

## GROMACS - Bug #1272

### Energy minimization with domain decomposition crashes

06/03/2013 06:31 AM - David van der Spoel

<b>Status:</b> Closed	
<b>Priority:</b> Normal	
<b>Assignee:</b> Berk Hess	
<b>Category:</b>	
<b>Target version:</b>	
<b>Affected version - extra info:</b>	<b>Difficulty:</b> uncategorized
<b>Affected version:</b> 4.6.2	

#### Description

A colleague asked me to look at a very large simulation he could not make work. A simple water slab of ~ 129 M atoms.

```
genconf -f spc216.gro -nbox 200 200 5
grompp -v -c out -f em -o em
mdrun -s em
```

My colleague tested on a big GPU system, I tested on a regular AMD 48 core system. it crashes with a DD error for me.

Steepest Descents:

Tolerance (Fmax) = 1.00000e-03

Number of steps = 25

Step 0:

The charge group starting at atom 7 moved more than the distance allowed by the domain decomposition (31.034332) in direction X

distance out of cell -341.377655

Old coordinates: 0.000 0.000 0.000

New coordinates: 0.000 0.000 0.000

Old cell boundaries in direction X: 341.378 372.412

New cell boundaries in direction X: 341.378 372.412

#### Associated revisions

##### Revision ec49c7fa - 06/05/2013 03:58 PM - Berk Hess

fixed DD internal state corruption with energy minimization

With energy minimization we need to reload old DD states after steps are rejected. There were two bookkeeping issues in the reload.

This could lead to all kinds of, but no silent, errors.

Fixes #1272

Change-Id: Ia44c3bad27f3efdee76fa93dd281690e44dde700

#### History

##### #1 - 06/03/2013 02:14 PM - Berk Hess

- Status changed from New to Accepted

- Assignee set to Berk Hess

I assume you were running the group scheme, whereas the GPUs will have to run the Verlet scheme?

Anyway, I get the same error with group, Verlet and also with RF iso PME.

But the error is different from yours. The energy and forces at step 0 seem to be ok:

Step= 0, Dmax= 1.0e-01 nm, Epot= -2.02225e+09 Fmax= 9.12274e+02, atom= 622087

But step 1 gives a settle error and -pforce 5000 shows large forces.

The coordinates are very strange: regular and identical for many particles. It seems that somewhere some coordinate buffers got messed up.

step 1 atom 97635271 x 279.397 248.914 1.022 force 5.47555e+03

step 1 atom 97635919 x 279.397 248.914 2.884 force 5.44673e+03  
step 1 atom 97636567 x 279.397 248.914 4.746 force 5.42407e+03  
step 1 atom 97637215 x 279.397 248.914 6.608 force 5.44247e+03  
step 1 atom 97634623 x 279.397 248.914 8.470 force 5.43415e+03  
step 1 atom 97638511 x 279.397 250.776 1.022 force 5.47089e+03  
step 1 atom 97639159 x 279.397 250.776 2.884 force 5.44283e+03  
step 1 atom 97639807 x 279.397 250.776 4.746 force 5.42009e+03  
step 1 atom 97640455 x 279.397 250.776 6.608 force 5.43953e+03  
step 1 atom 97637863 x 279.397 250.776 8.470 force 5.42750e+03  
step 1 atom 97641751 x 279.397 252.638 1.022 force 5.47208e+03  
step 1 atom 97642399 x 279.397 252.638 2.884 force 5.44174e+03  
step 1 atom 97643047 x 279.397 252.638 4.746 force 5.42176e+03  
step 1 atom 97643695 x 279.397 252.638 6.608 force 5.43861e+03  
step 1 atom 97641103 x 279.397 252.638 8.470 force 5.42884e+03  
step 1 atom 97644991 x 279.397 254.500 1.022 force 5.47362e+03

## #2 - 06/03/2013 04:46 PM - Berk Hess

- Status changed from Accepted to In Progress

MD seems to run fine, so that indicates the issue is somewhere in the energy minimization.

