

GROMACS - Feature #2888

Feature # 2816 (New): GPU offload / optimization for update&constraints, buffer ops and multi-gpu communication

CUDA Update and Constraints module

03/08/2019 04:22 PM - Artem Zhmurov

Status: New	
Priority: High	
Assignee: Artem Zhmurov	
Category:	
Target version: 2020	
Difficulty: uncategorized	
Description	
<ul style="list-style-type: none">• LINCS for non-water constraints.• SETTLE for water constraints.• Leap-frog integrator.• Merge of the three into single module.• Remove the scaffolding from LINCS, SETTLE and Leap-Frog:<ul style="list-style-type: none">◦ Coordinates, velocities, forces management.◦ PBC management.◦ Virial reduction.◦ Update tests.◦ Remove Impl.◦ Template computeVirial and updateVelocities◦ Remove the D2D copy for the coordinates (xp->x)	
Subtasks:	
Task # 3114: Possible improvements to update-constraints	New
Feature # 3162: Add virtual site support to GPU version of update-constraints.	New
Bug # 3163: gpuupdate / task assignment stabilization	New
Task # 3167: GPU update path user documentation	New
Feature # 3168: GPU update release notes	New
Bug # 3182: pressure coupling buggy with GPU update	New
Related issues:	
Related to GROMACS - Feature #2885: CUDA version of LINCS	New
Related to GROMACS - Feature #2886: CUDA version of SETTLE	New
Related to GROMACS - Feature #2887: CUDA version of Leap Frog algorithm	New
Related to GROMACS - Task #2936: introduce check that CPU-GPU transfers are m...	New
Related to GROMACS - Task #3171: schedule CPU H2D force contribution in separ...	New
Related to GROMACS - Task #3195: assess nightly master failures	New

Associated revisions

Revision 1c8eb7c5 - 06/18/2019 11:24 AM - Artem Zhmurov

Combine CUDA Leap-Frog, LINCS and SETTLE. I.

This is the first step in combining constraints and integrator into "UpdateAndConstraints" module. The initial merge does not imply any performance optimisation or code clean-up. Hence, this patch keeps all the temporary infrastructure that was built around SETTLE, LINCS and Leap-Frog to allow them to function as a separate units. In the following commits, this infrastructure will be removed and these three implementations will be more closely integrated. To enable, set GMX_UPDATE_CONSTRAIN_GPU environment variable. Note, that environment variables GMX_LINCS_GPU, GMX_SETTLE_GPU and GMX_INTEGRATE_GPU will no longer work.

Refs #2816, #2888

Change-Id: I8730aad0ecaa0230686fe89d1157b0da2f01f7bc

Revision fb7a59cd - 07/03/2019 09:42 AM - Artem Zhmurov

Combine CUDA Leap-Frog, LINCS and SETTLE. II.

Stand-alone CUDA implementations of Leap-Frog, LINCS and SETTLE required additional scaffolding for integration and testing. The most prominent part of this is the management of coordinates, velocities and forces, which is removed in this commit. Management of periodic boundary conditions and virial reduction will be removed in following commits.

Refs #2816, #2888

Change-Id: I4c65a6c7088fd8059f4e7fa3cb4637cb2af79ebc

Revision 747c371c - 07/03/2019 01:01 PM - Artem Zhmurov

Memory management fixes in CUDA version of LINCS

This fix is to prepare LINCS to run with DD.

1. The masses array size depends on the current number of atoms rather than on the number of constraints.
2. The size of other arrays should be based on the number of threads launched on the GPU, which include padding added to align coupled constraints with the thread blocks. Also renamed variable according to conventions.

Refs #2885 and #2888

Change-Id: I20cb53ebc6da6a1ff2ee1e385613b27c4a01d11f

Revision 1b64f6aa - 07/04/2019 08:23 AM - Artem Zhmurov

Use `reallocateDeviceBuffer(...)` in CUDA version of SETTLE

Refs #2886 and #2888

Change-Id: Ia45254a24eda8e6ad151b1f4c6583b1a2c926004

Revision 6385f296 - 07/05/2019 11:05 AM - Artem Zhmurov

Remove PImpl scaffolding from CUDA version of LINCS

The CUDA implementation of LINCS was initially introduced as a stand-alone feature. This required hiding CUDA-specific variables and subroutines into the private implementation subclass. Since the LINCS is not a part of Update and Constraints module, this is no longer required and can be removed.

