Multiscale methods for enabling scale-aware capability in CESM:

A SciDAC Climate Application

William Collins* (LBNL, PI) and

Peter Caldwell (LLNL), Bert Debusschere (SNL), Steve Ghan (PNNL), Don Lucas (LLNL), Lenny Oliker (LBNL), Todd Ringler (LANL), and Carol Woodward (LLNL)

Drivers for the project

- Many phenomena we want to simulate occur at the very small scales of clouds and ocean eddies.
- Two-way interactions are important for the organization and variability of the climate.
- Given cost of uniform ultra-high resolution, ***statically or dynamically refinable models*** could be a key experimental platform.

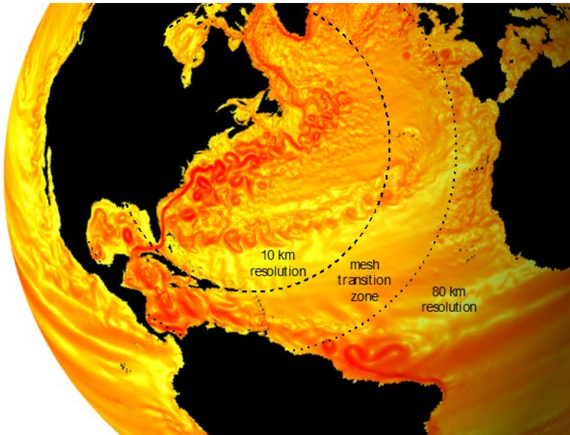

Goals of the project

- Develop, validate, and apply multiscale models of the climate system based upon atmospheric and oceanic components with variable resolution.
- Exploit new variable resolution unstructured grids based on finite element and finite volume formulations developed by DOE.
- Integrate advances in time-stepping methods, grid generation, and automated optimization methods for next-generation computer architectures.

