Lawrence Livermore National Laboratory

Progress on Advanced Dynamical Cores for the Community Atmosphere Model

June 2010

Art Mirin

Lawrence Livermore National Laboratory, P. O. Box 808, Livermore, CA 94551

This work performed under the auspices of the U.S. Department of Energy by
Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344

LLNL-PRES-438171

A primary motivator for advanced dycores is to overcome limitations of the lat-lon grid

- **Convergence of meridians at poles**
  - **limits timestep and provides unnecessarily high resolution in polar regions at expense of other locations throughout globe**
  - **limits scalability by inhibiting effective domain decomposition in longitude**
- **Several new dynamical cores are currently available in versions of CAM**
  - ***Homme* (spectral element, cubed sphere grid, cam trunk)**
  - ***Fvcubed* (finite volume, cubed sphere grid, cam branch)**
  - ***Mpas* [finite volume, icosahedral (hexagonal) grid, cam branch]**

AMWG session

- **The following three slides are for the AMWG summary**

OpenMP in Homme has been revived

- ***Collaborators: John Dennis (NCAR); Jim Edwards (IBM); Mark Taylor (Sandia)***
- **Homme can now utilize both distributed memory parallelism (MPI) and shared memory parallelism (OpenMP)**
- **Both forms of parallelism operate over finite elements**
  - **OpenMP does not provide additional parallelism but could provide more effective parallelism**
- **For NE16NP4 configuration on Cray XT5 (jaguarpf), OpenMP allows close to a 50% speedup (for same number of cpu's)**
  - **1536 MPI tasks, no OpenMP => 21.4 sim.years per comp. day**
  - **256 MPI tasks, 6-way OpenMP => 31.3 sim.years per comp. day**
- **Comprehensive scaling study in progress**

Further validation of fvcubed has been accomplished

- ***Collaborators: Will Sawyer (CSCS); Brian Eaton (NCAR); Christiane Jablonowski (U. Mich.)***
- **Dynamics validated using Jablonowski-Williamson test cases**
- **Diagnosed and corrected bugs relating to physics**
- **Evaluated and upgraded OpenMP implementation**
- **Compared physics tendencies on FV grid vs FVCUBED grid**
- **Currently being upgraded from cam3_6_57 to cam4_9_02**
- **About to be updated to latest GFDL dycore version**
- **Will carry out further validation and comprehensive performance evaluation**

MPAS dycore has been implemented and is undergoing validation

- ***Collaborators: Dan Bergmann, Jeff Painter, Mike Wickett (LLNL); Todd Ringler (LANL); Bill Skamarock, Michael Duda, Joe Klemp (NCAR)***
- **MPAS is amenable to local mesh refinement**
- **Compared physics tendencies on FV grid vs MPAS grid**
- **Compared MPAS driven by CAM (without physics) with MPAS-standalone for baroclinic wave tests**
- **Inclusion of cell boundary arrays in netcdf output allows visualization with CDAT**
- **Creating mapping files with corresponding lat-lon grid, for AMWG diagnostics package and operation with land model**
- **Running comprehensive aquaplanet comparison with other dycores**

SEWG session

- **The remaining slides are for the SEWG session**

OpenMP in Homme has been revived

- ***Collaborators: John Dennis (NCAR); Jim Edwards (IBM); Mark Taylor (Sandia)***
- **Homme can now utilize both distributed memory parallelism (MPI) and shared memory parallelism (OpenMP)**
- **Both forms of parallelism operate with respect to spectral elements**
  - **OpenMP does not provide additional parallelism but could provide more effective parallelism**
  - **potential benefits of shared memory over distributed memory parallelism include:**
    - **lower memory footprint**
    - **more favorable surface area to volume ratio**

