

GROMACS - Bug #1648

Atoms with zero LJ parameters but partial charges appear to cause crashes in free energy calculations when perturbed in the presence of other charges

11/26/2014 02:14 AM - David Mobley

Status:	Feedback wanted		
Priority:	Normal		
Assignee:	Michael Shirts		
Category:	core library		
Target version:			
Affected version - extra info:	5.0.4	Difficulty:	uncategorized
Affected version:	4.6.x		

Description

I've run several systems now with free energy calculations where I run into crashes whenever I perturb a molecule having partial charges on some atom which has zero LJ interactions and this atom is in the vicinity of other atoms having strong electrostatic interactions (as GAFF specifies zero LJ parameters for the hydroxyl hydrogen, this is not an unusual situation for me). The most recent example of this is in the calculation of solvation free energies of benzoic acid in the presence of other benzoic acid molecules, but it also happens when studying binding of polar molecules to a host containing a variety of carboxylic acids. I have never been able to find why this should cause crashes, nor have I been able to eliminate issues by reducing the timestep or other protocol modifications. The only thing I've tried which seems to avoid the issue is to make the hydroxyl hydrogen NOT have zero LJ parameters, but this isn't really a fix, since the force field specifies zero LJ parameters.

After correspondence on the mailing list, I suspect the problem is related to what Berk wrote on the developers list concerning how soft core is applied:

"A pair interaction is soft-cored when at least one of the two atoms is perturbed and at least one of the atoms has $c12=0$ at $\lambda=0$ or $\lambda=1$. So for atoms with LJ, interactions are only soft-cored when strictly necessary. For atoms without LJ, e.g. many hydrogens in many force fields, all perturbed interactions are soft-cored. For such atoms soft-coring is necessary when connected to a heavy atom that is decoupled. So the only cases where some interactions are soft-cored that are not strictly necessary is in a system where some atoms are decoupled and others are modified, then interactions with hydrogens connected to modified atoms would not need to be soft-cored."

I believe this is the issue in my case. Specifically, I am turning off the solute (carboxylic acid) and so its perturbed interactions with the REMAINING, non-perturbed solutes are soft-cored, meaning that charge sites can overlap. At least, that's my reading of what Berk said and it would perfectly explain the crashes I'm seeing.

Berk's response to this on list is that this should not apply in my situation, since sc-coul should ensure that the atoms in question also do not have partial charges (by virtue of soft-cored coulomb interactions) when their charge sites are overlapping. However, in my case, I have sc-coul set to its default value of zero, which means this (I think) does not avoid the problem.

Anyway - I'm not totally sure I understand what's going on here. However, what I DO know is that I have several systems (I am attaching my benzoic acid test system) that I can't get to run stably in free energy calculations when there are zero LJ parameters on a hydroxyl hydrogen, but which DO run stably in (a) conventional calculations (not free energy), or (b) when I put nonzero LJ parameters on the hydrogen. I strongly suspect the issue is related to how soft-core atoms are selected/handled.

In my benzoic acid test system, the first production simulation (prod.0) at $\lambda=0$ runs fine but the others crash at roughly progressively shorter and shorter times as λ increases. I'm assuming this relates to their likelihood of overlapping with another benzoic acid molecule or something similar.

Let me know if you want me to try repeating with sc-coul = yes to see if that avoids the issue (though this would not be a "fix" per se, since my approach should also be valid); I only recently became aware of this option.

History

#1 - 11/26/2014 02:19 AM - Michael Shirts

David Mobley wrote:

I've run several systems now with free energy calculations where I run into crashes whenever I perturb a molecule having partial charges on some atom

Do you have a log (and possibly .tpr) file available from one of the runs that crashes? Seeing the git hash (4.6.2-beta2 is a little vague) as well as knowing what the program itself would be useful to start diagnosing.

#2 - 11/26/2014 02:20 AM - Michael Shirts

Also, how long does it take to crash? Are there specific lambdas which crash?

#3 - 11/26/2014 05:24 AM - David Mobley

- File prod.4.tpr added

- File prod.4.log added

All lambdas crash except #0, though the amount of time is variable. All of them make it into production except one. Those which crashed soonest in this particular incident are #4 and #8 (the latter crashed in equilibration).