Clouds scales and model resolution

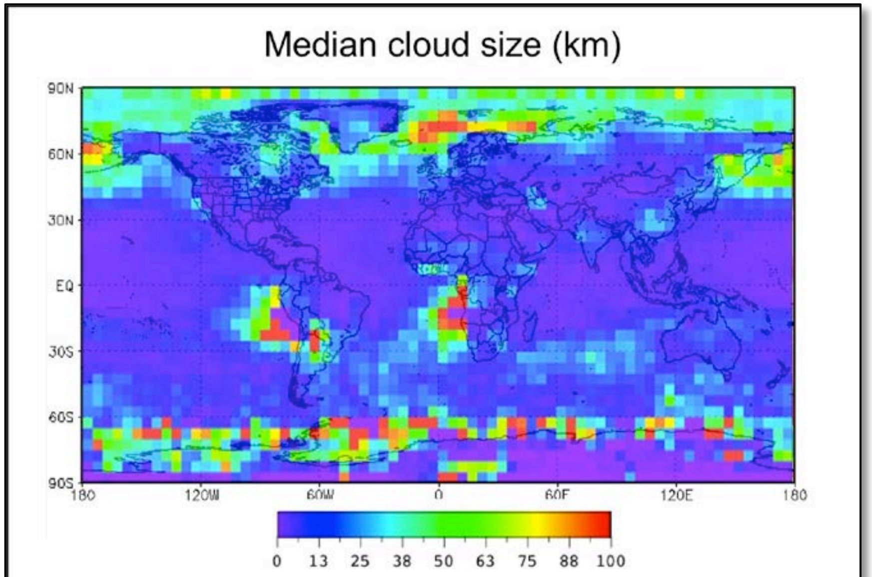

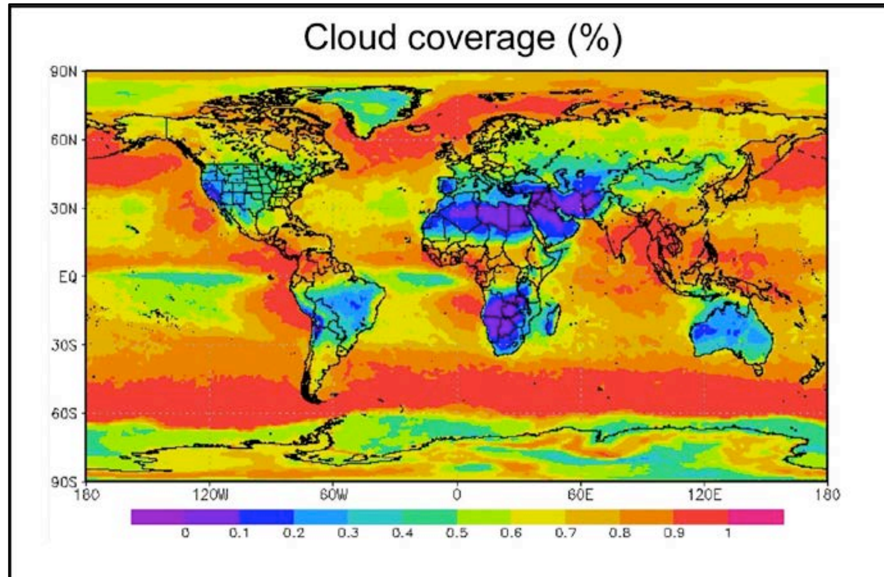

UHR AR5 AR4

Requirements for next generation physics

- Fidelity down to scales of key features of interest: cloud systems and ocean eddies
- Relaxation of usual parameterization assumptions:
  - *Ensembles:* Grid box may contain ~1 system, not >>1
  - *Scale separation:* Time steps and grid resolutions comparable to characteristic time and length scales of systems
  - *Equilibrium:* Due to scale "entanglement", sub-grid physics are not in equilibrium with boundary conditions imposed by resolved fields
  - *No memory:* Sub-grid systems will retain state across steps
  - *Scale awareness:* Physics needs to quasi-invariant to resolution
  - *Determinism:* Physics evolution is inherently stochastic.

