GROMACS - Bug #124

Distribution of instantaneous kinetic temperature generated by Nose-Hoover appears to be incorrect

01/07/2007 12:55 AM - John D.

Status:	Closed		
Priority:	High		
Assignee:	David van der Spoel		
Category:	mdrun		
Target version:	3.3.1		
Affected version - extra info:		**Difficulty:**	uncategorized
Affected version:			

Description

Hey guys!

Brief description: I compared all of the thermostats in gromacs 3.3.1 on a 500-particle LJ system to see whether they were self-consistent in generating correct distributions of the "instantaneous kinetic temperature", whose average should be the desired thermostat temperature, but the distribution of which is possible to compute analytically from the Maxwell-Boltzmann distribution. The only thermostat that appeared to pass this test is the Langevin integrator. Nose-Hoover appears to be broken in some way.

Detailed description of tests available at:

http://www.dillgroup.ucsf.edu/~jchodera/group/wiki/index.php/
Gromacs_3.3.1_thermostat_validation_tests

Please contact me for more information, or if you have trouble accessing the page, graphs, or code: John D. Chodera <jchodera@gmail.com>

Cheers!

History

#1 - 04/14/2007 08:46 PM - David van der Spoel

John,

I had a look at your tests, could you please upload a version that does not require any funky python or matlab to repdroduce it?

#2 - 04/14/2007 10:56 PM - John D.

(In reply to comment #1)

> John,
>
> I had a look at your tests, could you please upload a version that does not require any funky python or matlab to repdroduce it?

Hi David!

Do you mean you want a test case that doesn't use any Python imports (like numarray) or Matlab? Or do you mean something that doesn't use Python at all?

Thanks,

- John

#3 - 04/15/2007 08:56 AM - David van der Spoel

Python is OK, but I don't feel much like installing numarray for this test only (or matlab for that matter).

Is your analytical T distribution for 500 particles or for 1?

#4 - 04/15/2007 09:14 AM - John D.

(In reply to comment #3)

> Python is OK, but I don't feel much like installing numarray for this test only
> (or matlab for that matter).

Hi David,

Fair enough! Do you have another numerical python library installed, like numpy? Or are you hoping for pure Python here?

I will try to modify the Python files to spit out an .xvg file with the observed and analytical distributions, since you can then plot this on your own. Does that sound reasonable?

> Is your analytical T distribution for 500 particles or for 1?

The distribution of the instantaneous kinetic temperature is computed using the number of degrees of freedom as a parameter. The number of degrees of freedom is, of course, 3*(number of atoms) - (number of constraints). For this test, a box of 500 LJ spheres, the number of degrees of freedom used is 500*3 - 3, since I remove the COM translational momentum every timestep.

Also, as the system gets larger, the distribution of instantaneous kinetic temperature will get narrower, of course.

Thanks!

- John

#5 - 04/15/2007 09:29 AM - David van der Spoel

OK,

since we have a reference distribution now we only have to do the simulations
and then run g_energy and g_analyze -distr to get histograms. It doesn't take
more than a small C-shell or Perl (python) script to produce these.

#6 - 07/14/2007 02:42 PM - David van der Spoel

John,

any progress on a simpler test set?

#7 - 09/07/2007 11:46 AM - Berk Hess

I decided to do a test myself.
I used a 400 atom LJ system.
Both SD and Nose-Hoover give the correct chi2 distribution.
I tried Nose-Hoover with tau=2 and tau=10.
If anything tau=10 gives a slighlty too broad distribution,
but for be sure I would need to simulate much longer than
the 1 nanosecond I currently did.

My suspicion is that your switched potential switches too fast:
a switching range of 0.1 sigma is extremely short. This will
probably lead to inaccurate integration with dt=2 fs.

Also I would always use a shifted potential, since that minimizing
the switching effects.

My settings are (with LJ from SPC oxygen):
nstlist = 20
rlist = 0.9
vdwtype=shift
rvdw=0.7
rvdw_switch=0.5

Berk.

