Experiments

- **Held-Suarez**: Simplified physics, dry-dynamics
- **Doubly Periodic**: Full-model physics, no-land, no-ice, prescribed SST's
- **Aquaplanet**: Full-model physics, no-land, no-ice, prescribed zonally symmetric SST's
- **AMIP** (Atmospheric Model Intercomparison Project)**:** Full-model physics, land, prescribed SST's
- **Coupled:** Atmospheric (full-physics, chemistry, …, land, ocean, ice models…

Climate Studies

- **Purpose of study**: Provide fundamental understanding of role of clouds in dynamics of climate
- **Previous AMIP studies**:

  100km, 50km, 25km, and 12km resolutions
- **Ongoing AMIP studies**:

  3.5km resolution, 41 vertical Levels
- **Proposed Coupled studies**:

  3.5km atmosphere and 25km ocean

Cubed-Sphere Atmospheric Dynamics

- Fully compressible non-hydrostatic equations
- Mapping of cube to sphere
- Ideal [x,y] decomposition
- Finite-volume advection
- Grid nesting and stretching

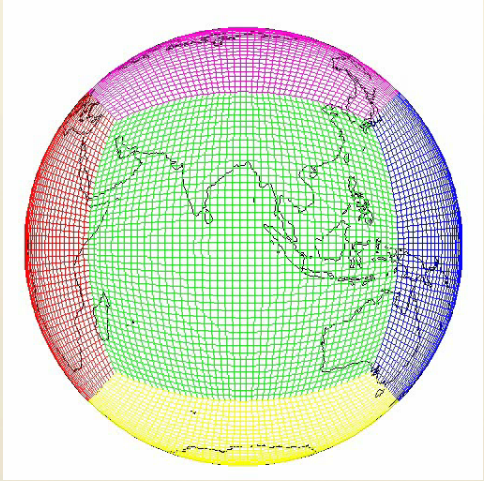

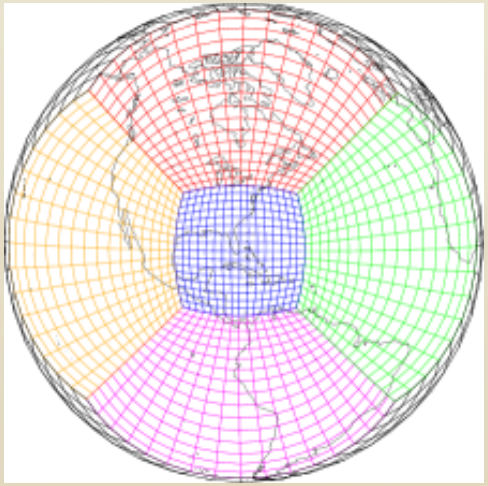

Status of Climate Experiments

- Experiments begun on May 1, 2005; run for several model months
- Tune model: **moisture** (precipitation), **heat** (outgoing longwave radiation), …, tropical cyclones
- Modified horizontal and vertical diffusion and micro-physics schemes

