Project Overview

- This SciDAC project will develop and test a global cloud resolving model (GCRM), i.e., a global atmospheric circulation model with a grid-cell spacing of approximately 3 km, capable of simulating the circulations associated with large convective clouds. The GCRM developed in this project will make use of a geodesic grid. It will use a non-hydrostatic dynamical core and parameterized cloud microphysics.

- Personnel:
  - Colorado State University - David A. Randall (PI), Joon-Hee Jung, Celal Konor, Ross Heikes, and Charlotte DeMott
  - Pacific Northwest National Laboratory - Karen Schuchardt, George Chin, Annette Koontz, Roger Marchand, Bruce Palmer
  - University of California at Los Angeles - Akio Arakawa (Co-PI)

Science Lesson

- Predicts the discretized 3D vorticity using finite-difference methods. Maintains the non-convergence of the 3D vorticity. Based on anelastic (expandable to pseudo-compressible) equations.

3-D view of the grid

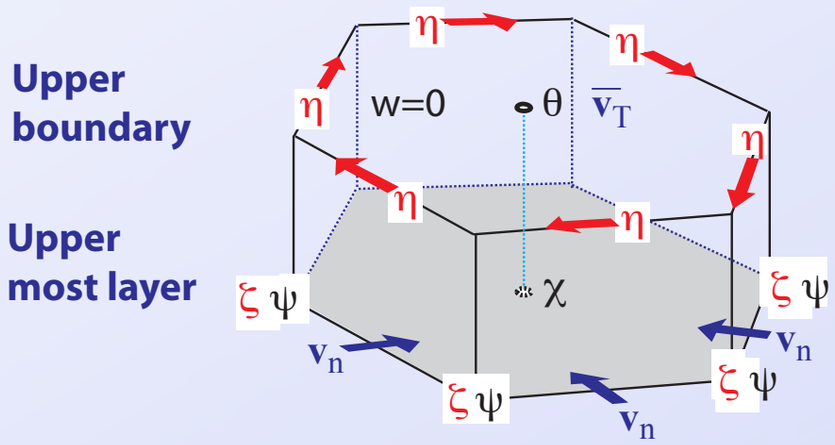

Charney-Phillips vertical grid

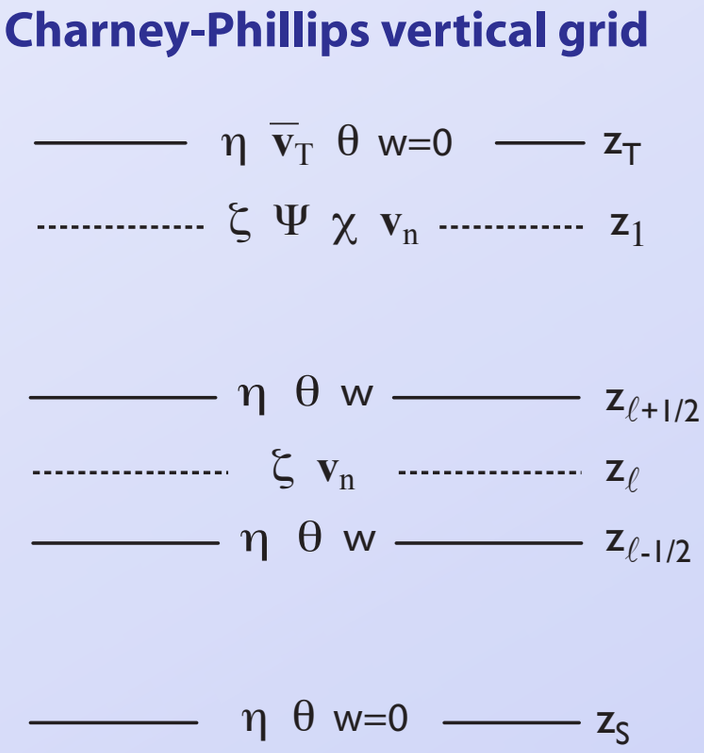

Science Lesson

- The 3D vector vorticity model borrows heavily from the existing barotropic and 3D full-physics hydrostatic model.
  - Geodesic grid structure
  - Data structure
  - Horizontal finite-difference operators, e.g. advection

