A Global Cloud Resolving Model

SciDAC
Scientific Discovery through Advanced Computing

Office of Science
U.S. Department of Energy

Goals

- Uniform global horizontal grid spacing of 4 km or better ("cloud permitting")
- 100 or more layers up to at least the stratopause
- Parameterizations of microphysics, turbulence (including small clouds), and radiation
- Execution speed of at least several simulated days per wall-clock day on immediately available systems
- Annual cycle simulation by end of 2011.

Motivations

- Parameterizations are still problematic.
- There are no spectral gaps.
- The equations themselves change at high resolution.
- GCRMs will be used for NWP within 10 years.
- GCRMs will be used for climate time-slices shortly thereafter.
- It’s going to take some time to learn how to do GCRMs well.

Scaling Science

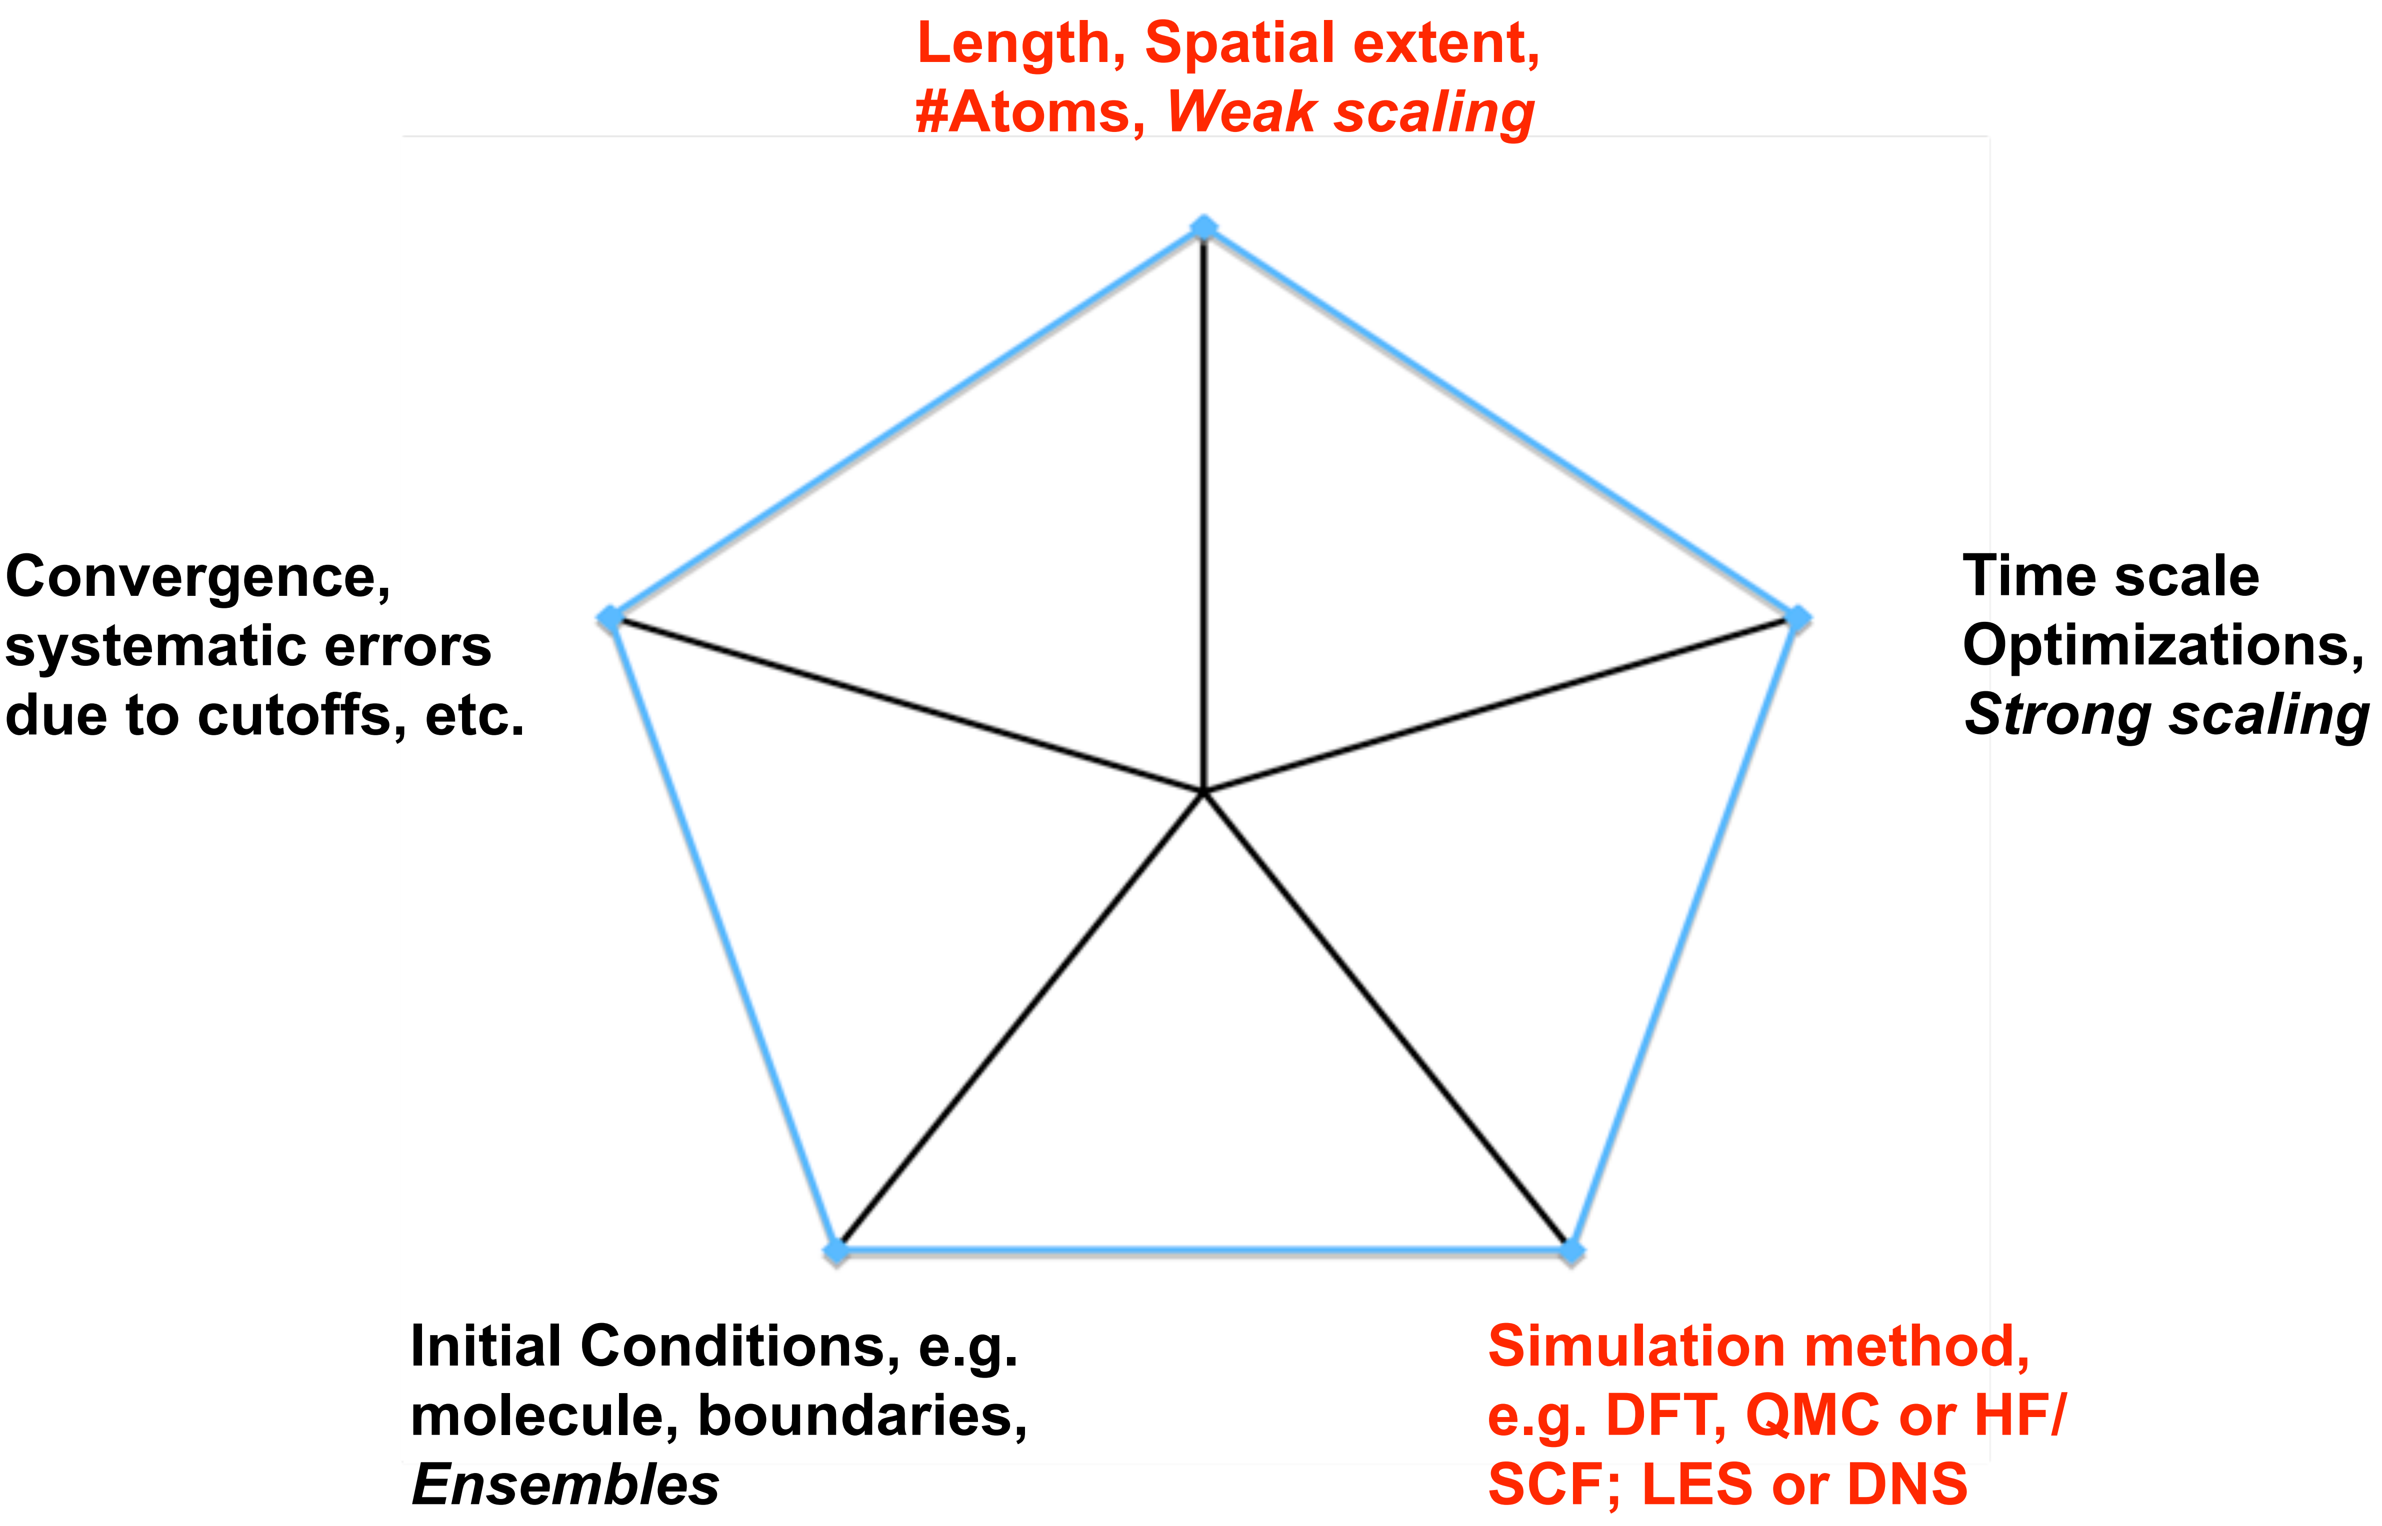

Basic design

- Non-hydrostatic
- Vertically propagating sound waves filtered
- Vorticity equation (instead of momentum equation)
- Mass and energy conserving
- Geodesic grid
- Z-coordinates (for now...)