Needs for scale aware parameterization

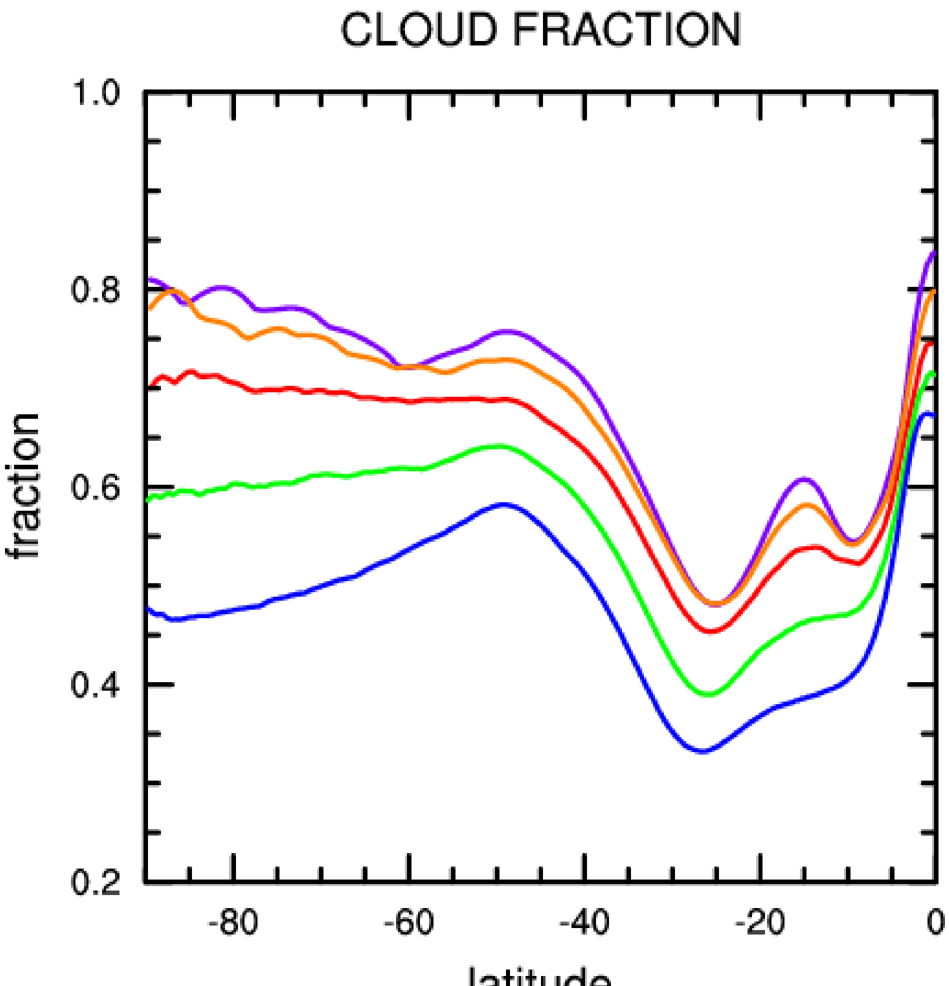

Taylor, 2011

CAM-SE with resolutions:
2.7° 1.9° 1.0° 0.5° 0.25°

- Monotonic decrease in LWCF also seen in CAM 3.1 Aqua Planet Experiments (Williamson, Tellus 2008)
- Cloud Fraction monotonically decreases with resolution
- CAM-SE simulations (shown) are similar.

