GROMACS - Bug #1101

fix & improve CPU oversubscription handling

01/02/2013 08:00 PM - Szilárd Páll

Status:	Closed		
Priority:	Normal		
Assignee:	Berk Hess		
Category:	mdrun		
Target version:	4.6		
Affected version - extra info:		**Difficulty:**	uncategorized
Affected version:			

Description

Oversubscription of the available CPU cores should be avoided in most (if not all) cases as it results in bad performance. This is made even worse by thread pinning. The current oversubscription check implemented in the gmx_omp_nthreads module is incorrect with separate PME nodes that use a different number of OpenMP threads than the PP nodes.

The following improvements are required:

- implementing correct check outside of gmx_omp_nthreads - it's not (only) OpenMP-related, it can happen with pure MPI/tMPI;
- turning off thread pinning when oversubscription is detected.

Associated revisions

Revision a1bd375a - 01/11/2013 07:05 PM - Erik Lindahl

Added basic CPU topology information to cpuid code

We can now detect the locality of hardware threads, cores,
and packages for Intel and AMD CPUs under Linux and Windows.
In particular, this provides an array with locality order
for logical processors that can be used to optimize placement.
Refs #1086, #1101.

Change-Id: I3f7985b1b67729376918c5a135b9157a9086235e

History

#1 - 01/07/2013 12:02 PM - Mark Abraham

Can we deal with this in the next fortnight or so for 4.6, or push it back to 4.6.1?

#2 - 01/09/2013 12:18 AM - Szilárd Páll

Mark Abraham wrote:

> Can we deal with this in the next fortnight or so for 4.6, or push it back to 4.6.1?

Depends what does "we" mean.

I have not strived to fixing the current code because it requires yet another splitting of the default communicator (the same is already done in 2-3 places) and doing a per-node thread enumeration (at least I don't know of a simpler way). However, as we discussed earlier (and now I'm realizing an issue for this is missing), there should instead be communicators set up and stored for inter- and intra-node communication which would remove current redundancy and enable implementing other features. This task is pretty simple, so I might as well just do it myself if nobody can pitch in.

#3 - 01/18/2013 03:59 PM - Erik Lindahl

- Status changed from New to Closed

Fixed by gerrit 2051, at least for what is realistic to expect in 4.6.