

GROMACS - Bug #315

append option continuing to write to previous trajectory crashed when dealing with a large traj.trr file

04/16/2009 11:26 PM - ckcumaa empty

Status:	Closed	
Priority:	Normal	
Assignee:	Erik Lindahl	
Category:	mdrun	
Target version:	4.0	
Affected version - extra info:		Difficulty: uncategorized
Affected version:		

Description

Overview:

"append" command to continue to write to the previous trajectory file crashes once the traj.trr file becomes larger than 2GB in Gromacs 4.0.4.

Steps to reproduce.

1. Run simulation with uploaded input files (3000 nitrogen gas system)
: mdrun -cpt 10
2. kill the job and just restart with a checkpoint file
3. use append option to continue to write a traj.trr file.
: mdrun -cpi state.cpt -cpt 10 -append yes

results.

restart with a checkpoint file and append option works only when the size of traj.trr file is less than 2GB. It always crashes when the traj.trr is larger than 2GB.

additional information:

I used a parallel running (4 cpu) and double precision gromacs so my actual restart command is
: mpirun -np 4 mdrun_d -cpi state.cpt -cpt 10 -append yes

Detailed error:

Once the size of traj.trr file becomes larger than 2GB then the "append" option does not work anymore and shows the following error message:

```
"Truncation of traj.trr file failed"
```

I've tried a lot of different systems but it always showed the same error message once the traj.trr became larger than 2GB.

At first, I thought it might be the problem of a compiler so I requested our local system administrator to look into this problem and he replied with the following:

```
The problem is the call to truncate() on line 1239 of checkpoint.c.
```

```
The problem is the value of outputfiles[i].offset for the trajectory file.
```

```
I put a quick modification in there to call strerror(errno) which on RHEL4 is set by truncate() when it fails.
```

```
The error was "invalid arguments".
```

I checked, nyx (the cluster) is correctly setting sizeof(outputfiles [i].offset) to 8 bytes (64b) thus the problem observed with a restart failing when a trajectory file is over 2Gbyte should not be happening.

```
The type: gmx_file_position_t
```

```
correctly uses the type off_t for the offset.
```

```
I did not look at the actual writing/reading of the checkpoint file.
```

```
As far as I can tell I don't see why this is failing, off_t is big enough to set offsets larger than 2Gbyte.
```

```
This issues happens with both PGI and GNU compilers.
```

We already report this issue to gmx user forum but we have not received any reply.

At this moment, we are out of ideas and decided to post this issue to bug report

History

#1 - 04/16/2009 11:29 PM - ckcumaa empty

Created an attachment (id=361)
input files for reproducing the error

#2 - 05/25/2009 06:03 PM - Erik Lindahl

Hi,

I seem to have problems reproducing this on my local system. How are you writing the file - is it possibly to a remote file server using a version of NFS that might not support large files?

Cheers,

Erik

#3 - 11/20/2009 03:29 PM - Ondrej Marsalek

This also happens to me at HLRN (<https://www.hlrn.de>) on the Lustre file system. Will provide more information, if needed, but at the moment I am not sure what exactly would be relevant.

#4 - 06/03/2010 11:16 AM - Rossen Apostolov

Hi Ondrej and Kyungchan,

can you still reproduce this with the current git master branch?

#5 - 06/11/2010 11:06 AM - David van der Spoel

I can not reproduce the error either, running in double precision over NFS to a file system which is 32 bit. The node I'm running on is 64 bit though. Please reopen if you can reproduce this still.

Files

example.zip	137 KB	04/16/2009	ckcumaa empty
-------------	--------	------------	---------------