Reverse Resolution Dependence in CAM

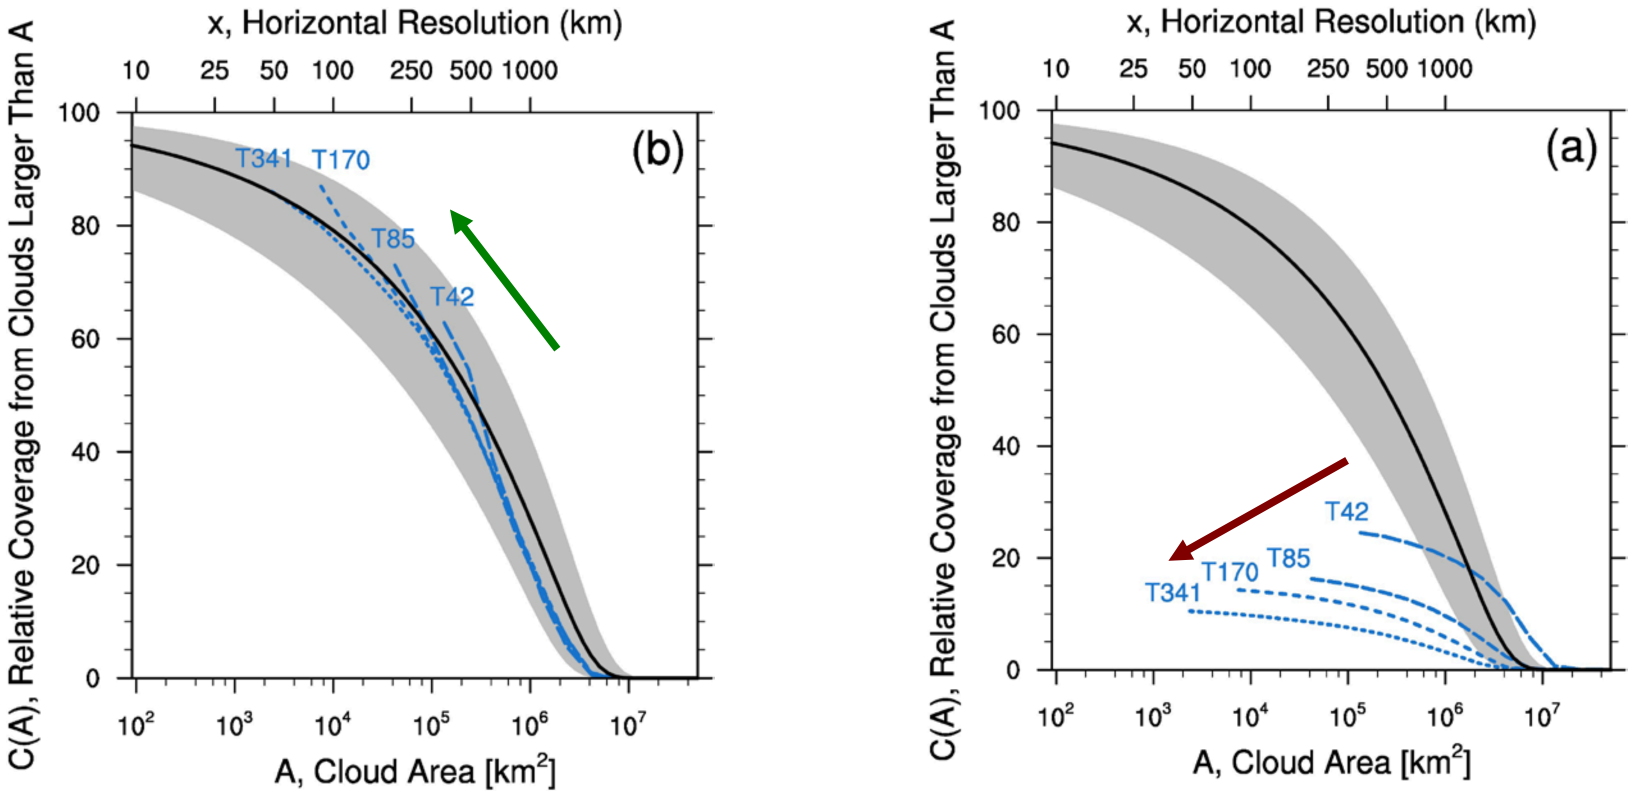

Expected Resolution Dependence

Actual Resolution Dependence

O'Brien et al, 2013

Project Elements

Variable Mesh Dycores

Model for Prediction Across Scales (MPAS)