The Homme dynamics advance is an OpenMP region

- **Many codes implement OpenMP over loops**
  - ***!$OMP parallel do***
  - ***do***
  - ***enddo***
- **In Homme, the time integration encompasses a single parallel region, with specific elements assigned to specific threads**
  - ***!$OMP parallel***
  - ***call dynamics_advance***
  - ***!$OMP end parallel***
  - **mpi communications and other key areas limited to master thread (*!$OMP master*)**
  - **synchronization accomplished through *!$OMP barrier***
  - **approach is potentially more efficient but more difficult to debug**

We encountered several issues in the OpenMP implementation

- **CAM reproducible sum (which references MPI) called from parallel region**
  - **restricted the *repro_sum* call to master thread**
  - **used temporary thread-shared buffer to transfer data between master and non-master threads**
    - ***global_buf(ie) = thread_private_mem(ie)***
    - ***!$OMP MASTER***
    - ***call repro_sum(global_buf, global_sum)***
    - ***!$OMP END MASTER***
- **Limited OpenMP support in CAM/Homme interface layer**
  - **enhanced OpenMP support**

We encountered several issues in the OpenMP implementation (cont.)

- **Indeterminate data clobbering during MPI communication**
  - **each thread packs data to be communicated into buffer accessible to master thread; master thread calls MPI; data is then extracted into thread-private structure**
  - **inserted *!$OMP barrier* calls between buffer unpacks and subsequent buffer packs**
- **Lack of OpenMP barrier after *timelevel_update* (which adjusts time level indices; called by master thread only) resulted in incorrect indices being used by non-master threads**
  - ***!$OMP MASTER***
  - ***call TimeLevel_update***
  - ***!$OMP END MASTER***
  - ***!$OMP BARRIER* [was not previously present]**
  - ***call prim_advance_exp* [uses new time level indices]**

OpenMP appears to pay off on Cray XT5

- **For NE16NP4 (~2-deg) configuration on Cray XT5 (jaguarpf), OpenMP results in a ~50% speedup (for same number of cpu's)**
  - **1536 MPI tasks, no OpenMP => 21.4 sim.years per comp. day**
  - **256 MPI tasks, 6-way OpenMP => 31.3 sim.years per comp. day**
- **Above result based on a single case (run multiple times); comprehensive scaling study in progress**

Effort to run CAM with fvcubed has been revived

- ***Collaborators: Will Sawyer (CSCS); Brian Eaton (NCAR); Christiane Jablonowski (U. Mich.)***
- **Dynamics validated using Jablonowski-Williamson test cases**
- **Initial performance tests encouraging**
  - **graph below from Cray XT4 (June 2009)**

