Geophysical Research Abstracts
Vol. 20, EGU2018-6463, 2018
EGU General Assembly 2018

How much precision is needed for accurate weather and climate simulations?

Matthew Chantry (1), Tobias Thornes (1), Peter Duben (2), and Tim Palmer (1)

(1) University of Oxford, Atmospheric, Oceanic and Planetary Physics, UK (matthew.chantry@physics.ox.ac.uk), (2) European Centre for Medium-Range Weather Forecasts, Reading, UK

With the plateau of CPU clock speeds, one cannot simply wait for higher resolution weather and climate simulations. Scalability of massively parallel codes is seen as the key to pushing towards 1km global models. Here we address a complementary approach. Predominately weather and climate models calculate using 64-bit arithmetic, a historic norm from machines which gave performance comparable to 32-bit arithmetic. Recently machine learning has moved in the direction of short precisions, 16-bits (half precision) and below, with the necessary linear algebra performed on accelerator hardware. The question is whether the weather and climate community can take advantage of these hardware developments. The inherent inaccuracies in current models, from large grid spacing to uncertain physical parameterisation schemes, suggest that lower than 64-bit precisions are feasible, but is half-precision achievable? Using software emulation we investigate kernels of ECMWF's OpenIFS model, focusing on the significand bits. On both weather and climate timescales large precision reductions are possible, even as resolution is increased towards operational levels. While uniform half-precision appears challenging we are able to reduce the majority of the investigated fields and calculations to this threshold. This will fuel further study into the dynamic ranges and exponent bits before hardware formulations can be tested.