Geodesic Grid

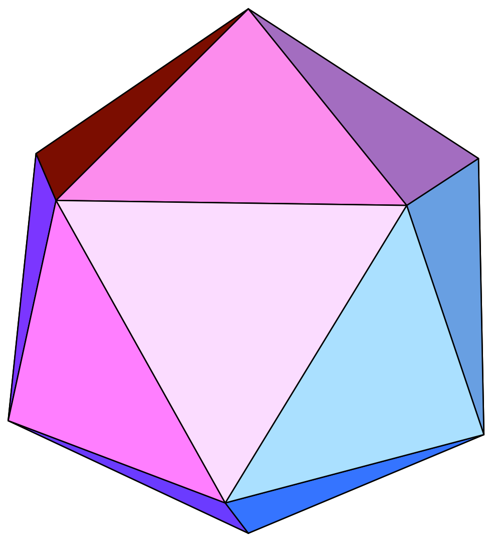

Icosahedron

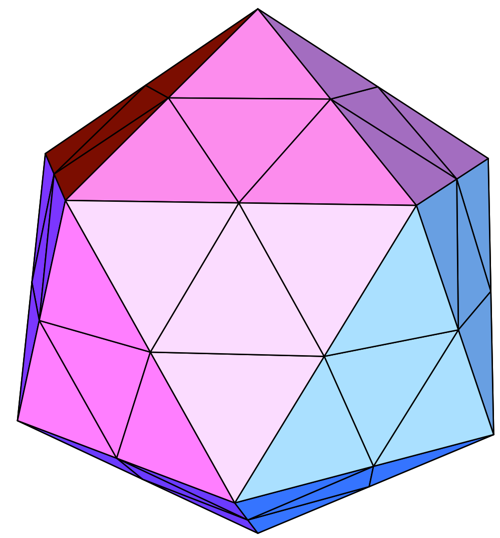

Bisect each edge and connect the dots

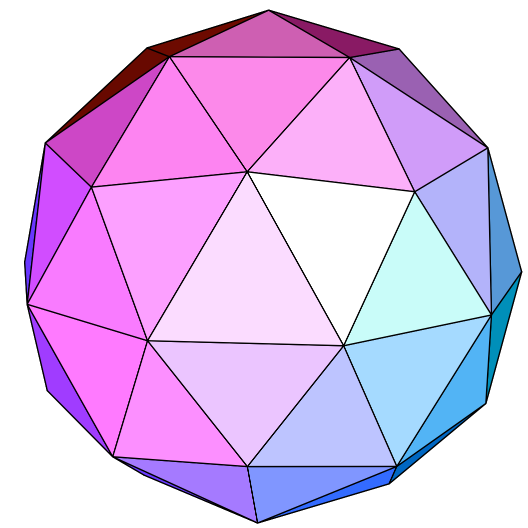

Pop out onto the unit sphere

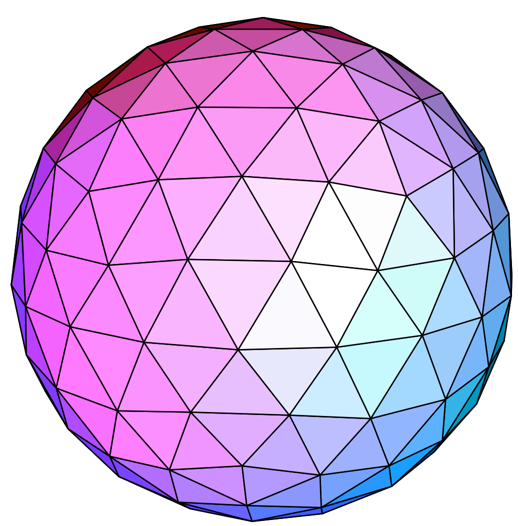

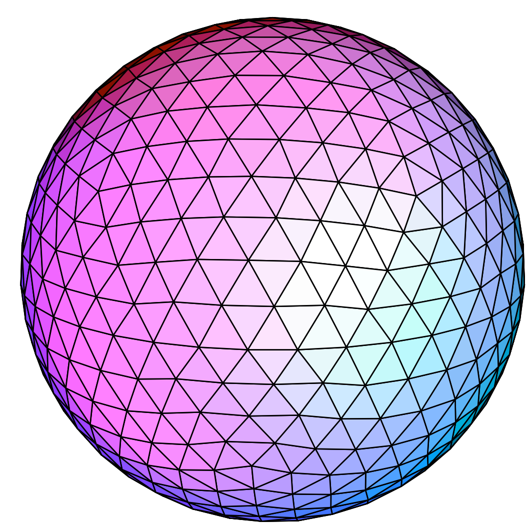

And so on, until we reach our target resolution...

Some grids of interest

Level of recursion	Number of grid columns	Distance between grid columns, km
9	2,621,442	15.64
10	10,485,762	7.819
11	41,943,042	3.909
12	167,772,162	1.955
13	671,088,642	0.977