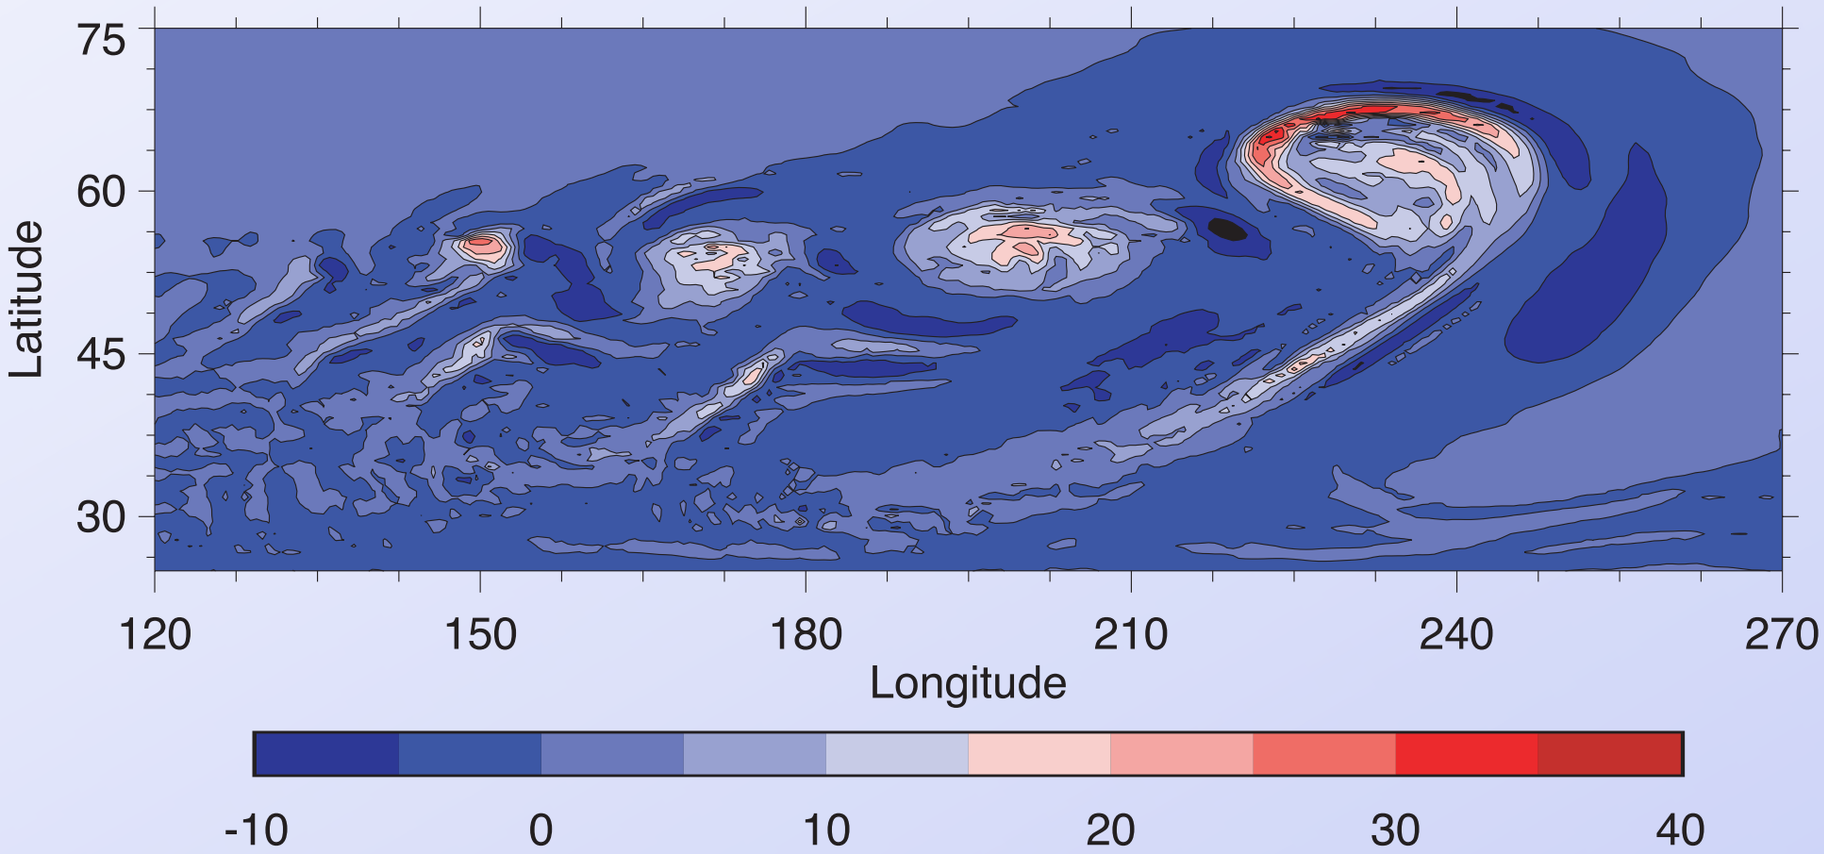

Parallel Programming Model

- The Geodesic Grid

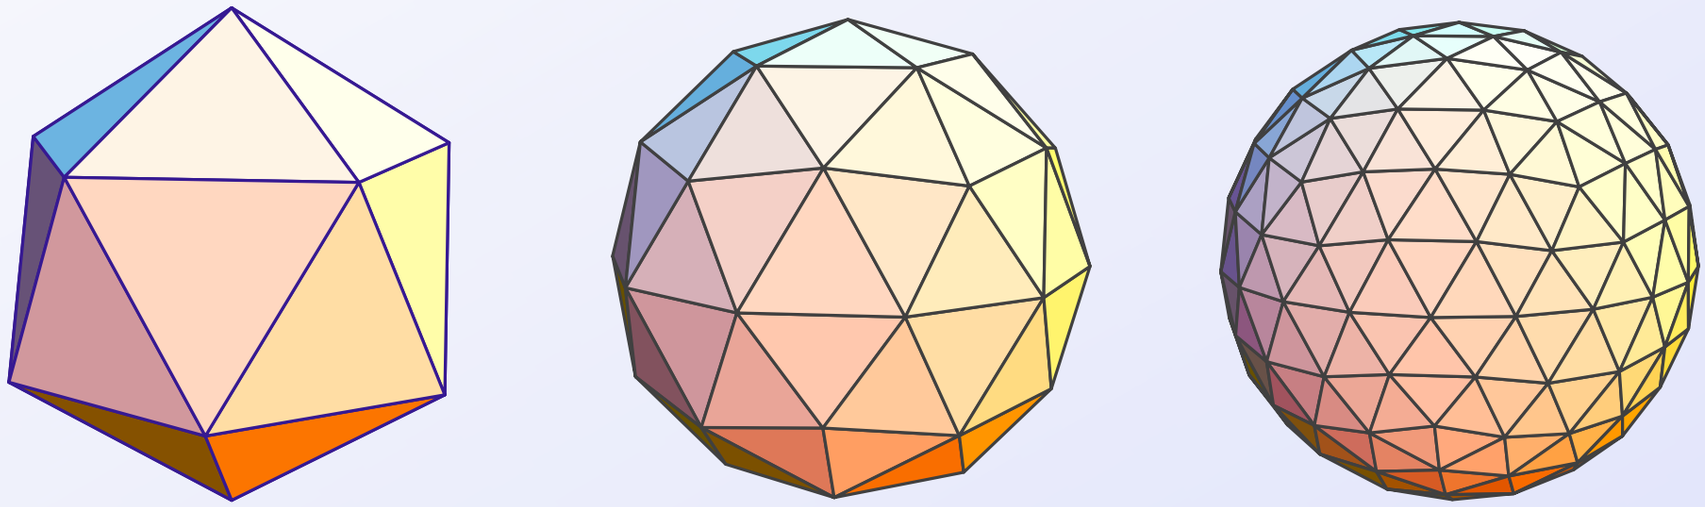

- Associate a Voronoi cell with each grid point

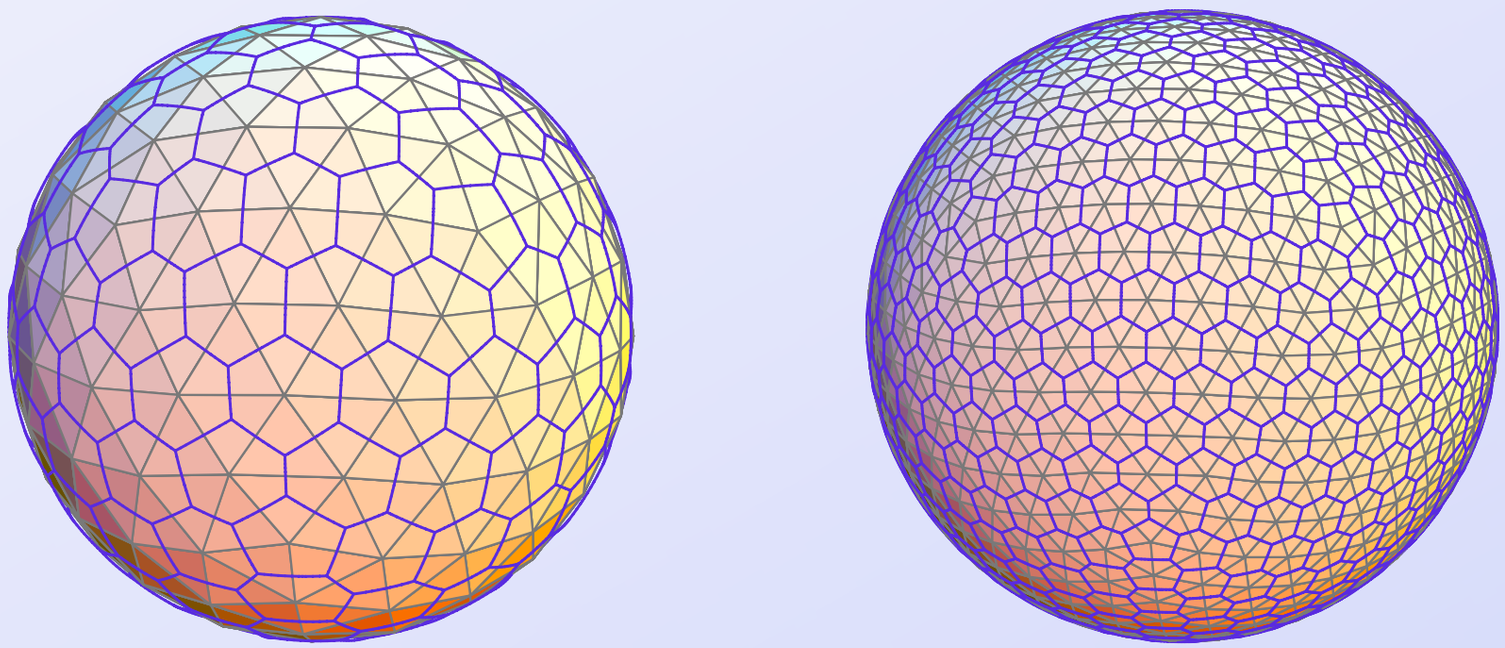

Parallel Programming Model

- Domain decomposition of the sphere
  - Concatenating the faces of the icosahedron allows the sphere to be partitioned into 10 quadrilateral regions
  - Further partition goes 40, 160, 640,...