#8 - 09/07/2007 12:10 PM - Berk Hess

You used Berendsen pressure coupling!!!

When I turned that on I got the peak.

A Berendsen thermo- and barostat do not provide a proper ensemble.
For most quantities this does not really matter,
but for temperature and pressure fluctuations it does.

Berk.

#9 - 10/04/2007 09:53 AM - John D.

(In reply to comment #8)

> You used Berendsen pressure coupling!!!
>
> When I turned that on I got the peak.
>
> A Berendsen thermo- and barostat do not provide a proper ensemble.
> For most quantities this does not really matter,
> but for temperature and pressure fluctuations it does.
>
> Berk.

Hi Berk,

Apologies for not responding earlier, but I've only had a chance to dig this up again and look at it.

No barostat was employed in the test set I reference and put online in the wiki link above. I used constant-volume simulations of LJ systems at a fixed density. No pressure coupling parameters were specified in the input files. The 'mdout.mdp' file that was generated says "Pcoupl = no". Unless I am misunderstanding, this means that "Berendsen pressure coupling" was not used.

One of the tests included a test of the Berendsen weak-coupling thermostat for temperature regulation. It is well-known to produce an ensemble somewhere in between the isokinetic ensemble (as tau -> 0) and the microcanonical ensemble (as tau -> infinity), behavior that was easy to see in my plots. As tau changed, the kinetic energy distribution changed drastically.

Provided there are no harmonic resonances (see reference [1] below) and the simulations are much longer than the coupling tiem, the Nose-Hoover thermostat should not have this same tau-dependent behavior. Unfortunately, my tests observed tau-dependent behavior of the kinetic energy histogram. I will test the suggestion above that this may be due to the short distance over which the cutoff was switched to zero (over a range of 0.1 sigma) to see if this abolishes the artifact.

Cheers,

- John

[1] B. L. Holian, A. F. Voter, and R. Ravelo. Thermostatted molecular dynamics: How to avoid the Toda demon hidden in Nosé-Hoover dynamics. Phys. Rev. E, 52(3):2338, 1995.

#10 - 10/04/2007 10:20 AM - Berk Hess

I looked again at the file I downloaded from your site
and now realized that the print in your python script
for the pressure coupling is commented out.

But I get perfect Ekin distribution with Berendsen pressure
coupling for my test LJ system.
Whereas with Berendsen pressure coupling I get exactly
the errors in the distribution that you show.

It would be nice if you could recheck stuff, before I try
to reproduce your results.

I think the switching function does not have a large influence
on the results.

Berk.

#11 - 10/04/2007 10:31 AM - John D.

(In reply to comment #10)

> I looked again at the file I downloaded from your site
> and now realized that the print in your python script
> for the pressure coupling is commented out.

Right -- I was eventually going to test the barostats too, but never got past testing the thermostats.

> But I get perfect Ekin distribution with Berendsen pressure
> coupling for my test LJ system.
> Whereas with Berendsen pressure coupling I get exactly
> the errors in the distribution that you show.

I'm sorry -- you said "Berendsen pressure coupling" twice. In which case (barostat or thermostat) did you get 'perfect Ekin distribution', and in which case did you get the tau-dependent behavior that I show in my data?

> It would be nice if you could recheck stuff, before I try
> to reproduce your results.

What precisely would you like me to re-check? I've adjusted the parameters to switch the LJ and Coulomb terms off smoothly between 4 sigma and 5 sigma now, which should certainly be smooth enough to avoid integration errors.

> I think the switching function does not have a large influence
> on the results.

I didn't think it would have, but I was happy to test if since you suggested you thought it might be a problem. :)

Cheers,

- John

#12 - 10/04/2007 10:41 AM - Berk Hess

Oops.
I meant without Berendsen pressure coupling, so NVT, I get correct results.
I did SD and Nose-Hoover with tau=2 and 10 ps.
All three Ekin distributions agree within the noise with the analytical
curve.

