ACME

Summary Version

1.0

Purpose of Benchmark

The purpose of this benchmark is to weak scale the ACME atmospheric dynamical core to the whole machine while efficiently using any accelerator devices and managing inter-node data transfers. The benchmark is developed to run on Titan.

Characteristics of Benchmark

The ACME model is a US Department Of Energy high-resolution Earth system model used for projecting future climate for impacts important to the DOE. Three scientific outcomes important for the ACME project are resolving current states and future climate changes regarding (1) the Earth's water cycle, (2) biogeochemistry interactions, and (3) the cryosphere. The dynamical core of the atmospheric model will be used for this RFP. This models gas-only stratified fluid flow on the rotating sphere as well as transporting chemical constituents used by physics routines along wind trajectories. It is discretized with the Spectral Element method on the cubed sphere. It uses MPI, OpenMP, and OpenACC for parallelization, threading, and accelerator vectorization.

Mechanics of Building Benchmark

A documented compile script located at homme/compile_scripts/titan/compile.cmake is provided that does everything necessary to build the benchmark on Titan. The job submission used on Titan is located at test/jw_baroclinic/openacc_work/openacc.job. This script would be submitted on Titan as "qsub openacc.job". Further, a README.CORAL2 file is placed in the root directory of the tarball with documentation of the relevant parameters that may be changed in the build and job submission scripts.

Mechanics of Running Benchmark

The benchmark is set up to be run in weak scaling mode at three different scales, each with 32 columns of elements per node:

- 304 x 304 x 6 ("ne304") on 17328 Titan nodes
- 216 x 216 x 6 ("ne216") on 8748 Titan nodes
- 152 x 152 x 6 ("ne152") on 4332 Titan nodes

Ideally, the number of time steps solved per unit of time will remain constant as the code is weak scaled to a larger node count.

CORAL-Class Problem

For the CORAL-class problem benchmarks, vendors are allowed to refactor the code in any way, so long as the following three criteria are met:

1. The three Titan benchmarks continue to have a relative L1-norm diff against the baseline files of 1e-10 or less after one day of simulation. This is simply to ensure the answer is still correct.
2. The number of time steps processed per second does not fall below 1.
3. The total number of columns of elements does not exceed 960 x 960 x 6 ("ne960"). This is important because hydrostatic assumption fails if the grid spacing becomes too small.

As for runtime parameters, the vendors may only change the "ne" parameter in openacc.job. To maintain a stable run at different values of "ne", the "tstep" and "nu" parameters must also be updated as functions of "ne" via the following relationships:

tstep = 9e3 * ne^-1
nu = 8.1e19 * ne^-3.322

Vendors may change the parallel decomposition in terms of processes, threading, work overlapping, and vectorization.

Verification of Results

The Figure Of Merit (FOM) is the number of time steps solved per second, and this number will be directly written out from the job submission script itself in a post processing step. Correct results can be verified via the NetCDF diff utility that computes relative difference norms between two NetCDF files. A fiducial baseline file is located at homme/baselines/ne304_baseline.nc to facilitate this test. The job submission script will compute this norm as a part of the job itself and give the output of the norm in the job output.