GROMACS - Bug #2867

abnormal box fluctuations on GPUs still there

02/22/2019 11:37 AM - Ramon Guixa-Gonzalez

Status:	Fix uploaded		
Priority:	High		
Assignee:	Berk Hess		
Category:	mdrun		
Target version:	2019.4		
Affected version - extra info:	Other vesions too	**Difficulty:**	uncategorized
Affected version:	2016.6		

Description

Hi there,

I am afraid that the fix applied to Bug #2845 has not resolved the issue of evident GPU-related box fluctuations. I tested the fix (http://manual.gromacs.org/documentation/2016.6/ReleaseNotes/release-notes.html#fix-incorrect-lj-repulsion-force-switching-on-gpus) on gromacs 2016.6 and, despite it took more that 300 ns, the box still fluctuates pretty abnormally.

I have attached the files from this test.

Ramon

Related issues:

Related to GROMACS - Bug #2845: critical box fluctuations when using GPUs	**Resolved**

Associated revisions

Revision 8fc3cc55 - 09/21/2019 12:01 PM - Berk Hess

Fix incorrect shift forces for CMAP

The shift force indices were inverted for the second and third
atom in the CMAP term, leading to incorrect virial and pressure
contributions when these atoms resided in different periodic images.

Fixes #2845 and #2867

Change-Id: I1946a1d375d6c62e4e6d23ee25b92b42a3d4a6f7

History

#1 - 02/22/2019 01:18 PM - Berk Hess

- Related to Bug #2845: critical box fluctuations when using GPUs added

#2 - 03/05/2019 10:23 PM - Ramon Guixa-Gonzalez

Hi, is there any other info I could provide to help? Do you have any hint on how to solve this issue?

#3 - 07/31/2019 02:47 PM - Alexander Vogel

I just saw that bug #2845 is labeled resolved so I'll repost my findings here and also continue updating here:

I can confirm this bug with GROMACS 2019.3. I have several simulations of a GPCR+Membrane system that get ripped apart after 100-200ns. But already before that unnaturally large fluctuations are seen that manage to recover however. After some time the system violently extends in x-y dimension (which is the membrane plane) and collapses in the z-dimension. I have two simulations of identical coordinates that were started with different random seeds but otherwise identical parameters and both get ripped apart. Interestingly, both systems have very similar box dimensions after the "event" although the disrupted membrane structure is very different between the two. Since I didn't notice the error immediately and the simulations did not crash, they continued for several hundert ns (first simulation for 450ns, second simulation for 900 ns). There still seem to be some pretty large fluctuations during that time but no violent changes as before.

I'll start a third version of the system today. A colleague of mine reported similar issues but I didn't have time to look at it yet.

I setup the simulation with charmm-gui and run them on an AMD 1700X + Nvidia RTX2080Ti. If more info is needed please tell me what you need.

Alexander Vogel

PS: The third identical copy started with a different random seed ripped apart in the same way after less than 100ns. The final box dimensions again are very similar to the first two cases. I'll now try to use charmm-gui with the same pdb structure of the protein and the same settings to create a different starting conformation (charmm-gui involves some randomness for placing lipids etc.). I'll report how that goes.

Anyway I don't really believe in a ForceField problem...because I can't imagine how even a big error would lead to such dramatic fluctuations in the box size. Also Ramon in Bug #2845 reported that the error does not occur with CPU. For that reason I also generate NAMD output for the new simulation so that I can verify if the problem occurs there too or not.

#4 - 07/31/2019 05:42 PM - Berk Hess

Ramon and I both saw the issue as well with CPU, so therefore it is not a GPU issue. We can't be sure it's a forcefield issue, but that is the most likely explanation.

#5 - 08/01/2019 10:08 AM - Alexander Vogel

Ok, sorry, I didn't see that you also observe it with CPU. I'll try to investigate further and see if I find something. What I find striking is that I also use POPC. I'll check its force field files if I find something.

Alexander

#6 - 08/01/2019 10:59 AM - Berk Hess

I don't think you can see much from the force field files. Lipid bilayer phases are always close to phase transitions. So the force field needs to be parameterized very accurately for systems to remain in the correct phase. Futhermore, dynamics in membranes is very slow, so it can take a very long time for anomalies to occur. Unless the force field has been tested in microsecond regime, there is always a chance of something going wrong.

#7 - 08/01/2019 11:55 AM - Alexander Vogel

- File jp101759q.pdf added

I'm running pure POPC simulations (with slightly different mdp parameters) for many microseconds in another project and nothing bad happens there:

1. 256 POPC: 8,05 microseconds
2. 256 POPC: 7,91 microseconds
3. 256 POPC: 7,22 microseconds
4. 1024 POPC: 1,65 microseconds
5. 4096 POPC: 301,3 nanoseconds

And the phase transition of POPC is at 277K which is pretty far from the 310K the collapsing simulation is running at (the pure POPC sims that run fine are at 298K). In addition the force field of POPC has been tested pretty rigorously with extensive comparison to experiment: "Update of the CHARMM All-Atom Additive Force Field for Lipids: Validation on Six Lipid Types" by Klauda et al. J. Phys. Chem. B 2010, 114, 7830–7843

POPC is one of the six lipids that were tested. I attached the paper for reference.

Therefore I doubt that the POPC force field is the issue unless there is a typo...that is what I want to check for. I took a first look at the files and they seemed fine however. Also the protein force field should be fine as should be water and ions or somebody else should have noticed already. And that is everything that is part of my simulation that gets ripped apart...there is nothing else in there.

Alexander

#8 - 08/01/2019 01:41 PM - Berk Hess

I don't know much about these CHARMM force fields. I assume that with the 277 K you mean the experimental transition temperature. I don't know what that of the model is and there could be other phases the model allows that are not observed experimentally.
I could also be that the exact cut-off parameters and treatment don't exactly match what the force field was parametrized with.

#9 - 08/01/2019 03:54 PM - Alexander Vogel

Yes, 277 K is the main phase transition that was determined experimentally for POPC. The simulations clearly are above that temperature and the membrane behavior in the simulations clearly corresponds to that of membranes above the main phase transition.

I'm a membrane biophysics person....so I would say I know pretty much about lipids and membranes. What is happening in the simulations is completely unphysical. The forces needed to rip a membrane apart like this must be huge. Lipid acyl chains hate water and during the "event" they get exposed to lots of it. Once I did simulations with lipids and water distributed randomly in a box. Within 1 ns they clustered to avoid as much water contacts as possible and within 50 ns they formed bilayers just by self-assembly. Here they do the opposite. They leave the bilayer and get exposed to lots of water within a few ns. This does not make any physical sense.

