U.S. NAVAL RESEARCH LABORATORY

Development and Testing of a Next Generation Spectral Element Model for the US Navy

Alex Reinecke[1], Kevin Viner[1], James Doyle[1], Sasa Gabersek[1], Matus Martini[2], John Mickalakes[3], Dave Ryglicki[4], Dave Flagg[1], and Francis Giraldo[5]

1 U.S. Naval Research Laboratory, Monterey, USA
2 DEVINE Consulting, Monterey, USA
3 UCAR, Boulder, CO
4 National Research Council, Monterey, USA
5 Naval Postgraduate School, Monterey, USA

17th Workshop on High Performance Computing
24-28 October 2016
European Center for Medium Range Weather Forecasting
Reading, UK

U.S. NAVAL RESEARCH LABORATORY

Overview

- Brief overview of the DoD HPC resources
- Overview of the NEPTUNE system as the US Navy's next generation modeling system
- Computational and I/O scaling
- Summary of work on physic package implementation and optimizations

DoD HPC Resources

System	Vendor / Model	Compute Nodes			Peak TFLOPS
		Core Type	Core Speed[1]	# of Cores	
Shepard	Cray XC30	Intel Xeon E5-2697v2	2.7	28,632	817
Armstrong	Cray XC30	Intel Xeon E5-2697v2	2.7	29,160	786
Thunder	SGI ICE X	Intel E5-2699v3	2.3	125,888	6
Topaz	SGI ICE X	Intel Xeon E5-2699v3 Haswell	2.3	124,416	5
Haise	IBM iDataPlex	Intel Xeon Sandy Bridge	2.6	18,880	435
Kilrain	IBM iDataPlex	Intel Xeon Sandy Bridge	2.6	18,880	435
Excalibur	Cray XC40	Intel Xeon E5-2698 v3	2.3	101,312	4
Riptide	IBM iDataPlex 1350	Intel Xeon Sandy Bridge	2.6	12,096	252
Conrad	Cray XC40	Intel Xeon E5-2698v3	2.3	49,248	2
Gordon	Cray XC40	Intel Xeon E5-2698v3	2.3	49,248	2
Spirit	SGI ICE X	Intel Xeon Sandy Bridge	2.6	73,440	1500
Garnet	Cray XE6	AMD Interlagos Opteron	2.5	150,912	1500
Copper	Cray XE6m	AMD Interlagos Opteron	2.3	14,976	135
Lightning	Cray XC30	Intel Xeon E5-2697v2	2.7	56,880	1281
Predator	SGI UV	Intel E5 Sandy Bridge	2.7	1,004	---

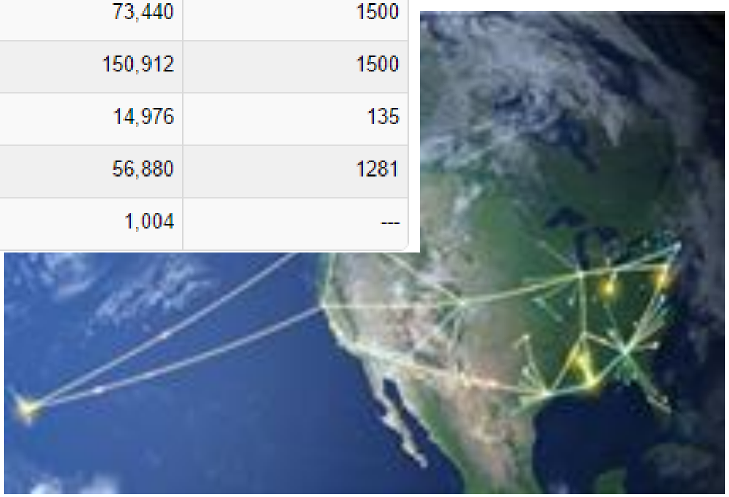

- 15 machines spread over 6 HPC centers
- Total 23 PetaFLOPS of peak computing
- Over 50 Petabytes of storage

U.S. NAVAL RESEARCH LABORATORY

NRL Marine Meteorology Division and Operational Partner FNMOC

- **Weather and seasonal climate**
  - Goal: 5 km deterministic global horizontal resolution
  - Transition ~2023
- **Global Forecasting:**
  - NAVGEM, SL-SI spectral transform
    - 4 times daily, currently 31km resolution
- **Limited Area (Mesoscale) Forecasting:**
  - COAMPS®, split-explicit FD
    - Over 70 unclassified COAMPS domains run daily down to 500m resolution
    - Domains are always in flux
- **Why invest in a new system?**
  - Neither system is positioned to exploit next generation architectures
  - Both systems are low order and non-conservative (though skillful nonetheless)
  - We are not a large laboratory, roughly 100 scientists, less than half working on model development

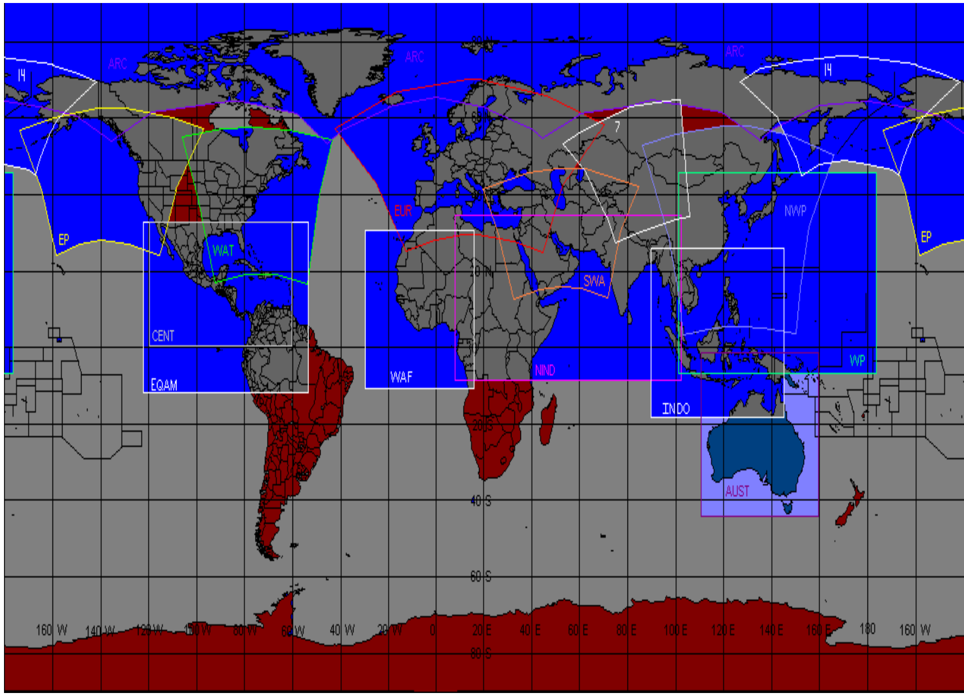

Next Generation NWP

- **Our goal is to develop next generation model, with capabilities:**
  - **Multi-scale non-hydrostatic dynamics (global-scale to mesoscale)**
  - **Highly scalable, flexible, accurate, and conservative numerical methods**
  - **Structured, unstructured or adaptive grids; scale-aware parameterizations**
- **A spectral element (SE) technique is the selected numerical method**
  - **The numerical solution is represented by a local polynomial expansion**
  - **Very small communication footprint implies excellent computational scaling**
- ***NEPTUNE* is the candidate with these attributes**

