

EGU2020-12904

<https://doi.org/10.5194/egusphere-egu2020-12904>

EGU General Assembly 2020

© Author(s) 2022. This work is distributed under the Creative Commons Attribution 4.0 License.



## Accelerated hydrologic modeling: ParFlow GPU implementation

Jaro Hokkanen<sup>1</sup>, Jiri Kraus<sup>2</sup>, Andreas Herten<sup>3</sup>, Dirk Pleiter<sup>3</sup>, and Stefan Kollet<sup>1</sup>

<sup>1</sup>FZ-Jülich, IBG-3, Germany

<sup>2</sup>NVIDIA GmbH, Germany

<sup>3</sup>FZ-Jülich, JSC, Germany

ParFlow is known as a numerical model that simulates the hydrologic cycle from the bedrock to the top of the plant canopy. The original codebase provides an embedded Domain-Specific Language (eDSL) for generic numerical implementations with support for supercomputer environments (distributed memory parallelism), on top of which the hydrologic numerical core has been built.

In ParFlow, the newly developed optional GPU acceleration is built directly into the eDSL headers such that, ideally, parallelizing all loops in a single source file requires only a new header file. This is possible because the eDSL API is used for looping, allocating memory, and accessing data structures. The decision to embed GPU acceleration directly into the eDSL layer resulted in a highly productive and minimally invasive implementation.

This eDSL implementation is based on C host language and the support for GPU acceleration is based on CUDA C++. CUDA C++ has been under intense development during the past years, and features such as Unified Memory and host-device lambdas were extensively leveraged in the ParFlow implementation in order to maximize productivity. Efficient intra- and inter-node data transfer between GPUs rests on a CUDA-aware MPI library and application side GPU-based data packing routines.

The current, moderately optimized ParFlow GPU version runs a representative model up to 20 times faster on a node with 2 Intel Skylake processors and 4 NVIDIA V100 GPUs compared to the original version of ParFlow, where the GPUs are not used. The eDSL approach and ParFlow GPU implementation may serve as a blueprint to tackle the challenges of heterogeneous HPC hardware architectures on the path to exascale.