Also I'm using standard cut-off parameters and I have done membrane simulations for many years using CHARMM and namd and never seen something even remotely similar. That is why I doubt that this is a force field issue. If there is some error in the force field it must be humongous. But the force fields predicts so many experimental values close to perfection...

Anyway...I'm still on the issue. I'm trying several things and I asked my co-workers to look into their simulations if they observe it as well. I'll also try to find a simulation that gets ripped apart where I have the namd input files to see it happens there as well. So I'm on the issue and will keep you updated if I find if it is a force field issue or a GROMACS issue.

#10 - 08/02/2019 08:25 AM - Berk Hess

Thanks for investigating this.

The force field can be compared by comparing the energies for a single frame between GROMACS and another code. Only the LJ energies could differ slightly due to a different LJ cutoff switch treatment.

Could you attach an mdp file?

#11 - 08/02/2019 02:39 PM - Alexander Vogel

- File mdout.mdp added

- File run.mdp added

I attached my mdp and a file that lists all mdp parameters set in GROMACS.

Alexander

#12 - 08/02/2019 04:30 PM - Berk Hess

One think I did not think of before is the possibility that the Parrinello-Rahman barostat with a coupling time of 10 ps might cause instabilities. What pressure coupling with what parameters are you using in CHARMM and NAMD?

I am now running Parrinello-Rahman with a coupling time of 20 ps to see if that changes things.

#13 - 08/05/2019 02:00 PM - Alexander Vogel

- File working_vs_3nonworking_sims.png added

- File run.mdp added

- File step7_production.mdp added

I think I'm closing in to the issue. I now have managed to run the exact same simulation of which 3 copies with different seeds were collapsing within less than 200 ns up to almost 400 ns without problems with GROMACS (see attached picture...the red line is the new simulation). I'll leave it running to confirm for longer time scales. What I did was running the simulation with the mdp that comes from charmm-gui instead of the modified mdp that our work group is usually using. As far as I know Ramon is using a similar setup as he is former member of the group and I think he came up with some of the changes.

I'm not an expert on the changes in the mdp, so could you please go over it and see if there are any differences that stand out (I attached both of them). The simulations with the run.mdp collapse while the one with step7_production.mdp seems to be running fine.

Thanks and best regards,

Alexander

#14 - 08/06/2019 05:43 PM - Berk Hess

I think your latest observation is just coincidence and you will see the same issue with more statistics.

I have been running the system provided by Ramon, which uses the step7_production settings, but with the V-rescale thermostat with tau-t=0.5. I changed to NH with tau-t=1 and see the same issue.

#15 - 08/13/2019 10:25 AM - Alexander Vogel

- File 1_microsecond_without_problems.png added

Now at 1 microsecond without problems.

#16 - 08/14/2019 01:04 PM - jose flores-canales

Michael Shirts has mentioned that a Monte-Carlo barostat is on the works in his lab. It will be great to see this and other barostats arriving on Gromacs.

#17 - 08/14/2019 05:56 PM - Ramon Guixa-Gonzalez

- File likeCharmm.jpeg added

Alexander Vogel wrote:

> Now at 1 microsecond without problems.

Alexander Vogel wrote:

Now at 1 microsecond without problems.

Alex, I also have the feeling new runs will show us otherwise, but let's see. At my end, I am running a test with a similar CHARMM-GUI-out-of-the-box mdp, namely using a longer tau-t (1.0) and a shorter tau-p (5.0). So far (~325 ns), no abnormal fluctuations present (pic attached). In the past, however, sudden fluctuations have often showed up around 400 ns and sometimes much later (e.g. 800 ns). I will report again beyond 1 micro.

#18 - 08/20/2019 12:40 PM - Alexander Vogel

- File update.png added

Another update. Now at 1.7 microseconds and still everything fine. I'll try to switch this simulation to the other mdp options...never did that before...have to see how it goes. If I get that working we will see if it shows up again.

Alexander

#19 - 08/20/2019 12:46 PM - Berk Hess

I would find it very surprising if only changing the T and p coupling settings "fixes" the issue. If anything, using Nose-Hoover and a shorter pressure coupling time should make things less stable. If this would really "fix" it, the only thing I can think of is that the system has some internal resonance which coincides with the 10 ps period of the barostat.

#20 - 08/20/2019 07:01 PM - Mayukh Chakrabarti

- File Box_fluctuations.png added

- File production.mdp added

I have been following this issue with some interest, as I have been experiencing the same issues with box fluctuations observed by others. I am currently using Gromacs 2018.7 with the CHARMM36 forcefield, July 2017 revision. My system contains a transporter protein embedded in a POPC bilayer. In multiple simulations of the same system with different random seeds, I observe large fluctuations in the X-Y dimensions, followed by an eventual collapse in the Z-dimension, usually after 100-200 ns in most instances. I have attached an image documenting these fluctuations in multiple runs. Equilibrating the system for a longer time does not fix the issue.

I too am a bit doubtful about specific mdp settings fixing this issue. I have also attached an example mdp file that I have been using. Other than the presence of two coupling groups as opposed to 3, it exactly follows the settings that are suggested by the CHARMM-GUI production.mdp file, and which Alexander has now tested for 1.7 microseconds (NH & Parinello-Rahman, tau-t = 1, tau-p = 5). In my case, I am observing the fluctuations using these settings.

#21 - 08/20/2019 10:16 PM - Berk Hess

Consensus is growing in the community that second order coupling like Nose-Hoover and Parrinello-Rahman causes issues (like instabilities).
Has someone tried using Berendsen pressure coupling?

#22 - 08/21/2019 10:12 AM - Alexander Vogel

- File collapse_after_mdp_change.png added

- File mdp_before_collapse.mdp added

- File changed_mdp_that_leads_to_collapse.mdp added

And BAM...there it is. Directly after changing the mdp the system starts to collapse. In the attached picture the dotted line marks the timepoint where I changed the mdp. The two different mdp files are also attached. I'll continue the simulation for one more day to see it fully collapse...just to be safe.

In the mdp several parameters were changed. The biggest difference is the change of the thermostat from Nose-Hoover to V-rescale. But also some other parameters change. If you want to me to change just a single one I can do that and see what happens...just tell me which one you would be most interested in.