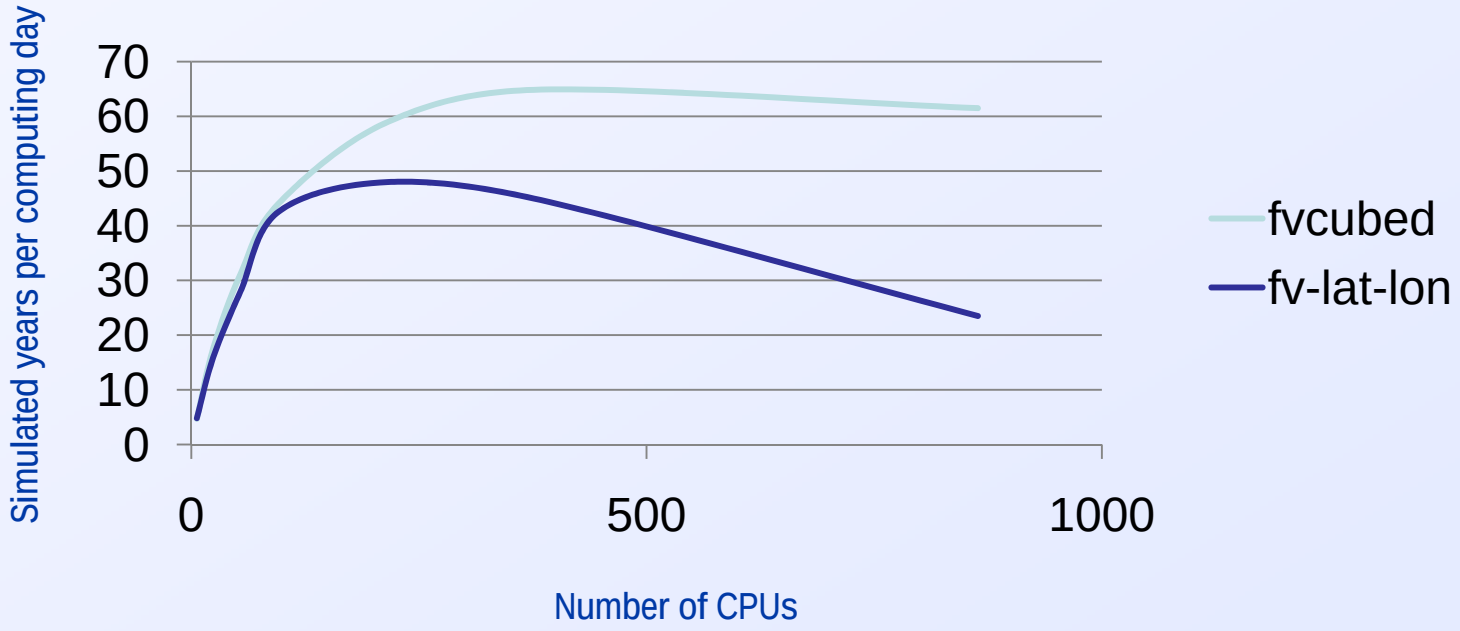

Further validation of fvcubed has been accomplished

- **Diagnosed and corrected bugs relating to physics**
  - **multiple chunks per dynamics block (phys_loadbalance=0) now supported**
  - **phys_loadbalance (>0) options requiring communication not yet implemented**
- **Evaluated and upgraded OpenMP implementation in both dycore itself and CAM/fvcubed interface layer**
- **Compared physics tendencies on FVCUBED grid vs FV grid**
  - **skip dynamics advance (for FV, need to convert potential temperature to virtual potential temperature to account for mismatch between input and output states to/from dyn_run)**

Temperature change over month (physics tendencies only, no dynamics)

