SAND2011-1619C

The Variable Resolution Spe… Dynamical Core in the Community Atmospheric Model (CAM)

Michael N. Levy, James R. Overfelt, and Mark A. Taylor

Sandia National Laboratories
Albuquerque, NM, USA

SIAM Conference on Mathematical & Computational Issues in the Geosciences
March 24, 2011

Sandia National Laboratories

Outline

HOMME (High Order Methods Modeling Environment)

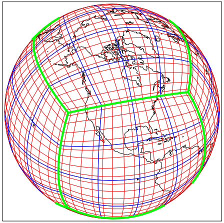

HOMME

- "A scalable and efficient spectral-element-based atmospheric dynamical core" (http://www.homme.ucar.edu)
- CAM-HOMME dynamical core available in CCSM
- Elements originally were squares on a cube, projected onto a sphere using gnomonic projection
  - Now able to use any quadrilaterals on cube or sphere (with great-circle edges)

Why Do We Refine?

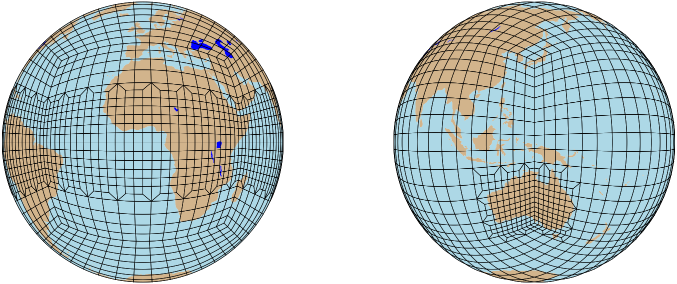

Benefits of Refinement

- High-res studies of specific areas (the tropics, Australia, etc)
  1. Refine over Atmospheric Radiation Measurement [ARM] sites
  2. Calibrate global parameters for high-res runs based on ARM data (significantly cheaper than tuning via 500+ global runs)
- Alternative to nested models for regional climate

How Do We Refine?

Refinement Options

1. Conforming or non-conforming?
2. Structured or unstructured?
3. Static or dynamic?

Constraint

The spectral element method, as implemented in HOMME, requires quadrilateral meshes tiling a sphere.

Spoiler: conforming unstructured static refinement

Conforming vs Non-Conforming [Quadrilateral] Refinement

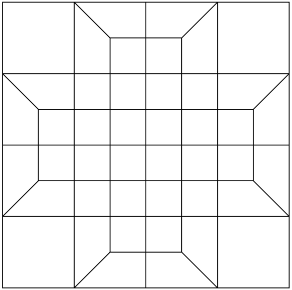

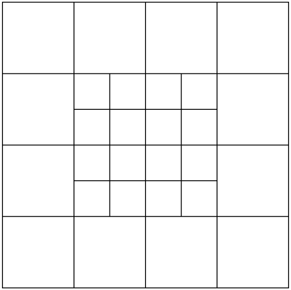

Conforming Refinement (left)

- **Focus of this talk**
- Every edge is shared by exactly two elements

Non-Conforming Refinement (right)

- Refine grid by splitting an edge
- Allows for “hanging nodes”

Structured vs Unstructured [2D] Meshes

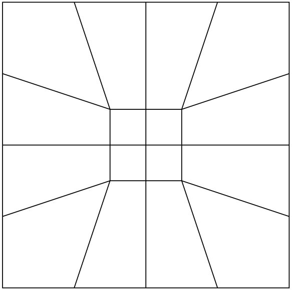

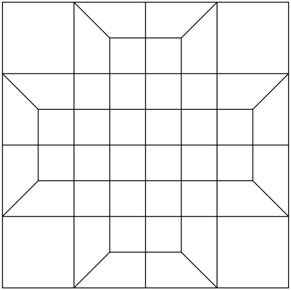

Structured Mesh (left)

- Domain is tiled by elements in such a way that elements can be numbered with (i, j) coordinates

Unstructured Mesh (right)

- **Focus of this talk**
- Domain is tiled arbitrarily

Static vs Dynamic Mesh Refinement

Static Refinement

- **Focus of this talk**
- Refine grid initially (based on topography, regional interests, etc), then run

Dynamic Refinement

- Refine grid continually throughout the run (based on gradients, mass, or some other user-defined criterion)
- Computationally more expensive, also far more complicated to implement

This project: conforming unstructured static refinement

More on Refinement Choice

Why Conforming Unstructured Static Refinement?