Alexander

#23 - 08/21/2019 03:25 PM - jose flores-canales

- File production_langevin.mdp added

I attach my mdp parameters file. It is basically the same from Charm-gui, except that I am using Langevin thermostat. For Parrinello-Rahman I am using Tau_p = 5 ps, the same system also failed for Tau_p = 20 ps.

#24 - 08/22/2019 09:05 AM - Berk Hess

I now tried Berendsen pressure coupling with tau-p both 5 and 10 ps. The box dimensions now behave much smoother and show less extreme fluctuations, but the issue is still there.

#25 - 08/22/2019 10:28 AM - Alexander Vogel

- File collapse_after_mdp_change_update.png added

Update from my side. So it didn't collapse completely yet...but it went much further than yesterday...see graph. After changing the mdp the problems clearly are back. I'm certain it is just a matter of time until it will collapse.

What is interesting is that Jose Flores-Canales also observes these problems after changing the thermostat from the charmm-gui standard parameters...so maybe it is the thermostat instead of the barostat...or maybe combinations of the two.

Also that idea of internal resonances becoming critical does not sound convincing at all to me for three reasons:

1. If there is one thing biomembranes in the liquid-crystalline phase are known for, then that they are extremely flexible. Having a critical resonance in a biomembrane would be like having one in bubble gum. The flexible nature of the material should dampen every oscillation effectively.

2. It doesn't look like oscillations becoming critical. In the simple idea I have of a critical oscillation, it should look like an oscillation that quickly becomes larger and larger. But there are no oscillations in any of the graphs...neither in mine nor in that of Ramon or jose flores-canales. Boxsize-Z always goes down...never up. That is not an osziallation to me.

3. It happens for different tau_p. E.g. see Jose Flores-Canales post above. I would expect tau_p to have an influence if it really were critical oscillations.

I know that I'm going to sound a bit ungrateful now but this issue is open for months already. There are three independent research labs that have reported it here (Ramon: Leonardo Pardo lab, Jose Flores-Canales: Aarhus University, myself: Hildebrand lab). A collegue of mine also postet the issue on twitter and several more people reported seeing the problem:

Wojciech Kopec (de Groot lab): https://twitter.com/wojciechkopec3/status/1158328579475877888
R. Thomas Ullmann (has his own lab): https://twitter.com/RTUllmann/status/1158409881432596482

Also the issue is devastating to any simulation where it happens. As far as I know Ramon does not use GPU anymore because of the bug. I don't know if I can rely on my simulations anymore.

A lot of non-GROMACS possible causes were brought up but none of them sound likely to me. It almost certainly is not the POPC forcefield. I have uploaded a papers that shows how good it is in comparison to experiment. I have GROMACS simulations running at the moment of pure POPC bilayers of various sizes (256 lipids, 1024 lipids, 4096 lipids (more than 1 million atoms)) and for a cumulative time of more than 25 microseconds and all of them have no problems (with standard charmm-gui mdp parameters). Also the critical oscillations do not sound convincing as a cause as stated above.

Changing the mdp parameters causes it or at least has a large impact on it. Too me that looks like the problem most likely has to do with GROMACS or one of its pressure/temperature coupling methods. I'm not an expert on these so I don't know what to look for. I try to help as much as possible by trying different parameters and running microsecond timescale simulations just for the bug. It is easily reproducible for my simulation...it should not be too hard to look at it in detail. Is there anybody really looking at it at the moment? What do I have to do, so that this is issue is taken seriously? Maybe it is...but the responses here don't really reflect that.

So now I sound ungrateful...which I didn't want to because I'm very grateful that GROMACS exists and that people take care of it and fix its bugs and make it better and faster all the time. It just puzzles me that this issue does not cause more (perceivable) concern.

Alexander

Edit: Sorry didn't see your post from 1h ago...probably ninja'd me while I was typing. So do I understand that right that it is probably not the barostat? Is there anything you want me to try with the mdp parameters of my simulation where it happens so reproducibly?

#26 - 08/22/2019 11:40 AM - Berk Hess

So you have seen no problems at all with the charmm-gui parameters on both CPU and GPU?
Have others used the same parameters and seen no issues either?

If that is the case we need to gradually change a problem mdp file into the charmm one and see when the issues disappear.

Note that we still don't know if this is a GROMACS or force-field issue. People tend to simulate much longer in GROMACS than in many other packages, so force field issues tend to appear first with GROMACS.
Also, details of the implementation of coupling algorithms matter. Some years ago I removed a 1 step delay of recording and using the temperature in the Nose-Hoover thermostat. This caused simulations of ice of a user to become unstable, but that was actually the correct simulation behavior. The delay added artificial damping. (I now a bilayer is not in a solid state, but this is just to show that subtle details can have a large impact on the behavior of simulations).

#27 - 08/22/2019 11:44 AM - Justin Lemkul

People have reported via the Twitter thread linked above that the standard CHARMM-GUI inputs can also lead to this problem.

I have validated this force field extensively, comparing forces and energies between GROMACS and CHARMM. They match almost perfectly for all molecules I've ever tested (encompassing amino acids and full proteins, nucleotides and DNA, lipids, small molecules, and carbohydrates). I have discussed this with Alex MacKerell and we agree that the odds of a force field problem are vanishingly small. Personally, I have several hundred ns of membrane simulations in OpenMM with C36 that are just fine.

We ran a cross-software comparison of membrane simulation outcomes in the CHARMM-GUI validation, but that was with version 5.1. At that point, everything matched quite well. We used CPU only then.

#28 - 08/22/2019 12:16 PM - Berk Hess

I was also thinking of the stability of the CHARMM forcefield itself (not the implementation in GROMACS).

In Stockholm we have many simulations with this forcefield up to a microsecond and we have never seen issues.

#29 - 08/22/2019 12:17 PM - Berk Hess

I would like to get the input for the CHARMM gui mdp parameter runs that failed, so I can check those myself.

#30 - 08/22/2019 01:59 PM - Berk Hess

Has someone observed if there first is a conformational rearrangement and then the box starts fluctuating or first the box starts to fluctuate and then the conformations change?

#31 - 08/22/2019 02:37 PM - Alexander Vogel

- File POPC_256x3.png added

- File POPC_1024(red)_4096(blue).png added

- File berkhess.tar.gz added

Many ansers to many different questions follow:

1. Yes with standard charmm-gui parameters there seem to be no issues in my case. The simulation ran for 1.7 microseconds without a hitch...look at the graphs I attached in former posts. However, problems might appear at even longer simulation times as Justin Lemkul suggests. So far I used only GPU.