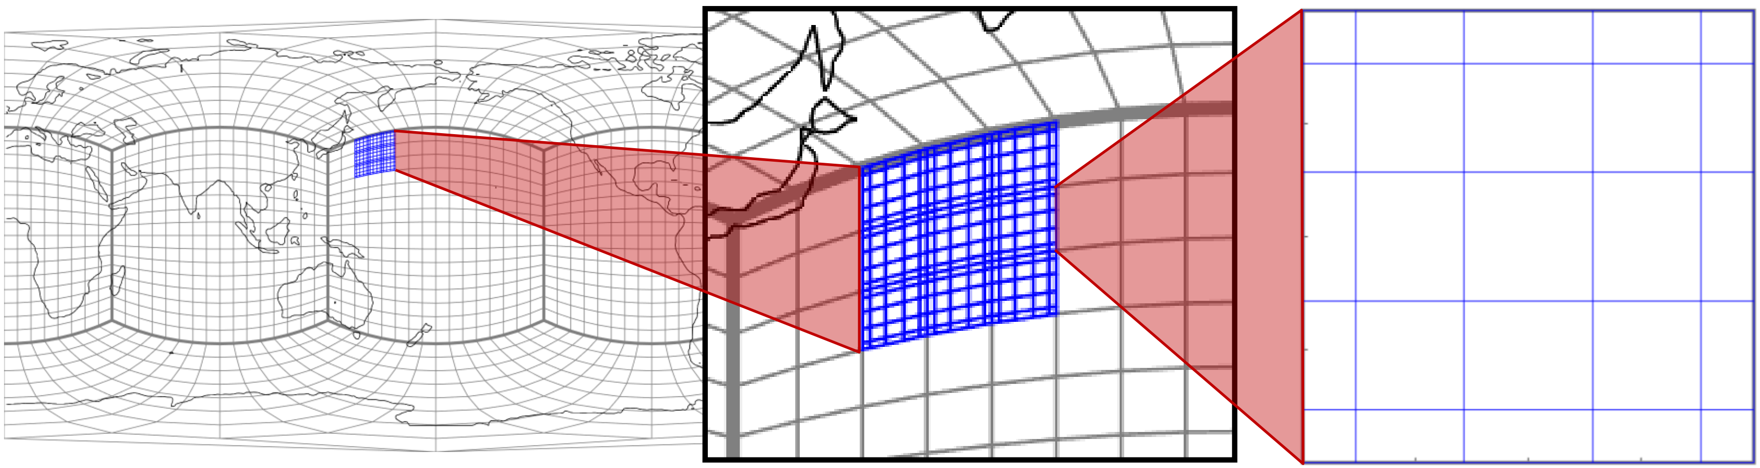

[1]***NEPTUNE*****:** *Navy Environmental Prediction sysTem Utilizing the NUMA[2] corE*

[2]***NUMA**:* *Nonhydrostatic Unified Model of the Atmosphere (F. Giraldo NPS)*

Why Spectral Elements?

- Small communication footprint
  - Only to transfer 1 row of data, regardless of numerical order
  - Excellent scaling potential
- Good potential for dense floating point computation
  - Projects well onto new architectures
- Downsides:
  - Physics implementation on irregular gird – doesn’t seem to be an issue
  - Tracer transport not ideal – new positive definite schemes

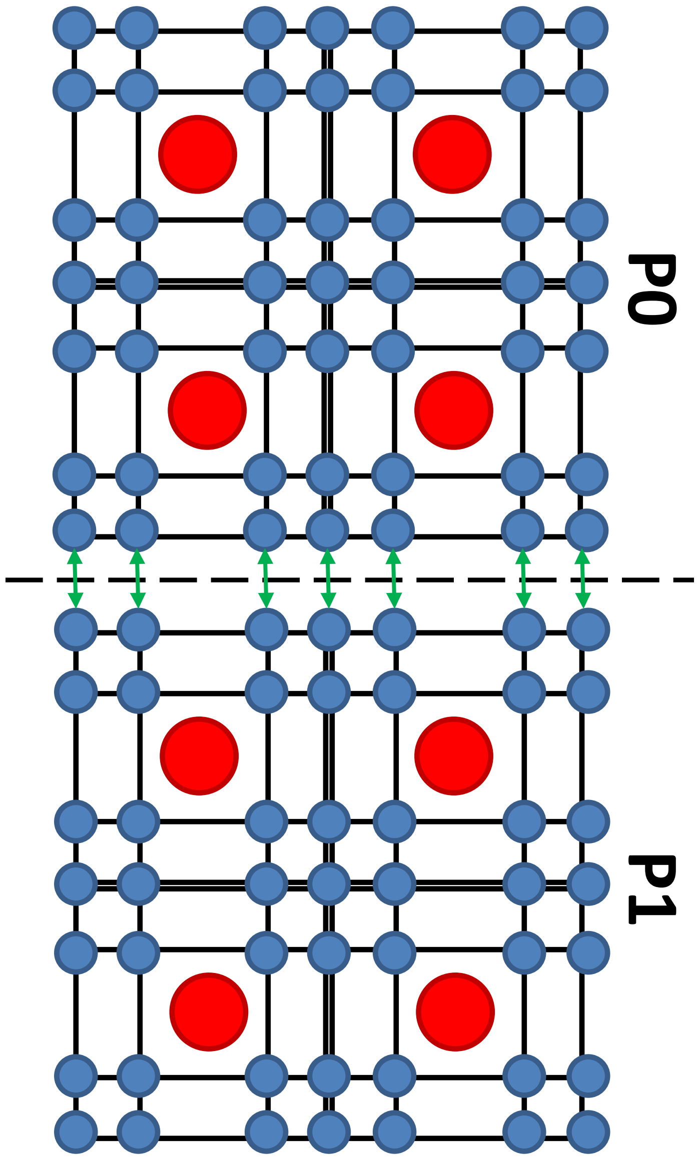

NEPTUNE – Overview and Layout

NEPTUNE is the US NAVY's next generation NWP system

NUMA is the dynamical core of NEPTUNE

NUMA uses a **three-dimensional** spectral element technique with a sphere centered Cartesian coordinate system on the cubed sphere