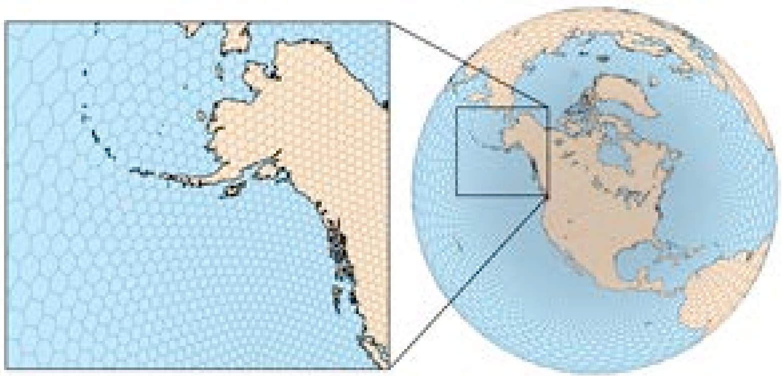

Spectral Element Dycore

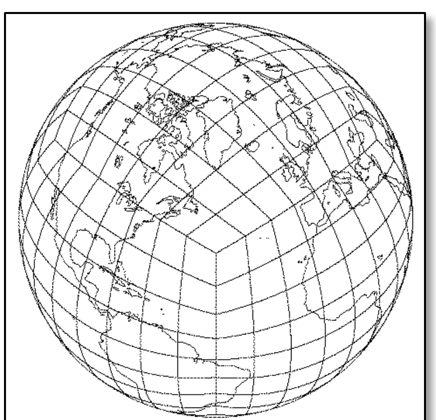

Physics-Dynamics Interface

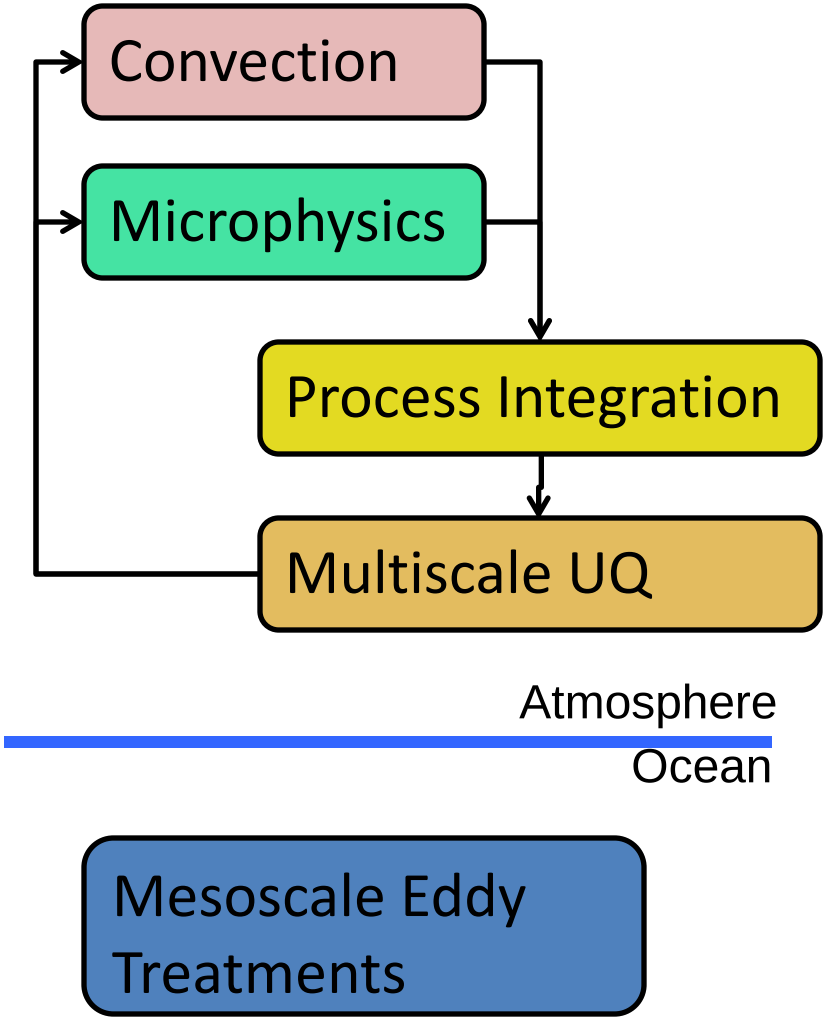

Ocean: MPAS-O Optimization

Variable Mesh Dycores

Model for Prediction Across Scales (MPAS)

