Geophysical Research Abstracts
Vol. 20, EGU2018-10474, 2018
EGU General Assembly 2018

Exascale computing in numerical weather prediction: massively parallel I/O in atmospheric models on conformal meshes

Dom Heinzeller (1,2,3), Michael Duda (4), Thomas Schwitalla (5), Harald Kunstmann (3,6)

(1) National Oceanic and Atmospheric Administration, Earth System Research Laboratory, Global Systems Division, Boulder, CO, USA, (2) University of Colorado Cooperative Institute for Research in Environmental Sciences, Boulder, CO, USA, (3) Karlsruhe Institute of Technology, Institute of Meteorology and Climate Research, Garmisch-Partenkirchen, Germany, (4) National Center for Atmospheric Research, Mesoscale and Microscale Meteorology Laboratory, Boulder, CO, USA, (5) University of Hohenheim, Institute of Physics and Meteorology, Stuttgart, Germany, (6) University of Augsburg, Department of Geography, Augsburg, Germany

With strong financial and political support from national and international initiatives, exascale computing is projected for the end of this decade. Leading operational weather forecasting centers such as ECMWF, NCEP and the UK Met Office are pushing towards convection-permitting, global simulations and large ensembles to improve their operational forecasting products. In order to fully exploit the capabilities of next-generation exascale computing systems, existing atmospheric modelling systems need to undergo significant development.

One challenge in this respect is the ability of the dynamical solver to scale out to orders of magnitude more cores than today. Traditional models on regular latitude-longitude grids as well as spectral models are challenged by this development as they require special treatment of certain grid cells or global communication at each model time step, both of which can seriously limit the parallel performance of the entire application. Models on conformal meshes like the Model for Prediction Across Scales (MPAS), on the other hand, can benefit to a much larger extent from modern computational architectures, since all mesh cells are treated equally and only local communication is required.

Another challenge for applications at extreme scale is the efficiency of the model's I/O layer. Traditional file formats such as NetCDF or GRIB are subject to severe scalability issues for large numbers of reading and writing tasks and for very large data sets, typically of the size of several terabytes in global, convection-permitting simulations. Model developers need to resort to specialized I/O libraries designed for massively parallel applications such as the SIONlib library, developed at Research Centre Jülich (FZJ), to avoid slowing down the model run.

In this contribution, we present the implementation of an additional I/O layer based on the SIONlib library in the MPAS model. Our implementation benefits from the SIONlib I/O layer in two ways as it reduces not only the read and write times by factors of 5 to 60 compared to parallel NetCDF formats, but also the time required to initialize the model by up to 90% by encoding task-dependent information, such as the horizontal mesh decomposition, in the SIONlib data.

We demonstrate that embracing the power of a conformal Voronoi mesh and a massively-parallel I/O layer enables us to conduct modelling experiments at extreme scale with the atmospheric component of MPAS, MPAS-A. In particular, we show results for the first application of MPAS on a global 1km mesh with more than 589 million horizontal grid columns.