GROMACS - Task #2185

add docs on MPI + CUDA w/wo MPS

05/23/2017 07:45 PM - Szilárd Páll

Status:	New	
Priority:	Normal	
Assignee:		
Category:	documentation	
Target version:		
Difficulty:	hard	

Description

We should document how, when, and why use thread-MPI, MPI, and MPI+MPS in GPU-accelerated runs.

Indirectly related: the documentation on running with MPI+CUDA, in particular in rank sharing scenarios is somewhat lacking -- especially since the port from the old wiki acceleration/parallelization page some of the info was left behind. The above would be a natural extension of some of that material which explains why is it useful to run multiple ranks per GPU device.

History

#1 - 05/23/2017 10:21 PM - Mark Abraham

Szilárd Páll wrote:

> We should document how, when, and why use thread-MPI, MPI, and MPI+MPS in GPU-accelerated runs.

https://redmine.gromacs.org/projects/gromacs/repository/revisions/master/entry/docs/user-guide/mdrun-performance.rst describes several examples. What is MPS?

> Indirectly related: the documentation on running with MPI+CUDA, in particular in rank sharing scenarios is somewhat lacking -- especially since the port from the old wiki acceleration/parallelization page some of the info was left behind. The above would be a natural extension of some of that material which explains why is it useful to run multiple ranks per GPU device.

Sure. The current organization has some multi-simulation examples at https://redmine.gromacs.org/projects/gromacs/repository/revisions/master/entry/docs/user-guide/mdrun-features.rst, but evolving somewhere to perhaps link to from one or other place could be good

#2 - 05/24/2017 05:17 PM - Szilárd Páll

Mark Abraham wrote:

> Szilárd Páll wrote:
>
> > We should document how, when, and why use thread-MPI, MPI, and MPI+MPS in GPU-accelerated runs.
>
> https://redmine.gromacs.org/projects/gromacs/repository/revisions/master/entry/docs/user-guide/mdrun-performance.rst describes several examples.

It does, however it also omits most of section 3 of the old acceleration & parallelization page page. Those few examples alone without much context are far less powerful IMO than linking use-cases and parallelization scenarios to the different mdrun launch configs; e.g. telling users when/why would you want to run #PP ranks > #GPUs, but also warn that #PP ranks == #cores is not ideal with GPUs.

> What is MPS?

https://docs.nvidia.com/deploy/pdf/CUDA_Multi_Process_Service_Overview.pdf

https://www.nvidia.com/object/running-jobs-in-gromacs.html

> Indirectly related: the documentation on running with MPI+CUDA, in particular in rank sharing scenarios is somewhat lacking -- especially since the port from the old wiki acceleration/parallelization page some of the info was left behind. The above would be a natural extension of some of that material which explains why is it useful to run multiple ranks per GPU device.

> Sure. The current organization has some multi-simulation examples at https://redmine.gromacs.org/projects/gromacs/repository/revisions/master/entry/docs/user-guide/mdrun-features.rst, but evolving somewhere to perhaps link to from one or other place could be good

Sure, but that is not what I was pointing out, ref above.