

GROMACS - Bug #1351

Simulations which are stable on Gromacs 4.5.3 are crashing on Gromacs 4.6.3.

10/04/2013 01:44 PM - Floris van Eerden

Status: Closed	
Priority: Normal	
Assignee: Berk Hess	
Category: mdrun	
Target version:	
Affected version - extra info:	Difficulty: uncategorized
Affected version: 4.6.3	
Description	
<p>I have systems which I can run without problem with GROMACS 4.5.5, but crash sooner or later with Gromacs 4.6.3. The system is a protein monomer (Photosystem II) embedded in a thylakoid membrane, of which some of the lipids contain constraints. Everything is in a hexagonal box. In this particular simulation all the charges of the protein are set to zero. But the crashes also occur when I try to simulate the dimer, with all charges at their normal value. I tried running it with domain decomposition, particle decomposition or on one thread, but in all cases i get sooner or later (sometimes the simulation runs quite some time) LINCS erros or a segmentation fault.</p> <p>Sincerely Floris</p>	

Associated revisions

Revision d5af279d - 10/16/2013 01:51 PM - Berk Hess

avoid division by zero in SIMD angles and dihedrals

The SIMD accelerated angle and dihedral code did not (correctly) check for dividing by zero, which can happen with aligned bonds. Fixes #1351

Change-Id: I326f90fca87ab5cca493204de4a58655465634ca

History

#1 - 10/04/2013 01:49 PM - Floris van Eerden

PS It is a Coarse Grained Martini system

#2 - 10/08/2013 01:50 PM - Mark Abraham

- Status changed from New to Accepted

Indeed, your startmd10tpr.tpr ran for 40ns with 4.5.7 and mdrun -rdd 1.8, and died with release-4-6 branch almost immediately.

I am trying a bisection to narrow down the point in the commit history that the crash first becomes reproducible, but that will take considerable time because success can take hours to measure! That said, there's a half dozen commits that are known offenders for latent bugs, so probably I will identify one of them and further diagnosis will be required.

Simpler examples of simulations that are stable or not (i.e. with respect to segfault) would be very useful - e.g. is your membrane stable (or not) on its own? Are any of the protein sub-units stable in water on their own? Is one of them fully representative of the complexity of the interactions present in any of them? I'm not worried if a protein sub-unit that isn't stable in the folding sense - just whether it is useful for probing what the numerical instability is.

#3 - 10/08/2013 03:55 PM - Floris van Eerden

Hello,

I am happy to read that you got the same error! So it is not just me ;)

About other systems:

The only other system which I ran on Gromacs 4.6.3 is the PSII dimer (with all charges normal) embedded in the thylakoid membrane (so an even bigger system) and that one was also not stable.

On Gromacs 4.5.5 the PSII dimer is stable at least up to 2 micro seconds, and the membrane itself also. But I do not have any of these simulations with Gromacs 4.6.3. I will try to run a patch of membrane overnight with Gromacs 4.6.3 and give you the result tomorrow.

Floris

#4 - 10/08/2013 04:44 PM - Mark Abraham

Thanks! I have tentatively identified the commit that introduced the Verlet cut-off scheme as the source of the problem. That commit is one of the serial offenders, and is relatively good news, because it should have had very little impact on your simulation. That might make the next stage easier.

#5 - 10/09/2013 08:36 PM - Floris van Eerden

- *File membrane added*

I did simulations of a membrane without protein, and also this system crashed really quick, see attachment. In the coming days I will try simulations of only the protein, without lipids.

#6 - 10/10/2013 02:02 AM - Mark Abraham

Thanks - that's already a very helpful simplification.

#7 - 10/10/2013 05:56 PM - Mark Abraham

Good news - running the first simulation with no SIMD at all also ran for 40ns (so the problem is with a SIMD path). That patch was nasty and introduced some SIMD to bonded interactions, as well as the Verlet cut-off scheme.

With non-bonded SIMD only: ran > 2ns.