fvcubed

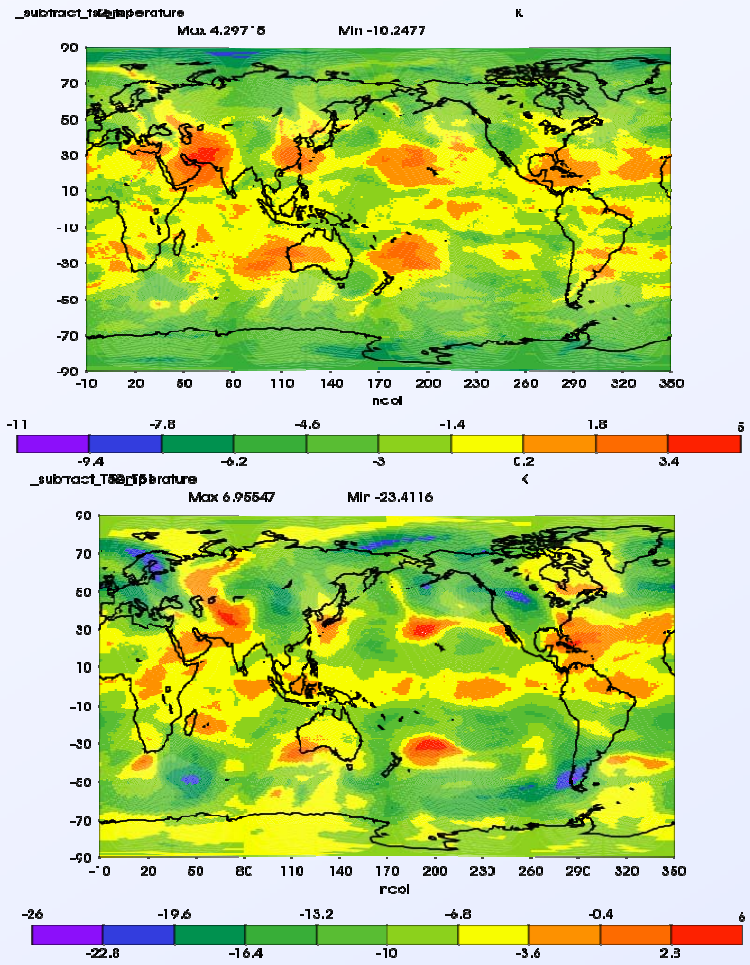

surface

500 hPa

fv

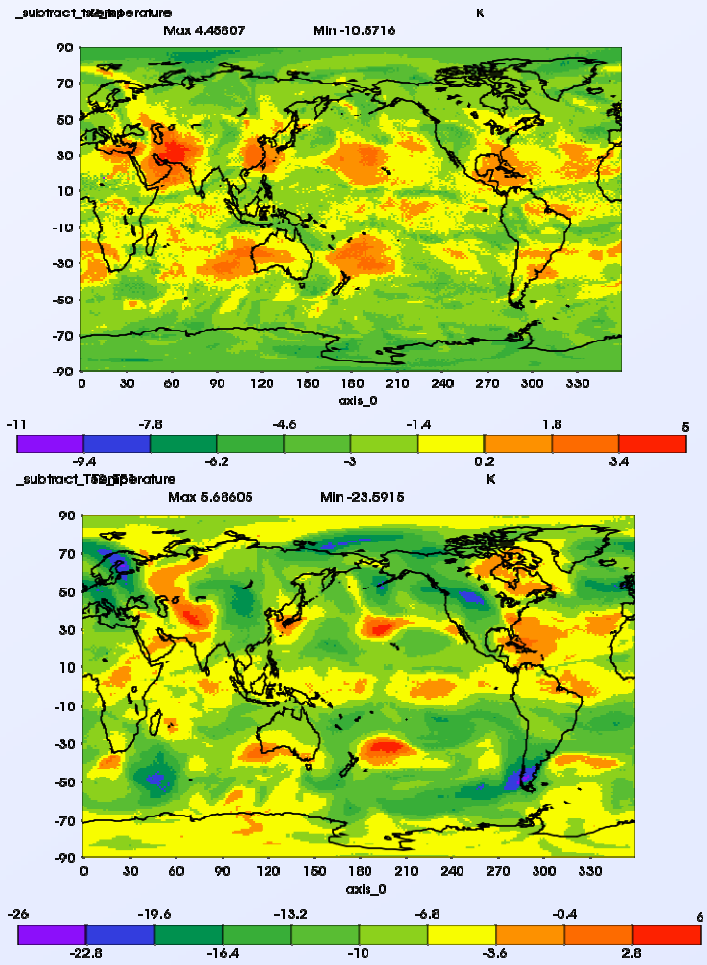

CAM/fvcubed implementation is being updated

- **Currently being upgraded from cam3_6_57 to cam4_9_02**
- **When complete, will update to latest GFDL dycore version**
  - **previously used NASA version of fvcubed, hence needed to extract added infrastructure (labor intensive)**
  - **have obtained direct access to GFDL archive (Balaji) and will use that version (provides easier path to future updates)**
- **Will then carry out further validation and comprehensive performance evaluation**

MPAS dycore has been implemented in CAM

- ***Collaborators: Dan Bergmann, Jeff Painter, Mike Wickett (LLNL); Todd Ringler (LANL); Bill Skamarock, Michael Duda, Joe Klemp (NCAR)***
- **MPAS (Model Prediction Across Scales) uses unstructured icosahedral (hexagonal) grid**
  - **finite-volume approach using C-grid (normal velocities at cell edges)**
  - **conformal grid (no hanging nodes) is amenable to local mesh refinement**
  - **dycores being developed for atmosphere and ocean**

