Geophysical Research Abstracts
Vol. 19, EGU2017-10729, 2017
EGU General Assembly 2017

Mixed Single/Double Precision in OpenIFS: A Detailed Study of Energy Savings, Scaling Effects, Architectural Effects, and Compilation Effects

Mike Fagan (1), Peter Dueben (2), Krishna Palem (3), Glenn Carver (4), Matthew Chantry (5), Tim Palmer (6), and Jeremy Schlacter (7)

(1) Rice University, Houston, Texas, United States (mfagan@cs.rice.edu), (2) ECMWF, Reading, UK (peter.dueben@ecmwf.int), (3) Rice University, Houston, Texas, United States (palem@rice.edu), (4) ECMWF, Reading, UK (Glenn.Carver@ecmwf.int), (5) University of Oxford, Oxford, UK (matthew.chantry@physics.ox.ac.uk), (6) University of Oxford, Oxford, UK (Tim.Palmer@physics.ox.ac.uk), (7) EPFL, Lausanne, Switzerland (jeremy.schlachter@epfl.ch)

It has been shown that a mixed precision approach that judiciously replaces double precision with single precision calculations can speed-up global simulations. In particular, a mixed precision variation of the Integrated Forecast System (IFS) of the European Centre for Medium-Range Weather Forecasts (ECMWF) showed virtually the same quality model results as the standard double precision version (Vana et al., Single precision in weather forecasting models: An evaluation with the IFS, Monthly Weather Review, in print). In this study, we perform detailed measurements of savings in computing time and energy using a mixed precision variation of the -OpenIFS- model. The mixed precision variation of OpenIFS is analogous to the IFS variation used in Vana et al. We (1) present results for energy measurements for simulations in single and double precision using Intel's RAPL technology, (2) conduct a -scaling- study to quantify the effects that increasing model resolution has on both energy dissipation and computing cycles, (3) analyze the differences between single core and multicore processing, and (4) compare the effects of different compiler technologies on the mixed precision OpenIFS code. In particular, we compare intel icc/ifort with gnu gcc/gfortran.