1. CAM-HOMME currently uses conservative SEM
   - Non-conforming refinement breaks conservation in SEM, would be better suited for DG (currently not part of CAM-HOMME)
   - Unstructured meshes allow more flexibility in refinement
2. Will be running CAM-HOMME with variable resolution by end of fiscal year
   - Dynamic refinement would take significantly longer to implement (and would restrict refinement options)

Changes Due to Refinement

Two Major Changes to HOMME

1. Implement hyperviscosity with variable viscosity coefficient, rather than static
2. Ability to read in mesh (Exodus file) rather than simply generate "uniform" meshes

And Some Minor Changes

- Map directly from element on sphere to reference element, bypassing cube (hope to move from "cubed sphere" to "global quads" in description of method)
- Change to resolution statistics (e.g. calculate average element length rather than average equatorial element length)

Changes Due to Refinement

Variable Viscosity

- Spectral Element core includes constant hyperviscosity

$$\frac{\partial \mathbf{u}}{\partial t} = -\nu \nabla^4 \mathbf{u},$$

implemented with auxiliary variable $\mathbf{f}$:

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} &= -\nu \left(\nabla(\nabla \cdot \mathbf{f}) - \nabla \times \hat{\mathbf{k}}(\nabla \times \mathbf{f}) \right) \\ \mathbf{f} &= \nabla(\nabla \cdot \mathbf{u}) - \nabla \times (\nabla \times \mathbf{u})\hat{\mathbf{k}} \end{aligned}$$

- Allowing ν to vary, hyperviscosity is

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} &= -\left(\nabla\sqrt{\nu}(\nabla \cdot \mathbf{f}) - \nabla \times \hat{\mathbf{k}}\sqrt{\nu}(\nabla \times \mathbf{f}) \right) \\ \mathbf{f} &= \nabla\sqrt{\nu}(\nabla \cdot \mathbf{u}) - \nabla \times \sqrt{\nu}(\nabla \times \mathbf{u})\hat{\mathbf{k}} \end{aligned}$$

Changes Due to Refinement

Variable Viscosity

- Spectral Element core includes constant hyperviscosity

$$\frac{\partial \mathbf{u}}{\partial t} = -\nu \nabla^4 \mathbf{u},$$

implemented with auxiliary variable $\mathbf{f}$:

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} &= -\nu \left(\nabla(\nabla \cdot \mathbf{f}) - \nabla \times \hat{\mathbf{k}}(\nabla \times \mathbf{f}) \right) \\ \mathbf{f} &= \nabla(\nabla \cdot \mathbf{u}) - \nabla \times (\nabla \times \mathbf{u})\hat{\mathbf{k}} \end{aligned}$$

- Allowing ν to vary, hyperviscosity is

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} &= -\left(\nabla(\nabla \cdot \sqrt{\nu}\mathbf{f}) - \nabla \times \hat{\mathbf{k}}(\nabla \times \sqrt{\nu}\mathbf{f}) \right) \\ \mathbf{f} &= \nabla(\nabla \cdot \sqrt{\nu}\mathbf{u}) - \nabla \times (\nabla \times \sqrt{\nu}\mathbf{u})\hat{\mathbf{k}} \end{aligned}$$

2 Questions about Variable Viscosity

1. Why vary ν?

Short Answer: Relationship between ν and effective resolution

- If ν is too large for fine mesh, results look similar to coarse mesh with same ν
- Alternately, if ν is too small for coarse mesh, results are noisy

