

GROMACS - Feature #3311

GPU infrastructure development

01/22/2020 08:22 AM - Artem Zhmurov

Status:	In Progress	
Priority:	Normal	
Assignee:		
Category:		
Target version:	2021-refactoring	
Difficulty:	uncategorized	
Description		
General goal is to develop platform-agnostic infrastructure for CUDA and OpenCL with perspective expansion to other technologies (SYCL?).		
Subtasks:		
Feature # 2967: GPU reallocateDeviceBuffer improvements		New
Task # 3312: Data type for coordinates, xyzq data, LJ parameters data to use for GPU bu...		In Progress
Bug # 3372: Re-enable RVec and float3 compatibility tests		Closed
Feature # 3313: Introduce and use opaque types for the DeviceStream and DeviceContext		Accepted
Task # 3314: Platform agnostic DeviceStream		Resolved
Task # 3315: Platform agnostic DeviceContext		Resolved
Task # 3316: Context and Stream manager		Accepted
Task # 3317: Improve testing of the GPU code		Resolved
Feature # 3318: Use wrappers for the GPU buffer copy/allocations		In Progress
Task # 3319: Use DeviceBuffer instead of native GPU types in NBNXM		In Progress
Task # 3320: Remove duplicating D2H/H2D wrappers in NBNXM		In Progress
Task # 3321: Add D2D wrapper		Accepted
Task # 3322: Add reallocate(...) function that does not care about the contents of the ...		Accepted
Task # 3323: Rework the StatePropagatorDataGpu		In Progress
Task # 3324: Rework CMake handling of GPU code		New
Related issues:		
Related to GROMACS - Task #2936: introduce check that CPU-GPU transfers/assign...		New

Associated revisions

Revision 6197aaed - 01/29/2020 03:31 PM - Artem Zhmurov

Split and move the checkDeviceBuffer(...) function from PME

Resolving a TODO.

Also fixed the formatting in neighboring comment.

Refs. #3311.

Change-Id: I1687981cc80e2388714cbbb3113f37e34582e31c

Revision ca9c6942 - 02/06/2020 12:24 AM - Artem Zhmurov

Make OpenCL DeviceVendor into enum class and move to GPU traits

The device context in OpenCL requires the information on vendor when constructed. To prepare for opaque DeviceContext, the vendor enum was moved into OpenCL traits.

Refs. #3311, needed for #3315.

Change-Id: Iec22ff17543b6a99407048de6e0cd82bb7218fb0

Revision e742ad10 - 02/06/2020 03:49 AM - Artem Zhmurov

Move DeviceInfo into GPU traits

The DeviceInfo is needed upon construction of DeviceContext. To prepare for opaque DeviceContext type, it is moved to GPU traits and renamed according to the common naming scheme.

Refs. #3311, needed for #3315.

Change-Id: I2a9f1d932f142d645df75901521a734d208de509

Revision 84e5a0e6 - 03/10/2020 08:02 AM - Artem Zhmurov

Use init(..) function to build DeviceContext

This patch unifies the logic of OpenCL context creation in PME and NBNXM by using the same init(..) function for the DeviceContext object.

Also, the DeviceInfo is now de-referenced directly after the check on the pointer validity and passed along as a const reference, which improves the clarity of the code.

Refs. #3315, #3311.

Change-Id: I5ba0f530918f3340fa1a5ad3e8d60fe4e0967dab

Revision 6975fbfd - 03/11/2020 03:59 PM - Artem Zhmurov

Take over management of OpenCL context from PME and NBNXM

This patch set creates the DeviceContext in runner and passes it to the consumers (PME and NBNXM). This removes unnecessary management code duplication, makes the device buffers in two modules compatible.

Fixes #2522

Fixes #3315

Refs #3311

Change-Id: I10358cfaced5b5c7dbddd95679c9a9703f3a2c0

Revision 8eadec22 - 03/13/2020 08:36 AM - Artem Zhmurov

Make DeviceStream into a class

Refs #3314

Refs #3311

Change-Id: I270864f0e82af63f91a91c9951bf678795680fa

Revision e3d904d0 - 03/13/2020 08:36 AM - Artem Zhmurov

Use DeviceStream init(...) function to create streams

Change the stream creation procedures from direct calls to CUDA and OpenCL API to using pre-defined init(...) method of the DeviceStream class.