2. About the force-field...if Alex MacKerell says that the force-field is ok, I believe him. Also as stated multiple times I do have three simulations of pure POPC bilayers of very long timescale: 8.7 microseconds + 8.6 microseconds + 7.8 microseconds. I just checked them again...they are fine...no fluctuations whatsoever...see attached graph. I also checked my two larger simulations (1k and 4k lipids)...they are much shorter (385 ns + 183 ns) but they also show no fluctuations...see attached other graph. I can't think of any more convincing way to show that it is not the POPC force-field that is causing the issue.

3. I attached the input files for my various collapsing systems. The tar file contains the following folders:

3_collapsing_runs: This folder contains the three replicas I showed at the beginning that collapsed in less than 200ns.

charmm-gui_run_no_problems: This folder contains the start of 1.7 microsecond run that had no problems.

charmm-gui_run_changed_mdp: This folder contains the end of 1.7 microsecond run that had no problems. I used that to start a run with changed mdp options that showed problems within 100ns or so.

Thanks for looking into it!

Alexander

#32 - 08/22/2019 02:48 PM - Alexander Vogel

Your question about conformations:

The membrane stays intact pretty long. As I said...it is pretty flexible. Just after the simulation box has deformed considerably the membrane ruptures. I hope that is what you were asking for.

Alexander

#33 - 08/22/2019 04:01 PM - Berk Hess

What I only noticed or realized now is that there is a switch from from 3 center of mass motion removal groups in the system to 1. In general there can be issues that the solvent starts sliding over the membrane and therefore people use separate COMM removal groups. I don't know if this correlates with the issue.
I'm running a check run with velocity output to see if there if momentum is building up between membrane and solvent.

#34 - 08/22/2019 04:02 PM - Berk Hess

PS Thanks Alexander for all the detailed information!

#35 - 08/22/2019 10:46 PM - Ramon Guixa-Gonzalez

- File pic1.jpeg added

- File pic2.png added

Hey all, well I checked this run I started a few days ago and, as I expected, the system collapse after almost 600 ns (check pic1 attached). The most interesting thing here is that I used almost (see below) exactly the same parameters CHARMM-GUI provides in its output files.

The ONLY difference between a CHARMM-GUI-like mdp file and the one I show here is the thermostat: V-rescale (me) versus Nose-Hoover (CHARMM-GUI) (see pic2 attached). I will now re-run the same simulation using an mdp file where I will only change V-rescale into Nose-Hoover in order to simulate an exact copy of the CHARMM-GUI parameters. I am not entirely sure, but I think, however, that I already simulated one of my systems literally using the mdp production file provided by CHARMM-GUI and observed similar fluctuations including the collapse of the system. May I remind one thing that I already reported from the beginning of this issue: none of the simulations showing these dramatic effects (likely more than 25 or 30 different runs now) end up crashing...??? Berk, regardless of the fluctuations, wouldn't you expect mdrun to crash upon the collapse of a protein membrane system?

#36 - 08/22/2019 11:09 PM - Berk Hess

I haven't seen crashes yet.

I think I tried all combinations of thermostats and barostats.
My current working hypothesis is that is could be due to the use of 3 T-coupling groups combined with COMM removal for the whole system. I have a run running, but it needs to run longer.

#37 - 08/23/2019 08:07 AM - Berk Hess

I also got box fluctuations with the COMM groups equal to the T-coupling groups.
Now I have no clue at all what it could be. All combinations of thermostats and barostats I tried can fail. But it still seems that some combination of mdp parameters do not fail, or?

#38 - 08/23/2019 09:13 AM - Alexander Vogel

The question is: When do you call a simulation "not failed"? How many microseconds? I got mine running up to 1.7 microseconds without problems (barostat: Parrinello-Rahman, thermostat: Nose-Hoover) but there are people that report seeing it after 6 microseconds.

When the system is compressed in Z-direction (and therefore expanded in the X-Y-plane) the forces must be crazy high because the membrane surely should not like to be stretched like that...in particular just before rupture. The hydrophobix effect should be very strong trying to avoid any water molecules in between lipids. Could one not just analyze the forces individually for such a snapshot and see what is causing the system to expand even further in the X-Y-plane although there should be very high forces from the lipids and waters close to them to keep the system together?

I would do it but I'm not an expert in these kinds of things and don't know how...

Alexander

#39 - 08/23/2019 09:14 AM - Berk Hess

I now looked at the potential energy and that looks unaffected by the large changes in box shape and membrane area. Unfortunately we don't know the free-energy differences between these states. But I am back to my original hypothesis that these stretched out membrane states are very similar in free-energy to the wanted state. If there are separated by a relatively small free-energy barrier, differences in coupling algorithms could provoke or prevent transitions.

I don't know of a way to assess the free-energy difference between such states. One check would be taking an extended state, maybe not too extreme, and run that in NAMD, CHARMM or OpenMM to see if it goes back to the original state.

#40 - 08/23/2019 09:23 AM - Berk Hess

Alexander, you are arguing based on the real world. But we are simulating models, not the real world. I do not doubt that the developers of this force field have done a good a careful job, but membrane force fields are particularly tricky. There is a delicate balance between hydrophilic interactions, hydrophobic interactions, dihedral potentials in the tails and entropy. You might get the right area per lipid and good looking conformations, but with a slightly incorrect balance. Or even nastier, you could get all the balance correct in the wanted state, but there are unforeseen free-energy minima in unwanted, previously unexplored conformations. We have seen this before in e.g. DNA force fields.

As the potential energy does not change, there will likely be no large forces in these extended states. But maybe such states should be disfavored due to lower entropy. That should still lead to a pressure contracting the membrane again though, and we do not see that.

#41 - 08/23/2019 11:40 AM - Justin Lemkul

Berk Hess wrote:

> I now looked at the potential energy and that looks unaffected by the large changes in box shape and membrane area. Unfortunately we don't know the free-energy differences between these states. But I am back to my original hypothesis that these stretched out membrane states are very similar in free-energy to the wanted state. If there are separated by a relatively small free-energy barrier, differences in coupling algorithms could provoke or prevent transitions.
>
> I don't know of a way to assess the free-energy difference between such states. One check would be taking an extended state, maybe not too extreme, and run that in NAMD, CHARMM or OpenMM to see if it goes back to the original state.