Complicated Answer: Locally, want to dissipate near the grid scale

- This implies we want $\nu = \nu(\Delta x)$

2. What's the best way to vary ν?

Dissipation rate of mode $k = 2\Delta x$ is $\nu/\Delta x^4$

- Keeping dissipation rate constant $\Rightarrow \nu = C\Delta x^4$

2 Questions about Variable Viscosity

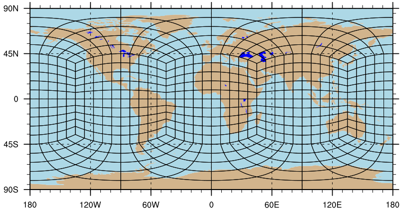

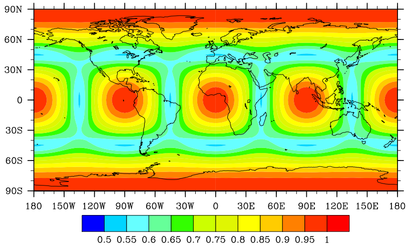

More on Varying ν

Since this is 2D code, vary ν by square of element area:

- Define area of element $\Omega_e := A_e$ and $\Omega_i :=$ largest element.
- For any element Ω_j, $\nu = \nu_0 \left(\frac{A_j}{A_i}\right)^2 \Rightarrow 0 < \nu \leq \nu_0$
- Continuity is enforced by averaging over element corners then using bilinear interpolation for element edges / interior nodes

SWTC 1

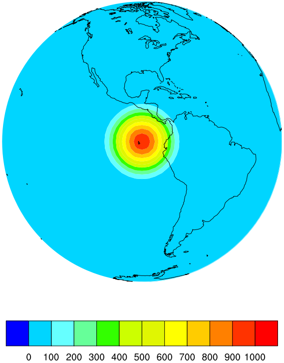

Williamson et al. – Test 1

- Advect a cosine bell around the globe
- A great test for refinement: refine the path of the bell

Two Refinements

Refinement Scheme #1

Start with 6×6 uniform grid on all faces, refine to 24×24 on the four equatorial faces ($1 \rightarrow 16$ splitting)

Refinement Scheme #2

Start with 12×12 uniform grid on all faces, refine to 24×24 on the four equatorial faces ($1 \rightarrow 4$ splitting)

Notes:

- The transition from coarse to fine occurs on equatorial faces
- Compare numerically to 24×24 uniform grid
- "Improve" grid by smoothing

SWTC 1 – Refinement #1

Not Smoothed

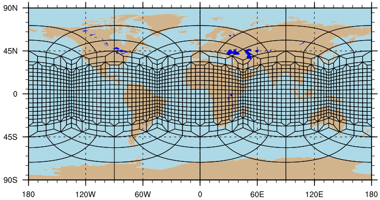

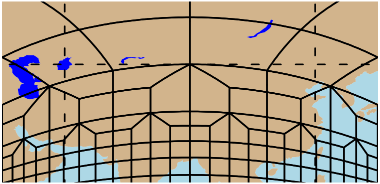

Smoothed (Mean Ratio technique)

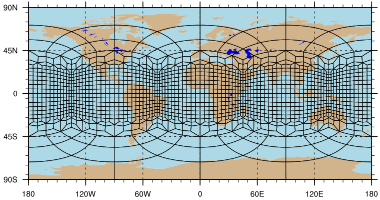

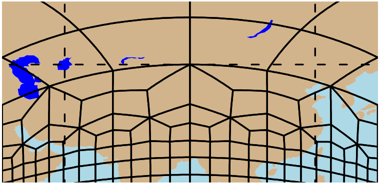

SWTC 1 – Refinement #2

Not Smoothed

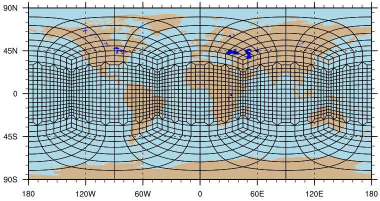

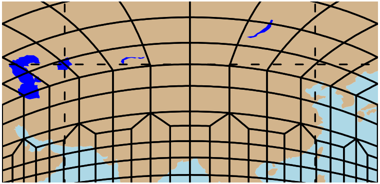

Smoothed (Mean Ratio technique)

