

GROMACS - Feature #2054

PME on GPU

09/27/2016 11:22 AM - Aleksei lupinov

Status:	Accepted	
Priority:	High	
Assignee:	Aleksei lupinov	
Category:	mdrun	
Target version:	2020	
Difficulty:	hard	
Description		
<p>This is a general issue to discuss and keep track of the PME GPU implementation progress.</p> <p>PME for CUDA is in Gromacs 2018. The current task is to implement PME for OpenCL 1.2, and unify bunch of easily unifiable PME/NB CUDA/OpenCL code on the side. The same original PME CUDA restrictions will apply to the first OpenCL implementation:</p> <ol style="list-style-type: none">1) PME order of 4 only - mostly a programming convenience and a sane default (even though it would be fun to change spread/gather kernel assumptions and logic, to try out order of 8).2) No PME decomposition (only a single process can run the whole PME GPU task, either with or without NB) - can be changed in a separate project.3) No free energy (~no multiple grids - not a difficult thing to implement).4) No Lennard-Jones PME (~no multiple grids + no LJ solver).5) Single precision only (pretty much a given with GPUs and the approximate computation that is PME method). <p>Additionally, OpenCL-specific implementation will at first have the warp size being fixed at 32 (while it should be trivial to relax it to 16 or multiples of 32).</p> <p>There should also be a checklist of CUDA/OpenCL wrapper changes here, that will facilitate porting the code without too much duplication.</p> <p>Other broad issues that are relevant and connected to PME, to keep in mind:</p> <ol style="list-style-type: none">1) Reworking the GPU/device assignment.2) Rethinking the GPU task scheduling.3) The input/output data formats and providers both on CPU and GPU (common GPU data framework, conversion kernel for NB, pagelocked allocator for host pointers...).		
Subtasks:		
Bug # 2208: cuFFT linking		New
Task # 2240: GPU emulation mode support for PME missing		Accepted
Bug # 2303: PME GPU opt-in mixed mode is broken with PME tuning		Closed
Task # 2453: PME OpenCL porting effort		Resolved
Task # 2498: OpenCL memory pinning/mapping		New
Task # 2500: detect and allow linking external cIFFT, or no cIFFT		Closed
Task # 2514: PME OpenCL reductions with intrinsics		New
Task # 2515: cIFFT RocM compatibility problem		Closed
Task # 2516: Support PME OpenCL execution width < 16		New
Task # 2519: Improve/remove PME OpenCL kernel barriers		New
Task # 2520: Treat OpenCL kernel width more diligently		New
Task # 2521: Implement alternating PME/NB wait for OpenCL		New
Task # 2522: OpenCL context duplication		New
Task # 2527: Rename GpuEventSynchronizer to something more fitting (after mergin PME Op...		New
Task # 2529: Improve test timeouts handling		Closed
Task # 2531: Consider optimizing tabulated data access on GPU		New
Task # 2532: enable queue priorities in OpenCL		New
Task # 2535: consider compiling opencl fft kernels once		New
Bug # 2536: cIFFT execution not timed in PME		Closed

Task # 2537: Simplify PME solve reduction	New
Task # 2538: organize more of the PME GPU code along task-specific lines	New
Task # 2696: ensure PME queue is flushed	In Progress
Related issues:	
Related to GROMACS - Task #2092: Tests running on GPU, and hardware assignment	New
Related to GROMACS - Task #2124: PME GPU user interface suggestions	Closed
Related to GROMACS - Task #2053: refine notation in GPU code	New
Related to GROMACS - Task #2524: struct alignment/packing for OpenCL host & d...	New
Related to GROMACS - Task #3031: evaluate the impact of particle order on PME	New

Associated revisions

Revision 76c7a1a4 - 10/11/2017 01:49 PM - Aleksei lupinov

PME spline+spread CUDA kernel and unit tests

The CUDA implementation of PME spline computation and charge spreading for PME order 4 is added in pme-spread.cu.

The unit tests for PME CPU spline/spread stages (e8cf7c0) are also extended to work with the PME CUDA kernel, using the same reference data. The tests iterate over all CUDA GPUs which are compatible with Gromacs.

Refs #2054, #2092.

Change-Id: I5ec49f030b9b94395db28fa454ea25c3efb05d1

Revision 4231cc37 - 10/11/2017 03:11 PM - Aleksei lupinov

PME force gathering - CUDA kernel + unit tests

The CUDA implementation of PME force gathering for PME order 4 is added in pme-gather.cu. The unit tests for PME CPU force gathering (d20a5d36) are extended to work with the CUDA kernel, using the same reference data. The tests iterate over all Gromacs-compatible CUDA GPUs.