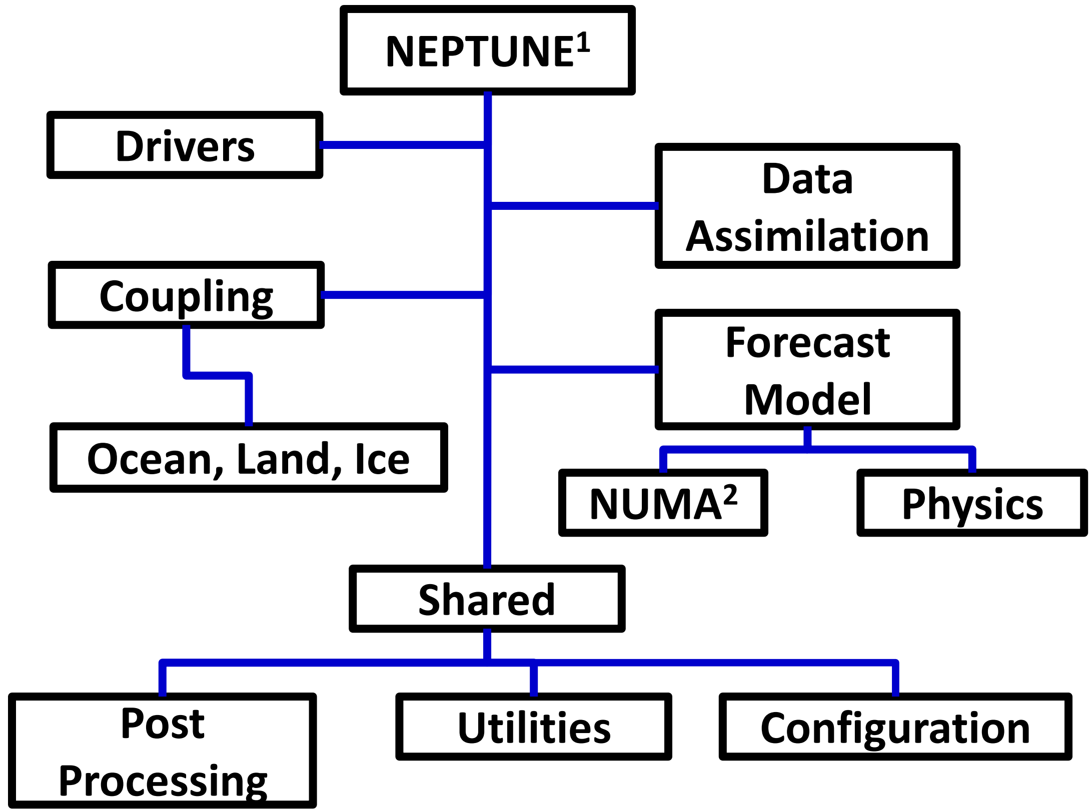

[1]***NEPTUNE:*** *Navy Environmental Prediction sysTem Utilizing the NUMA[2] corE*
[2]***NUMA:*** *Nonhydrostatic Unified Model of the Atmosphere (Giraldo et. al. 2013)*

Dynamical Core Overview

- Deep atmosphere equations
- Non-hydrostatic
- Unified 3D spectral element (CG/DG) dynamical core
- Q-array stored in (k,i) order, element boundary points share memory
- Cartesian coordinate system for limited area
- Sphere centered Cartesian coordinate system for global prediction
- Inexact integration (mass lumping)
- 1D/3D IMEX time integrators and split-explicit RK3 integration for efficiency

numa

Non-hydrostatic Unified Model of the Atmosphere

Kelly and Giraldo 2012 JCP, Giraldo et. al J. Sci. Comput. 2013, Kopera and Giraldo 2014 JCP

Global Model Applications

<Δx> = 5km global average horizontal resolution

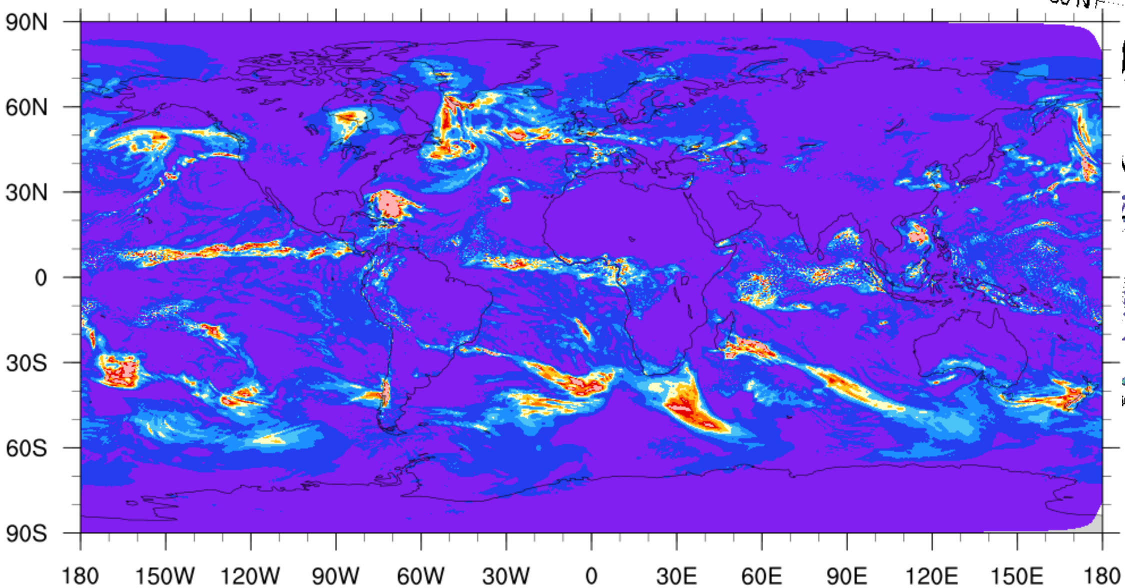

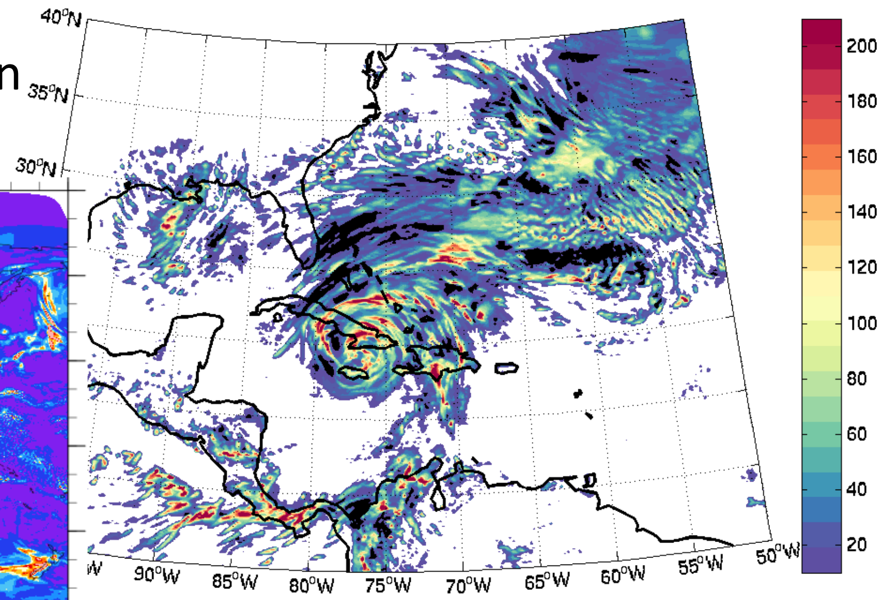

- Cold start with GFS initial conditions for Hurricane Sandy simulation
- Realistic rain-band structure in early stages of Sandy, good position forecast at later stages
- Correct location of precipitation in ITCZ and mid-latitude storm belt

