Geophysical Research Abstracts
Vol. 21, EGU2019-8129-1, 2019
EGU General Assembly 2019

Smart Routing Network Decomposition for the Massive Parallelization of Global Hydrological Models

Maren Kaluza (1,2), Luis Samaniego (1), Stephan Thober (1), Robert Schweppe (1,2), Rohini Kumar (1), and Oldrich Rakovec (1)

(1) UFZ, CHS, Leipzig, Germany (maren.kaluza@ufz.de), (2) University of Potsdam

Parameter estimation of high-resolution, continental-scale hydrological models requires powerful supercomputers. Therefore efficient parallelization of the model source code and the usage of extremely efficient algorithms are essential. Improved performance of large-domain hydrological models will take us a step closer to more realistic description of the global hydrology. This is in accordance with the ESM project (1), to which this work contributes. One of the main challenges to achieve this goal is the parallelization of the global streamflow routing network. While the parallelization of grid related processes is straightforward because there is no dependence with fluxes estimated in faraway cells, the parallelization of the streamflow routing along a dense streamflow network is a hard computational problem because of the dependency of streamflow estimated in every cell with that estimated upstream.

An optimal river network decomposition is necessary to preserve the routing order and gain speedup. To achieve this goal, a hybrid parallelization approach is adopted (MPI-OpenMP). We recursively cut off approximately equal sized subdomains to gain computational load balancing and distribute them among all available computational nodes using several effective scheduling algorithms for tree structured data. The corresponding problem in complexity theory is the partitioning and scheduling of forest data structures. The objective of finding a near optimal decomposition is a hard and by far a not entirely explored problem. This first step is carried out with MPI for large sub-river networks and inside these subdomains, OpenMP is applied for grid processes and routing to profit from the shared memory available for each node. The resulting routing schedule is static and can be used for each optimization step.

In this study, we apply this parallelization strategy to the open-source mHM (www.ufz.de/mhm) and test the speedup on the EDgE data (Samaniego et al. 2019 BAMS under review) representing all of the European rivers. This test was performed on the high-performance supercomputer JUWELS at Jülich Supercomputing Center (JSC).

Here, we present and discuss the impact of the parallelization parameters on the speedup curve: 1) the size of the subdomains, 2) the balance between MPI and OpenMP processes, 3) the buffer size of collected data messages between the subdomains, and 4) the impact of the chosen scheduling algorithm. We show, that buffering arrays of dimension 1000 lead to a communication overhead with MPI if the size of the subdomains is smaller than 20 cells. Here it is also shown that the slope of the speedup curve starts steeper without OpenMP but in the Hybrid Version the speedup overtakes the version without OpenMP. This insights allowed us to conclude that we are able to gain a nearly linear speedup with the hybrid version using up to 1500 CPUs. Finally, we discuss the limiting factors of the possible speedup.

1. https://www.esm-project.net/