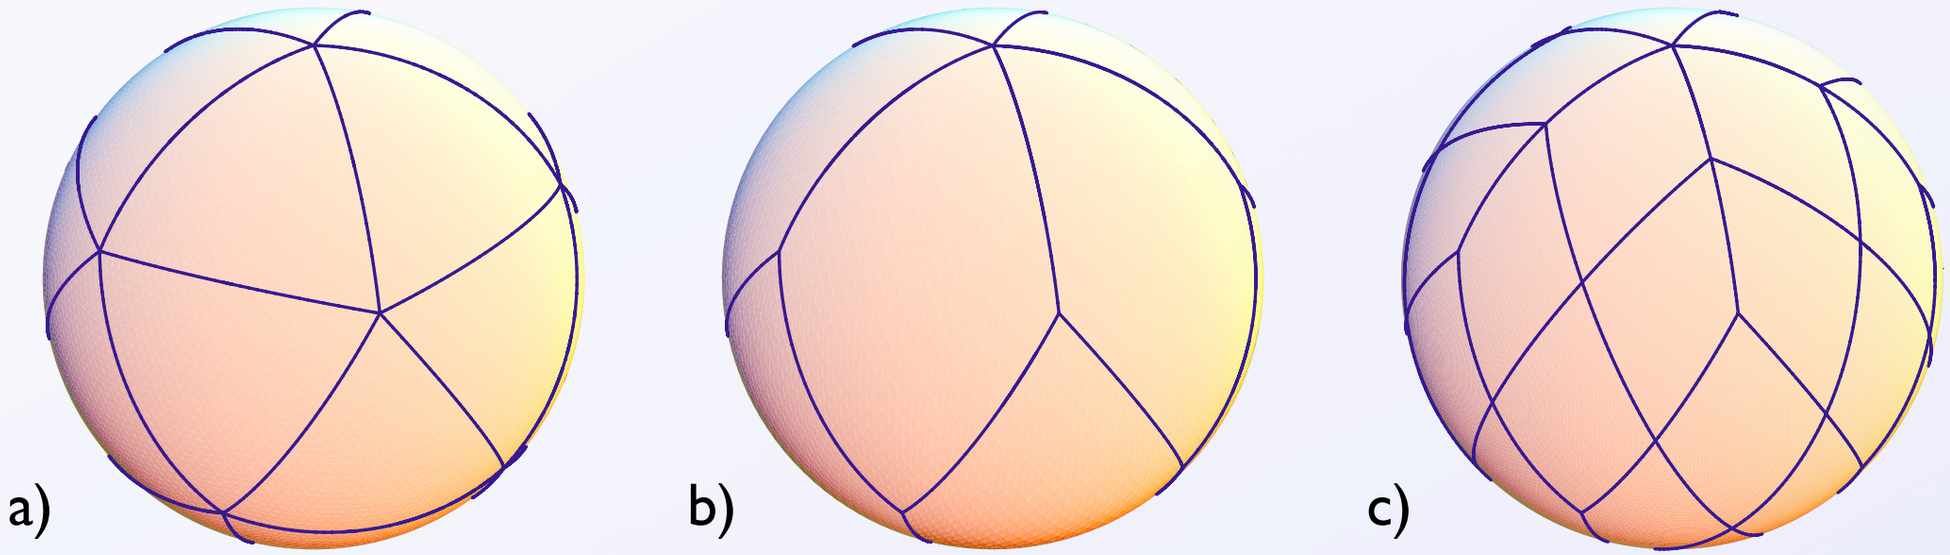

Parallel Programming Model

- The subdomain blocks are numbered as follows:

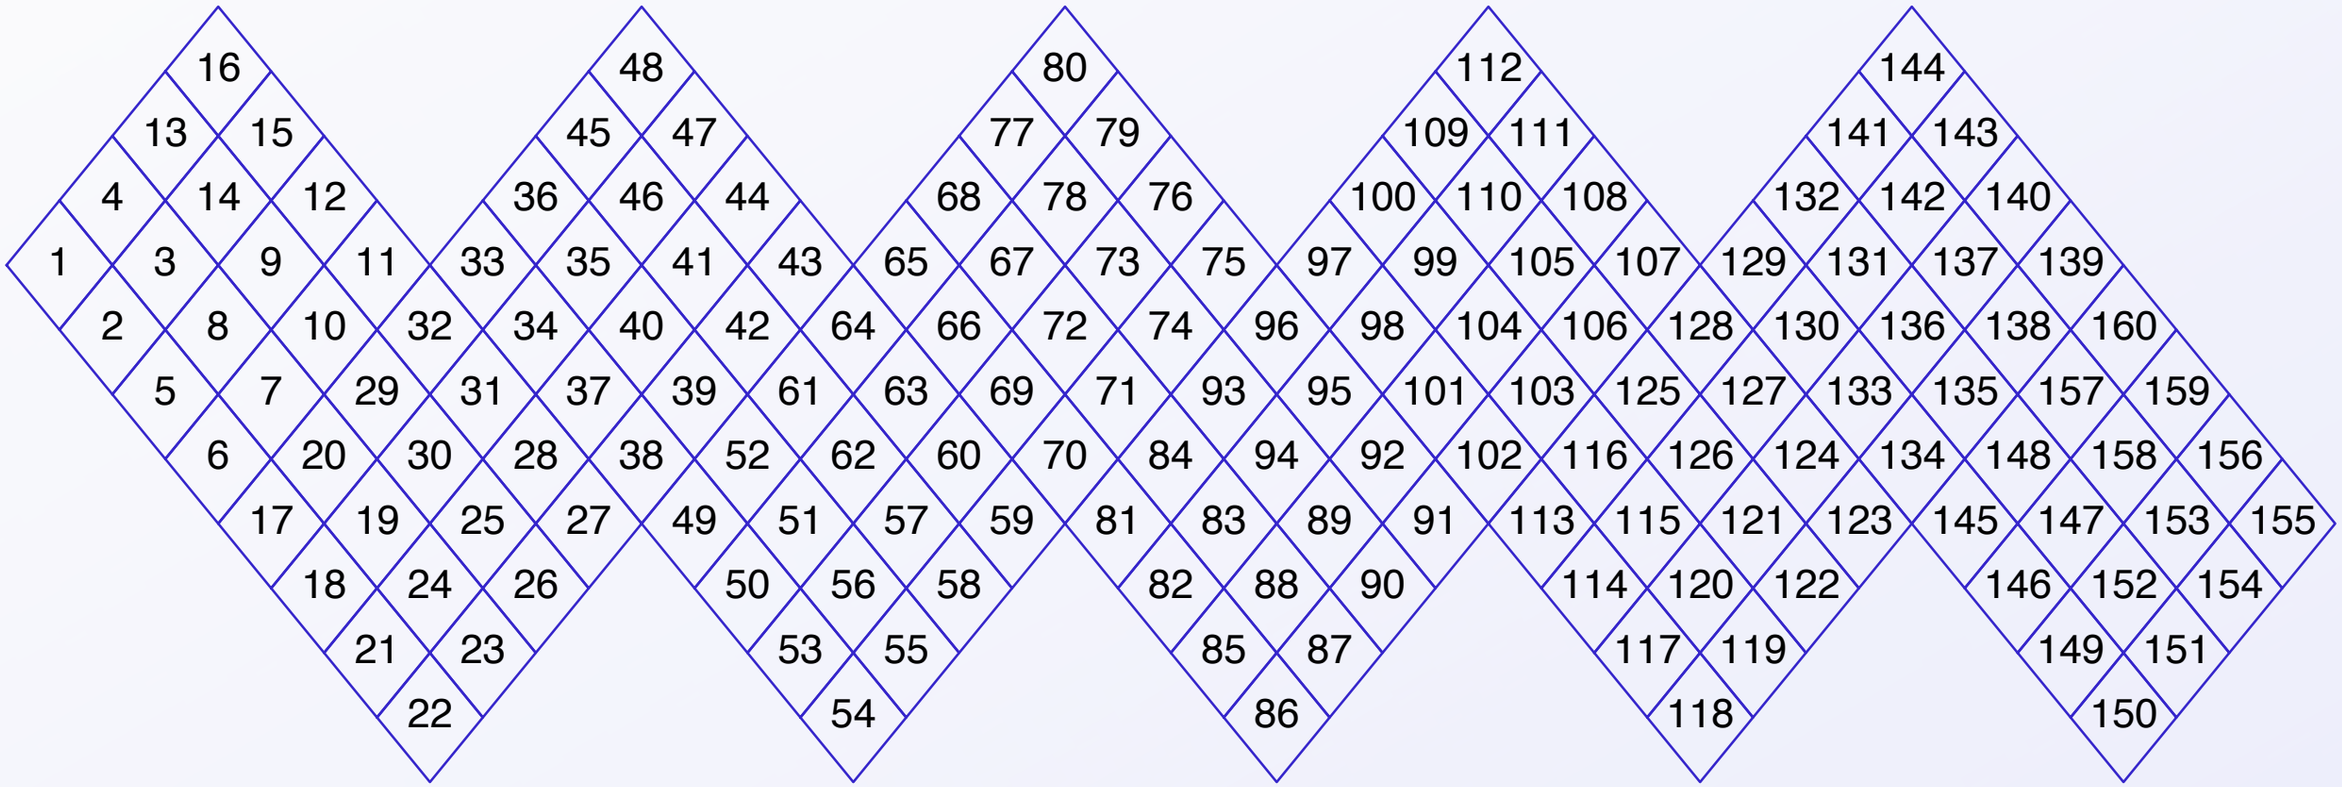

- Nonblocking MPI is used for communication between subdomain blocks.
- Currently runs on...
  - IBM SP RS/6000 Seaborg NERSC
  - IBM BlueGene/L Argonne
  - Apple XServe G5 104-processor CSU

I/O Patterns and Strategy

- Processes are grouped so that only one group member writes to disk.
  - The remaining processes within a group communicate with the I/O process via message passing
  - The I/O processes read/write their own files
  - The I/O processes could be exclusivily dedicated to I/O

