

## GROMACS - Task #3093

Task # 3370 (New): Further improvements to GPU Buffer Ops and Comms

### rework GPU direct halo-exchange related force reduction complexities

09/18/2019 02:53 AM - Szilárd Páll

<b>Status:</b>	In Progress
<b>Priority:</b>	High
<b>Assignee:</b>	
<b>Category:</b>	mdrun
<b>Target version:</b>	2021-infrastructure-stable
<b>Difficulty:</b>	uncategorized
<b>Description</b>	
Force reduction is now done in two stages, if there is halo exchange CPU contribution is already reduced with communicated data early, while in other cases the transfer of force this happens later. The current mechanism also relies on position-dependent code with leading to implicit dependencies rather than explicit event-based sync with a record closely succeeding the producer and an eventWait enqueue at preceding the consumer task.	

#### Associated revisions

##### Revision 54c24729 - 09/18/2019 03:36 PM - Alan Gray

GPU Force Halo Exchange

Activate with `GMX_GPU_DD_COMMS` environment variable.

Extends GPU Halo exchange feature to provide GPU Force halo exchange functionality. Does not yet support virial steps, which require an extra shift force reduction - these are currently performed on the non-buffer ops / non direct-comm path. Also has same limitations as coordinate halo exchange.

Performs part of #2890. Future work to improve synchronization towards a more one-sided scheme (#3092) and to make dependencies more explicit (#3093)

Change-Id: `lfc23cc8db2655f7258e68b34e7cdc7b71994e1e8`

##### Revision 8a0d4d97 - 10/13/2019 01:27 PM - Szilárd Páll

Enable `StatePropagatorGpuData` for force transfers

Force transfers have been switched to use `StatePropagatorGpuData` already before. This change updates the synchronization mechanisms as:

- replaces the previous stream sync after GPU buffer/ops reduction with a `waitForcesReadyOnHost` call;
- removes the barriers in `copyForces[From|To]Gpu()` as dependencies are now satisfied: most dependencies are intra-stream and therefore implicit, the exception being the halo exchange that uses its own mechanism to sync H2D in the local stream with the nonlocal stream (which is yet to be replaced Refs #3093).

Refs. #3126.

Change-Id: `l8bfd39f79c87f20492c4ae287d6f19261724f806`

##### Revision 7073b54d - 11/27/2019 06:48 AM - Alan Gray

Event-based Dependency for GPU Force Halo Exchange

Introduces new event recorded when exchanged forces are ready on GPU, and passes this into force buffer ops using `dependencyList`. Removes previous mechanism of forcing local stream to wait on non-local stream.

Addresses part of #3093

Refs #3194

Change-Id: `l768898839e5c6a653894d5eb80354f0e423e06ed`

## History

---

### #1 - 09/18/2019 12:39 PM - Szilárd Páll

- Private changed from Yes to No

### #2 - 10/08/2019 02:18 PM - Szilárd Páll

- Subject changed from *rework GPU direct halo-exchange related force reduction complexities* to *rework GPU direct halo-exchange related force reduction complexities*

### #3 - 10/11/2019 05:53 PM - Szilárd Páll

- Priority changed from Normal to High

what is the

### #4 - 11/01/2019 11:10 AM - Alan Gray

- Status changed from New to In Progress

- Target version changed from 2020 to 2020-beta3

Addressed in pending changes

<https://gerrit.gromacs.org/c/gromacs/+/13863>

and

<https://gerrit.gromacs.org/c/gromacs/+/13885>

### #5 - 11/01/2019 08:58 PM - Szilárd Páll

Alan Gray wrote:

Addressed in pending changes

<https://gerrit.gromacs.org/c/gromacs/+/13863>

and

<https://gerrit.gromacs.org/c/gromacs/+/13885>

Thanks for the update. I've just flagged the former which shows 3 additional tests failing on the gpucomm matrix compared to previous triggers.

I think we need to fix correctness issues of the code before we can really move forward with new changes.

### #6 - 12/02/2019 01:14 PM - Paul Bauer

- Target version changed from 2020-beta3 to 2020-rc1

has this here been resolved?

### #7 - 12/02/2019 02:17 PM - Szilárd Páll

Paul Bauer wrote:

has this here been resolved?

Partly, but not fully. We still have a conditionality of when do we upload local forces to the GPU based on the code-path, which I think is undesired code complexity. There is however no room for this in the release branch. Should be bumped to later (but preferably not to "infrastructure-stable").

### #8 - 12/20/2019 08:19 AM - Paul Bauer

- Target version changed from 2020-rc1 to 2021-infrastructure-stable

bumped to a 2021 target

### #9 - 02/07/2020 01:25 PM - Alan Gray

- Status changed from In Progress to Closed

Moved to umbrella task <https://redmine.gromacs.org/issues/3370>

### #10 - 02/14/2020 11:35 AM - Alan Gray

- Status changed from Closed to In Progress

- Parent task changed from #2890 to #3370

Re-opening and moving to subtask of [#3370](#), so we don't lose the discussion.