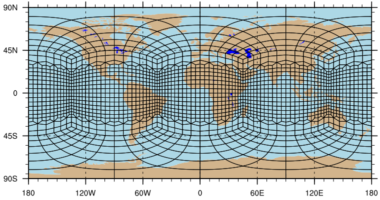

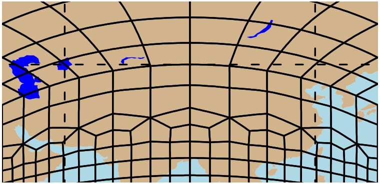

SWTC 1 – Computational Efficiency (Tabular)

Grid	# Elem	tstep (s)	Work Units	L^2 error
12×12	864	720	1	$5.784 \cdot 10^{-2}$
18×18	1944	480	3.38	$2.333 \cdot 10^{-2}$
24×24	3456	360	8	$1.225 \cdot 10^{-2}$
Refine 1	1656	360	3.83	$1.343 \cdot 10^{-2}$
Refine 2	1920	360	4.44	$1.261 \cdot 10^{-2}$
Smooth 1	1656	360	3.83	$1.228 \cdot 10^{-2}$
Smooth 2	1920	360	4.44	$1.227 \cdot 10^{-2}$

Table Details

- 1 work unit = computation (time) to run coarsest uniform grid

  As resolution increases, time step decreases

- Refined grids are slightly more work than 1944 element uniform grid, error is comparable to 3456 element uniform grid

  Even better results w/ smoothing

SWTC 1 – Computational Efficiency (Graphical)

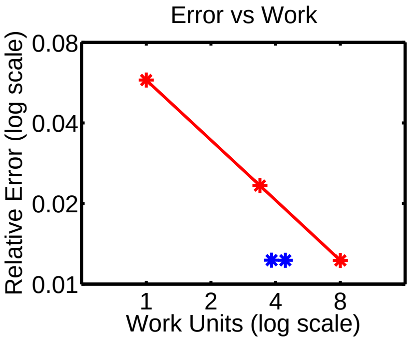

Graph Details

- Same data as presented in table on previous slide
- Red line = uniform meshes, blue points = refined / smoothed
- Not shown: advecting through low-res region of refined mesh (over poles) results in same error as global low-res mesh

SWTC 5

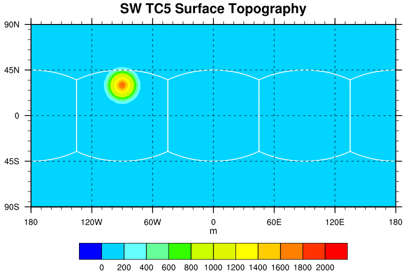

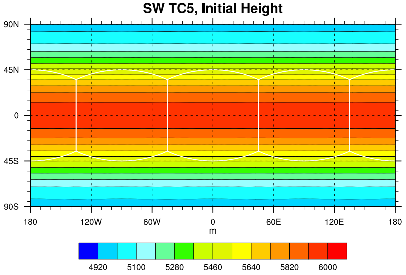

Williamson et al. – Test 5

- Flow around an isolated mountain
- Another good test for refinement: refine around the mountain

Experiment

Mountain has radius of 20°, refine area w/ radius 30°
Compare meshes based on coarsest elements

Notation

Grid: `N20_x4_s9`

- **N20** Begin with uniform grid based on 20×20 elements on each face of cubed sphere
- **x4** Refine such that edge length in coarse region is 4 times the length of that in fine
- **s9** Apply smoothing operator to grid 9 times

Source

Exploring a Multi-Resolution Modeling Approach within the Shallow-Water Equations
T. D. Ringler, D. Jacobsen, M. Gunzburger, L. Ju, M. Duda, W. Skamarock
Submitted to Monthly Weather Review