Ref #2054

Change-Id: I162e3a14cb9aa8ddeac17c5ad1ca709df72b8986

Revision 7bec7e1f - 10/11/2017 04:35 PM - Aleksei lupinov

PME solving - CUDA kernel + unit tests

The CUDA implementation of PME solving is added in pme-solve.cu. The unit tests for PME CPU solving are extended to work with the CUDA kernel, using the same reference data.

The CUDA solver supports 2 grid dimension orders: YZX and XYZ (unlike the CPU one which only supports YZX). This is also tested.

Lennard-Jones solving is not implemented. The tests iterate over all Gromacs-compatible CUDA GPUs.

Refs #2054

Change-Id: I610e7f077f39a64089dd9b80df9905094b10459

Revision 2747fc48 - 10/21/2017 07:18 PM - Aleksei lupinov

Add calls to the PME GPU stages

This adds the inactive calls to PME GPU stages both for PP+PME and PME-only ranks.

Ref #2054

Change-Id: I5af2ab95cedff422c39592255f01205d42fc7eb7

Revision 4c2fe1e6 - 09/07/2018 09:28 AM - Magnus Lundborg

Check q perturbation when PME on GPU is tested

If charges are not perturbed allow running PME on the GPU in FE simulations.

Refs #2054.

Change-Id: Ibc610cb63afaadf4aa97608b8e03b6906fe2d026

History

#1 - 11/17/2016 03:48 PM - Szilárd Páll

I'd suggest creating subtasks to track progress of what needs to be done.

#2 - 12/13/2016 02:33 PM - Aleksei lupinov

There are a couple of PME CPU unit test patches sitting idle in Gerrit.

These are <https://gerrit.gromacs.org/6251/> and <https://gerrit.gromacs.org/6337>.

I would like to get these in sooner rather than later,

as the GPU spline computation/spreading patch involves the unit test which builds both on those and on the main PME GPU patch <https://gerrit.gromacs.org/6212> as well.

Don't be discouraged by their sizes - most of that is just generated reference data, the actual code is ~500 lines in each.

#3 - 01/16/2017 11:59 AM - Gerrit Code Review Bot

Gerrit received a related patchset '10' for Issue [#2054](#).

Uploader: Aleksei lupinov (a.yupinov@gmail.com)

Change-Id: gromacs~master~1f5ec49f030b9b94395db28fa454ea25c3efb05d1

Gerrit URL: <https://gerrit.gromacs.org/6357>

#4 - 01/30/2017 12:00 PM - Gerrit Code Review Bot

Gerrit received a related DRAFT patchset '2' for Issue [#2054](#).

Uploader: Aleksei lupinov (a.yupinov@gmail.com)

Change-Id: gromacs~master~1162e3a14cb9aa8ddeac17c5ad1ca709df72b8986

Gerrit URL: <https://gerrit.gromacs.org/6437>

#5 - 02/10/2017 06:18 PM - Szilárd Páll

- *Difficulty hard added*

- *Difficulty deleted (uncategorized)*

Additional tasks that are blockers for the release:

- PME-GPU user-interface (command line, manual device assignment, log reporting, etc.)
- user documentation + examples
- testing on multiple generation of devices (CC 2.0?)
- testing with multiple CUDA releases
- performance evaluation (at least to determine the range of use-cases where it makes sense to use a GPU for PME)

#6 - 02/13/2017 12:21 PM - Gerrit Code Review Bot

Gerrit received a related patchset '4' for Issue [#2054](#).

Uploader: Aleksei lupinov (a.yupinov@gmail.com)

Change-Id: gromacs~master~1c610e7f077f39a64089dd9b80df9905094b10459

Gerrit URL: <https://gerrit.gromacs.org/6459>

#7 - 03/09/2017 04:02 PM - Aleksei lupinov

- *Related to Task #2092: Tests running on GPU, and hardware assignment added*

#8 - 03/09/2017 04:02 PM - Aleksei lupinov

- *Related to Task #2124: PME GPU user interface suggestions added*

#9 - 03/20/2017 05:48 PM - Aleksei lupinov

- *Related to Task #2053: refine notation in GPU code added*

#10 - 04/03/2017 12:12 PM - Gerrit Code Review Bot

Gerrit received a related DRAFT patchset '16' for Issue [#2054](#).