With Berendsen pressure and Nose-Hoover temperature coupling I get exactly
the tau dependent error in Ekin you show.

But as I understand, you already rechecked and confirm that you see
the Ekin errors without pressure coupling.

Berk.

#13 - 10/04/2007 10:44 AM - Berk Hess

Could you attach a top, gro and mdp file for say the tau=10 ps case
where it goes very wrong, so I can easily check your results?

Berk.

#14 - 10/17/2007 10:06 AM - John D.

Created an attachment (id=251)
Histogram of instantaneous kinetic temperatures for Langevin integrator and Nose-Hoover thermostat for 500-atom Lennard-Jones system.

This is the latest results for the 500-atom Lennard-Jones system using 10 ns of simulation (with the first 5 ns discarded to equilibration). I will attach the input files and analysis scripts used to generate this in a subsequent attachment.

The only feature of note is that the behavior of Nose-Hoover at tau = 10 ps seems to significantly deviate from the expected distribution. I'm not certain if this is a 'bug' or some odd sort of resonance artifact.

These results were produced with gromacs 3.3.1.

I didn't have Matlab handy, so there are no error bars on this plot.

#15 - 10/17/2007 10:13 AM - John D.

Created an attachment (id=252)
Test setup for the 500 particle LJ thermostat test.

Berk,

Sorry for taking so long to respond to your last email -- I've been distracted with other projects and hadn't had time to double-check my latest test

results from the cluster.

Here is the archive of my test before generating any data. In the data/ subdirectory are separate subdirectories for each test. You can 'tcsh run.sh' batch script to generate the data. The data/nose-hoover-10.00ps/ directory is the one that is behaving suspiciously according to the energy histograms.

There is a script in the base directory called 'process_energies.py' that compiles histograms from the .edr files, though you can apparently do this with your own gromacs tools as well.

I didn't include any of the actual data I had generated with gromacs 3.3.1 because it ended up being something like 600 MB of data -- too much to post here. The test doesn't take long to generate 20 ns of data for the argon system, though -- just a few hours.

Cheers,

- John

#16 - 10/17/2007 05:03 PM - Berk Hess

Sorry, but all this python stuff is too complicated for me
(and slightly tricky).
To be really sure we are running the same simulations,
could you just attach the top, gro and mdp file used
to generate the tpr for the problematic simulation?

Thanks,

Berk.

#17 - 10/17/2007 07:24 PM - John D.

Hi Berk,

It's in there! In data/, you can find the .mdp, .gro, and .top files for each test in the test subdirectories, as well as a two-line shell script to create the .tpr and run it with mdrun.

The especially problematic one is

data/nose-hoover-10.00ps

Apologies if that was unclear from my last email.

- John

#18 - 10/18/2007 10:26 AM - Berk Hess

I have reproduced your problem.
I now also realize what the difference between your and my system is.
I have a liquid, whereas you have a dilute gas.

I think you have an ergodicity problem.
Your Nose-Hoover oscillation time is 10 ps,
whereas the autocorrelation time of the velocities is longer: 12 ps.
I think this will lead to ergodicity issues.
I don't know if this means the Ekin distribution could be incorrect
or that you need to simulate extremely long to get the right distribution.

Anyhow, I think this is a general Nose-Hoover ergodicity issue,
not a Gromacs problem.

Berk.

#19 - 10/18/2007 03:29 PM - Berk Hess

I should rephrase my previous comments.
I guess one should not compare the velocity autocorrelation time
and the Nose-Hoover tau.

But I still think the problem is caused by ergodicity issues
due to the very slow energy transfer in the system, due to particles
only feeling each other on average on times of 10 ps.

I have now simulated 100 ns and the kinetic energy distribution
is far from converged. This is most obvious from the envelope
of the kinetic energy profile as a function of time.
Only twice this envelope reaches +-30 K. Thus one needs at least

1000 ns and maybe much longer to obtain convergence.
I don't know if it would finally reach the proper distribution or not.