Refs #2816, #2888

Change-Id: I9698224d4702dfb8d99106999335c62e83a511df

Revision b1150eee - 07/26/2019 10:20 AM - Artem Zhmurov

Remove PImpl scaffolding from CUDA version of SETTLE

GPU version of SETTLE was implemented as a class with private implementation so it will be possible to initialize on non-CUDA hosts. Now, the implementation can be hidden inside the Update and Constraints PImpl so that the CUDA specific types and calls can be exposed in SETTLE and private implementation is no longer needed there.

Refs #2816, #2888

Change-Id: I4c78f2629be34b42bb5f4f7d34970c3e41515691

Revision 1bfc9ba5 - 08/06/2019 03:52 PM - Artem Zhmurov

Remove PImpl scaffolding from CUDA version of Leap-Frog

Private implementation in CUDA version of Leap-Frog was used to introduce this integrator as a stand-alone unit.

Now it is merged with constraints, PImpl is no longer needed.

Refs #2816, #2888

Change-Id: lea82abef016b7e15b9be44a0e1b446e12e582d3c

Revision b1be1e72 - 08/08/2019 11:13 AM - Artem Zhmurov

Refactor Leap-Frog tests and connect them to CPU version

This introduces test data object and runners to the Leap-Frog tests, which are now connected to the CPU version of Leap-Frog. This also makes possible to include tests based on the reference values, which are needed to make sure that the temperature and(or) pressure control works fine in new implementations.

Refs. #2816, #2888.

Change-Id: ld2d934c43138889ad178a94126cab4da2895bb5a

Revision d1f2302e - 08/09/2019 05:35 PM - Artem Zhmurov

Refactoring of the SETTLE tests

Current version of tests for CUDA version of SETTLE was a quick addition to the old tests, with direct comparison of the GPU implementation with the old original CPU-based implementation. This commit rearranges the test structure, making it possible to apply the same set of tests to both implementations. There is no changes to the tests themselves. Currently, comparison tests will run twice and will dry-run on CUDA builds without CUDA-capable devices.

TODO: Add comparison with pre-computed values for coordinates, velocities and virial. Remove the CPU vs GPU comparison tests.

Refs #2886, #2888.

Change-Id: lfc6af9af6c93787b919b785348f9f4547b6c267

Revision 0cd72f2b - 08/09/2019 07:59 PM - Artem Zhmurov

Prepare Update and Constraints for Domain Decomposition

Initial GPU-based version of the update and constraints was not designed to run with the Domain decomposition. This introduces a couple of fixes to the memory management that should allow the module to work with the DD enabled. The memory buffers are now re-allocated at the set(...) stage, if so needed.

Refs. #2816, #2888.

Change-Id: l155884f5797252cf048a6400a2dd7b042d355b7e

Revision 7bd1c817 - 08/13/2019 04:13 PM - Artem Zhmurov

Make use of reference data in SETTLE tests

As a temporary measure, the CPU and GPU versions of SETTLE were tested against each other. Making use of the reference data framework allows to test them against precomputed values. Now, the final positions, velocities and virial are properly tested in CPU and, if available, in GPU versions.

Refs. #2886, #2888.

Change-Id: l8e54e1a741263b8bf9774a21141c527f58130fa9

Revision 1fbaf8ff - 08/16/2019 11:29 AM - Artem Zhmurov

Remove PImpl scaffolding from CUDA version of SETTLE

GPU version of SETTLE was implemented as a class with private

implementation so it will be possible to initialize on non-CUDA hosts. Now, the implementation can be hidden inside the Update and Constraints PImpl so that the CUDA specific types and calls can be exposed in SETTLE and private implementation is no longer needed there.

Refs #2816, #2888

Change-Id: I4c78f2629be34b42bb5f4f7d34970c3e41515691

Revision 3d35e919 - 08/16/2019 11:29 AM - Artem Zhmurov

Remove PImpl scaffolding from CUDA version of Leap-Frog

Private implementation in CUDA version of Leap-Frog was used to introduce this integrator as a stand-alone unit. Now it is merged with constraints, PImpl is no longer needed.

Refs #2816, #2888

Change-Id: lea82abef016b7e15b9be44a0e1b446e12e582d3c

Revision 039709b7 - 08/16/2019 11:29 AM - Artem Zhmurov

Prepare Update and Constraints for Domain Decomposition

Initial GPU-based version of the update and constraints was not designed to run with the Domain decomposition. This introduces a couple of fixes to the memory management that should allow the module to work with the DD enabled. The memory buffers are now re-allocated at the set(...) stage, if so needed.

Refs. #2816, #2888.

Change-Id: I155884f5797252cf048a6400a2dd7b042d355b7e

Revision 4b07f76f - 08/26/2019 03:07 PM - Artem Zhmurov

Add temperature coupling to CUDA Leap-Frog integrator

Velocity rescaling added to the integrator. None, one, or multiple rescaling groups are supported. Tests that include temperature coupling are added.

