GROMACS - Bug #499

mdrun_mpi fails randomly

08/10/2010 01:49 PM - Timofey Kushnir

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

Description

Created an attachment (id=509)
liquid cell test model

Every time I run mdrun_mpi using 80 processes with command line

`which mpirun` --hostfile hostfile --mca btl openib,self -np 80 `which g_mdrun_mpi` -v

mdrun_mpi ends with an error similar to this

...
[node08:15172] * **An error occurred in MPI_Waitall**
[node08:15172] on communicator MPI_COMM_WORLD
[node08:15172] **MPI_ERR_TRUNCATE: message truncated**
[node08:15172] * MPI_ERRORS_ARE_FATAL (your MPI job will now abort)
...

With 64 processes mdrun_mpi fails only sometimes.
With number of processes grater 64 mdrun_mpi fails almost always.

Is it a bug in mdrun_mpi, or limitation of the model?

History

#1 - 08/11/2010 09:43 AM - Berk Hess

My guess would rather be that there is a problem with your mpi installation/configuration.
Do you get any warnings or errors before the crash in stdout, stderr or the log file?

PS 80 nodes is far beyond the efficient scaling limit for this system, which would be around 10 nodes.

Berk

#2 - 08/11/2010 01:44 PM - Timofey Kushnir

(In reply to comment #1)

> My guess would rather be that there is a problem with your mpi installation/configuration.

Unfortunately, no. I've tried all MPI realizations we have (including commercial ones) and use all available compiler suites (including commercial ones). The result is the same.

> Do you get any warnings or errors before the crash in stdout, stderr or the log file?

No, the failed run stdout trailing lines look like this:

...
vol 0.54! imb F 31% pme/F 0.73 step 276800, remaining runtime: 26 s
vol 0.53! imb F 25% pme/F 0.76 step 276900, remaining runtime: 26 s

vol 0.58! imb F 23% pme/F 0.77 step 277000, remaining runtime: 26 s
vol 0.60! imb F 23% pme/F 0.74 step 277100, remaining runtime: 26 s
vol 0.57! imb F 28% pme/F 0.71 step 277200, remaining runtime: 26 s
[node09:10284] * **An error occurred in MPI_Waitall**
[node09:10284] on communicator MPI_COMM_WORLD
[node09:10284] **MPI_ERR_TRUNCATE: message truncated**
[node09:10284] * MPI_ERRORS_ARE_FATAL (your MPI job will now abort)
----
mpirun has exited due to process rank 8 with PID 10284 on
node node09 exiting without calling "finalize". This may
have caused other processes in the application to be
terminated by signals sent by mpirun (as reported here).

...

> PS 80 nodes is far beyond the efficient scaling limit for this system, which
> would be around 10 nodes.

OK, but on 10 nodes it takes about 25 minutes to finish while on 72 nodes about 7 minutes, the time difference is substantional.

Well, I found that for 72 nodes with this commandline

`which mpirun` --hostfile hostfile --mca btl openib,self -np 72 `which g_mdrun_mpi` -v -dlb yes -npme 8 -dd 4 4 4

the run finishes correctly in about 20% cases showing good run time.

Unfortunately, no core files are dropped so I cannot distinguish where the error is. Could you please help me in further investigations?

#3 - 08/11/2010 02:48 PM - Berk Hess

Have you also tried openmpi?

You might need to set the right limit in your shell to enable core dumps.

Berk

#4 - 08/11/2010 03:00 PM - Timofey Kushnir

(In reply to comment #3)

> Have you also tried openmpi?

These errors are produced using OpenMPI 1.4.1

> You might need to set the right limit in your shell to enable core dumps.

It's already done. On all nodes `ulimit -c' shows `unlimited'.

#5 - 08/22/2010 10:47 AM - Berk Hess

I assume there nothing new to report about this issue?

What version of Gromacs are you using exactly?

Berk

#6 - 08/23/2010 10:00 AM - Berk Hess

Created an attachment (id=523)
fixed src/mdlib/pme_pp.c for Gromacs 4.0

#7 - 08/23/2010 10:02 AM - Berk Hess

There was a bug when a direct-space node had to send 0 charges to a pme-only node.
I have fixed it for the next 4.5 beta and have attached a fixed file for 4.0.

PS I would like to mention again that this is far above the scaling limit,
having 0 charges on a node is not ideal.

Thanks for reporting this,

Berk

Files

topol.tpr	86.1 KB	08/10/2010	Timofey Kushnir
pme_pp.c	12.8 KB	08/23/2010	Berk Hess