Comparing three grids

N10_x1.g

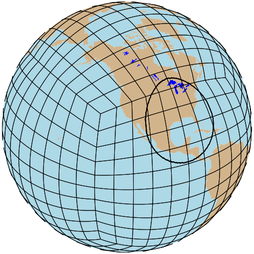

N20_x1.g

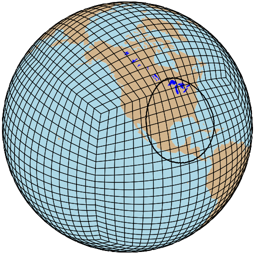

N10_x2_s6.g

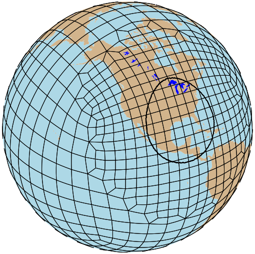

First wave of results

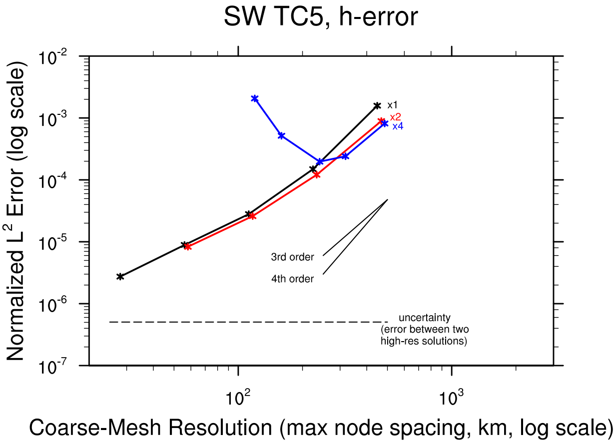

Problem in the *x4* Grids

N10_x4_s6_old.g

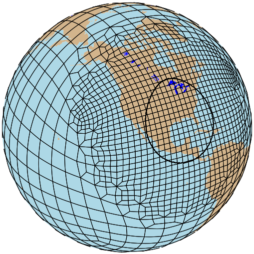

N10_x4_s6.g

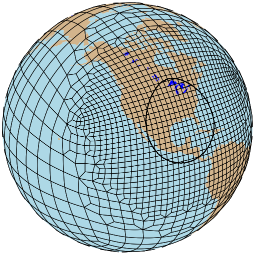

Transition Region

We kept the size of the fine mesh the same, but enlarged the transition region. This fixed the low-res `x4` grids, but still had a problem around `N40`.

Second Wave of Results

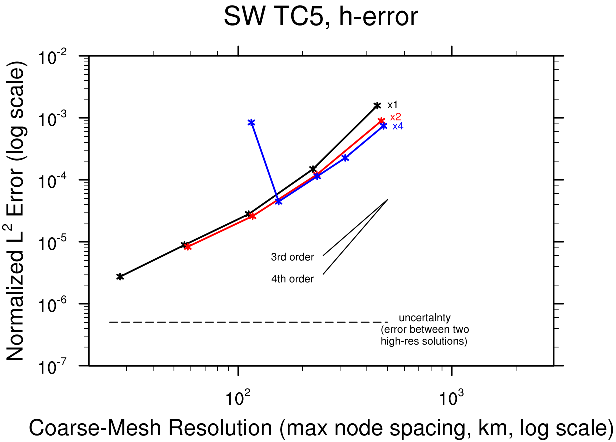

Future Work

Still to Come

1. Improve grid construction ("sizing function" determines how elements are located, Anderson et al.)
    - Though transition region improved look of `x4` results, didn't help much with `x8` grids
2. Initial 3D runs: start w/ aquaplanet
3. Full 3D: look into vertical dissipation (Tribbia and Temam report)
4. Get mesh refinement working with DG core of HOMME

Acknowledgements

Sandia is a multiprogram laboratory operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin, for the U.S. Department of Energy's National Nuclear Security Administration.