Limited Area Model Application

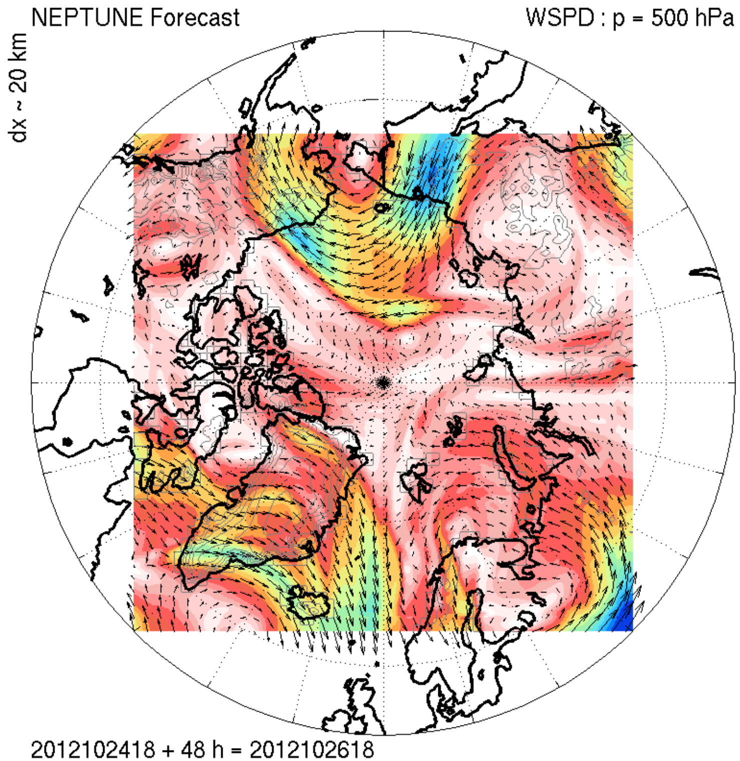

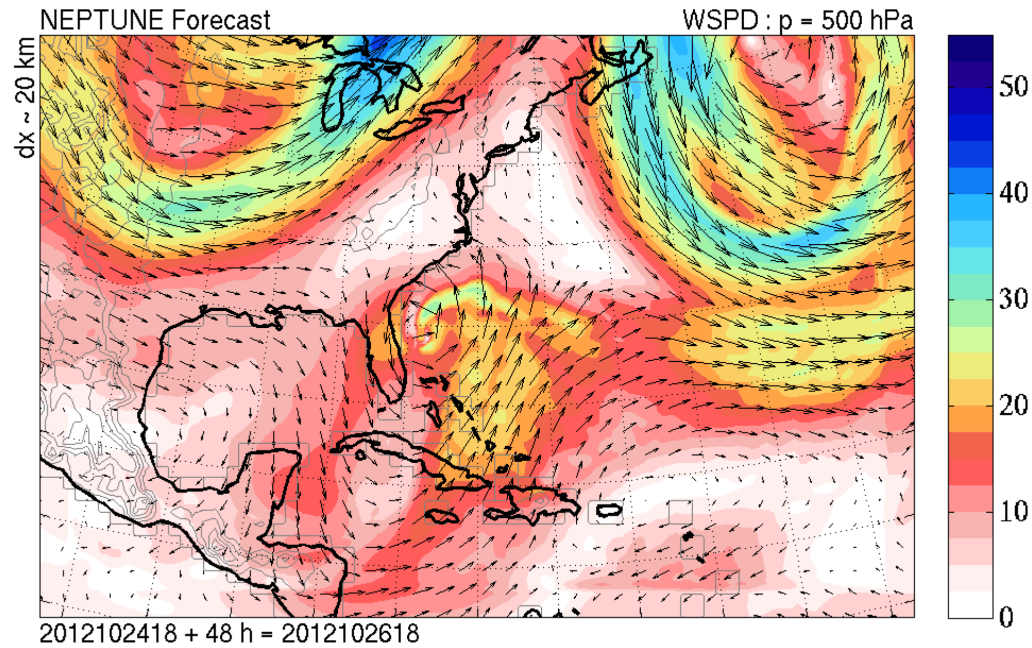

- Developed a limited area version of NEPTUNE with set of projections
  - Lambert Conformal, Mercator, Stereographic
  - Solution prescribed at lateral boundaries using Davies BC's

NEPTUNE Grid Generation

- Cubed sphere grid with sphere centered Cartesian coordinate system
  – P4est for grid generation and MPI decomposition
  – Parallel and scalable, can build large grids on the fly
- Limited area domain structured as a cube, called “the cube” domain
  – Map projections are built into the Jacobian
  – P4est used for parallel grid generation

