

GROMACS - Feature #2885

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

CUDA version of LINCS

03/08/2019 04:19 PM - Artem Zhmurov

Status:	Closed	
Priority:	Normal	
Assignee:	Artem Zhmurov	
Category:	core library	
Target version:	2021-infrastructure-stable	
Difficulty:	uncategorized	
Description		
Adapt the LINCS constraints to work efficiently on CUDA-enabled GPUs.		
TODO:		
<ul style="list-style-type: none">A separate class that contains the logic.Reduction for the virial using shuffle.Many GPU version.Free energy.		
Ideas for kernel improvement:		
<ul style="list-style-type: none">Use analytical solution for matrix A inversion (for small matrices of H-bonds constraints), inverted matrix itself can be reused rather than recomputed.Move more data to local/shared memory and try to get rid of atomics (at least on the device level).Use locality of coupled constraints better (maybe go from block-sync to warp-sync)Introduce mapping of thread id to both single constraint and single atom, thus designating Nth threads to deal with Nat <= Nth coupled atoms and Nc <= Nth coupled constraints.		
Testing:		
<ul style="list-style-type: none">Initial integration to the constraints test.Add bigger systems to test virial reduction and overall redistribution of constraints among threads.Generalization of tests for different platforms.		
Related issues:		
Related to GROMACS - Feature #2886: CUDA version of SETTLE		Closed
Related to GROMACS - Feature #2887: CUDA version of Leap Frog algorithm		Closed
Related to GROMACS - Feature #2888: CUDA Update and Constraints module		Closed

Associated revisions

Revision 0a1aae78 - 04/28/2019 05:29 PM - Artem Zhmurov

CUDA version of LINCS constraints.

Implementation of the LINCS constraints for NVIDIA GPUs. Currently works isolated from the other parts of the code: coordinates and velocities are copied to and from GPU on every integration timestep. Part of the GPU-only loop. Loosely based on change 9162 by Alan Gray. To enable, set the environmental variable GMX_LINCS_GPU.

Limitations:

- Works only if the constraints can be split in short uncoupled groups (currently < 256, designed for H-bonds constraints).
- Does not change the matrix inversion order for constraints triangles.
- Does not support free energy computations.
- Assumes no communications between domains (i.e. assumes that there is no constraints connecting atoms from two different domains).

5. Number of thread per blocks should be a power of 2 for reduction of virial to work.

TODOs:

1. Move more data from the global memory to local.
2. Change .at() to []
3. Add sorting by the number of coupled constraints to decrease warp divergencies.
4. numAtoms should be changeable (for multi-GPU case).

Refs #2816, #2885

Change-Id: I3c975cf898053b7467bcd30459e60ce2c8852be6

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 af1e0e7e - 01/20/2020 03:15 PM - Artem Zhmurov

Rename LincsCuda into LincsGpu

This is to follow general naming conventions across the code.

Refs #2885, #2888.

Change-Id: Ifa7e3febeff1d958155ed02daa97d26e828e8381

History

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

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

#2 - 03/08/2019 05:06 PM - Artem Zhmurov

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

#3 - 03/08/2019 05:06 PM - Artem Zhmurov

- Related to Feature #2888: CUDA Update and Constraints module added

#4 - 03/08/2019 05:11 PM - Artem Zhmurov

- Description updated

#5 - 03/12/2019 01:33 PM - Artem Zhmurov

- Description updated

#6 - 12/02/2019 11:32 AM - Artem Zhmurov

- Description updated

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

Most of the features are done for 2020, the rest is bumped to 2021

#7 - 01/22/2020 11:49 AM - Artem Zhmurov

- Status changed from New to Closed

Initial implementation and integration are done. Possible improvements moved to [#3114](#).