EMS Annual Meeting Abstracts
Vol. 14, EMS2017-615, 2017

ESCAPE: Scaling Weather and Climate for Exascale

Alastair McKinstry, Oisin Robinson, and Michael Lysaght
NUI Galway, Irish Centre for High-End Computing, Galway, Ireland (alastair.mckinstry@ichec.ie)

ESCAPE is the Horizon 2020 project, Energy-efficient Scalable Algorithms for Weather Prediction at Exascale. To reach the cloud-resolving scale of 1km, we need both Exscale (1018 flops/sec) compute power, but also to stay within power budget of < 10-20 MW. To do this, ESCAPE addresses the ETP4HPC Strategic Research Agenda 'Energy and resiliency' priority topic, developing a holistic understanding of energy-efficiency for extreme-scale applications using heterogeneous architectures, accelerators and special compute units.

To do this, we define fundamental algorithmic building blocks ('Weather & Climate Dwarfs') underlying weather and climate services. These are then used for subsequent co-design, optimization, and adaptation efforts: building performance models to understand the limitations and requirements of these dwarves, optimising them for multiple “accelerator” architectures, and exploring the requirements of new alternative algorithmic approaches within this performance analysis framework.
ESCAPE combines the resources of multiple agencies: the European Centre for Medium-Range Weather Forecasts, Danmarks Meteorologiske Institut, Deutscher Wetterdienst, l'Institut Royal Météorologique de Belgique, Météo-France and MeteoSchweiz; and on the compute side, Instytut Chemii Bioorganicznej Polskiej Akademii Nauk; Loughborough University, The Irish Centre for High-End Computing, Bull SAS, NVIDIA Corporation; and Optalysys Ltd.

ICHEC has been working with porting and analysing the performance of a number of dwarves on accelerators, in particular the Intel Xeon Phi accelerators and nVidia GPU processors, compared to traditional CPU architectures. We explore the requirements needed for speedups on these platforms and the bottlenecks and limitations possible. To date we have optimised the advection schemes LAITRI (Semi-Lagrangian) and MPDATA (multidimensional positive definite advection algorithm (for finite volume methods) and the BiFourier and Spherical Harmonics dwarves.

We show that speedups relative to traditional Xeon CPUs are possible but depend heavily on data layout for optimal vectorisation. Within ESCAPE, the “ATLAS” data abstraction library and the “Grid Tools” domain specific language system have been developed by ECMWF and Meteo Swiss to abstract data layout details from the algorithms. We show the changes that need to be made to these components for maximal performance, and show the performance limits of the system produced as seen using “Roofline” models approaches the bandwidth limits of the hardware.