Outgoing Longwave Radiation: August 17, 2005

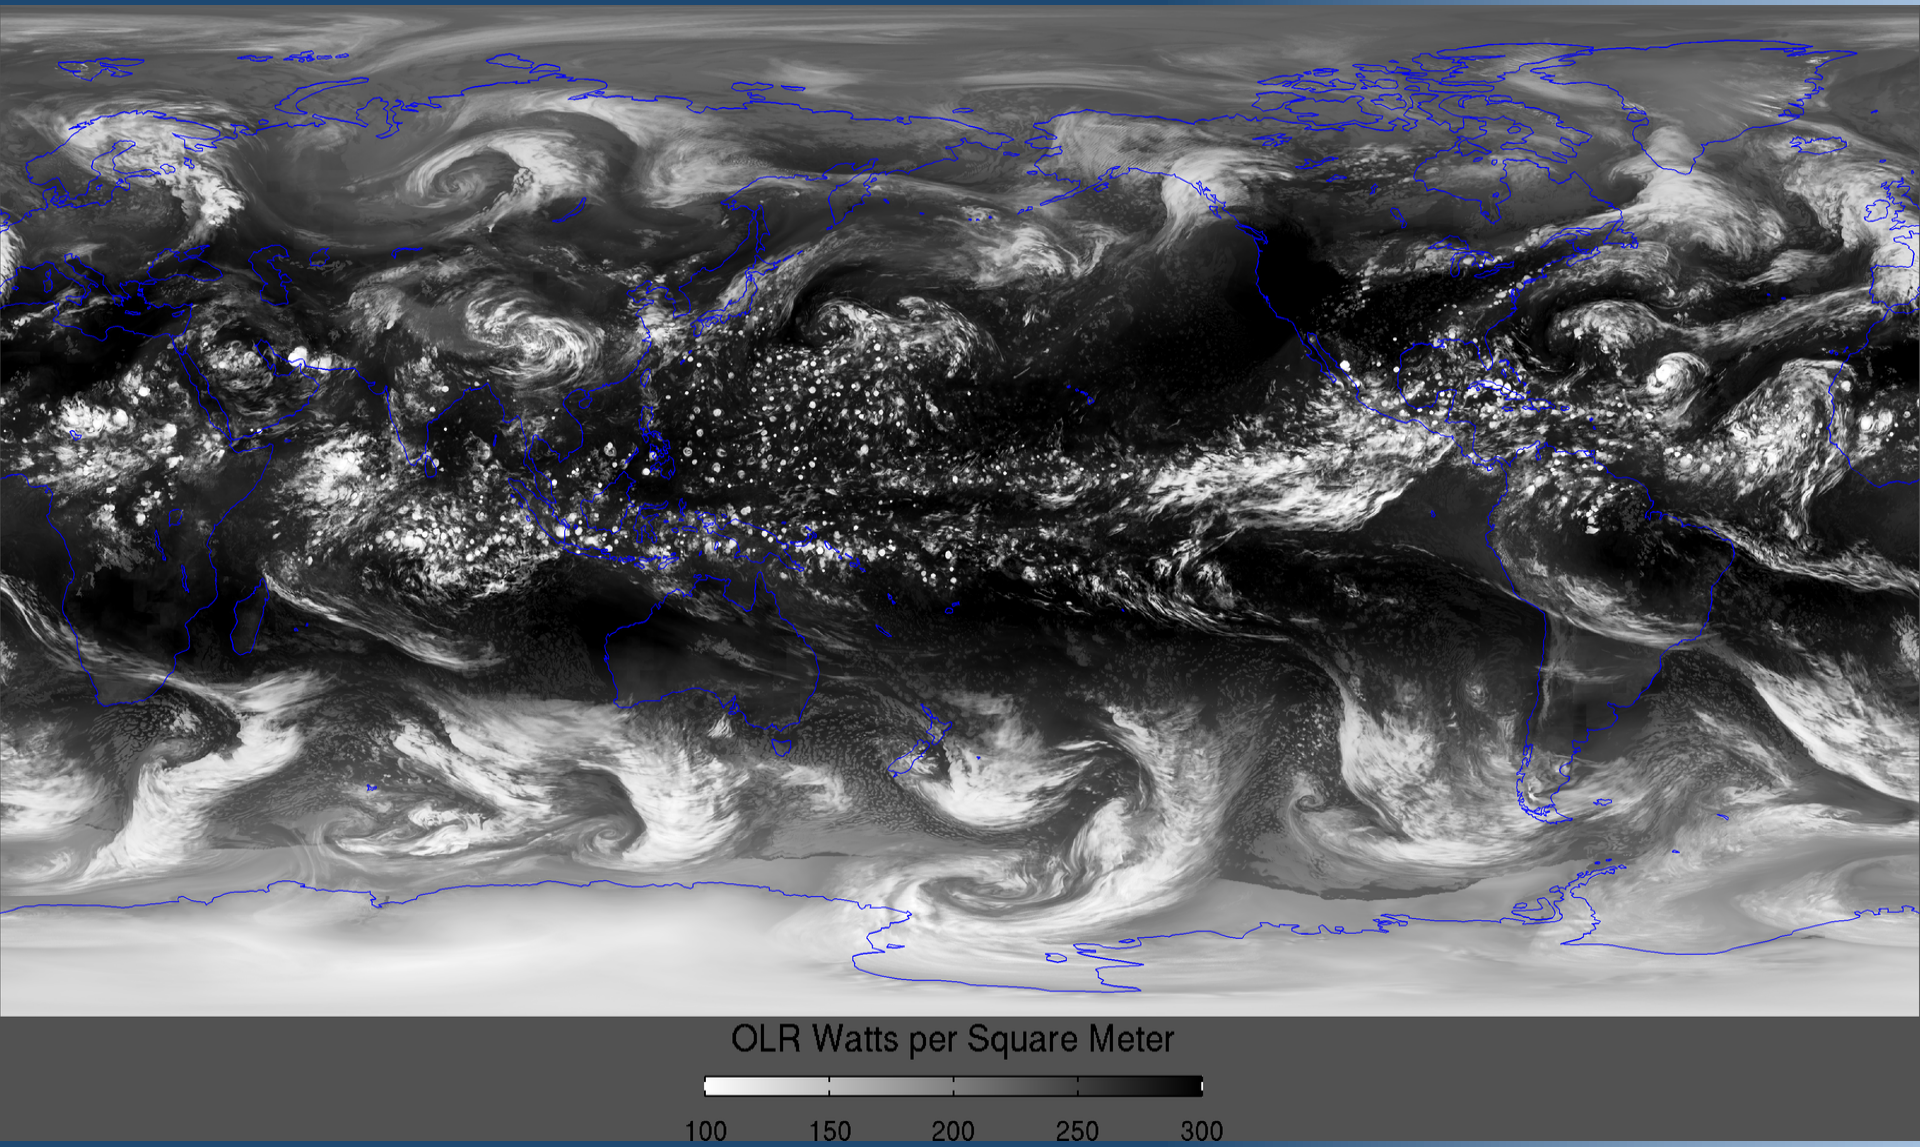

Computational Infrastructure

- Code written by GFDL scientists in FORTRAN 90 with MPI, and NetCDF
- 0.75M lines of executable code
- Software improvements over last 5-years: BGP and BGQ
  - OpenMP at high-level: atmosphere, land, …
  - Eliminate global domain arrays
  - Improve communication and I/O scaling

Computational Performance

- 3.5km executes on BGQ on 15,360 MPI ranks: (8-MPI ranks/node and 8-OpenMP threads/rank)
- 20 model days per wall clock day
- 8-times improvement per node BG/P to BG/Q
- Generates 250GB history files/model day
- History files downloaded and post-processed at GFDL
- 0.25M lines of post-processing scripts

