

GROMACS - Task #2520

Feature # 2054 (Accepted): PME on GPU

Task # 2453 (Resolved): PME OpenCL porting effort

Treat OpenCL kernel width more diligently

05/24/2018 06:21 PM - Aleksei lupinov

Status:	New
Priority:	Normal
Assignee:	
Category:	
Target version:	
Difficulty:	uncategorized
Description	
As discussed on https://gerrit.gromacs.org/#/c/7924/ , we're currently treating OpenCL kernel execution width with CUDA-centric world view. Function getWarpSize() queries CL_KERNEL_PREFERRED_WORK_GROUP_SIZE_MULTIPLE for a test kernel and returns 64 for AMD, 32 for NVidia accordingly. We should likely be using vendor extensions (CL_DEVICE_WAVEFRONT_WIDTH_AMD from cl_amd_device_attribute_query, CL_DEVICE_WARP_SIZE_NV from cl_nv_device_attribute_query, CL_KERNEL_MAX_SUB_GROUP_SIZE_FOR_NDRANGE from cl_intel_subgroups) and only use CL_KERNEL_PREFERRED_WORK_GROUP_SIZE_MULTIPLE as a fallback. We should also be querying this "warp" size for each kernel separately and storing it.	
Related issues:	
Related to GROMACS - Task #2516: Support PME OpenCL execution width < 16	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

History

#1 - 11/01/2018 08:39 AM - Mark Abraham

- Related to Task #2516: Support PME OpenCL execution width < 16 added