Berk.

#20 - 10/20/2007 08:11 AM - John D.

Berk,

Thanks for tracking down the cause of the deviation! I'd agree that a weakly-interacting dilute gas could possess very slow timescales for thermalization of kinetic energy, and could lead to the observed behavior with a single Nose-Hoover coordinate (rather than a Nose-Hoover chain or stochastic thermostat like Langevin or Andersen).

I'm rather surprised that the system corresponds to a dilute gas, because I thought I had picked a temperature and density in the liquid part of the phase diagram. My numbers originally came from the NIST site set up to validate Monte Carlo codes (see detailed information from my README below). I had to add a switching function to get reasonable energy conservation (which was not done in the Monte Carlo test system) but it is possible I messed up a conversion and ended up with a system in the gas phase rather than the liquid phase, or even worse, along the gas-liquid phase coexistence curve. I'll try to figure out what happened.

Thanks,

- John

(FROM README)

DESCRIPTION

To validate a thermostat, we want to ensure that it samples from the correct thermodynamic ensemble. The simplest way to check this is to compare the potential and kinetic energy distributions to either an analytically-computable distribution or a reference simulation from a code that is known to be correct. This test is a simple internal consistency check, where the distributions generated by the various thermostat options available within gromacs are compare to each other.

The test system consists of a box of 500 LJ molecules at reduced temperature T* = kB T / epsilon = 0.85 and reduced density rho* = rho sigma^3 = 0.009, in the
liquid phase. The radius sigma and well depth epsilon were chosen to be the same as the oxygen in TIP3P water (sigma = 1.7682 A, epsilon = 0.1521 kcal/mol) with a mass of 18 amu so as to yield timescales that roughly mimic systems of the sort we simulate. A switching function was employed to smoothly switch the LJ interactions to zero over 3 to 3.1 sigma. The volume was held constant. Center of mass velocity was removed every timestep, and the pairlists were rebuilt every timestep.

10 ns simulations were run under NVT conditions for each of the thermostats in gromacs, discarding the initial 1 ns to equilibration. Several choices of tau_t (10 ps, 1 ps, 0.1 ps) were used for each thermostat, but Brownian dynamics would hang for 10 ps and 1 ps, so this data is not present in the plots below. A 2 fs timestep was used, energies written every 0.1 ps, and the single-precision version of mdrun was used to run the simulations. The potential energy and instantaneous kinetic temperature ranges sampled by all simulations were each divided into 40 histograms, and the histogram values and uncertainties computed and plotted.

The original idea was to compare to data computed at NIST here, but they use straight cutoffs, which would cause energy conservation problems:

See http://www.cstl.nist.gov/srs/LJ_PURE/mc.htm

> But I still think the problem is caused by ergodicity issues
> due to the very slow energy transfer in the system, due to particles
> only feeling each other on average on times of 10 ps.
>
> I have now simulated 100 ns and the kinetic energy distribution
> is far from converged. This is most obvious from the envelope
> of the kinetic energy profile as a function of time.
> Only twice this envelope reaches +-30 K. Thus one needs at least
> 1000 ns and maybe much longer to obtain convergence.
> I don't know if it would finally reach the proper distribution or not.
>
> Berk.

#21 - 10/20/2007 08:32 AM - David van der Spoel

John,

the coordinates you have given are probably intended to be in angstrom. So you probably want scale them down by a factor of 10.

#22 - 10/22/2007 10:48 AM - Berk Hess

I now have 1000 ns with tau=10 ps for your dilute gas
and Ekin distribution looks close to what it should be.
So there does not seem to be a problem with Nose-Hoover in Gromacs.

Berk.

Files

temperature_histogram.pdf	4.2 KB	10/17/2007	John D.
LJ-500-particles-thermostat-test.tgz	74.2 KB	10/17/2007	John D.