I'm fairly certain this crash also happens in 4.6.3 and 4.6.4 - I believe I had the student repeat in actual versioned GROMACS and he had the same issues. I certainly ran into the same issue in my host-guest system in 4.6.2. Can you remind me though how I would retrieve the git hash?

I could test it in a more recent GROMACS version, and I could also try removing everything but the benzoic acid molecules (i.e. removing all of the solvent). But this would take me a few days at least - the recent upgrade of my cluster was a substantial one and we basically have to recompile/reinstall everything, and I have to figure out how to use the new queueing system, etc. Let me know.

I'm uploading a log and tpr, for one which crashed after about 60 ps of production. Naively I would expect this to be one of those "sporadic crash" issues where if you run a job which crashes on a different architecture it might run longer, etc. Though we've run this benzoic acid system in several different solvents and always the vast majority of the runs crash regardless of what we do -- though it can take a couple nanoseconds before it happens.

Thanks.

#4 - 11/26/2014 09:38 AM - Berk Hess

If I understood your mdp setup correctly, you first decouple Coulomb and then LJ. At lambda index 1 you should then have full LJ, right? So that should not crash due to any (non)soft-core issue, since you only have weakened charges and the original, full LJ interactions that prevent overlap. Or do I misinterpret the lambda state settings Michael implemented?

#5 - 11/26/2014 05:30 PM - Michael Shirts

David Mobley wrote:

Can you remind me though how I would retrieve the git hash?

It's now printed in the logfile (since 4.6)

All lambdas crash except #0, though the amount of time is variable. All of them make it into production except one. Those which crashed soonest in this particular incident are #4 and #8 (the latter crashed in equilibration).

OK, testing things out now. Haven't gotten a crash yet with 5.0, but only 100 ps or so.

I'm uploading a log and tpr, for one which crashed after about 60 ps of production. Naively I would expect this to be one of those "sporadic crash" issues where if you run a job which crashes on a different architecture it might run longer, etc. Though we've run this benzoic acid system in several different solvents and always the vast majority of the runs crash regardless of what we do -- though it can take a couple nanoseconds before it happens.

Thanks, this helps understand what the computer thinks the internal settings are -- though I don't see anything off at first glance.

#6 - 11/26/2014 06:33 PM - David Mobley

Berk - yes, I should have full LJ, which is why it's so peculiar it's crashing and that addition of LJ parameters to the hydrogen atom eliminates crashes. I am NOT actually 100% sure that this particular possible issue is the cause of the crashes in my case, I just know that I've now run into a variety of different instances where having zero LJ parameters on a hydroxyl hydrogen in a free energy calculation seems to cause a crash. And based on what you said about how soft-core atoms are being selected, it seemed as though that COULD be the explanation.

Michael - the git hash is 8d37c3e0da5ad2925a456d6c4200ecbf6c01e5d3, though since you have the log file you presumably had that now.

I'll work on getting mdrun going on our cluster again so I can do some more testing myself.

This all came up because a collaborator asked about how soft-core atoms are selected in GROMACS, and the answer suddenly made me think it could be an explanation for these puzzling crashes I've had on my list of things to investigate.

#7 - 11/26/2014 06:48 PM - Michael Shirts

Thanks! I'm going to dig a little more into this with the debugger now, since I've run out to 300 ps with 4.6.5 and not gotten any crashes -- so replicating will be hard todo. I'll look at exactly what the interactions are at different lambda for atoms with no LJ (as opposed to what we think they are).

#8 - 02/02/2015 07:05 PM - David Mobley

I have verified all of these issues still persist for me both in the git has referenced above, and with the released version of GROMACS 4.6.7. I'm seeing roughly 50% of jobs failing by 2.5 ns, though most make it at least a few hundred picoseconds before crashing. This is on a fresh compilation of 4.6.7 on our new cluster setup.

The failures do seem sporadic, in that I don't see the **same** run necessarily always crash again and again, nor always at the same spot. These are typical "large force" type crashes - i.e. lots of LINCS warnings (including possibly "table limit" warnings) followed by a segmentation fault. Though I should note that in previous testing we tried removing bond constraints and reducing the timestep and still had crashes, just preceded by slightly different symptoms.

