

GROMACS - Bug #2345

Avoid duplicate over-subscription note+warning

12/12/2017 07:39 AM - Mark Abraham

Status: Closed	
Priority: Normal	
Assignee: Mark Abraham	
Category: mdrun	
Target version: 2018-beta3	
Affected version - extra info:	Difficulty: uncategorized
Affected version: 2018-beta2	
Description	
The thread-affinity and OpenMP management code both warn about oversubscription:	
<pre>gmx mdrun -s reference_s -deffnm test -nt 5</pre>	
...	
<pre>GROMACS: gmx mdrun, version 2018-beta2-dev-20171211-2e91fcf47 Executable: /home/mabraham/git/r2018/build-cmake-clang-debug/install/bin/gmx Data prefix: /home/mabraham/git/r2018/build-cmake-clang-debug/install Working dir: /home/mabraham/git/regressiontests/complex/nbnxn_rf Command line: gmx mdrun -s reference_s -deffnm test -nt 5</pre>	
<pre>Reading file reference_s.tpr, VERSION 2016-beta1-dev-20160524-3a9e67d (single precision) Note: file tpx version 110, software tpx version 112 Changing nstlist from 10 to 25, rlist from 0.944 to 1.068</pre>	
<pre>Using 5 MPI threads</pre>	
<pre>WARNING: Oversubscribing the available 4 logical CPU cores with 5 thread-MPI threads. This will cause considerable performance loss!</pre>	
<pre>NOTE: Oversubscribing the CPU, will not pin threads</pre>	
<pre>NOTE: Thread affinity was not set.</pre>	
The one in the OpenMP code only works when there's no separate PME ranks, and not on nodes with only PME ranks, assumes that each PP rank has the same number of threads (which might be wrong if the user has set up MPI creatively), and has an existing todo to move it elsewhere and support PME-only ranks, so probably the best approach is to remove it.	
Related issues:	
Related to GROMACS - Bug #2342: Avoid yelling about thread pinning twice	Closed

Associated revisions

Revision 7607e626 - 12/19/2017 02:02 AM - Berk Hess

Make oversubscription warning consistent

The hardware thread oversubscription warning was only issued with OpenMP and without separate PME ranks. Now it actually reduces the thread count over the physical node.
Also moved the thread affinity up to the earliest possible point.

Refs #2345

Change-Id: Ifdf62c723fd87b0ddaab0df1e2f1bf36b461ea33

History

#1 - 12/12/2017 07:39 AM - Mark Abraham

- Related to Bug #2342: Avoid yelling about thread pinning twice added

#2 - 12/12/2017 08:14 AM - Gerrit Code Review Bot

Gerrit received a related patchset '1' for Issue [#2345](#).
Uploader: Mark Abraham (mark.j.abraham@gmail.com)
Change-Id: gromacs~release-2018~lb9f91a5e4db73b399060a25ab7e927a785231fdc
Gerrit URL: <https://gerrit.gromacs.org/7332>

#3 - 12/12/2017 08:36 AM - Berk Hess

- Status changed from New to Resolved
- Assignee changed from Mark Abraham to Berk Hess

#4 - 12/12/2017 08:42 AM - Mark Abraham

- Status changed from Resolved to Fix uploaded

The fix has just a +1 right now :-)

#5 - 12/12/2017 08:51 AM - Berk Hess

- Assignee changed from Berk Hess to Mark Abraham

I misinterpreted this issue. The double affinity note should be fixed and is fixed with my fix. But the oversubscription warning has nothing to do with affinities (although it also triggers the affinity note). The oversubscription performance warning should certainly stay. We should also note that we do not set affinities.

#6 - 12/12/2017 09:01 AM - Mark Abraham

Berk Hess wrote:

I misinterpreted this issue. The double affinity note should be fixed and is fixed with my fix. But the oversubscription warning has nothing to do with affinities (although it also triggers the affinity note). The oversubscription performance warning should certainly stay. We should also note that we do not set affinities.

This issue is orthogonal to the affinity issues. I am happy to reword the remaining warning to mention the performance loss.

#7 - 12/12/2017 02:18 PM - Erik Lindahl

Let's have **one** warning/note about what is happening, not three...

#8 - 12/12/2017 02:22 PM - Erik Lindahl

Suggestion:

WARNING: Oversubscribing the available 4 logical CPU cores with 5 thread-MPI threads.
This is bad for performance, and prevents us from pinning threads.

In general, what is the reason we have to echo a "note" in every case we are not pinning threads? Since we already warn about cases we consider severe, I see no reason why we just must tell the user about an internal decision we made.

#9 - 12/15/2017 04:14 PM - Gerrit Code Review Bot

Gerrit received a related patchset '1' for Issue [#2345](#).
Uploader: Berk Hess (hess@kth.se)
Change-Id: gromacs~release-2018~lfd62c723fd87b0ddaab0df1e2f1bf36b461ea33
Gerrit URL: <https://gerrit.gromacs.org/7356>

#10 - 12/19/2017 02:35 AM - Mark Abraham

Erik Lindahl wrote:

Suggestion:

WARNING: Oversubscribing the available 4 logical CPU cores with 5 thread-MPI threads.
This is bad for performance, and prevents us from pinning threads.

In general, what is the reason we have to echo a "note" in every case we are not pinning threads? Since we already warn about cases we consider severe, I see no reason why we just must tell the user about an internal decision we made.

I agree such minor things don't belong on stderr.

#11 - 12/19/2017 02:37 AM - Mark Abraham

- Status changed from Fix uploaded to Closed