



Contribution ID: 135

Type: **Poster**

Porting the hydrologic model ParFlow to different accelerator architectures using eDSL and Kokkos

Wednesday 26 February 2025 19:40 (20 minutes)

The ParFlow hydrologic model is an integrated variably saturated groundwater, surface water flow simulator that incorporates subsurface energy transport and land surface processes through the integration of the Common Land Model (CLM) as a module. In addition ParFlow has been coupled to atmospheric models, such as WRF, COSMO and ICON. ParFlow is also integrated in the German climate and weather prediction ICON(-Land) software ecosystem as part of the WarmWorld project and it is an important component of the Terrestrial Systems Modeling Platform (TSMP), which enables integrated simulations from the bedrock, across the land surface to the top of the atmosphere with the coupled ICON/COSMO-CLM-ParFlow modeling system.

ParFlow is written in C (with additional Fortran 90 parts (especially if CLM is enabled)) and uses the parallelization methods MPI, OpenMP, native CUDA support and the programming model Kokkos (with the backend CUDA, HIP or OpenMP). As a matter of fact, the parallelism in ParFlow has been abstracted early on in what is called an embedded Domain Specific Language (eDSL) which is leading to a best-practice separation-of-concerns, which means the domain scientist/developer does not see e.g. a single MPI call when programming in ParFlow.

Since future hardware will be characterized by varying architectures there was a demand to also enable HIP for AMD architectures. For example, the pre-exascale HPC system LUMI of EuroHPC is based on an accelerator consisting of AMD GPUs. To implement HIP for ParFlow via the eDSL and Kokkos (Cuda was already implemented) the porting work started in December 2022 in a sprint of the national Earth System Modelling Initiative (natESM) and was continued and finished at IBG3 (Forschungszentrum Jülich). Based on the eDSL the AMD porting was done using Kokkos because the Kokkos ecosystem already includes an implementation of the HIP programming model for AMD GPUs and which resulted for ParFlow in a high degree of performance portability. Performance and scalability have been demonstrated on JUWELS Booster (Jülich Supercomputing Centre) for Nvidia GPUs (Nvidia A100) and also on the LUMI supercomputer at CSC (Finland) with AMD MI250X accelerators.

In this poster we will present the eDSL of ParFlow and also how Kokkos (with the Cuda and HIP backend) is included in its eDSL and to allow ParFlow to reach performance portability on the basis of the eDSL and Kokkos, especially for Nvidia (Cuda) and on AMD platforms (HIP) with the change of only a limited amount of lines. We will also present scaling plots for different machines and accelerators and how to enable HIP as a Kokkos backend of ParFlow using the eDSL of ParFlow.

I want to participate in the youngRSE prize

no

Primary authors: BENKE, Joerg; Dr FAHAD, Muhammad (Forschungszentrum Jülich)

Co-authors: Dr HERTEN, Andreas (Forschungszentrum Jülich); Dr CAVIEDES-VOULLIEME, Daniel (Forschungszentrum Jülich); Prof. KOLLET, Stefan (Forschungszentrum Jülich)

Presenter: BENKE, Joerg

Session Classification: Poster and Demo Session together with Reception

Track Classification: Research Software: domain-specific languages