

CP2K Developers Meeting

December 8th 2021, 15:00-16:00 CET




CP2K Developers Meeting

1. Introduction
2. CP2K Release (all)
3. Double-Hybrid Gradients (Frederick Stein)
4. Updates
5. Update from Alfio
6. GPU Acceleration (all)
7. Licensing of Performance critical Code under BSD (Ole Schütt)
8. RI-HFX forces (Augustin Bussy)
- 9.
10. Next CP2K developers meeting (all)




Introduction




CP2K Release (all)

- time table
 - between christmas and new year
- open issues
 - [possible gpu issue](#) (multigpu)
 - general issue probably within DBCSR
 - dashboard should be green
 - HIP disabled
- Arch files updated



Double-Hybrid Gradients (Frederick Stein)

- stress tensors work now
- good speedup from ADMM and new implementation
- paper in preparation
- further optimization from memory usage and communication
- crash in grid code
 - not enough shared memory in GPUs
 -



Update from Ole

My TODO list for the next ~2 months:

- X-mas release
- Refactor Dockerfiles, upload production images to DockerHub.
 - /cp2k name
 - NVIDIA seems comfortable with CUDA in Containers that also have GPL-stuff
- Merge new sparse matrix code (DBM) for tensors (DBT).
 - merge after the release
 - fork dbcsr-tensor, move back to cp2k-code
- Improve collocate/integrate GPU kernels for large basis sets, ie. metals.
 -
- Look into GPU migration of GEEP to speed up QM/MM.



Update from Alfio Lazzaro via mail

As part of HPE, I'm working on the DBCSR porting to AMD GPUs, as part of the LUMI effort. This is a joint collaboration between HPE, AMD and (partially) CSCS.

For the time being, we are considering only two tests:

- a. RI-MP2 (H2O-128), which requires to update the COSMA library to make better use of the multi-gpu (COSMA is a CSCS library, that's where the CSCS comes in play)
- b. H2O-DFT-LS, which requires to update the DBCSR library

For a. there is no direct relation to the CP2K release (COSMA is an external dependency). For b, DBCSR is (still) integrated in CP2K, so it needs an update for the next CP2K release.

AMD did a recent PR in DBCSR, which is supposed to fix an issue with HIP and threading. There are other problems with the multi-gpu environment, which I will try to fix next week. As of now, the plan is to have a new DBCSR release in two weeks time.

Notes:

- Pilot project on LUMI in April, big chunk/full of LUMI, 2-3 days
- double hybrid, stress tensor, ~400 atoms
- contact HPE because of submatrix method → showcases




Paderborn

- Nersc allocation for exaflop submatrix method on A100s
- Extension from submatrix+xtb to submatrix+dft
- Pull request of the former is in preparation
- BMBF-Project (with ELPA@RZG, FHI) application for GPU acceleration with CP2K and DFTB+
 - submatrix/CP2K as science case
- Hossam: periodic multipoles (based on Phys. Rev. B 100, 245135 (2019)) as a pull request before Christmas/Release
- D4 dispersion correction after the release as a pull request
- ERI on FPGA
 - mostly algorithm design right now
- FFT-acceleration with FPGAs
 - general effort, prime-number ffts
 - e.g. 64^3 -ffts supporting all ffts up to 64



Update from Augustion Bussy

- forces for RI-HFX
- GPU acceleration of RHO and MO flavours on Daint
- sparse tensor contraction not accelerated on GPUs
- production quality code by Christmas
- Ole: need test case for single Daint node
- next step: more optimizations



Licensing of Performance critical Code under BSD (Ole Schütt)

<https://github.com/cp2k/cp2k/tree/master/tools/toolchain#licenses>


- already implemented
- NVIDIA can only help if they don't have to touch GPL code
- all relevant contributors have already agreed
- now BSD for these parts



GPU Acceleration (all)

Current Efforts:

- Ole Schütt: collocate and integrate in grid code, Nvidia already looked at it
 - plan: extend to higher angular momentum, solve resource usage balance
- Ole Schütt: tensor code on GPUs, fairly simple and generic kernel with “tensor cores” in GPUs, problems are tensors with small block sizes
- Ole Schütt: GEEP acceleration on GPUs (in January), maybe use grid-code
- Alfio Lazzaro: usage on AMD GPUs (linear-scaling, RI-MP2)
- Matt Watkins: ERIs in periodic systems (collaboration with CRYSTAL code guys) via libcint library, exploratory phase
- Paderborn: Submatix with xTB, extension DFT
- Joost (CSCS): fft, spFFT acceleration on GPUs
- Joost (CSCS): eigensolver on GPUs



Next CP2K developers meeting (all)

- monthly or **bi-monthly**?
 - 1h?
- four times per year, up to two hours