Again, there is no reason I'm aware of that there should be abnormally large forces in this case unless there are issues with the implementation of soft core potentials that are specific to this type of case. The ONLY TIME we have run into these issues in all of our work on solvation is when we have hydroxyl groups (with no LJ parameters on the hydroxyl hydrogen) in close proximity. In this particular case we've got more than one benzoic acid molecule in solution so we think this is happening when the hydroxyl groups approach one another (hence, at least in part, the sporadic nature of the issue).

Please let me know if you have specific things it would be helpful for me to look at to help troubleshooting. Otherwise plan is to next:

- Test if the problem persists in GROMACS 5
- Re-verify that adding LJ parameters to the hydrogen avoids the problem (my student did this but I was going to check for myself as well)

If it would help I could also try and strip the system down to something more minimal which would reproduce the problem - i.e. maybe I could reproduce it for just a dimer in the gas phase or something similar...?

#9 - 02/09/2015 06:35 PM - David Mobley

- *File bug.tar.gz added*

I have now reproduced this problem in 5.0.4 with a somewhat simpler protocol. In this case I'm running without constraints with a 1 fs timesteps and my runs typically crash with an error something like:

WARNING: Listed nonbonded interaction between particles 1 and 10 at distance 2.407 which is larger than the table limit 2.200 nm.

which I assume is just standard "blowing up". It's worth noting that now, many of my runs are crashing about halfway through (100 ps) NPT equilibration. The failure rate is now perhaps slightly lower - I'm seeing approximately five of my 20 runs crash.

I'm uploading the input files. Next I will attempt to re-verify that the issue goes away if I add LJ parameters to the hydroxyl hydrogen (though as noted before, this shouldn't matter).

Thanks.

#10 - 02/11/2015 12:20 AM - Mark Abraham

- *Affected version - extra info set to 5.0.4*

David Mobley noted at [#1684](#) that the issues is also present in 5.0.4. Thanks!

#11 - 02/11/2015 07:47 PM - David Mobley

I have reproduced my student's work and verified that the issue disappears if the hydroxyl hydrogen LJ parameters are set to be nonzero (in 5.0.4; my student checked this in 4.6.4 with the same conclusion).

So, it seems safe to say that soft core is working incorrectly in this case and resulting in overlap of hydroxyl hydrogens with other atoms (probably the nearby oxygens) when they ought not to be able to overlap - that is to say, when they are still carrying partial charges. (In this protocol, we turn the partial charges linearly to zero before beginning to modify LJ parameters, so it should never be possible for the hydroxyl hydrogens to overlap with other atoms while they still carry a charge).

Please let me know if there are any other tests we can do which would be helpful.

#12 - 02/15/2015 02:42 AM - Michael Shirts

I'll start debugging, but if you know which ones blow up (or if it's totally stochastic) that could help.

which I assume is just standard "blowing up". It's worth noting that now, many of my runs are crashing about halfway through (100 ps) NPT equilibration. The failure rate is now perhaps slightly lower - I'm seeing approximately five of my 20 runs crash.

#13 - 02/15/2015 02:45 AM - Michael Shirts

I have a .tpr that causes the problem (it's the .tpr from [#1648](#), which is where I gave up debugging this when I couldn't reproduce the result from run to run) , but I want to get an example running that fails without free_energy on, as I think that will be easier to debug.

#22 - 06/25/2015 06:07 AM - Michael Shirts

- File *runfail_in.gro* added
- File *runfail.top* added
- File *runfail_nonzero.top* added
- File *runfail.mdp* added

OK, I have some files that reproduce the problem here.

With *runfail_nonzero.top*, it should run to completion (~5 min). With *runfail.top*, it should fail in the first 1000 steps or so. Basically, if the acid H's have no LJ terms, it fails - if there is any LJ interaction, it works.

Run with -nt 1 for best reproducibility.

Files

benzoic_acid.tar.gz	127 KB	11/26/2014	David Mobley
prod.4.tpr	216 KB	11/26/2014	David Mobley
prod.4.log	227 KB	11/26/2014	David Mobley
bug.tar.gz	120 KB	02/09/2015	David Mobley
runone.save2.tpr	215 KB	06/24/2015	Michael Shirts
runfail_in.gro	3.08 KB	06/25/2015	Michael Shirts
runfail.top	27.8 KB	06/25/2015	Michael Shirts
runfail_nonzero.top	27.9 KB	06/25/2015	Michael Shirts
runfail.mdp	3.4 KB	06/25/2015	Michael Shirts