High Resolution Climate Modeling with Geodesic Grids

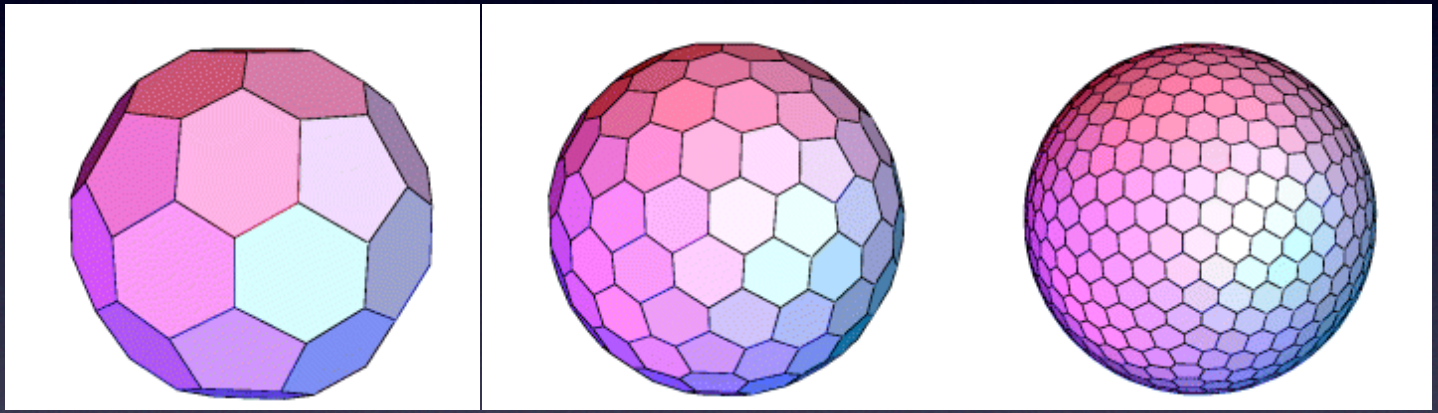

Don Dazlich
David Randall, PI
Colorado State University

Data Layout

- There is an underlying logically rectangular order in the grid.