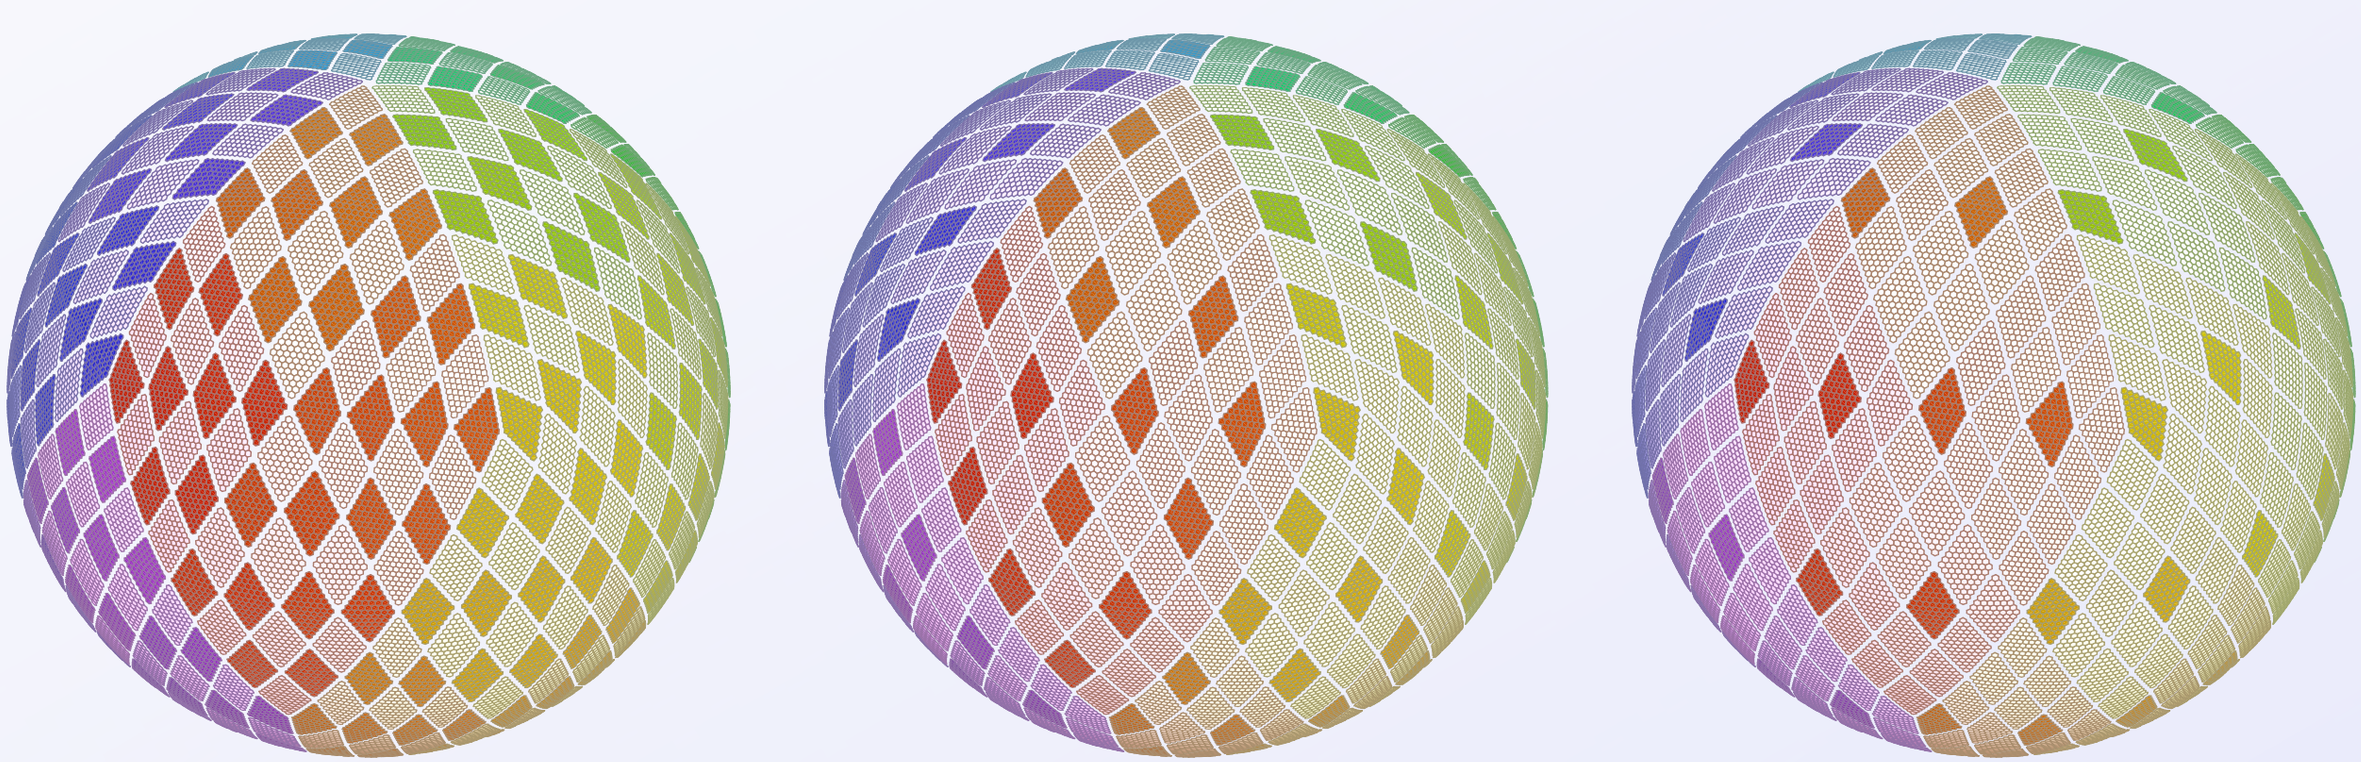

Groups of 2 Groups of 4 Groups of 8

I/O Patterns and Strategy

- Storage requirements for a restart file
  - Assume the following fields:
    1. Vertical component of vorticity at cell centers
    2. Horizontal component of vorticity at cell edges (X3)
    3. Potential temperature at cell centers
    4. 6 species of water at cell centers
  - One prognostic time level and three time levels of tendencies
  - Total bytes = (11 fields) X (4 AB3) X (8 bytes per word) X (number of cells) X (number of layers)

number of bytes (1e9)		number of cells (1e6)			
		10.5 (8 km)	41.9 (4 km)	167.8 (2 km)	671.1 (1 km)
number of layers	64	236	944	3780	15118
	96	354	1417	5669	22677

Memory requirements per process

- Assume the following 3D fields (prognostic and diagnostic):
  1. 11 prognostic variables
  2. vertical velocity at cell centers
  3. normal velocity at cell edges (×3)
  4. radiation fluxes at cell centers
  5. 16 other fluxes and other stuff
- 32 total 3D (96 layer) fields

number of bytes (1e9)		number of cells (1e6)				
		10.5 (8 km)	41.9 (4 km)	167.8 (2 km)	671.1 (1 km)	2684 (0.5 km)
number of processes	1	258	1031	4123	16492	65971
	1024	0.25	1.01	4.03	16.1	64.4
	2048		0.50	2.01	8.05	32.2
	4096			1.01	4.03	16.1
	8192			0.50	2.01	8.05