With SIMD everywhere except proper dihedrals: run > 0.5ns

With SIMD everywhere except force-only proper dihedrals: run > 1ns.

Otherwise, crash.

This suggests that with your large time steps, something like that some atoms are going close enough to co-linear that the SIMD routines are returning a garbage value. We do have some checks for this kind of thing, but it would seem we have missed one.

#8 - 10/11/2013 11:43 AM - Berk Hess

Mark, it seems you didn't check the SIMD angles, or did you?

My guess was that the problem is in there, since in Martini angles could be 180 degrees.

Could one of you try the change below, which should fix this issue. If it works I will submit a proper, clean fix.

```

--- a/src/gmplib/bondfree.c
+++ b/src/gmplib/bondfree.c
@@ -1141,6 +1141,14 @@ angles_noener_simd(int nbonds,

    cos_S      = gmx_mul_pr(rij_rkj_S, gmx_mul_pr(nrij_1_S, nrkj_1_S));

+   {
+       gmx_mm_pr minone_S;
+       minone_S = gmx_set1_pr(-1.0);
+       cos_S = gmx_max_pr(cos_S, minone_S);
+   }
+   theta_S    = gmx_acos_pr(cos_S);

    sin_S      = gmx_invsqrt_pr(gmx_max_pr(gmx_sub_pr(one_S, gmx_mul_pr(cos_S, cos_S)),

```

#9 - 10/11/2013 12:22 PM - Berk Hess

The angles are not the problem, but rereading Mark's comment I think he means to say that the issues are in the SIMD dihedral code.

But the dihedral code looks very robust. There are is only one division, which involves the distance between atoms j and k. But in Martini atoms j and k can not end up on top of each other, can they?

#10 - 10/11/2013 12:39 PM - Mark Abraham

By "SIMD everywhere" I meant including SIMD angles (ie. by hacking the `#ifdef SIMD_BONDED` suitably), but I will try your hack anyway.

#11 - 10/11/2013 12:49 PM - Mark Abraham

Yes - my diagnosis so far is that the issue is shown by `pdihs_noener_simd()`, because if the only change I make to the code is to call its fallback routine `pdihs_noener()`, then I get the > 1ns run I reported earlier.

#12 - 10/11/2013 12:50 PM - Mark Abraham

Tried Berk's angle hack, still crashed around 24000 steps.

#13 - 10/11/2013 03:48 PM - Berk Hess

- *Status changed from Accepted to Resolved*

The SIMD dihedral handled 180 degree bond angles incorrectly. I uploaded a fix:

<https://gerrit.gromacs.org/#/c/2697/>

#14 - 10/11/2013 03:56 PM - Floris van Eerden

- *File NVT_noMembrane added*

Hello,
I am happy that you found the problem.
As I said I would, I am also doing a simulation without the membrane (the protein has however some co-crystallized lipids which I left), with NVT. And now the simulation already ran for 329800 steps. Would anybody have an idea while it does not crash in this case?

#15 - 10/11/2013 04:12 PM - Berk Hess

It will only crash when two adjacent bonds in a dihedral are (nearly) co-linear. The chance of this happening is very small and it could be that that only happens in Martini lipids.

#16 - 10/11/2013 04:27 PM - Floris van Eerden

Thanks for the reply.
It seems indeed the case, because the previous simulation also seems to work well with pressure coupling.

#17 - 10/11/2013 04:48 PM - Mark Abraham

I have the protein+membrane stable with Berk's fix on gerrit for > 4ns so far. Will run the full 40 ns

#18 - 10/16/2013 01:55 PM - Berk Hess

- % Done changed from 0 to 100

Applied in changeset [d5af279d407572fe0f8f09aacf4e773961222691](#).

#19 - 12/03/2013 04:23 PM - Rossen Apostolov

- Status changed from Resolved to Closed

Files

runfiles	9.53 MB	10/04/2013	Floris van Eerden
membrane	2.2 MB	10/09/2013	Floris van Eerden
NVT_noMembrane	3.61 MB	10/11/2013	Floris van Eerden