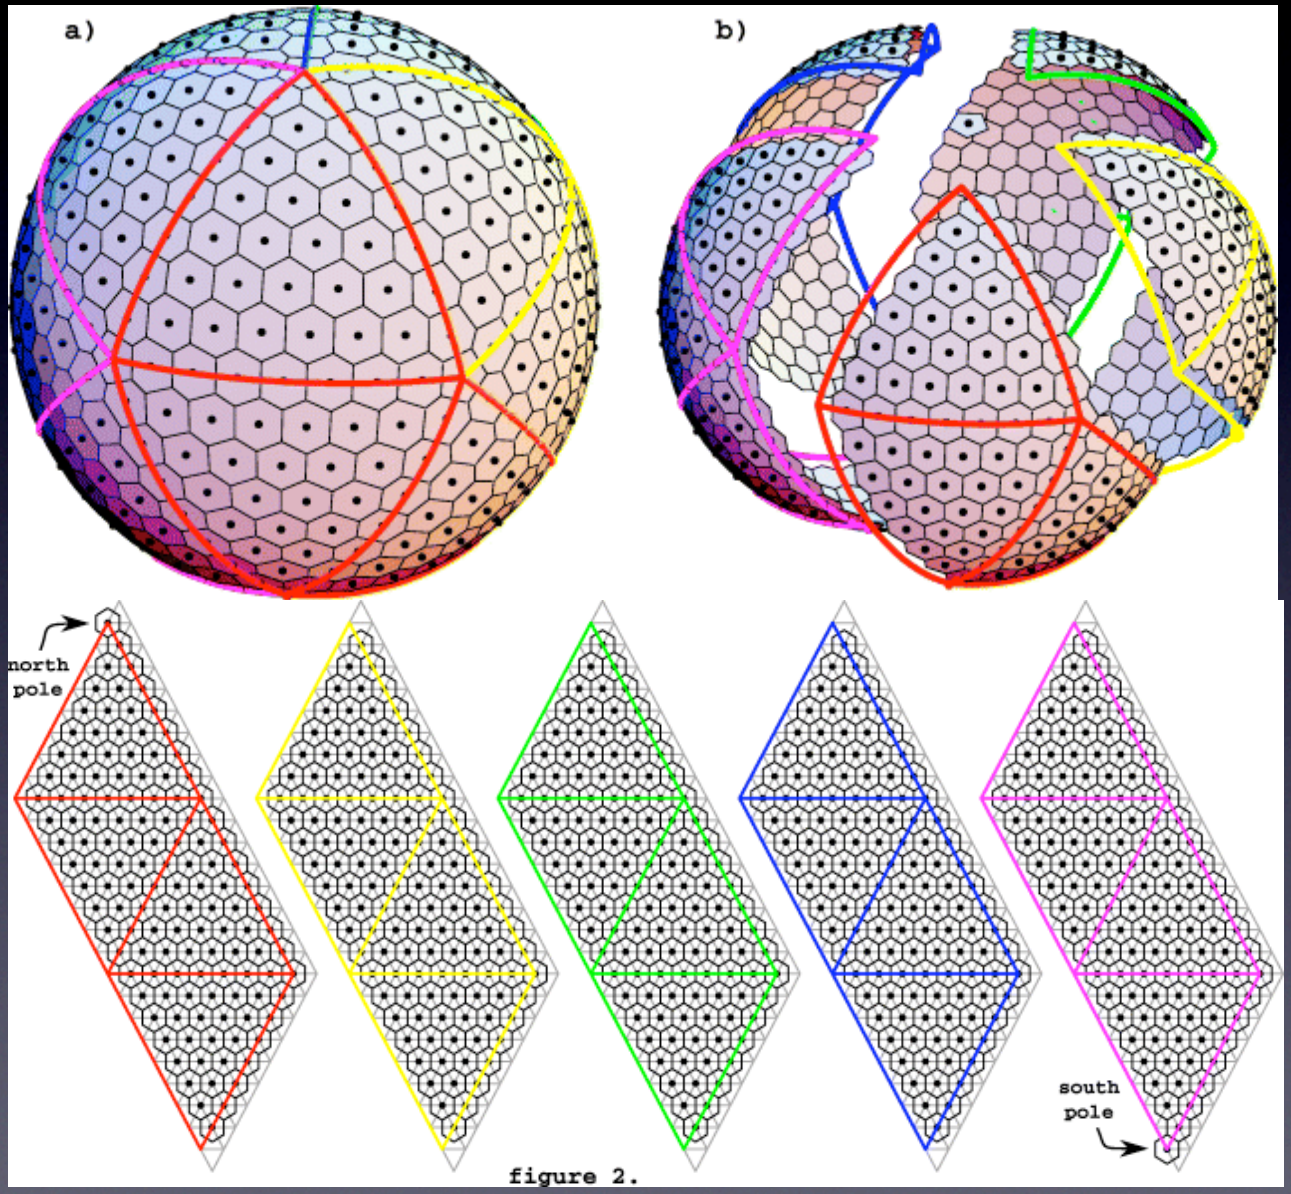

Parallelism

- All parallelism is MPI. The panels are subdivided and the subpanels are parceled out to the processes.
- Looking to implement hybrid MPI/OpenMP.

Models 1

- Atmosphere - 10242 gridcells, convective parameterization in each cell replaced with 2D cloud resolving model. Have run up to 1280 processes - load balance issues.
- Instead of a 4km global cloud resolving model at 100,000 times the cost, we simulate some cloud scale processes at roughly 200 times the cost of traditional GCM.