U.S. NAVAL RESEARCH LABORATORY

Cubed Sphere Grid

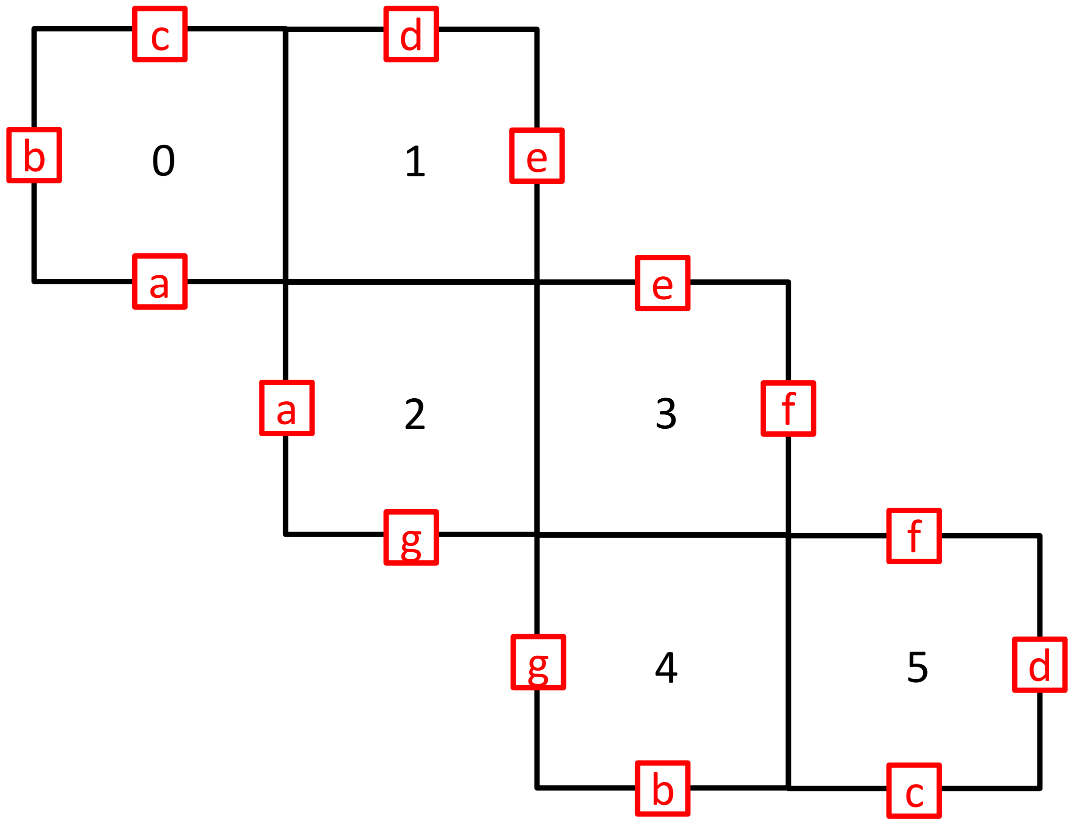

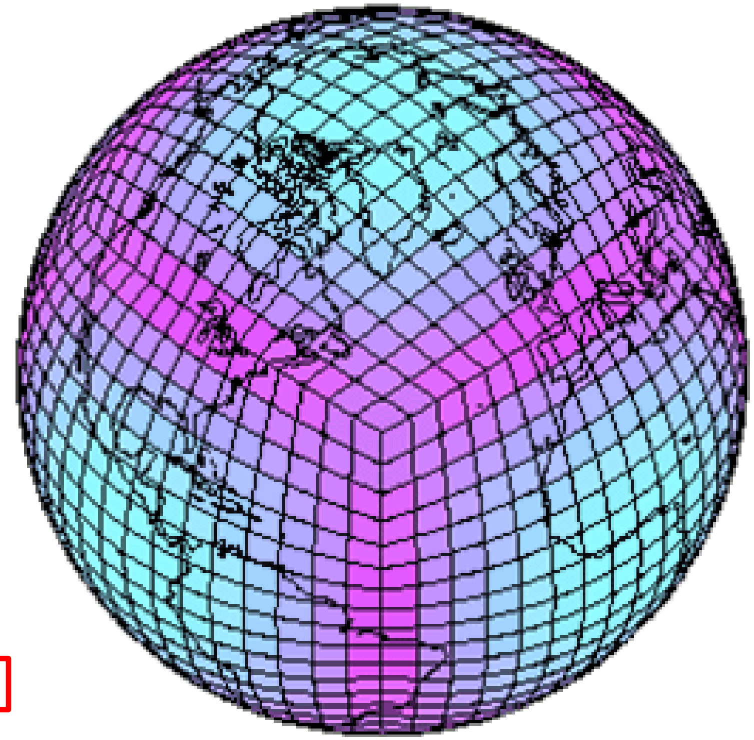

- Traditional cube sphere (i.e. FV3, HOMME) stitch 6 cubes together
  - Matching conditions along edges, requires conditional logic for edges
- NEPTUNE uses a Cartesian, sphere centered coordinate system
  - No matching condition needed, no conditionals for edge conditions

Sphere Centered Coordinate System

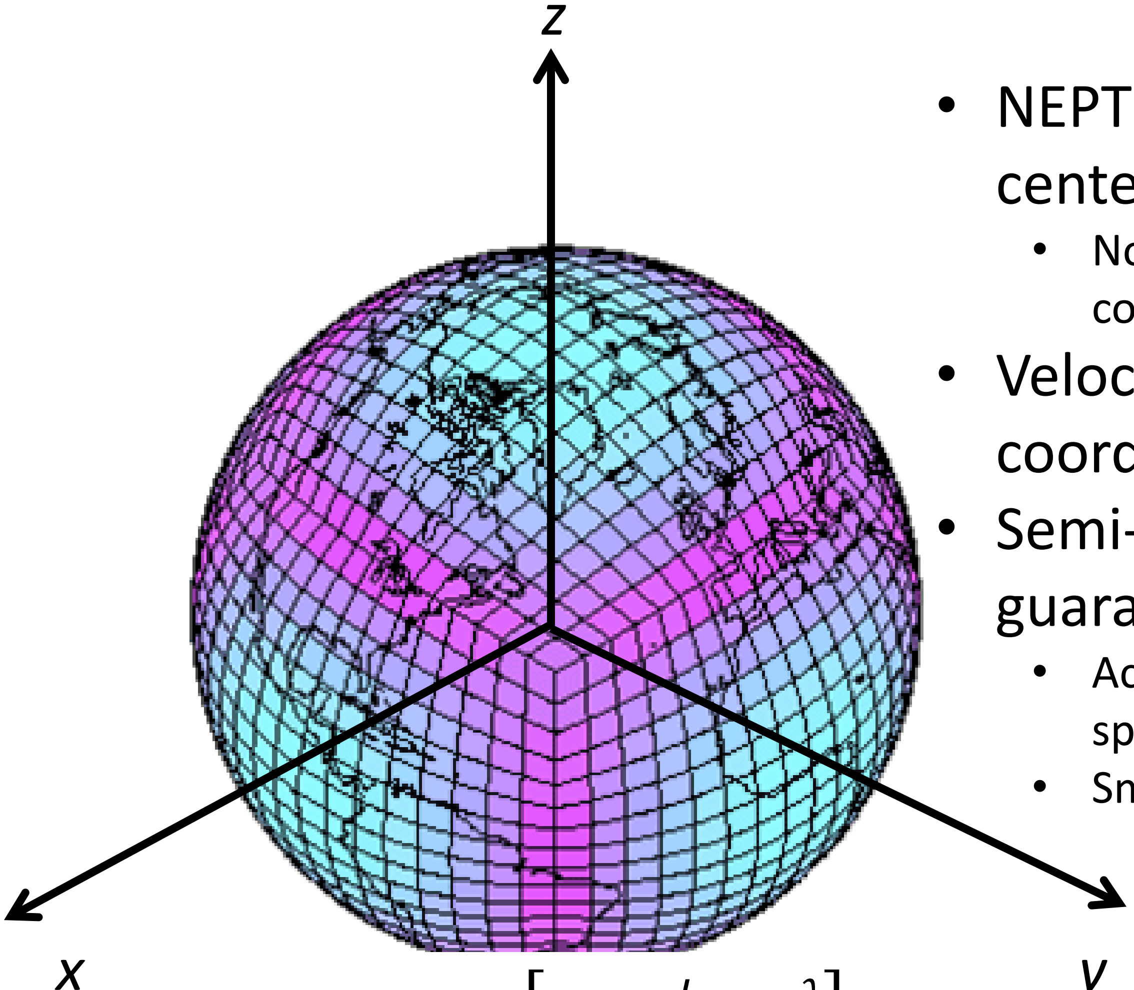

$$\vec{x} = \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} r\cos\phi\cos\lambda \\ r\cos\phi\sin\lambda \\ r\sin\phi \end{bmatrix}$$