I could do this pretty easily if someone wants - I have inputs readily available that I know work.

Another test I would like to perform is to compute the single-point energy of a crazy coordinate file. If someone (Alexander?) can post a .gro/.pdb file of a snapshot that is torn apart (extreme point of a box fluctuation, for instance) and also report the energies that GROMACS reports at that snapshot, I will compare energies in CHARMM and forces so we can do an apples-to-apples comparison of the software on the same snapshot.

#42 - 08/23/2019 11:46 AM - Berk Hess

I trust the energies in GROMACS and they look fine.

But a check I just did is start from a stretched conformation and run a "relaxation"-run with fixed box x/y dimensions. Now I get negative pressures of around -40 bar along x and y. So the system actually seems to want to contract back. I guess this means that some momentum is built up that causes the pressure and box fluctuations. It does not seem to be due to the conformation of the bilayer, although I can't 100% exclude that now.

#43 - 08/23/2019 04:52 PM - jose flores-canales

I am running MD with AMBER18/pmemd.cuda using a restarting file before my system underwent transformation.

#44 - 08/23/2019 05:26 PM - Ramon Guixa-Gonzalez

Berk Hess wrote:

> I haven't seen crashes yet.
>
> I think I tried all combinations of thermostats and barostats.
> My current working hypothesis is that is could be due to the use of 3 T-coupling groups combined with COMM removal for the whole system. I have a run running, but it needs to run longer.

Berk, as I just reported (see comment #34 above), there is so far not any combination of parameters that do not fail other than changing the thermostat from V-rescale to Nose-Hoover,which is what I am testing now. In Alexander's hand, it also seems to be this switch what is causing the problem.

#45 - 08/25/2019 12:49 PM - Berk Hess

- File largearea_t232ns.pdb added

I now ran a longer and slightly more careful fixed area run of Ramon's system at 13% larger x/y area. This gives a surface tension of 0.014 +-0.006 N/m. This is about a factor 2.5 lower than one would expect from the area compressibility modulus of the force field.

This seems to indicate that there are conformational changes which lower the surface tension and might enable the observed fluctuations, if there is a significant initial fluctuation to overcome some barrier.

I have attached a frame so people can look at the bilayer.

#46 - 08/26/2019 09:24 AM - Alexander Vogel

- File now_collapsed.png added

The simulation where I switched from Nose-Hover to V-rescale took somewhat longer to totally collapse but it happend over the weekend...see attached screenshot. Most striking is how the fluctuation appear immediately after switching to V-rescale and never go away although the system manages to recover in between.

Berk Hess wrote:

> As the potential energy does not change, there will likely be no large forces in these extended states. But maybe such states should be disfavored due to lower entropy. That should still lead to a pressure contracting the membrane again though, and we do not see that.

I'm not suprised that the potential energy of the lipids is not affected that much by the extended state...lipids are very flexible after all. I would expect the largest forces coming from entropy because the hydrophobic effect is entropic in nature. When the membrane is stretched the water molecules should have more access to the hydrophobic tails that are not as well shielded by the lipid head groups anymore. This should cause a considerable entropic effect.

Alexander

#47 - 08/26/2019 05:01 PM - Berk Hess

As I wrote, I see about a factor 2.5 lower surface tension then expected for 13% area increase. So it seems that configurations are possible that, probably, have lower free-energy than in reality, although likely still higher than the free-energy of the wanted area.

#48 - 08/26/2019 10:32 PM - Berk Hess

I now have 140 ns of 13% larger area run and the surface tension is 36 +-68 bar*nm. So it's zero within the statistical error.
It is about time that someone starts a run with the pdb I uploaded in a different package to check what surface tension that gives. Then we should be able to see if this is a GROMACS or a force field issue.

#49 - 08/27/2019 01:39 AM - Justin Lemkul

Berk Hess wrote:

> I now have 140 ns of 13% larger area run and the surface tension is 36 +-68 bar*nm. So it's zero within the statistical error.
> It is about time that someone starts a run with the pdb I uploaded in a different package to check what surface tension that gives. Then we should be able to see if this is a GROMACS or a force field issue.

I'd be happy to run a simulation in OpenMM (and perhaps NAMD) but I will need the topology of NEC to generate a CHARMM PSF.

#50 - 08/27/2019 09:50 AM - Berk Hess

Justin Lemkul wrote:

> Berk Hess wrote:
>
> > I now have 140 ns of 13% larger area run and the surface tension is 36 +-68 bar*nm. So it's zero within the statistical error.
> > It is about time that someone starts a run with the pdb I uploaded in a different package to check what surface tension that gives. Then we should be able to see if this is a GROMACS or a force field issue.
>
> I'd be happy to run a simulation in OpenMM (and perhaps NAMD) but I will need the topology of NEC to generate a CHARMM PSF.

Ramon should then provide the input you ask for.

I can also run Alexander's system, but that will take some time.

#51 - 08/27/2019 02:07 PM - Alexander Vogel

- File run3_snapshots.png added

- File run3_20.0.trr added

- File run3_46.2.trr added

- File run3_70.0.trr added

- File eq7_3.gro added

- File summary.txt added

Hello everybody,

I prepared a set of files you might find useful. To keep the file size managable I chose run3 from my original set as it collapsed after ~60ns and the whole trajectory is just ~100ns long. The first attachment is a graph of the Boxsize-Z over time. In this graph there are three dotted lines at 20ns, 46.2ns and 70ns at which I extracted the snapshots run3_20.0.trr, run3_46.2.trr, and run3_70.0.trr. There also is the eq7_3.gro that you probably need to see the snapshots. I also extracted everything I could from the edr files at these 3 timepoints which is summarized in summary.txt.

I also compressed my whole run including the two equilibration steps. The first equilibration step is the actual equilibration, the second equilibration step is just a very short run with random seed to generate different replicas of the same system (e.g. run1, run2, run3 in my case where I only gave you run3 because that collapsed fastest). This can be downloaded there: https://www.dropbox.com/s/bzdn36io9orgh24/run3.tar.gz?dl=0

Tell me if you need anything else.

Alexander

#52 - 08/27/2019 05:12 PM - Berk Hess

Thanks Alexander, but Justin (only) needs "the topology of NEC to generate a CHARMM PSF".

#53 - 08/27/2019 05:13 PM - Justin Lemkul

Berk Hess wrote:

> Thanks Alexander, but Justin (only) needs "the topology of NEC to generate a CHARMM PSF".