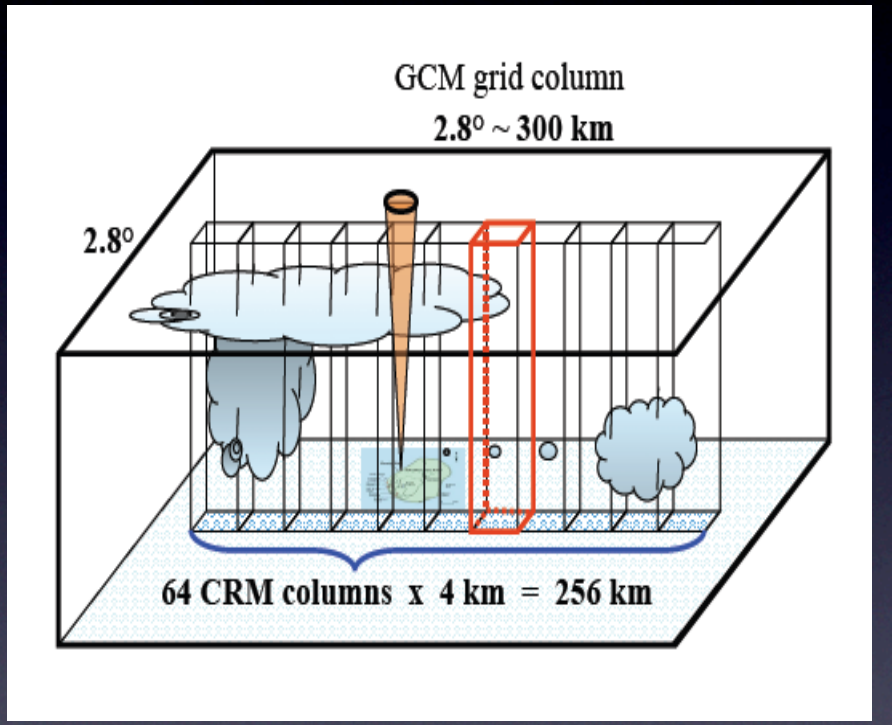

Models: 11

- Ocean model - currently running 655362 cells (1/4 deg). Target is 2621442 (1/8 deg) to resolve eddies.
- Scaling is excellent out to 320 procs. Code bottleneck here (some structures are O(nprocs) and we are transplanting some operators that avoid this.

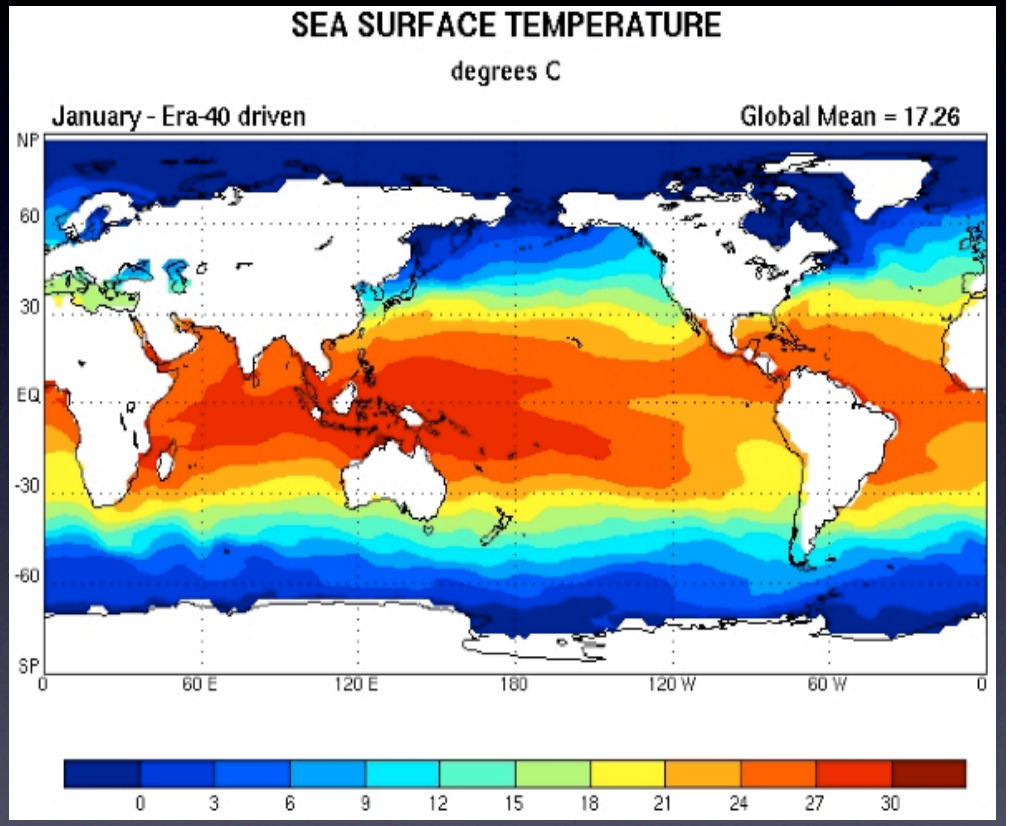

I/O

- Diagnostic output (netcdf) is written one file per process. Post-processing joins these files. We are considering one file per n processes at higher process counts.
- Input and checkpoint files are direct-access. Each record corresponds to one grid cell. This complicates visualization (transpose). We are looking at MPI-IO.