



## **Adaptation of the anelastic solver EULAG to high performance computing architectures.**

Damian Wójcik (1), Miłosz Ciżnicki (2), Piotr Kopta (2), Michał Kulczewski (2), Krzysztof Kurowski (2), Zbigniew Piotrowski (1), Krzysztof Rojek (3), Bogdan Rosa (1), Łukasz Szustak (3), and Roman Wyrzykowski (3)

(1) Institute of Meteorology and Water Management - National Research Institute, Warszawa, Poland (damian.wojcikimgw.pl), (2) Poznań Supercomputing and Networking Center, Poznań, Poland, (3) Częstochowa University of Technology, Częstochowa, Poland

In recent years there has been widespread interest in employing heterogeneous and hybrid supercomputing architectures for geophysical research. Especially promising application for the modern supercomputing architectures is the numerical weather prediction (NWP). Adopting traditional NWP codes to the new machines based on multi- and many-core processors, such as GPUs allows to increase computational efficiency and decrease energy consumption. This offers unique opportunity to develop simulations with finer grid resolutions and computational domains larger than ever before. Further, it enables to extend the range of scales represented in the model so that the accuracy of representation of the simulated atmospheric processes can be improved. Consequently, it allows to improve quality of weather forecasts.

Coalition of Polish scientific institutions launched a project aimed at adopting EULAG fluid solver for future high-performance computing platforms. EULAG is currently being implemented as a new dynamical core of COSMO Consortium weather prediction framework. The solver code combines features of a stencil and point wise computations. Its communication scheme consists of both halo exchange subroutines and global reduction functions. Within the project, two main modules of EULAG, namely MPDATA advection and iterative GCR elliptic solver are analyzed and optimized. Relevant techniques have been chosen and applied to accelerate code execution on modern HPC architectures: stencil decomposition, block decomposition (with weighting analysis between computation and communication), reduction of inter-cache communication by partitioning of cores into independent teams, cache reusing and vectorization. Experiments with matching computational domain topology to cluster topology are performed as well. The parallel formulation was extended from pure MPI to hybrid MPI - OpenMP approach. Porting to GPU using CUDA directives is in progress.

Preliminary results of performance of the new implementation show promising increase of computational efficiency. The details depend on the particular setup of the numerical experiments. The gained experience will be a base for a complete EULAG version for emerging HPC architectures.