I asked for some files separately in a Twitter thread that another user had started a few days ago. I want to do some more thorough comparisons in CHARMM. I do still need the NEC topology for the system you are working on.

#54 - 08/27/2019 06:02 PM - Alexander Vogel

Sorry...forgot to also upload the whole charmm-gui.tgz which contains all psf files you should need:
https://www.dropbox.com/s/jrb85zgh2cvdhsh/charmm-gui.tgz?dl=0

Alexander

#55 - 08/28/2019 02:58 PM - Alexander Vogel

- File step5_production_changed_to_collapsing.mdp added

- File waterbox.png added

So I tried a stupid thing: I setup a waterbox of the same size as my collapsing protein+membrane+water+ions system and simulated it with the "problematic" mdp. I just deleted the PROT and MEMB groups from the tc-grps and the corresponding values because they don't exist anymore. I attached the mdp for reference. I know...the mdp parameters probably are not the best for a waterbox because of the semi-isotropic pressure coupling...but the result still seems way out of what I would call normal. We start at a Boxsize-Z of ~12.5 and within 33ns we reach a Boxsize-Z of ~3000 (plot of Boxsize-Z is attached) when the simulation crashes:

Fatal error:
Step 16814700: The total potential energy is nan, which is not finite. The LJ and electrostatic contributions to the energy are 0 and 0, respectively. A non-finite potential energy can be caused by overlapping interactions in bonded interactions or very large or Nan coordinate values. Usually this is caused by a badly- or non-equilibrated initial configuration, incorrect interactions or parameters in the topology.

At this point there is lots of vacuum in between slabs of water and even some individual water molecules in vacuum. At the end the box even extends in the X-Y-plane while still extending in the Z-direction. I took a video...see it on twitter: https://twitter.com/MelbourneLE/status/1166695203572527106

Now this could be because the mdp parameters are not really suited for a waterbox...as I said I'm not an expert on such things. But it does not seem reasonable to me at all. If this really is not because the mdp parameters are not suitable I would say we finally can exclude that the problem has anything to do with the force field of the lipids, the protein, or even the ions. Maybe the water force field but I really would expect a bug in GROAMCS...maybe something to do with the semiisotropic pressure coupling?

Alexander

#56 - 08/28/2019 04:01 PM - jose flores-canales

I think this is expected. Once I forgot to choose isotropic NPT in a protein-solvent system. Instead the simulation run with anisotropic NPT and the protein end-up interacting with its images along one axis in a very thin box.

#57 - 08/28/2019 05:08 PM - Alexander Vogel

jose flores-canales wrote:

> I think this is expected. Once I forgot to choose isotropic NPT in a protein-solvent system. Instead the simulation run with anisotropic NPT and the protein end-up interacting with its images along one axis in a very thin box.

Really? Did you see the video? There is lots of vacuum in the box after a few ns...can't believe that this is expected...

Vacuum is the most hydrophobic thing I can think of...and water is very polar...

Alexander

#58 - 08/28/2019 05:09 PM - Justin Lemkul

jose flores-canales wrote:

> I think this is expected. Once I forgot to choose isotropic NPT in a protein-solvent system. Instead the simulation run with anisotropic NPT and the protein end-up interacting with its images along one axis in a very thin box.

Anisotropic coupling distorts any box - water, membrane, whatever.

Semiisotropic should not and certainly (I would think) should not be so dramatic as what Alexander is showing. A drift over time, perhaps, but not so extreme.

#59 - 08/29/2019 12:36 PM - Alexander Vogel

- File iso_vs_semiiso.png added

So in the twitter discussion, Justin Lemkul suggested to run my membrane simulation with istropic pressure coupling instead of semiisotropic: https://twitter.com/JustinLemkulVT/status/1166696918963163137

I did that for run3 (because that one collapsed fastest) and although I'm only at 78 ns yet it already looks much better. All other 3 replicas of the system had already collapsed (run3) or at least shown significant spikes in Boxsize-Z (run1 and run2) by that time. With isotropic pressure coupling none of that is seen. In addition the fluctuations in general are much less...see the attached graph...it is almost a straight line vs. the others that fluctuate a lot more even whem not showing extreme spikes.

Alexander

PS: Could we now please get some confirmation that we all agree that this is not because of the lipid force field and that it is likely a bug in the

GROMACS pressure coupling...maybe in the semiisotropic part...and that somebody is looking at it...that would be really nice after all the work I put into this, trying to convince people that it is worth looking into.

#60 - 08/29/2019 03:11 PM - jose flores-canales

- File box_dimensions.png added

- File box_dimension_gromacs.png added

I don't have enough statistical sampling, but the box is stable with pmemd.cuda engine. Following Berk's suggestion, I took a checkpoint early when the Z-box dimension begins to shrink and converted the gromacs topology parameters to amber compatible topology file with gromber/parmed3. So far I have got 150 ns and the box rapidly recovers its Z-dimension (see the first graph attached). This is maintained near its previous average dimension ~ 122 Angstroms obtained with GMX (see graph attached, 1000 ns and dimensions in nm). Force-field is the same Charmm36 for all simulations. Charmm36 parameter files were obtained directly from the MacKerell's website.

pmemd.cuda is running with Langevin thermostat and Monte-carlo barostat, while gromacs was ran with Langevin integrator and PR barostat. Semiisotropic NPT conditions and 12/10 force-switch were used for both MD engines.

#61 - 08/29/2019 05:05 PM - Alexander Vogel

jose flores-canales wrote:

> I don't have enough statistical sampling, but the box is stable with pmemd.cuda engine. Following Berk's suggestion, I took a checkpoint early when the Z-box dimension begins to shrink and converted the gromacs topology parameters to amber compatible topology file with gromber/parmed3. So far I have got 150 ns and the box rapidly recovers its Z-dimension (see the first graph attached). This is maintained near its previous average dimension ~ 122 Angstroms obtained with GMX (see graph attached, 1000 ns and dimensions in nm). Force-field is the same Charmm36 for all simulations. Charmm36 parameter files were obtained directly from the MacKerell's website.
>
> pmemd.cuda is running with Langevin thermostat and Monte-carlo barostat, while gromacs was ran with Langevin integrator and PR barostat. Semiisotropic NPT conditions and 12/10 force-switch were used for both MD engines.

Which simulation is that? GPCR + membrane from me or something else? If it contains lipids this would further disprove the hypothesis, that it is the lipid force field (or any other force field of stuff that is part of your simulation) that is causing the issue. Thanks for following up on this!