Jablonowski Test Case

- 2621442 cells (15.64km) on 640 cores of franklin
- 850 hPa relative vorticity

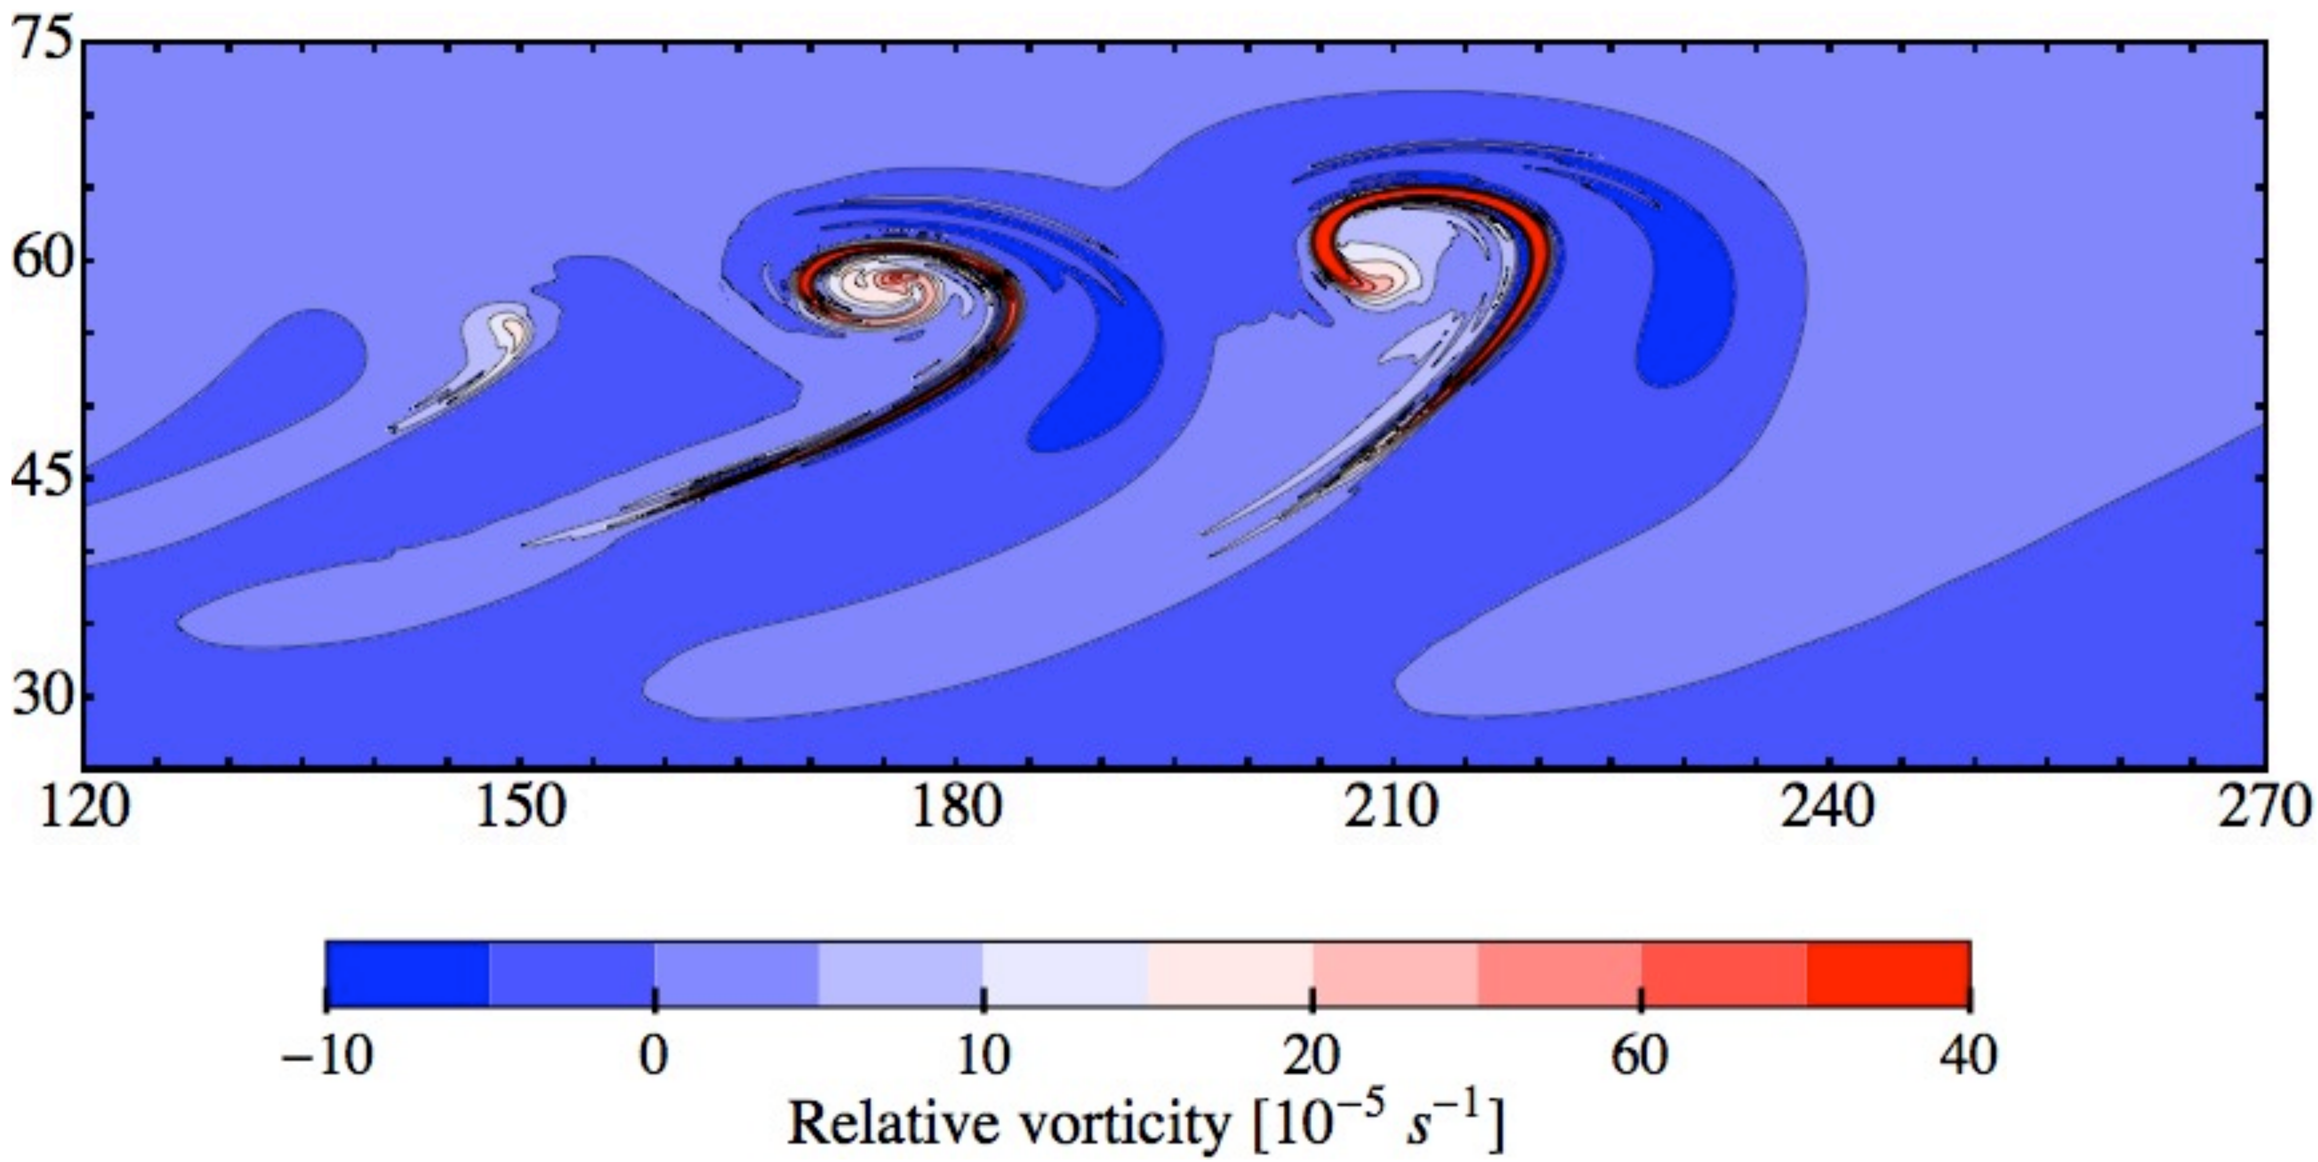

Scaling test of 3D-multigrid on Franklin

Time (s)		Number of cores			
		2560	5120	10240	20480
Grid resolution	41,943,042 **(11)** (3.909km)	19.57	10.96	5.56	2.87
	167,088,642 **(12)** (1.955km)	85.76	39.37	21.91	10.84

Scaling test of 3D-multigrid on Jaguar

- The **NCCS Cray XT5** with 181,00 cores
- 20 V-cycles
- 80 layers

Time (s)	Number of cores		
	2560	10240	40960
167,088,642 **(12)** (1.955km)	80.123	18.381	4.768

Full dynamical core on Franklin

Grid	PEs (Nodes)	GFlop /sec	Sec/ day
5	40 (10)	5.4	26
6	160 (40)	17.70	66
7	640 (160)	57.5	130
8	2560 (640)	168.30	355
9	2560 (640)	339.7	1403
10	5120 (1280)	638.3	5495
11	10240 (2560)	1366.4	20139

We think this can speed up by about a factor of two.

Key (rough) numbers

- >~ 40 million grid columns
- >~ 100 layers
- >~ 10 3D prognostic fields
- >~ 10 3D diagnostic fields
- >~0.4 TB per full write
- Time step ~ 10 seconds -- not just a stability issue
- Can use at least 20 K processors on XT5 -- probably 40 K
- Will produce about 5 simulated days per wall-clock day on 20 K processors with a 4 km grid spacing
- ~50000 processor hours/simulated day on Grid 11

Computational challenges

- **Efficient execution on a very large number of processors**
- **Parallel I/O (especially O)**
- **Management and distribution of the voluminous model output**
- **Analysis and visualization**

These are "infrastructure" issues that will be faced by anyone using a GCRM.

API Design

(Karen Schuchardt and Bruce Palmer, PNNL)

The API can be configured to allocate n nodes to serve as IO Aggregators.