Computational Performance

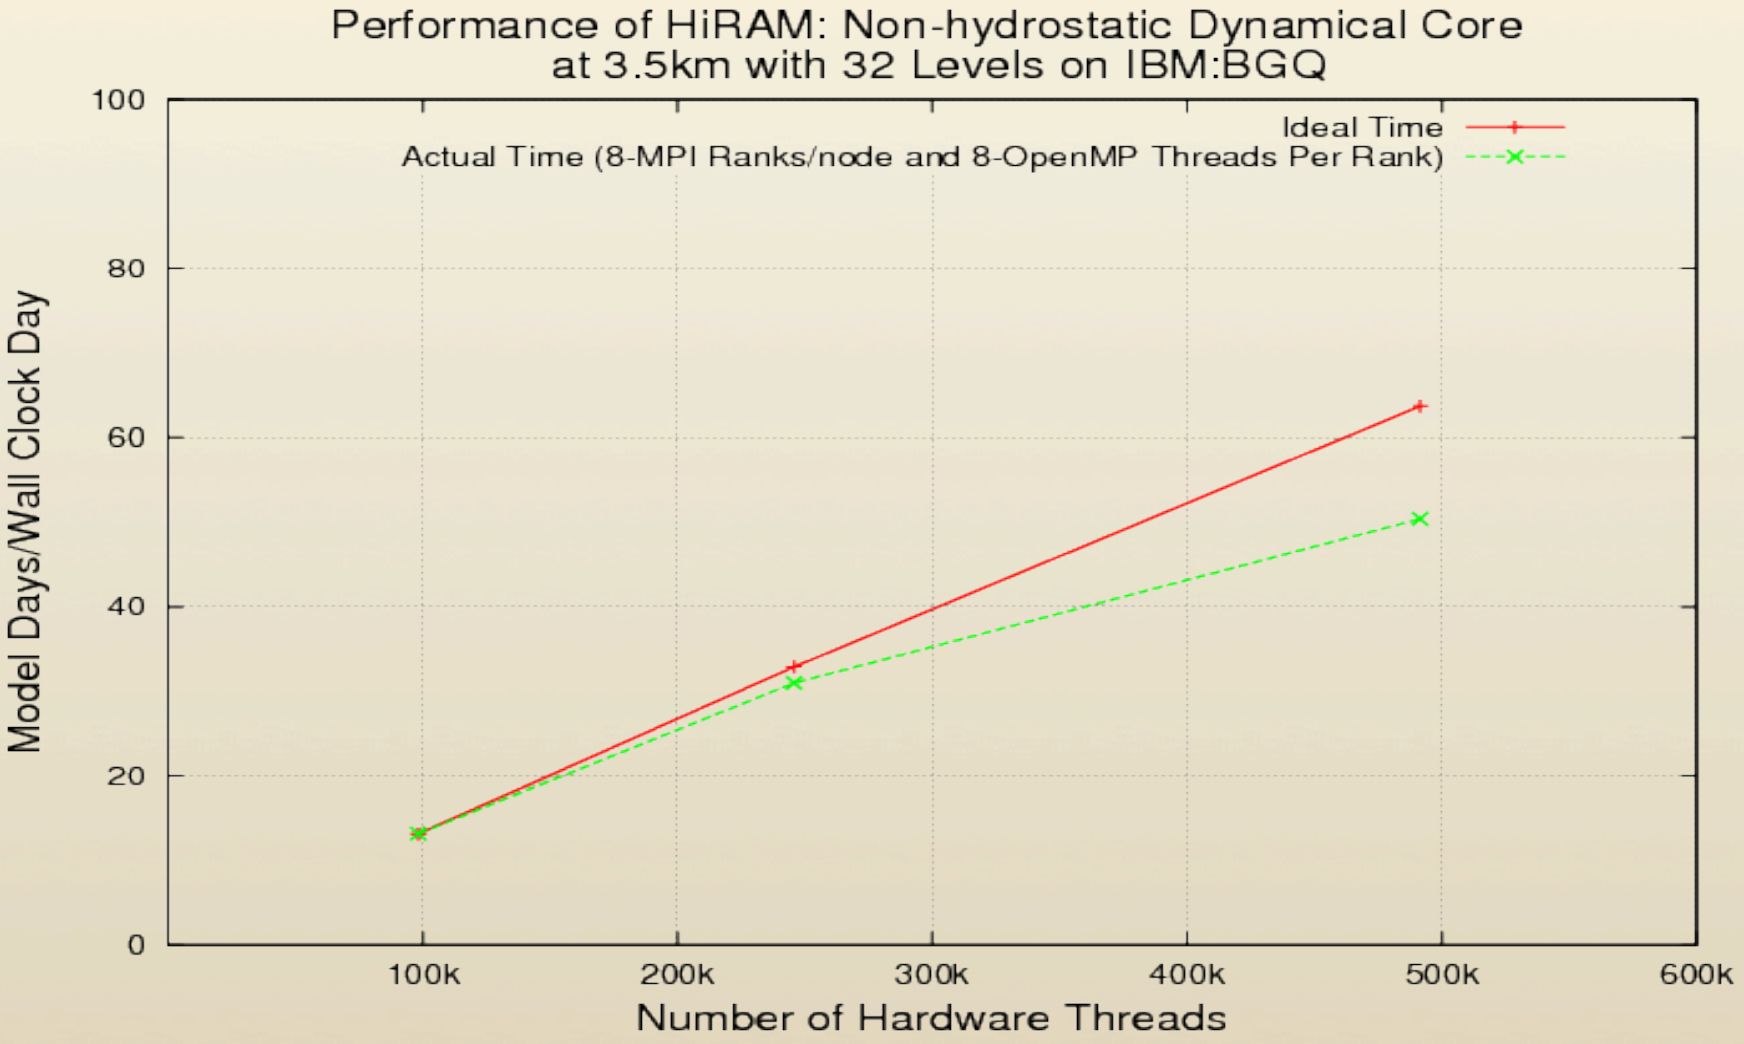

Computational Challenges

- 10-year development cycle
- Infrastructure: (land-sea,.. masks, initial conditions)
- Issues with performing climate science at these scales:
  - Significant development time (2-day segments) months to equilibrate model
  - Dedicated resources needed for climate experiments
- Storage: 250GB/day

Computational Findings

- Improve performance on BG/Q to guide development
- Standalone versions of: atmosphere, physics, land, coupler, I/O, communications,… for performance testing
- Code now scales to 1M hardware threads
- Initialization was significant issue now takes 10% of time

Computational Findings

- Model timer counters: Atmosphere=75%, Physics=15%, land and coupler=10%. (85% compute and 15% communication)
- HPM metrics, general characteristics:

  40% Floating-point (FPX) and 60% integer (FXU). 10B/ cycle (not bandwidth limited)
- Line and function-level timing data to identify:
  - Loop-level-vectorization
  - Regions for additional OpenMP threading
  - Data copies across subroutine boundaries

Computational Findings

- MPI-profile data used to improve: MPI and I/O performance
- Increase parallelism: blocking (decomposing in “i,j,k” and “i,j”) and task (separating radiation from other physics-packages)
- Deal with storage/post-processing: 2PB of data per scenario (20-years)

Computational Directions

- How should we proceed with infrastructure development?
  - Language: FORTRAN-90
  - Parallelism: Directives (OpenMP, OpenACC)
- Future compute platforms
- Moved to different compute platforms: Titan and Stampede

Climate and Computational Challenges

- Coupled model with 3.5km atmosphere and 25km ocean started experiments
- Compute requirement: **3 model-years per wall clock day. Getting 20 model-days per wall clock day**
- **Fifty-times** improvement to achieve required throughput

Questions?

chris.kerr@noaa.gov