- NEPTUNE uses a Cartesian, sphere centered coordinate system
  - No matching condition needed, no conditionals for edge conditions
- Velocity is defined relative to the coordinate system
- Semi-analytic Jacobian terms guarantee:
  - Accurate representation of the cubed sphere
  - Smooth transitions across face edges

U.S. NAVAL RESEARCH LABORATORY

Data Layout in NEPTUNE

Q-array layout: beneficial for 1D IMEX and physics because vertical varies the fastest

HOWEVER

There may be significant room for improvement

2D Slice of q array

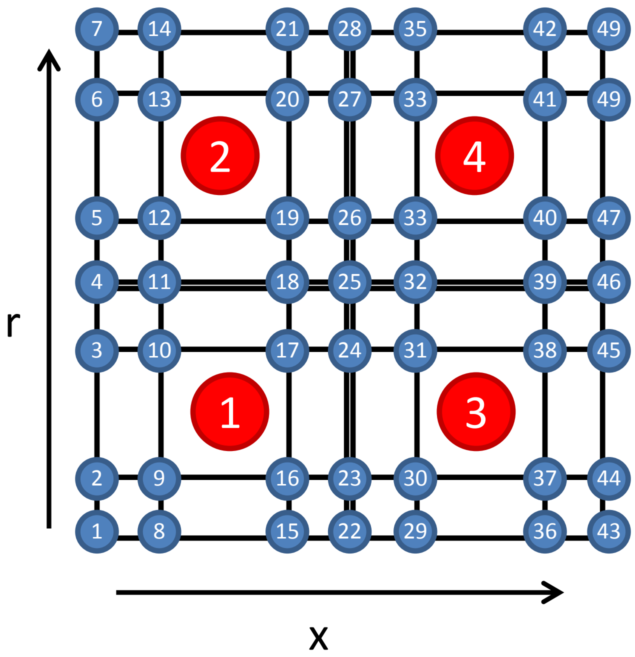

For each SE operation data is reordered to contiguous (i,j,k) order on the element

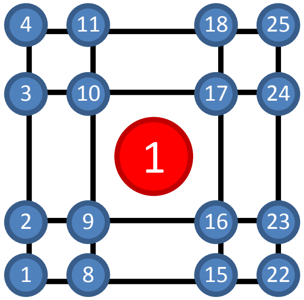

Occurs 1x for calls to create_rhs
Occurs 2x for calls to create_lap
create_lap is called 1x per create_rhs

create_rhs per time step depend on the time integrator:
(at least) 1x per stage of a RK integrator

NEPUTNE Scaling Tests

DOD HPC CAP Program

- *Capabilities and Acceptance Program*
  - *30 days to port and run scaling tests on new systems prior to general use*
  - *Access to 3 systems in June 2015*
  - *CRAY-XC40 and 2 SGI ICE X machines*
- *3D dry dynamical core test, baroclinic instability (Jablonowski 2006)*
  - *12-km fixed horizontal resolution -> increase nodes*
  - *Fixed # of elements per node -> increase nodes*

U.S. NAVAL RESEARCH LABORATORY

Weak Scaling Test

12-km Baroclinic Wave (Jablonowski 2006)

Cray XC-40 (Conrad--NAVO)

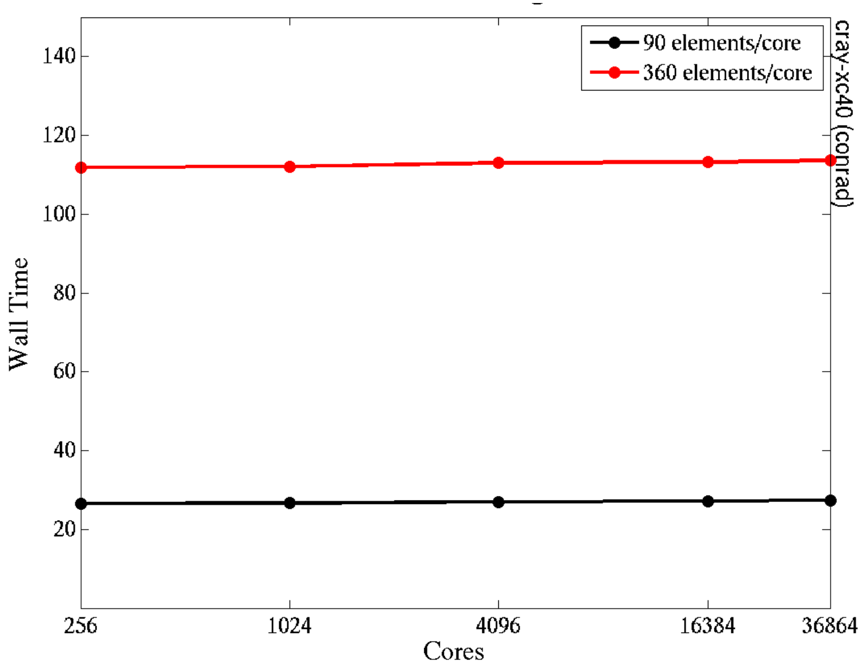

SGI ICE X (Thunder--AFRL)

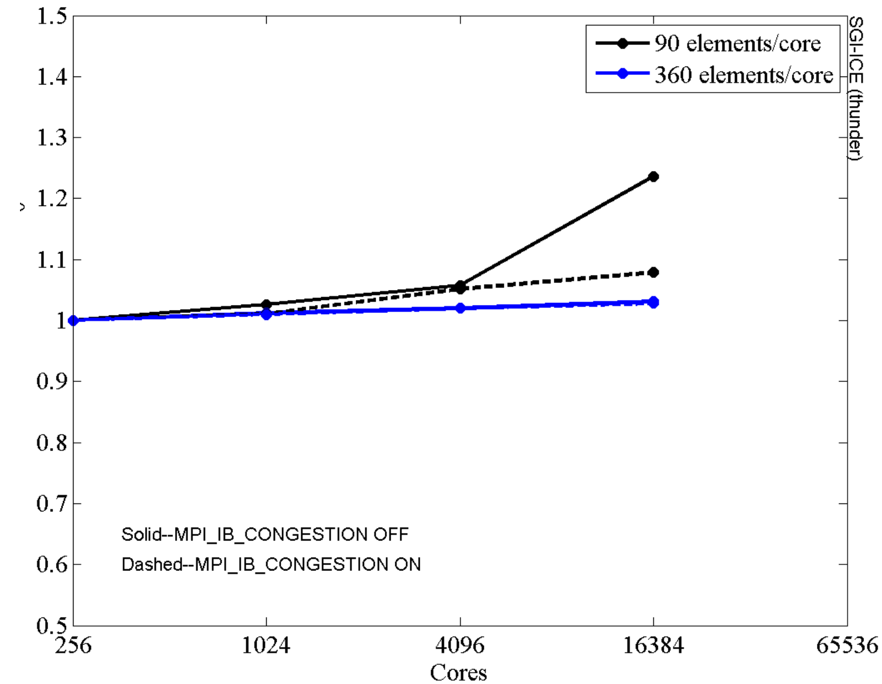

- XC40 – code compiled and ran out of the box
- ICE X required a minor amount of modification to the environment:
  - Set MPI_IB_CONGESTION=1