Alexander

#62 - 08/29/2019 05:23 PM - jose flores-canales

It is my own system, a GPCR on POPC lipids bilayer and TIP3P. I have struggled with this for some months. For the time being, I am working in other projects using the Amber engine.

#63 - 08/30/2019 09:54 AM - Alexander Vogel

- File iso_vs_semiiso_update1.png added

jose flores-canales wrote:

> It is my own system, a GPCR on POPC lipids bilayer and TIP3P. I have struggled with this for some months. For the time being, I am working in other projects using the Amber engine.

You can try running this sim with isotropic pressure coupling instead of semiisotropic in GROMACS. So far it totally fixes the problem for me. My isotropic sim is now at 160ns and doesn't show a hitch...see attached graph.

Alexander

#64 - 08/30/2019 11:01 AM - jose flores-canales

I have been thinking about it, but I think I will settle down with initial semiisotropic NPT equilibration (~50 ns) followed by NVT production runs. I will probably take the configuration with the box volume closest to the average for NVT.

#65 - 09/02/2019 09:10 AM - Alexander Vogel

- File iso_vs_semiiso_update2.png added

Another update for the system with isotropic pressure coupling...the final one for some weeks as I will be on vacation. Now at 450 ns and still no problems with isotropic pressure coupling. To me this looks as if the semiisotropic pressure coupling is causing the issues...because that is the only difference between the three runs with semiisotropic that collapse reproducibly and the isotropic pressure coupling of the exact same system with exactly the same mdp parameters besides isotropic vs. semiisotropic.

I'll keep it running over my vacation so that we should several microseconds then to be sure...

Alexander

#66 - 09/02/2019 09:20 AM - Berk Hess

Of course the box cannot deform with isotropic pressure coupling or NVT. The deformation is only possible when x/y is coupled independently from z. That does not tell us anything about the source of the issue.

#67 - 09/02/2019 09:53 AM - Alexander Vogel

- File correct_loop_sims_V-rescale_semiisotropic.png added

Berk Hess wrote:

> Of course the box cannot deform with isotropic pressure coupling or NVT. The deformation is only possible when x/y is coupled independently from z. That does not tell us anything about the source of the issue.

Thanks for the update...wasn't aware of that. As I said I'm no expert on the mdp parameters. I always use semiisotropic as I'm only simulating membranes. Anything else you want me to run for this system while I'm on vacation?

Also I just checked three new simulations I started before the weekend. They are a new charmm-gui setup of a very similar system (replaced a few amino acids of the receptor)...same thing happens...again to all three replicas. So with the same mdp it basically happens to me always. I'll now change them to Nose-Hoover instead of V-rescale. We'll see if that fixes it as before after my vacation.

Alexander

#68 - 09/03/2019 08:04 PM - Berk Hess

A question about other MD packages. Do other packages have a coupling option equivalent to "semi-isotropic" in GROMACS? If so, have people tested a system which shows issues in GROMACS with that in another package?

#69 - 09/03/2019 09:16 PM - Justin Lemkul

Berk Hess wrote:

> A question about other MD packages. Do other packages have a coupling option equivalent to "semi-isotropic" in GROMACS? If so, have people tested a system which shows issues in GROMACS with that in another package?

OpenMM has a semiisotropic MC barostat. I plan to set up a run from Alexander's coordinates sometime soon, I just haven't had a moment to get it going yet.

#70 - 09/05/2019 02:47 PM - jose flores-canales

- File box_all.png added

- File distances_all.png added

- File box_dimensions_distance_amber.png added

Berk Hess wrote:

> A question about other MD packages. Do other packages have a coupling option equivalent to "semi-isotropic" in GROMACS? If so, have people tested a system which shows issues in GROMACS with that in another package?

Semiisotropic NPT with MC barostat for PMEMD/AMBER (see Z-box dimension in box_dimensions_distance_amber.png). I described the simulation parameters in message #60.

I have been comparing the GPCR protein structure fluctuations between the successful and unsuccessful GMX 2018.6 runs. I noticed that protein fluctuations (shrinking distances between the extra- and intra-cellular terminus of two contiguous helices), along the Z axis resulted in lipid bilayer deformation (distances_all.png). At least for my GPCR system on POPC lipids.

To test that protein fluctuations might disrupt the lipid bilayer, I ran two different sets of GMX trajectories starting from the same configuration (obtained before the lipid bilayer starts collapsing). The only difference consists on: one set with protein positional restraints and the other one with no restraints. See the results in box_all.png. So based on these results, I think the problem is the P-R barostat integrator and notorious with some combinations of thermostat and barostat integrators, such that when a protein has an specific type of fluctuations the lipid bilayer thickness shrinks. This could be specific to GMX simulations of GPCR proteins, which have great flexibility in the membrane (it should also be the case in MD simulations).

All GMX runs were produced with Langevin thermostat and semiisotropic P-R barostat. C36FF was used for all simulations. AMBER MD simulations are from a slightly different system (protein+ligand), but it shows that the protein has large fluctuations (compared to GMX trajectories), while the lipid bilayer maintains in average its dimensions (see box_dimensions_distance_amber.png).

#71 - 09/17/2019 06:02 PM - Berk Hess

Someone in our group had a contracting instead of expanding bilayer area. He noticed that the initial bilayer structure produced by CHARMM-gui had

non-level head group planes and a region that looked like some sort of cavity inside the bilayer. He noticed similar issues in the initial structures attached at this issue here and those persist up to the time when the box starts to deform. So we have another hypothesis: no GROMACS or force field issue, but an issue in CHARMM-gui which produces bad initial bilayer conformation under certain circumstances (but we don't know which).

#72 - 09/18/2019 03:06 PM - jose flores-canales

- File step5_charmm2gmx.pdb added

- File top_view_highlighted.png added

- File Lateral_view_HG.png added

I haven't found irregularities on the head group planes and an obvious cavity is difficult to find. This membrane is asymmetric (two more lipids in one leaflet). One of the graphs highlights two lipids on the protein border (top_view_highlithted.png), while the space on the other leaflet is occupied by protein helices. I attach the initial charmm-gui structure.

#73 - 09/21/2019 11:28 AM - Berk Hess

Thanks to great investigation work of one of our PhD students I found the source of the issue: an index bug in the shift force correction for the CMAP term. The pressure is incorrect when a the second and third atom in a CMAP dihedral reside in different periodic images. This can happen when a protein is positioned over one of the box edges.

