventional force field simulations/CG simulations with tabulated potentials before with A/B, and all of them seemed to be fine if I didn't use nolist=-1 with tabulated potentials.</p> <p>Can anyone give me any suggestions?</p> <p>Thanks, Lanyuan</p>	

History

#1 - 10/30/2009 05:49 PM - Berk Hess

Isn't your time-step simply too large?
Use seem to be using LJ particles with the LJ parameters of water, but with a mass of 1. A time step of 2 fs seems too large for stably integrating such a system. A rough estimate would be that this is equivalent to a $2 \cdot \sqrt{16/1} = 8$ fs in water, which would become unstable (this mass of oxygen is 16).

Berk

#2 - 11/03/2009 10:53 PM - Lanyuan Lu

Dear Berk,
Isn't the mass just 18.0 as that I defined in cg.top?
Thanks,
Lanyuan

#3 - 11/04/2009 11:55 AM - Berk Hess

Indeed, your mass is correct.

Looking once again at your mdp file, I noticed that you set table-extension to zero. This is the source of the problem, it should be large enough to account for diffusion, as the manual states.

Also note that rvdw has no effect with user vdwtype (unless it is larger than rlist).

Berk

#4 - 11/04/2009 11:58 PM - Lanyuan Lu

Dear Berk,
You're right. The problem is due to table-extension.
What did you mean by "Also note that rvdw has no effect with user vdwtype (unless it is larger than rlist)"? Aren't rvdw/rlist the short/long cutoffs in the Verlet list, for nblast=-1?
BTW, in the manual, it says that nblast=-1 is not sufficient for large systems, since sometimes only a few atoms passed the buffer and the update is required. I think the method in LAMMPS can be applied to avoid this problem. In LAMMPS, you can set the neighbor list check only performed every n steps.
Thanks very much and I appreciate your kind help.
Lanyuan

#5 - 11/05/2009 09:38 AM - Berk Hess

The cut-off only affects the functional form of the interactions.
With a user table the functional form is completely dictated by the table and the cut-off has no effect (but rlist does).

nblast=-1 is always "sufficient", it might just be somewhat more costly than have a fixed nblast. Checking only every n steps defeats the purpose of nblast=-1, then you might as well set it to a fixed value.

Berk

Files

input.tar	270 KB	10/23/2009	Lanyuan Lu
-----------	--------	------------	------------