Strong Scaling Test

12-km Baroclinic Wave (Jablonowski 2006)

Cray XC-40 (Conrad--NAVO)

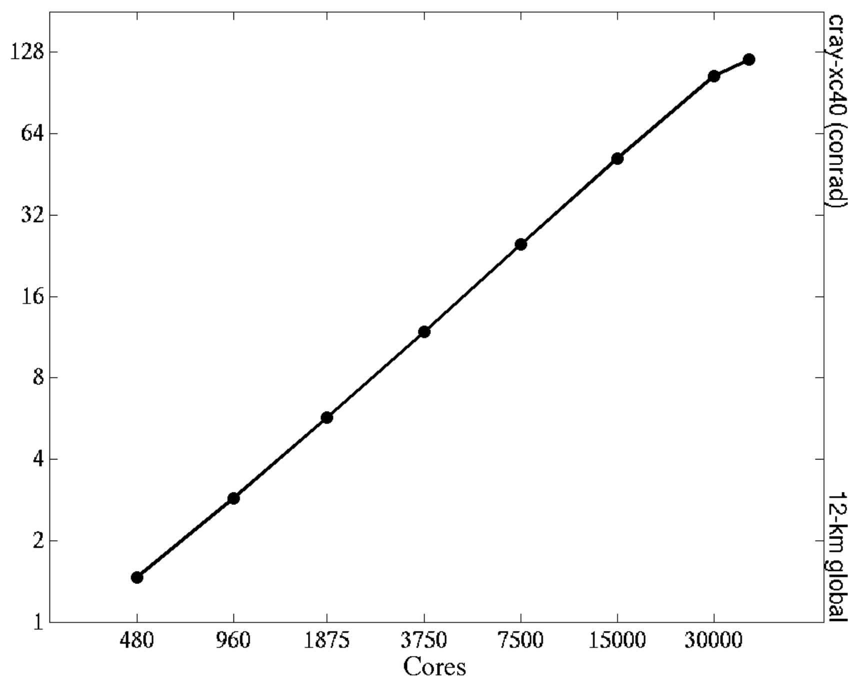

SGI ICE X (Thunder--AFRL)

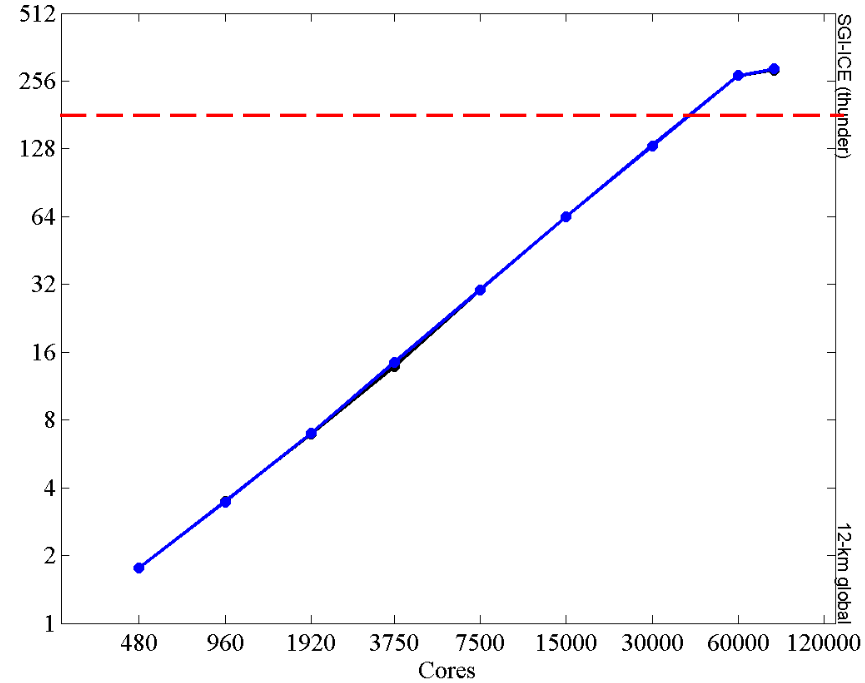

- Near linear scaling on both machines up to 1 column of elements per compute core
- Number of cores to reach operational requirement is high
  - Version of code tested was immature, very little effort to optimize
  - Significant room for improvement