The fix is swapping indices one on line:

```
index cbf89ad42..515e186bd 100644
--- a/src/gromacs/listed-forces/bonded.cpp
+++ b/src/gromacs/listed-forces/bonded.cpp
@ -3267,7 +3267,7 @ cmap_dihs(int nbonds,
rvec_inc(fshift[t21], f1_k);
rvec_inc(fshift[t31], f1_l);

-       rvec_inc(fshift[t21], f2_i);
+       rvec_inc(fshift[t12], f2_i);
rvec_inc(fshift[CENTRAL], f2_j);
rvec_inc(fshift[t22], f2_k);
rvec_inc(fshift[t32], f2_l);
```

#74 - 09/21/2019 11:35 AM - Berk Hess

- Category set to mdrun

- Status changed from New to Fix uploaded

- Target version set to 2019.4

#75 - 09/21/2019 11:39 AM - Berk Hess

PS: To be clear, this is unrelated to (non) use of GPUs. CMAP is always computed on the CPU.

Also the forces and energy are always correct and the pressure is correct when atoms involved in CMAP are in the same periodic image (which is why we did not observe this issue earlier).

#76 - 09/21/2019 12:37 PM - Justin Lemkul

Outstanding - thanks for the hard work in tracking this down. That was a tricky issue!

#77 - 09/21/2019 12:58 PM - Berk Hess

Thanks everyone for their help!

This was one of the hardest issue to track down. What put me off track is that I somehow thought that this also occurred in a pure membrane system without a protein.

At least the good news is that all results of simulations with the protein in the middle, which is the most common situation, are correct.

#78 - 09/21/2019 02:03 PM - Ramon Guixa-Gonzalez

Berk Hess wrote:

> Thanks everyone for their help!
>
> This was one of the hardest issue to track down. What put me off track is that I somehow thought that this also occurred in a pure membrane system without a protein.
>
> At least the good news is that all results of simulations with the protein in the middle, which is the most common situation, are correct.

Oh my god, these is so good news. The idea of not being able to use GROMACS was just killing me...
Many many thanks to all!!

#79 - 09/21/2019 02:21 PM - jose flores-canales

Amazing news! Thanks a lot Berk and your team for this detective work. I will test the patched GMX2018 ASAP.

#80 - 09/23/2019 02:50 PM - Berk Hess

I can also add that the error in the pressure is on the order of 1 bar, so very small. Therefore this likely went unnoticed with isotropic pressure coupling and likely even the effect on relevant results was negligible. On the other hand, with anisotropic pressure coupling there is no hydrostatic pressure preventing large box deformations, so a continuous error of 1 bar can completely deform a box.

Files

2016_6.zip	43.4 MB	02/22/2019	Ramon Guixa-Gonzalez
jp101759q.pdf	4.36 MB	08/01/2019	Alexander Vogel
mdout.mdp	10.7 KB	08/02/2019	Alexander Vogel
run.mdp	1.76 KB	08/02/2019	Alexander Vogel
working_vs_3nonworking_sims.png	26.5 KB	08/05/2019	Alexander Vogel
run.mdp	1.76 KB	08/05/2019	Alexander Vogel
step7_production.mdp	1.16 KB	08/05/2019	Alexander Vogel
1_microsecond_without_problems.png	23.2 KB	08/13/2019	Alexander Vogel
likeCharmm.jpeg	758 KB	08/14/2019	Ramon Guixa-Gonzalez
update.png	22.9 KB	08/20/2019	Alexander Vogel
Box_fluctuations.png	249 KB	08/20/2019	Mayukh Chakrabarti
production.mdp	2.5 KB	08/20/2019	Mayukh Chakrabarti
collapse_after_mdp_change.png	27 KB	08/21/2019	Alexander Vogel
mdp_before_collapse.mdp	1.16 KB	08/21/2019	Alexander Vogel
changed_mdp_that_leads_to_collapse.mdp	1.8 KB	08/21/2019	Alexander Vogel
production_langevin.mdp	1.3 KB	08/21/2019	jose flores-canales
collapse_after_mdp_change_update.png	27 KB	08/22/2019	Alexander Vogel
POPC_256x3.png	16.6 KB	08/22/2019	Alexander Vogel
POPC_1024(red)_4096(blue).png	14.3 KB	08/22/2019	Alexander Vogel
berkhess.tar.gz	24 MB	08/22/2019	Alexander Vogel
pic1.jpeg	433 KB	08/22/2019	Ramon Guixa-Gonzalez
pic2.png	91.7 KB	08/22/2019	Ramon Guixa-Gonzalez
largearea_t232ns.pdb	5.43 MB	08/25/2019	Berk Hess
now_collapsed.png	25.9 KB	08/26/2019	Alexander Vogel
run3_snapshots.png	17.7 KB	08/27/2019	Alexander Vogel
run3_46.2.trr	1.09 MB	08/27/2019	Alexander Vogel
run3_20.0.trr	1.09 MB	08/27/2019	Alexander Vogel
run3_70.0.trr	1.09 MB	08/27/2019	Alexander Vogel
eq7_3.gro	6.29 MB	08/27/2019	Alexander Vogel
summary.txt	4.05 KB	08/27/2019	Alexander Vogel
step5_production_changed_to_collapsing.mdp	1.76 KB	08/28/2019	Alexander Vogel
waterbox.png	15.5 KB	08/28/2019	Alexander Vogel
iso_vs_semiiso.png	28.9 KB	08/29/2019	Alexander Vogel
box_dimensions.png	7.42 KB	08/29/2019	jose flores-canales
box_dimension_gromacs.png	44.9 KB	08/29/2019	jose flores-canales
iso_vs_semiiso_update1.png	28.9 KB	08/30/2019	Alexander Vogel
iso_vs_semiiso_update2.png	27.7 KB	09/02/2019	Alexander Vogel
correct_loop_sims_V-rescale_semiisotropic.png	27.2 KB	09/02/2019	Alexander Vogel
box_all.png	10.4 KB	09/05/2019	jose flores-canales
distances_all.png	14.7 KB	09/05/2019	jose flores-canales
box_dimensions_distance_amber.png	10.2 KB	09/05/2019	jose flores-canales
step5_charmm2gmx.pdb	9.98 MB	09/18/2019	jose flores-canales

top_view_highlighted.png	1.06 MB	09/18/2019	jose flores-canales
Lateral_view_HG.png	194 KB	09/18/2019	jose flores-canales