GROMACS - Task #2092

Tests running on GPU, and hardware assignment

12/19/2016 10:37 AM - Aleksei Iupinov

Status:	New	
Priority:	Normal	
Assignee:		
Category:	testing	
Target version:	future	
Difficulty:	uncategorized	

Description

How should the hardware assignment be treated both outside and inside unit tests?
I would like to discuss CUDA in this case (now that the PME CUDA spreading kernel and its unit tests are almost ready for review), but maybe there are other relevant areas.

Currently the PME CUDA spreading unit test works in a default CUDA context. That would mean always using "first" visible CUDA-capable GPU.
What should the behavior of any GPU unit test be like?
1) Keep it context-unaware - single test run in a single context with no control over it.
2) Iterate over all available GPU contexts - supporting testing multiple contexts per single run, at least.
3) Same as 2, but with more built-in "smartness", e.g. skipping multiple devices with same compute capability.

In CUDA case, the visibility of devices to the executable can be controlled with an environment variable CUDA_VISIBLE_DEVICES, so all the "smartness" can be deferred to Jenkins/tester.
This is why I'm leaning towards the option 2, but would like to read opinions about the general approach here.

Related issues:

Related to GROMACS - Feature #2054: PME on GPU	**Accepted**
Related to GROMACS - Task #2355: update bundled googletest	**Closed**

Associated revisions

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

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: If5ec49f030b9b94395db28fa454ea25c3efb05d1

Revision 37969a6d - 10/06/2020 11:37 AM - Artem Zhmurov

Use new GPU infrastructure in MDLib tests

This make use of common device testing infrastructure in MDLib tests,
where both GPU and CPU implementations are tested. The GPU runners
will now be executed on all the detected devices, not only on the
default one. Also, this will allow to use the MDLib tests in OpenCL
and SYCL, where proper device context object is needed.

Closes #3317

Closes #2254

Related #2092

History

#1 - 01/16/2017 11:59 AM - Gerrit Code Review Bot

Gerrit received a related patchset '10' for Issue #2092.
Uploader: Aleksei Iupinov (a.yupinov@gmail.com)
Change-Id: gromacs~master~If5ec49f030b9b94395db28fa454ea25c3efb05d1
Gerrit URL: https://gerrit.gromacs.org/6357

#2 - 03/09/2017 04:02 PM - Aleksei Iupinov

- Related to Feature #2054: PME on GPU added

#3 - 03/24/2017 03:52 PM - Szilárd Páll

Based on offline discussion:

- we should have a separate test that runs the device enumeration and initialization only
- we need to report PME CPU and GPU tests separately so it's clear what tests have been run (not just in case of failure dumping the error which might indicate whether CPU or GPU test failed)

#4 - 11/28/2017 06:06 PM - Mark Abraham

- Target version set to 2018

What remains to consider here?

#5 - 11/29/2017 09:49 AM - Aleksei Iupinov

Well, one small thing would be to consider printing GPU info in the beginning (end?) of the test run. Trivial within a single binary, I had a draft for that somewhere. Not sure about propagating that to general make check output. Currently I keep relying at heuristic "hmm, some of the solve/gather/spread tests took 20x longer than the other, I assume it means the GPU kernels ran" ;-)

#6 - 12/19/2017 03:26 AM - Mark Abraham

Aleksei Iupinov wrote:

> Well, one small thing would be to consider printing GPU info in the beginning (end?) of the test run. Trivial within a single binary, I had a draft for that somewhere. Not sure about propagating that to general make check output. Currently I keep relying at heuristic "hmm, some of the solve/gather/spread tests took 20x longer than the other, I assume it means the GPU kernels ran" ;-)

In master, #2355 may help with some aspects, e.g. that the test runner might write a string that includes "GPU" rather than just a number for the instances of the parmeterized test fixture.

#7 - 12/19/2017 05:27 PM - Aleksei Iupinov

- Tracker changed from Feature to Task

- Priority changed from High to Normal

- Target version changed from 2018 to future

I think it's fair to retarget this then, assuming https://gerrit.gromacs.org/#/c/7349/ is completed.

#8 - 12/19/2017 05:28 PM - Aleksei Iupinov

- Related to Task #2355: update bundled googletest added