Uploader: Aleksei lupinov (a.yupinov@gmail.com)

Change-Id: gromacs~master~19e705b86d5aa07d59544de68234cdd6242ad1194

Gerrit URL: <https://gerrit.gromacs.org/6472>

#11 - 05/19/2017 01:23 PM - Aleksei lupinov

- Blocked by Task #2183: GPU-accessed memory page-locking and page sizes added

#12 - 05/26/2017 07:32 PM - Aleksei lupinov

I would like to urge everyone to review the low-level PME GPU building blocks:

<https://gerrit.gromacs.org/#/c/6357/> (spreading kernel)

<https://gerrit.gromacs.org/#/c/6459/> (solving kernel)

<https://gerrit.gromacs.org/#/c/6437/> (gathering kernel)

<https://gerrit.gromacs.org/#/c/6212/> (the data structures, their management, cuFFT calls) - this one is large and already has some renaming/cleanup TODOs, which would be much easier to resolve when all these 4 changes are merged in.

There is more work to do if you look at <https://gerrit.gromacs.org/#/q/topic:pme>, so I suggest that these components are reviewed first - they have been sitting there for a while, and I think they would do more good being tested by users of the master branch with included unit-tests.

Note that there is a GPU-task assignment change <https://gerrit.gromacs.org/#/c/6205/> which sits at the bottom of the PME GPU branch (and in hindsight probably should have started higher), so I would appreciate more reviews on that as well. Otherwise, rebasing the core PME GPU changes once they're reviewed to skip it should be trivial.

#13 - 06/01/2017 03:22 PM - Gerrit Code Review Bot

Gerrit received a related DRAFT patchset '2' for Issue [#2054](#).

Uploader: Aleksei lupinov (a.yupinov@gmail.com)

Change-Id: gromacs~master~15af2ab95cedff422c39592255f01205d42fc7eb7

Gerrit URL: <https://gerrit.gromacs.org/6670>

#14 - 11/28/2017 06:01 PM - Mark Abraham

- Status changed from New to Resolved

This is now implemented as intended

#15 - 11/28/2017 06:06 PM - Mark Abraham

- Blocked by deleted (Task #2183: GPU-accessed memory page-locking and page sizes)

#16 - 11/28/2017 07:34 PM - Szilárd Páll

Mark Abraham wrote:

This is now implemented as intended

Yeah, suggest closing.

#17 - 11/29/2017 09:19 AM - Aleksei lupinov

OK, I'm just not familiar with issue tracking logic - if we implement e.g. coordinate conversion kernel, or PME GPU decomposition, and make an issue for that, is it alright for it to have a closed parent?

#18 - 11/29/2017 12:07 PM - Szilárd Páll

Aleksei lupinov wrote:

OK, I'm just not familiar with issue tracking logic - if we implement e.g. coordinate conversion kernel, or PME GPU decomposition, and make an issue for that, is it alright for it to have a closed parent?

We targeted this feature for the next release. While some of the subtask did not materialize (but overall the feature was implemented), it might be cleaner to close this issue and continue with the few smaller remaining tasks. [#2208](#) and [#2240](#) should be possible to resolve on way or another.

#19 - 12/19/2017 03:01 AM - Mark Abraham

yeah we'll make new tasks (that can refer to this one) when we decide to do future work to add more functionality.

#20 - 12/19/2017 03:03 AM - Mark Abraham

- Status changed from Resolved to Accepted

- Target version changed from 2018 to 2019

can't close this while sub tasks remain open, so retargeting

#21 - 03/05/2018 02:41 PM - Aleksei lupinov

- *Description updated*

#22 - 05/26/2018 03:17 PM - Szilárd Páll

- *Related to Task #2524: struct alignment/packing for OpenCL host & device code added*

#23 - 09/03/2018 09:38 AM - Magnus Lundborg

Does anyone have any rough suggestions where to start implementing PME with free energy on GPU?

#24 - 09/04/2018 12:03 PM - Gerrit Code Review Bot

Gerrit received a related patchset '1' for Issue [#2054](#).

Uploader: Magnus Lundborg (magnus.lundborg@scilifelab.se)

Change-Id: gromacs~master~lbc610cb63afaadf4aa97608b8e03b6906fe2d026

Gerrit URL: <https://gerrit.gromacs.org/8305>

#25 - 11/29/2018 11:22 AM - Mark Abraham

- *Target version changed from 2019 to 2020*

#26 - 07/10/2019 11:44 AM - Szilárd Páll

- *Related to Task #3031: evaluate the impact of particle order on PME added*