## #3 - 06/03/2013 05:38 PM - Berk Hess

I now have a 21600 water system, genconf -nbox 10 10 1, which gives DD constraint missing atom errors when run with -dd 4 4 1 or -dd 8 2 1 or -dd 2 8 1, but not -dd 16 1 1. Less than 16 ranks seems to work fine.

So this seems to have nothing to do with system size.  
(and it only seems to happen with energy minimization)

## #4 - 06/03/2013 09:22 PM - Berk Hess

- Subject changed from Large simulation (129 Matoms) crashes to Energy minimization with domain decomposition crashes

- Status changed from In Progress to Fix uploaded

I changed the subject, as the bug was not much related to large systems, but rather to energy minimization.

## #5 - 06/05/2013 04:00 PM - Berk Hess

- Status changed from Fix uploaded to Resolved

- % Done changed from 0 to 100

Applied in changeset [ec49c7fa35c551ec7e9ef936ad48d3f2dd3efacd](#).

## #6 - 06/06/2013 06:50 PM - David van der Spoel

mdrun crashes for me with SEGV due to a big box anyway, even when running algorithm=md rather than steep:

```
genconf -f spc216.gro -nbox 400 400 2
```

could it be out of memory? But that would give another crash wouldn't it?

## #7 - 06/06/2013 06:50 PM - David van der Spoel

The nodes have 48 cores and 128 Gb RAM.

## #8 - 06/14/2013 05:00 PM - David van der Spoel

- Status changed from Resolved to Accepted

Reopening, because large systems do crash as indicated previously.

## #9 - 05/16/2014 12:20 PM - Rossen Apostolov

I'm trying to reproduce this but on a 32core, 32GB machine. The input is made with

```
genconf -f spc216.gro -nbox 100 200 2,
```

bigger systems give out of memory crashes.

The above system runs with:

- default mdrun with 24PP (8x3x1) and 8PME nodes

- -dd 8 4 1
- -dd 16 2 1

It eventually crashes with:

- -dd 4 4 2 , message "A charge group moved too far between two domain decomposition steps  
This usually means that your system is not well equilibrated") and then "The charge group starting at atom 19830606 moved more than the distance allowed by the domain decomposition (1.862060) in direction Z"
- -dd 8 2 2 , same

All runs give plenty of e.g. "Water molecule starting at atom 23512618 can not be settled".

**#10 - 06/12/2014 02:11 AM - Erik Lindahl**

But does it work great when only using a single thread (i.e., without SETTLE errors)? Otherwise it's an equilibration problem.

**#11 - 06/12/2014 06:18 PM - Berk Hess**

Could you run it with reaction-field as well, to see if the issue might be in PME?  
We have run systems this large, but mostly (only?) without PME.

I first thought is might be an indexing issue, since 129M is close to all 31-bits in an integer, but 100x100x2 should be small enough to run fine. You can run with GMX\_DD\_DEBUG=1 if we suspect there is a DD indexing error.

spc216.gro is well equilibrated, so it can't be an equilibration issue.

**#12 - 06/25/2014 11:34 AM - Rossen Apostolov**

Strangely I couldn't reproduce my own tests, i.e. the system doesn't crash despite the numerous settle errors. Using a single thread clears most of the settle errors, in fact only 1 of the water molecules is reported. RF also produces a lot of settle errors. This is with 4.6.5.

**#13 - 06/16/2015 11:50 PM - Erik Lindahl**

- *Status changed from Accepted to Resolved*

Original problem was resolved, and since Rossen couldn't reproduce the problem he saw I'll assume this is working as intended.

**#14 - 06/16/2015 11:50 PM - Erik Lindahl**

- *Status changed from Resolved to Closed*

**Files**

em.mdp	8.58 KB	06/03/2013	David van der Spoel
topol.top	154 Bytes	06/03/2013	David van der Spoel
spc216.gro	28.6 KB	06/03/2013	David van der Spoel