MPAS uses conformal, variable resolution grid

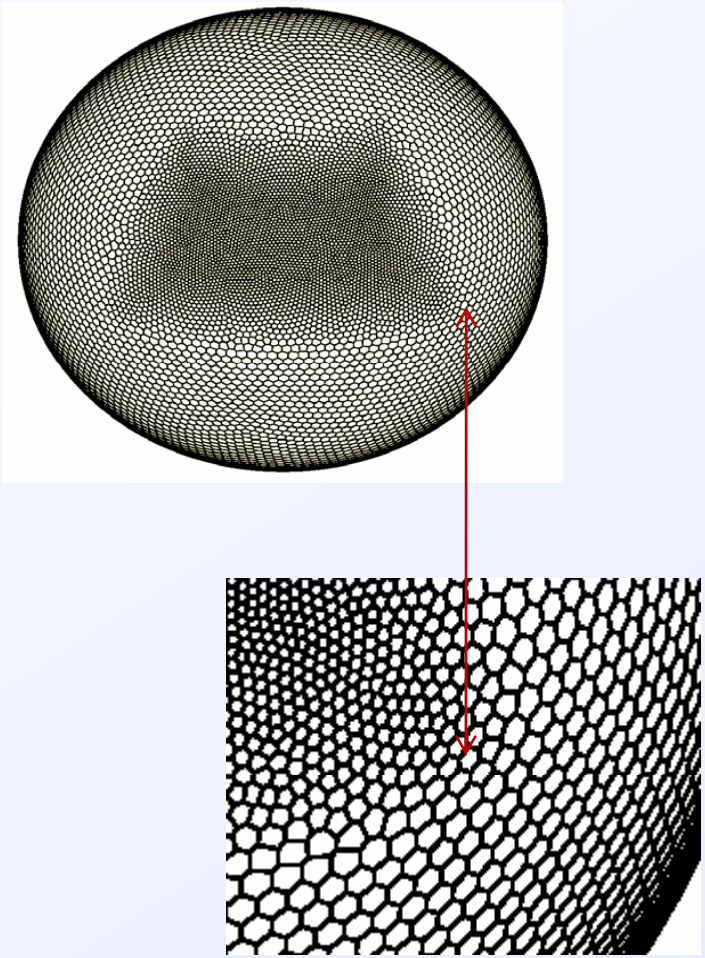

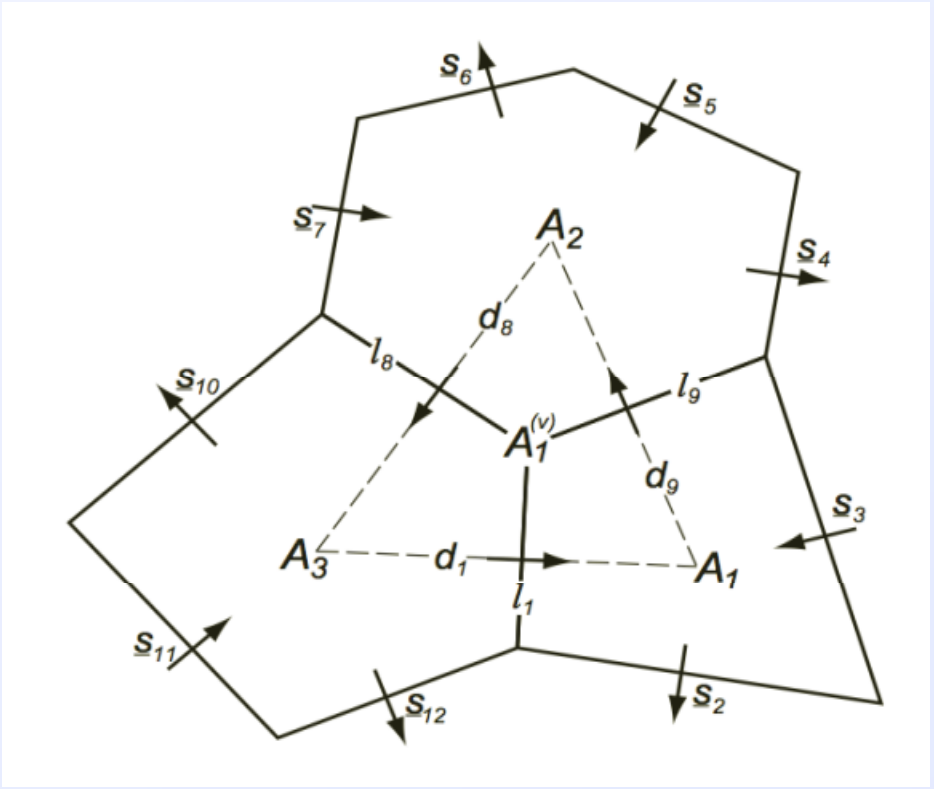

MPAS uses vertical coordinate different from those of other CAM dycores

- **MPAS uses an eta-based vertical coordinate, but with dry pressure instead of total pressure**
  - **Pd(x,k) = A(k)*(p0-pt) + B(k)*(psd(x)-pt) + pt**
  - **reduces to same functional form as other dycores if pt=0**
- **Important instances of A(k),B(k) in CAM replaced by pressure state variable or reference pressure**
- **Reference pressure (used for parameterizations) supplied by reference pressure module**
  - **Pref(k) = A(k)*(p0-pt) + B(k)*(p0-pt) + pt**

CAM/MPAS is undergoing validation

- **Compared physics tendencies on FV grid vs MPAS grid**
- **Inclusion of cell boundary arrays in netcdf output allows visualization with CDAT**
- **Compared MPAS driven by CAM (without physics) with MPAS-standalone for baroclinic wave tests, including advection of passive tracers**
- **Codes run on CRAY XT5 (jaguar) and Opteron-Infiniband system (atlas)**

Zonal velocity change over month (physics tendencies only, no dynamics)