The API is designed to support multiple parallel (or serial) IO layers

pnetcdf, netcdf4, netcdf3…

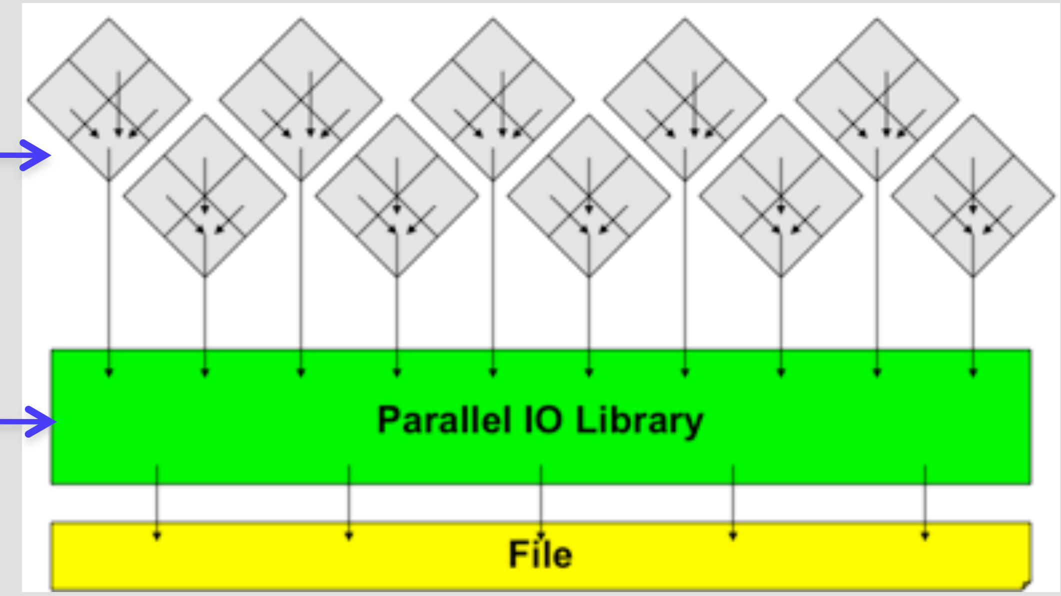

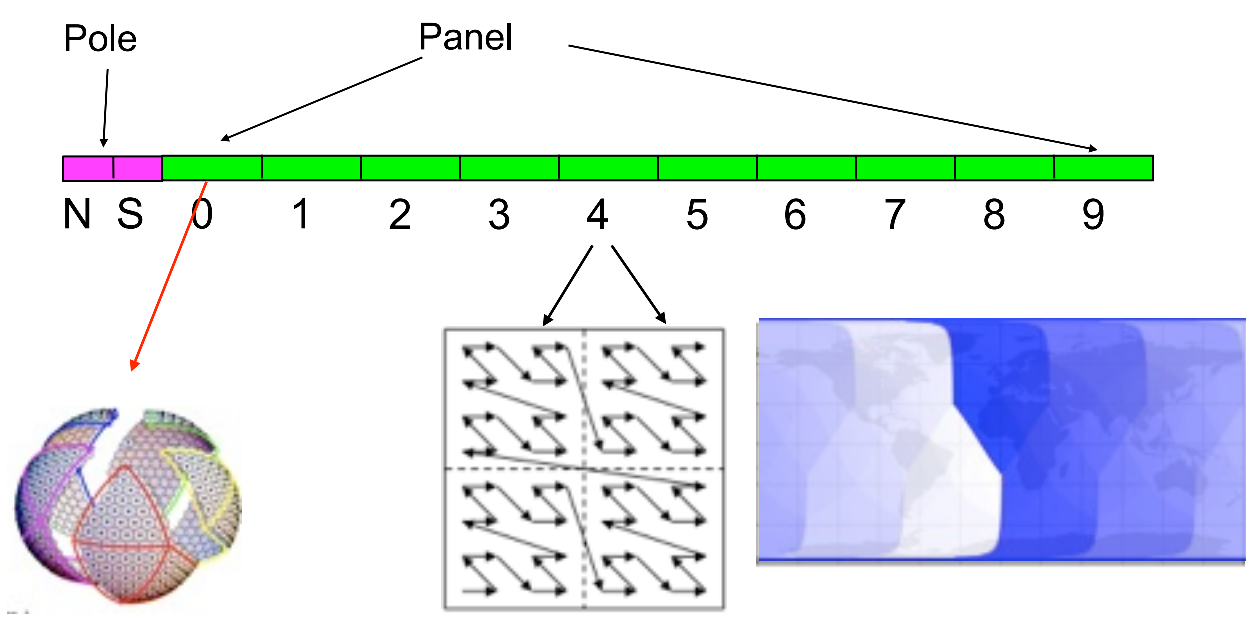

Grid and associated data linearized so that the sequence of grid cells follows a self-similar space-filling two-dimensional curve

- •Blocks within panels can be written as contiguous blocks
- •Order not dependent on number of processors
- •For parallel analysis, achieves good locality without special handling

Summary

- Qualitatively different
- Just barely feasible now
- Weak scaling and new "simulation method"
- Output volume huge but controllable
- Analysis and visualization challenges