Refs #3314

Refs #3311

Change-Id: I96a0ca41f251b9925ef9bed77c4f355939b65c6d

Revision c26d93da - 03/20/2020 12:26 PM - Artem Zhmurov

Small fixes to the DeviceStream

1. Fix compiler warning on having const modifier for bool in function declaration.
2. Fix comments.
3. Introduce isValid(...) method and use it inside the class.

Refs #3314

Refs #3311

Change-Id: I482aee831461f6b170c5fbf90f3f3e978282d226

Revision 99f4253d - 03/23/2020 04:46 AM - Artem Zhmurov

Introduce DeviceStreamManager

Make a separate object that will be handling the creation, management and destruction of the GPU context and streams. It is detached from the rest of the code in this patch, but will be attached in the follow-up.

Refs #3316
Refs #3311

Change-Id: I2c59b930ac266d89f9e0172b83f07e9858f0b

Revision db2c0d2e - 03/25/2020 07:47 AM - Artem Zhmurov

Make use of the DeviceStreamManager

Use the DeviceStreamManager throughout the code. The manager is owned by the runner and created when GPU is active. The consumers get the context and streams if needed.

TODOs:

1. Make builders and move the selection on whether the stream should be created there. The builders should take the manager and pass the context and the stream to the consumer. Builders should have the option to create a stream.
2. Makefile in ewald tests uses old infrastructure. Also, the device context management should be lifted from there and utilized in all the tests that can run on GPU hardware.

Refs #3316
Refs #3311

Change-Id: I0d08adbe1dee19c1890e55f0e0cf79cea97d39bd

Revision 1ced5fb7 - 04/16/2020 01:38 PM - Artem Zhmurov

Unify CUDA and OpenCL lookup-table creation

In CUDA code, textures are used for the lookup-tables, whereas in OpenCL they are created as a read-only buffers. This commit hides these differences behind a unified wrapper.

Refs #3318
Refs #3311

Change-Id: I003e0c982c2452a2753e331b46fc59f0b7e1b711

Revision 5f8899ba - 05/04/2020 07:44 PM - Artem Zhmurov

Unify CUDA and OpenCL lookup-table creation

In CUDA code, textures are used for the lookup-tables, whereas in OpenCL they are created as a read-only buffers. This commit hides these differences behind a unified wrapper.

Refs #3318
Refs #3311

Change-Id: I003e0c982c2452a2753e331b46fc59f0b7e1b711

Revision c048437f - 05/05/2020 08:37 AM - Artem Zhmurov

Unify CUDA and OpenCL lookup-table creation

In CUDA code, textures are used for the lookup-tables, whereas in OpenCL they are created as a read-only buffers. This commit hides these differences behind a unified wrapper.

Refs #3318
Refs #3311

Change-Id: I003e0c982c2452a2753e331b46fc59f0b7e1b711

Revision 986b2bb1 - 05/11/2020 10:31 AM - Artem Zhmurov

Unify CUDA and OpenCL lookup-table creation

In CUDA code, textures are used for the lookup-tables, whereas in OpenCL they are created as a read-only buffers. This commit hides these differences behind a unified wrapper.

Refs #3318

Refs #3311

Change-Id: I003e0c982c2452a2753e331b46fc59f0b7e1b711

Revision d3ce8501 - 05/15/2020 06:38 PM - Artem Zhmurov

Unify CUDA and OpenCL lookup-table creation

In CUDA code, textures are used for the lookup-tables, whereas in OpenCL they are created as a read-only buffers. This commit hides these differences behind a unified wrapper.

Refs #3318

Refs #3311

Change-Id: I003e0c982c2452a2753e331b46fc59f0b7e1b711

History

#1 - 01/23/2020 12:00 PM - Artem Zhmurov

- *Tracker changed from Task to Feature*

#2 - 01/23/2020 12:12 PM - Artem Zhmurov

- *Target version set to 2021-refactoring*

#3 - 02/05/2020 03:43 PM - Szilárd Páll

- *Related to Task #2936: introduce check that CPU-GPU transfers/assignments are made between compatible types added*