fv

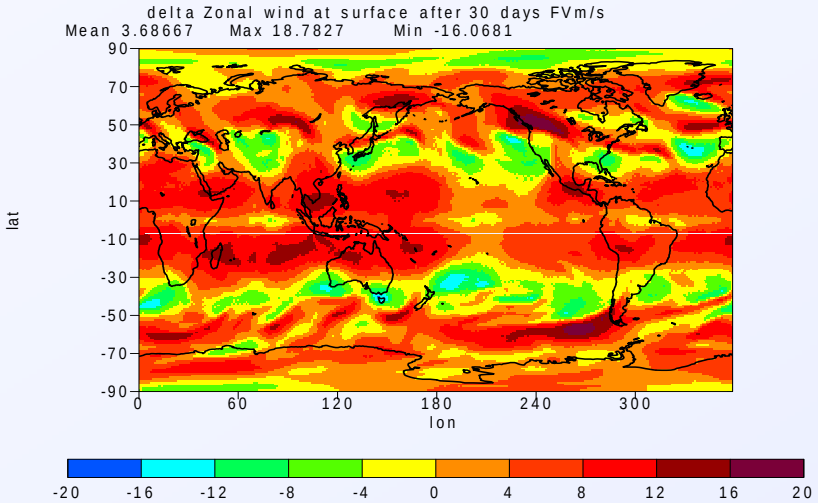

surface

mpas

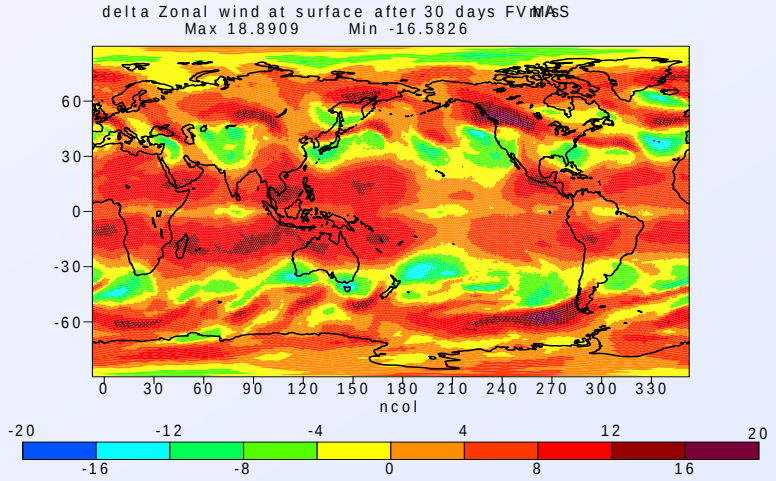

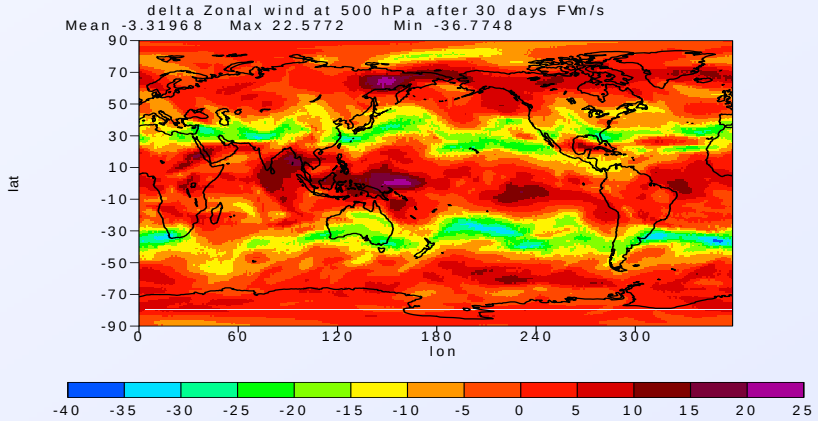

500 hPa

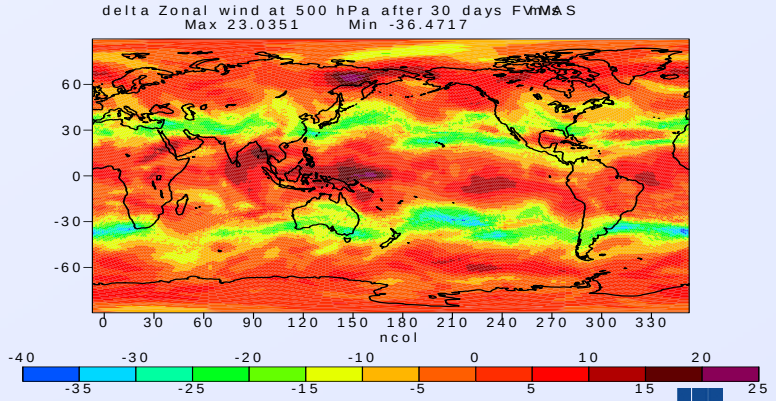

Baroclinic wave test (Jablonowski and Williamson)

- **Initial velocity field is zonally symmetric and contains mid-latitude northern and southern hemisphere jets**
- **Apply zonal wind perturbation**
- **Evaluate using 40962-cell grid (nominally 1-deg)**
- **Compare with literature**
- **Compare CAM/MPAS with MPAS-standalone**

Results of baroclinic wave test (10 days)