Post Simulation Data Service

- Our partners at PNNL are developing a web portal with the following features:
  - Enable users to search and locate datasets produced by the GCRM. A semantic model for the *metadata* will enable users to discover data related to their original searches or to browse based on selected criteria. The metadata services will make use of existing metadata standards and promote the processing, comparison, and sharing of data.
  - Hosts data reduction and data analysis. To accommodate the large data sets, we expect to build parallel tools for standard mathematical operations such as averaging, and expand on the capabilities of the NetCDF Operators.
  - Visualization services

Visualization and Analysis

- We are using Google Earth to visualize the 2D data on the earth. Capability to zoom and rotate.
- Jeff Daily and Karen Schuchardt and at PNNL
- For a demo see http://climate.pnl.gov/resources/gedemo.shtml

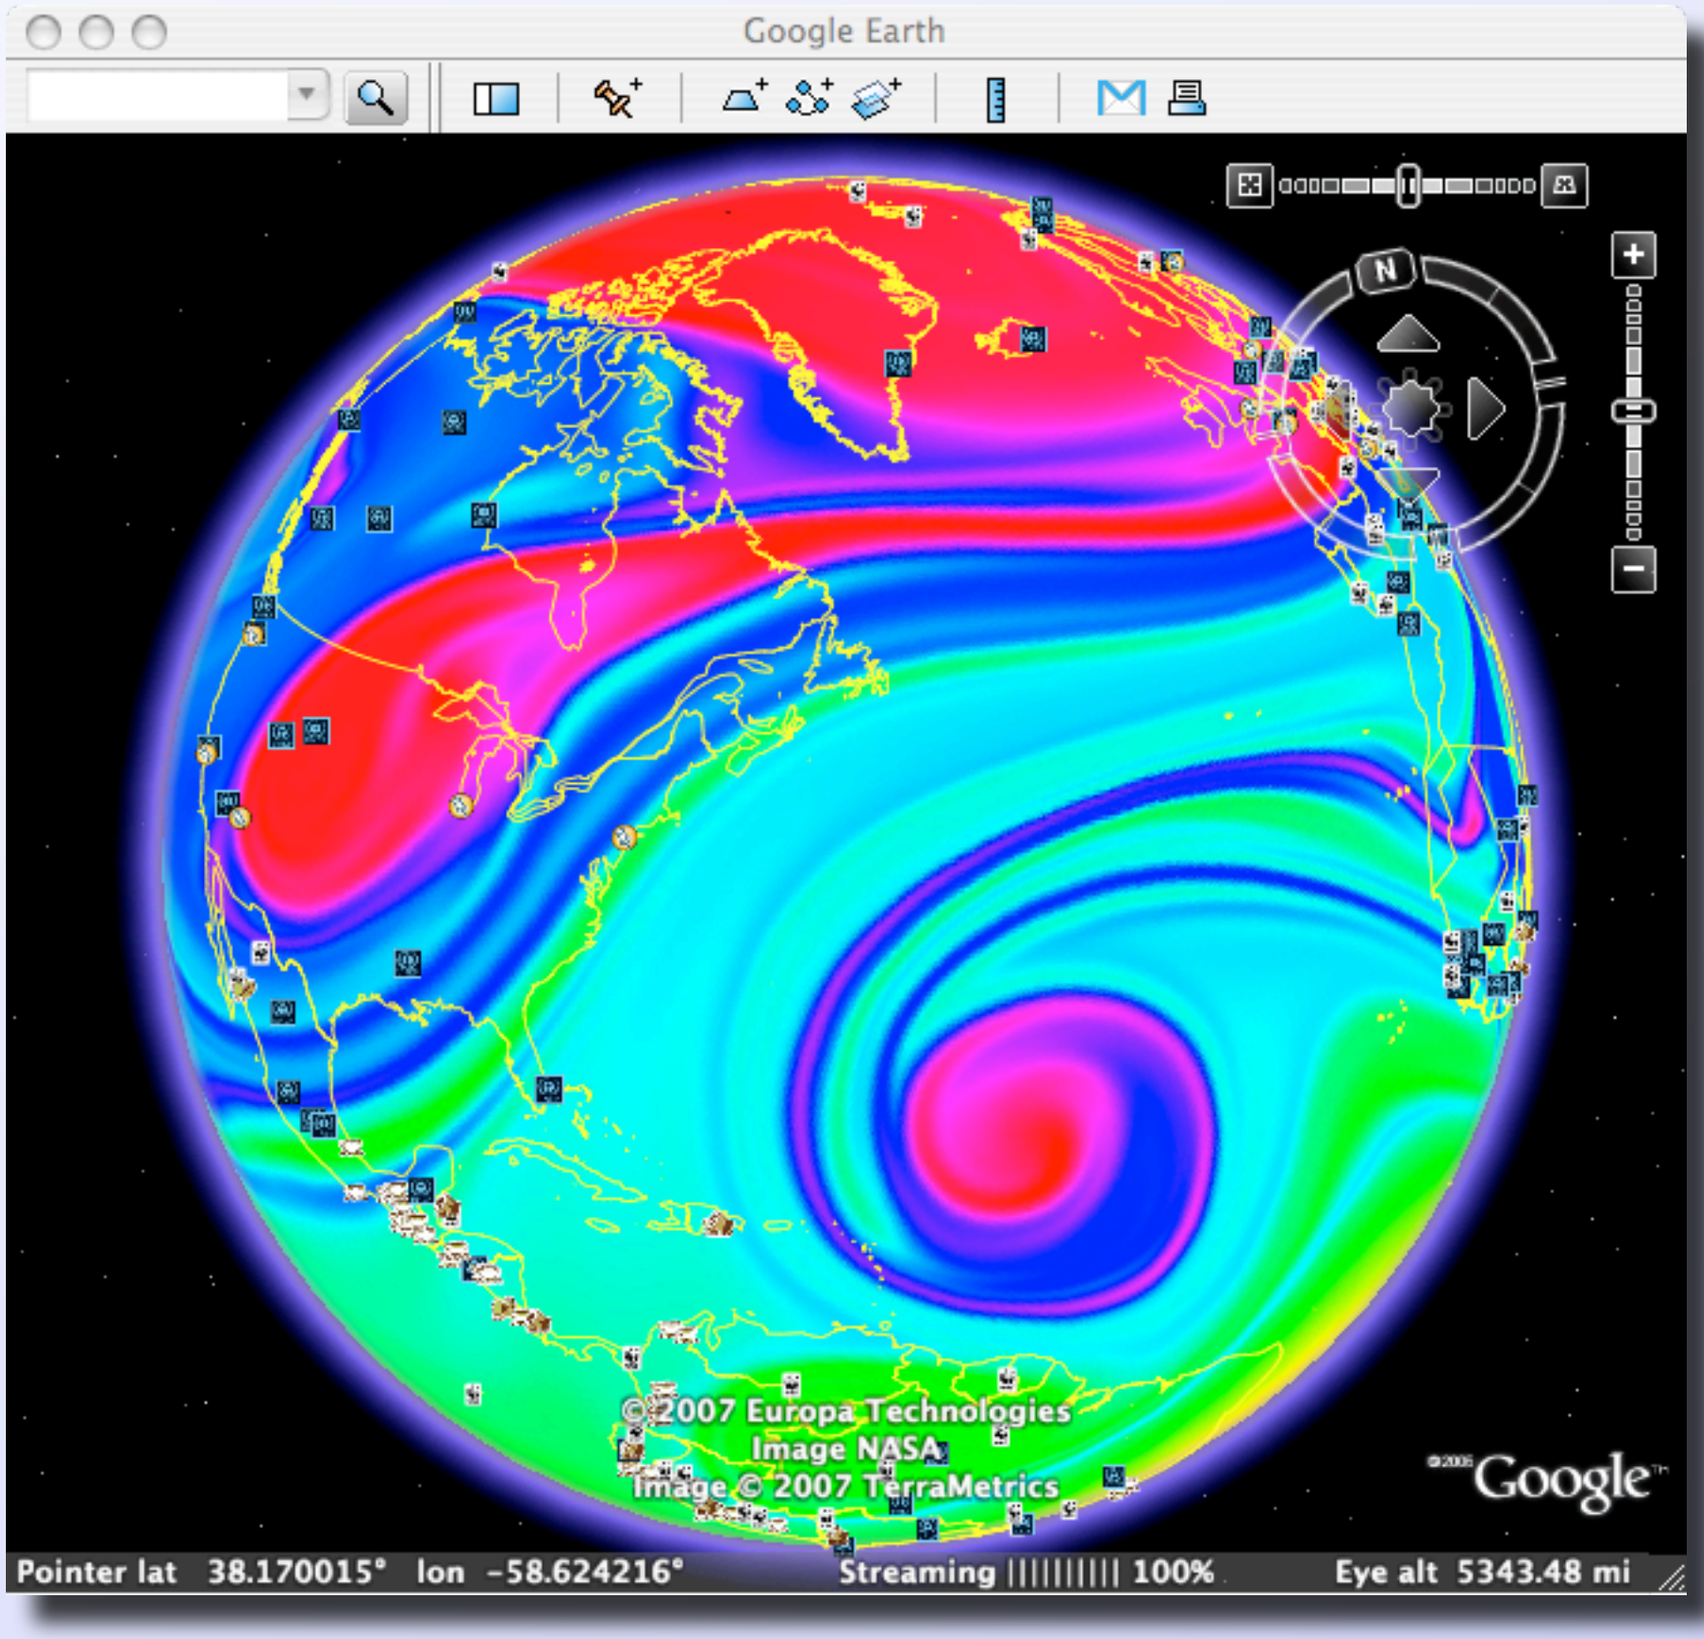

Status and Scalability

- The multigrid solver is an important component of the hydrostatic model. Data motion within the multigrid solver is representative of the model as a whole.
- Time to perform 10 multigrid V-cycles on 16 layers

Seaborg Time (s) (cells per process)		number of cells (1e6)			
		2.62 (16 km)	10.5 (8 km)	41.9 (4 km)	167.8 (2 km)
number of processes	80	5.65 (32K)	25.13 (131K)		
	160	3.27 (16K)	11.87 (65K)		
	320	2.14 (8K)	5.48 (32K)	25.42 (131K)	
	640		3.82 (16K)	12.62 (65K)	
	1280		2.87 (8K)	5.92 (32K)	25.89 (131K)

Roadmap

- Almost all the communication between processes is for the ghost-cell updates. With this in mind further optimize the code:
  - Assign subdomains to processes to take advantage of the pseodo-Cartesian topology. For example, experiment with MPI_CART_CREATE.
  - The finite-difference operators have small stencils. Overlap communication and computations.