NoseHoover temperature control is not implemented.

Refs. #2887, #2888.

Change-Id: I1e6850eeb74de829554260fba6a6e6c1c63ceb46

Revision 7ddb7204 - 09/03/2019 06:49 AM - Artem Zhmurov

Add Parrinello-Rahman pressure coupling to CUDA Leap-Frog integrator

Parrinello-Rahman isotropic pressure coupling scheme is now added to CUDA version of the Leap-Frog integrator. The Leap-Frog tests are updated to check the new addition.

Refs. #2887, #2888.

Change-Id: Icf42667621c16a994e68baf5158ea4abac387928

Revision 79aab161 - 10/02/2019 07:07 PM - Artem Zhmurov

Eliminate D2D copy in update constraints

The intermediate coordinates (x' or x_p) are only needed inside the update-constraints module (for the constraints algorithms) and never used outside. Hence, the x_p variable can be used to save the coordinates before update, while x stores the final coordinates. This way, there is no need to make a D2D $x_p \rightarrow x$ copy after applying the constraints, since x will have the correct data.

Refs. #2888, #3114.

Change-Id: I363b633976a236a8e2bf2137c21d3bf0a765cb06

Revision f310be38 - 10/16/2019 01:54 AM - Szilárd Páll

Trigger synchronizer when local forces are ready

The synchronizer is created and managed in StatePropagatorDataGpu and is passed to the nonbonded module at the force buffer ops init.

Refs #2888 #3126

Change-Id: Ie9bf0b6cd8511fe282e377e48f3940e591db214c

Revision 7bbfb57c - 10/16/2019 10:21 AM - Artem Zhmurov

Link GPU force producer and consumer tasks

The GPU event synchronizer that indicates that forces are ready for a consumption is now passed to the GPU update-constraints. The update-constraints enqueue a wait on the event in the update stream before performing numerical integration and constraining. Note that the event is conditionally returned by the StatePropagatorDataGpu and indicates that either the reduction of forces on the GPU or the H2D copy is done, depending on offload scenario on a current timestep.

Refs. #2816, #2888, #3126.

Change-Id: Ie12b0c55b75ec5f0c31ce500a2760fb4d5cf3b91

History

#1 - 03/08/2019 05:06 PM - Artem Zhmurov

- Related to Feature #2885: CUDA version of LINC3 added

#2 - 03/11/2019 11:54 AM - Artem Zhmurov

- Related to Feature #2886: CUDA version of SETTLE added

#3 - 03/13/2019 01:37 PM - Artem Zhmurov

- Related to Feature #2887: CUDA version of Leap Frog algorithm added

#4 - 03/13/2019 07:10 PM - Artem Zhmurov

- Description updated

I have LINC3 with some tests for it, SETTLE with some tests for it and Leap-Frog integrator with some tests for it. Now I combine them into one "Update and Constrain" module. Any ideas for the test that the merge was successful?

#5 - 03/13/2019 07:12 PM - Artem Zhmurov

- Subject changed from CUDA GPU-only loop to CUDA Update and Constraints module

- Description updated

#6 - 03/21/2019 08:43 PM - Gerrit Code Review Bot

Gerrit received a related patchset '4' for Issue [#2888](#).

Uploader: Artem Zhmurov (zhmurov@gmail.com)

Change-Id: gromacs~master~I8730aad0ecaa0230686fe89d1157b0da2f01f7bc

Gerrit URL: <https://gerrit.gromacs.org/9329>

#7 - 03/21/2019 10:35 PM - Artem Zhmurov

- Description updated

#8 - 03/22/2019 02:10 PM - Artem Zhmurov

- Description updated

#9 - 04/24/2019 11:32 PM - Artem Zhmurov

- Description updated

#10 - 04/24/2019 11:32 PM - Artem Zhmurov

- Description updated

#11 - 04/29/2019 11:10 AM - Szilárd Páll

- Related to Task #2936: introduce check that CPU-GPU transfers are made between arrays of compatible types added

#12 - 08/08/2019 10:47 AM - Artem Zhmurov

- Description updated

#13 - 09/19/2019 02:13 PM - Artem Zhmurov

- Description updated

#14 - 10/08/2019 03:09 PM - Artem Zhmurov

- Description updated

#15 - 10/22/2019 12:20 AM - Szilárd Páll

- Related to Task #3171: schedule CPU H2D force contribution in separate stream added

#16 - 11/01/2019 09:22 PM - Szilárd Páll

- Related to Task #3195: assess nightly master failures added