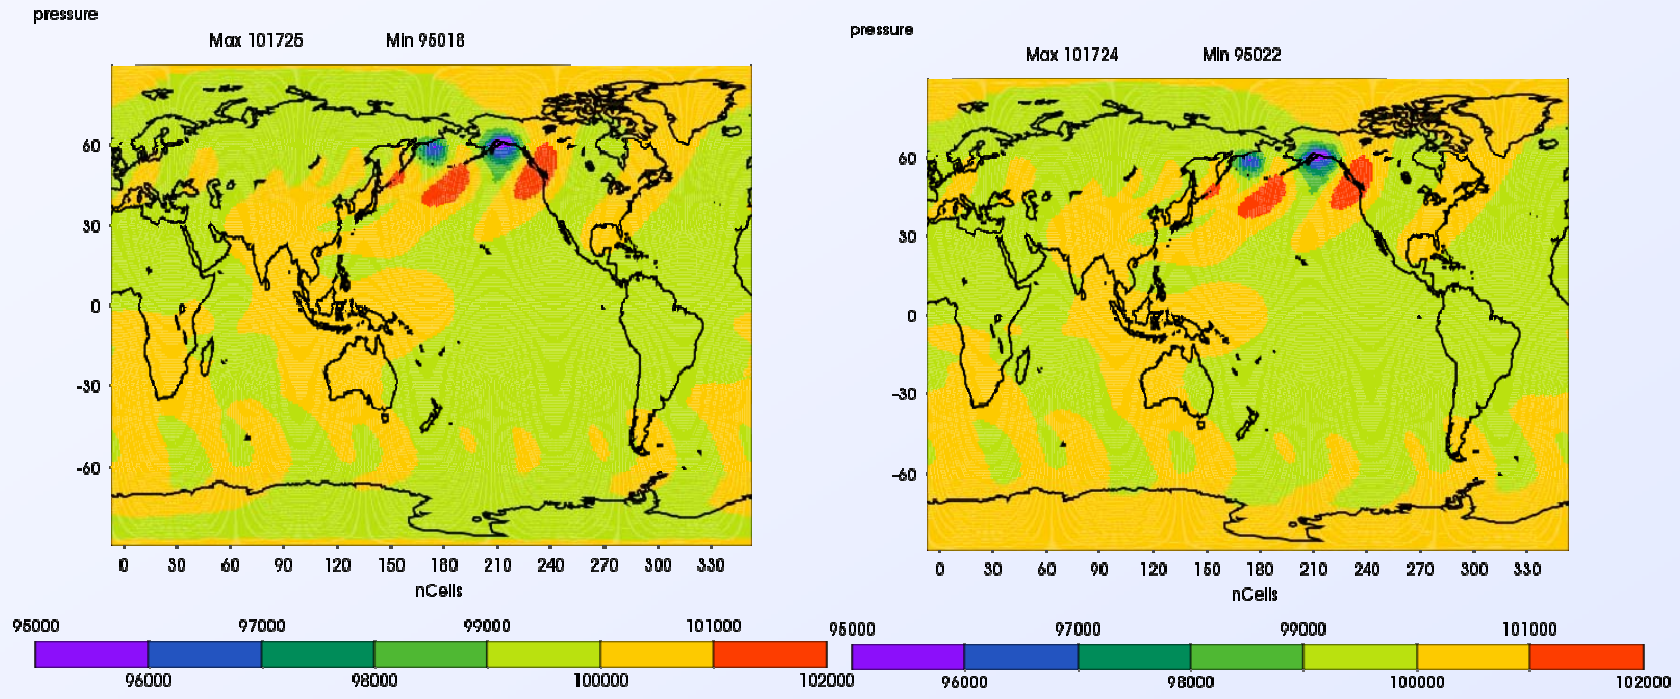

CAM/MPAS

MPAS-standalone

Current activities and future plans

- **Using newly re-written *Scrip* to create mapping files to convert history output to latitude-longitude grid**
  - **will enable utilization of AMWG diagnostics package**
  - **will enable operation with land model**
- **Initiating comprehensive aquaplanet comparison with other dycores**
- **Carrying out performance evaluation across resolutions and process counts**
- **Plan to carry out cases with land model**
- **Plan to investigate simulation using locally refined grid over key geographic region**