I/O Optimization

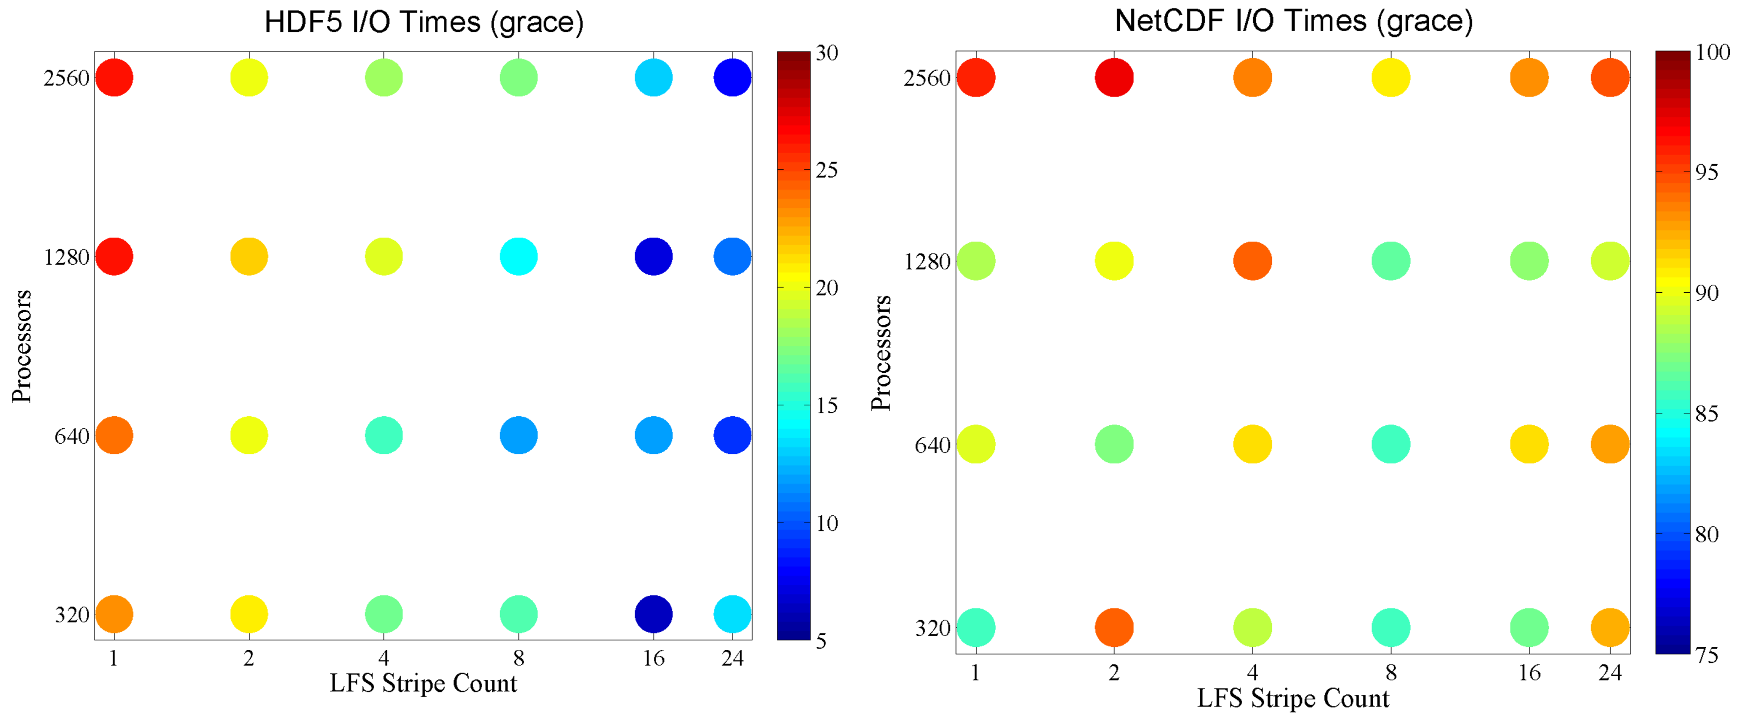

- Test problem is ~30 km global resolution with 32 vertical levels (E80P4L32)
- 14 three-dimensional variables are output 3 times
- Plan to incorporate asynchronous I/O through ESMF interface (leveraging NAVGEM effort)

Physics Implementation

- Physics packages implemented through the NUOPC unified physics driver
  - Data order is switched from (k,i) layout to (i,k) layout
- Special attention is given to the irregular vertical structure associated with vertical spectral elements
  - Testing different mappings to a regular vertical grid
- Initial full package is GFS hydrostatic physics
  - Working on optimizing this package for KNL through collaboration with the University of Utah

U.S. NAVAL RESEARCH LABORATORY

Costal Mesoscale Modeling

Next Generation Modeling -- NEPTUNE

- Development and implementation of physics parametrizations

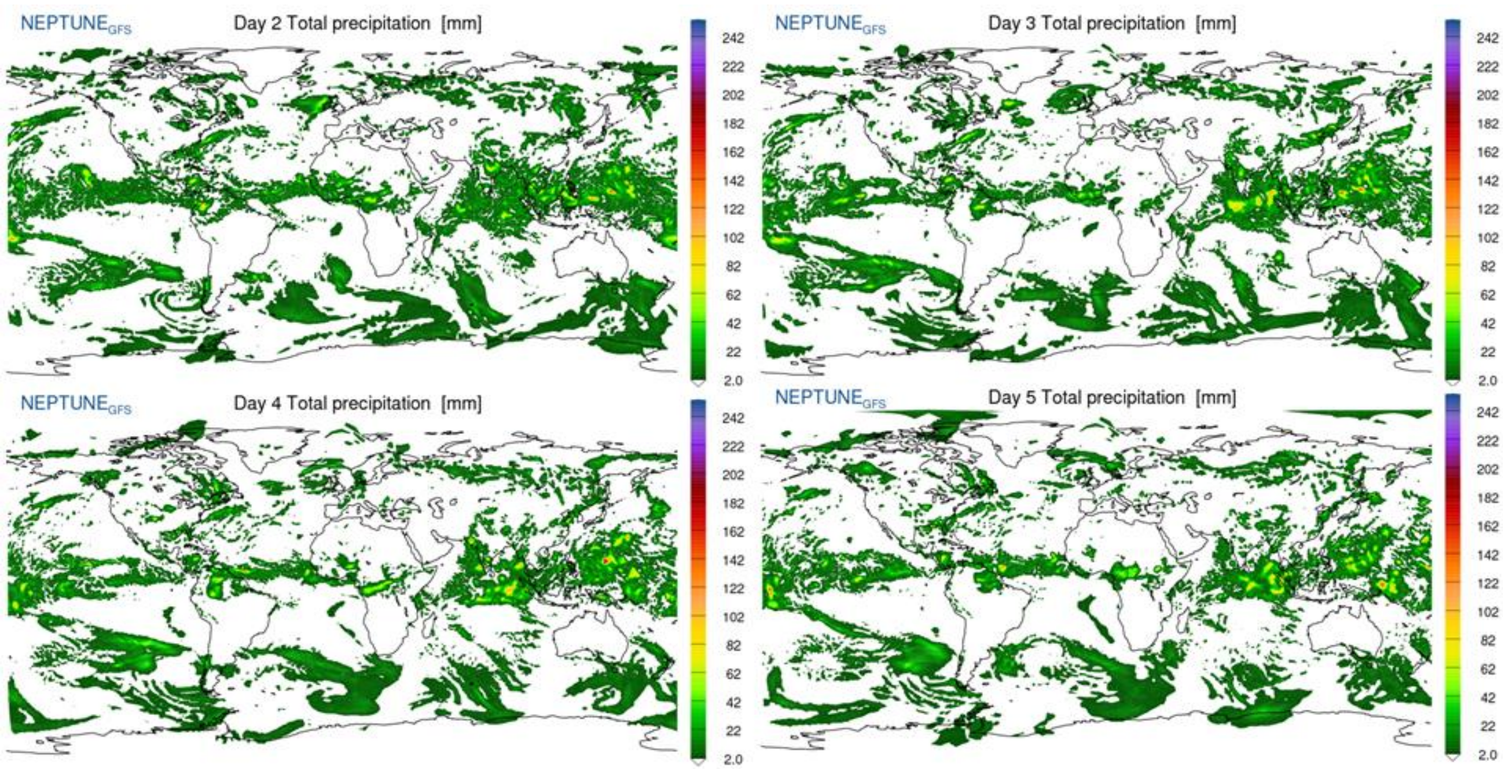

- Goal is to run a year worth of cold start simulations with 15 km resolution

Summary

- NRL is developing NEPTUNE, a three-dimensional spectral element NWP system
  - Unify limited area and global modeling
  - Operational transition in the next decade
- Testing shows excellent strong scaling, however, need to improve overall runtime by optimizing the data layout and incorporated openMP

U.S. NAVAL RESEARCH LABORATORY

Thank-you