GROMACS - Task #2516

Feature # 2054 (Accepted): PME on GPU

Task # 2453 (Resolved): PME OpenCL porting effort

Support PME OpenCL execution width < 16

05/23/2018 12:29 PM - Aleksei Iupinov

Status:	New	
Priority:	Low	
Assignee:	Aleksei Iupinov	
Category:		
Target version:		
Difficulty:	hard	

Description

PME CUDA/OpenCL code is implemented with the hardcoded assumption of 16 threads per atom (PME_SPREADGATHER_THREADS_PER_ATOM).
This corresponds to spreading/gathering in 2 dimensions - one can search for assignments of ithy and ithz in the spread and gather kernel files.
This logic has to be changed to only use 1 dimension to support execution widths < 16, e.g. on Intel.
Changing assignments/loop code themselves should be easy, but expect more pitfalls :-)

Related issues:

Related to GROMACS - Task #2030: make the OpenCL nobonded kernels work on Int...	**Closed**
Related to GROMACS - Task #2519: Improve/remove PME OpenCL kernel barriers	**New**
Related to GROMACS - Task #2520: Treat OpenCL kernel width more diligently	**New**

Associated revisions

Revision 691f1d0e - 11/07/2018 12:55 PM - Szilárd Páll

Ensure minimum exec width of the PME OpenCL kernels

This change adds checks to make sure that we don't execute incorrect
kernels in the case of the rare event if the Intel OpenCL compiler
decides to generate spread or gather kernels for 8-wide execution.

Refs #2516 #2520

Change-Id: I7ab33accebe908a56eb194e8245dfcfa6f817324

Revision a19dd7d5 - 08/15/2019 09:46 PM - Szilárd Páll

Fix OpenCL gather reduction

On >=16-wide execution it is correct (narrower is checked and excluded
during compilation).

TODO: Consider changing the default on NVIDIA & Intel where offloading
PME is generally not advantageous to performance.

Addresses part of #2519
Refs #2453 #2516

Change-Id: I24beaaeea096954ba32b3a80251945a9d82a3c05

History

#1 - 05/23/2018 04:44 PM - Aleksei Iupinov

- Related to Task #2030: make the OpenCL nobonded kernels work on Intel iGPU added

#2 - 05/24/2018 06:05 PM - Aleksei Iupinov

- Related to Task #2519: Improve/remove PME OpenCL kernel barriers added

#3 - 05/24/2018 06:08 PM - Aleksei Iupinov

As discussed with Roland, changing the code might be not needed at all.
One might get away with inserting additional synchronisation points as needed, and maybe still treating warp_size on host and device as multiple of 16.

#4 - 10/31/2018 09:12 PM - Gerrit Code Review Bot

Gerrit received a related patchset '1' for Issue #2516.
Uploader: Szilárd Páll (pall.szilard@gmail.com)
Change-Id: gromacs~release-2019~I7ab33accebe908a56eb194e8245dfcfa6f817324
Gerrit URL: https://gerrit.gromacs.org/8635

#5 - 11/01/2018 08:39 AM - Mark Abraham

- Related to Task #2520: Treat OpenCL kernel width more diligently added

#6 - 06/05/2019 05:32 PM - Szilárd Páll

- Priority changed from Normal to Low

Dropping priority as we do not expect to work on this anytime soon nor do we expect hardware that would need it.