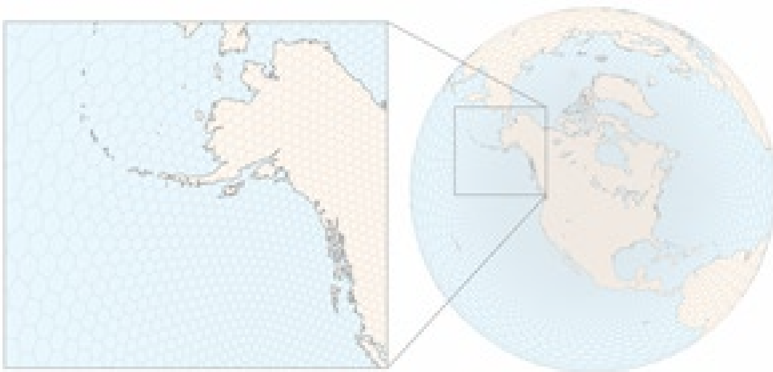

Spectral Element Dycore

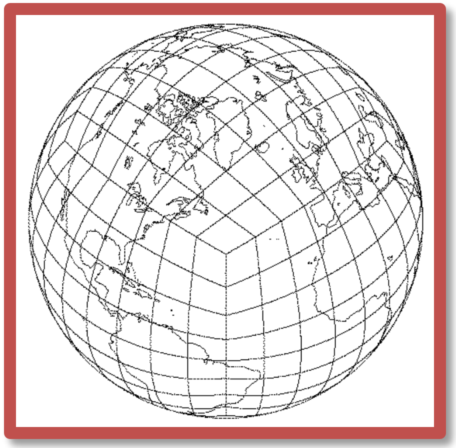

Physics-Dynamics Interface

Convection

Microphysics

Process Integration

Multiscale UQ

Atmosphere

Ocean

Mesoscale Eddy Treatments

Spectral Element Method

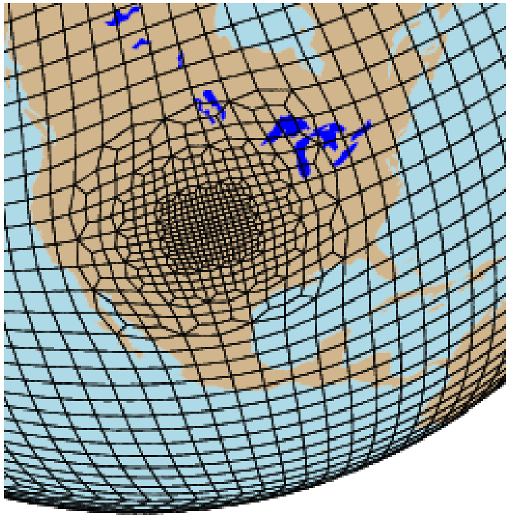

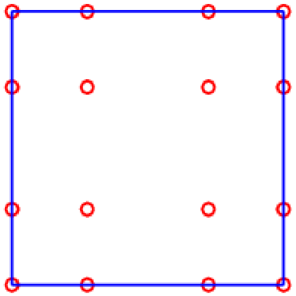

Taylor, 2011

- Spectral Elements: A Continuous Galerkin Finite Element Method
- All inter-element communication is isolated to one step in method, providing a clean decoupling of computation & communication
- High-order (4th) discretization
- Excellent local conservation: mass, energy, potential temperature, 2D PV
- All properties preserved on fully unstructured grids

CAM-SE Scalability

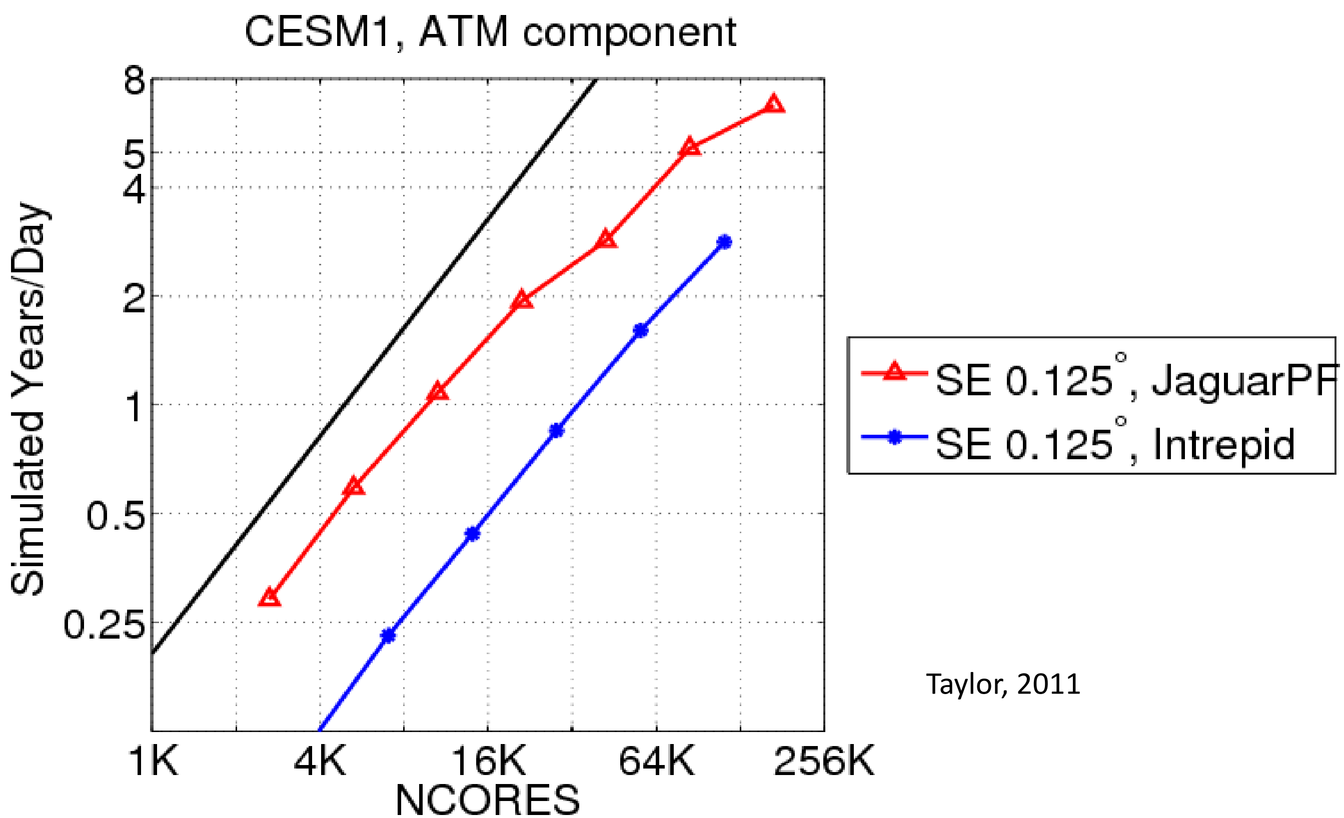

Taylor, 2011

- Excellent scaling to near full machine on both LCFs:

- **Intrepid:** Excellent scalability, peak performance at 115K cores, 2.8 SYPD.
- **JaguarPF:** Good scalability, peak performance at 173K cores, 6.8 SYPD.

Development of fully implicit solver for SE

SciDAC Multiscale

Goal: Maintain time step size with spatial refinement

Application of solver to a refined grid (2^{o} refined to 1^{o}) near a mountain feature

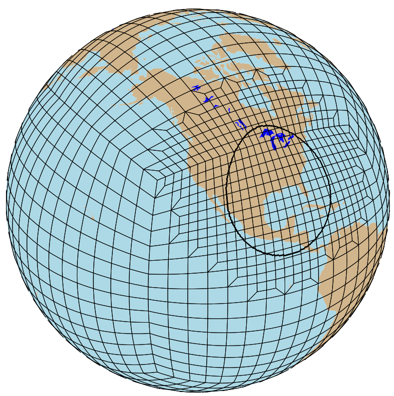

SW Test: Implicit 1200s time step converged successfully with an L_2 norm of 3.1e-4 for 1 day compared to explicit CFL limited 120s time step

Implementation: Extend implicit solver to hydrostatic dycore

Coarse baroclinic test case: Implicit backward Euler solves for a 1800s time step

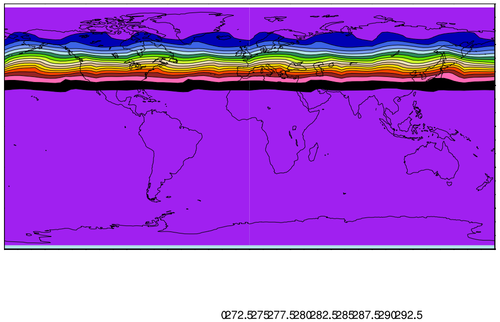

J. Ribbeck, K. Evans, C. Baker, R. Archibald, M. Taylor, A. Salinger

Preconditioning to accelerate implicit solves

Efficiency gains from implicit methods are not fully realized without an effective preconditioner:

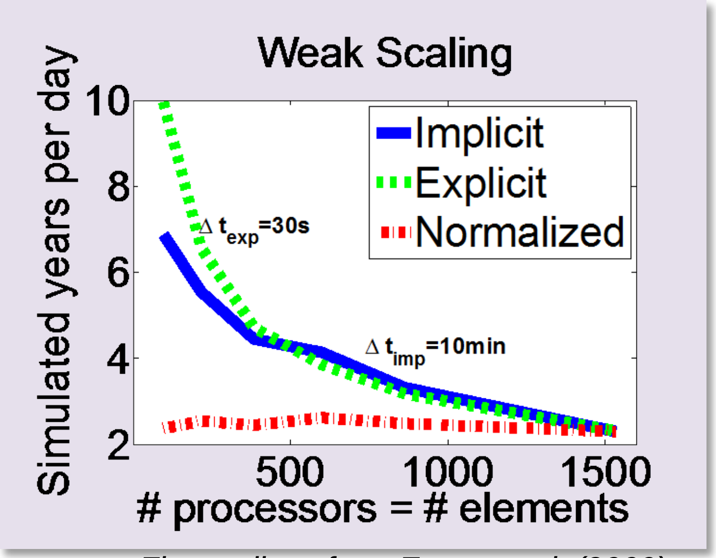

Fig: scalings from Evans et al. (2009)

1. Cast in vector formulation

$$\frac{dy}{dt} = f\left(y^{n+1}\right) \approx \frac{y^{n+1} - y^{n}}{\Delta t}$$

2. Apply Newton-Krylov method w/ preconditioner *P:*

$$F\left(y_{k+1}\right) \approx F\left(y_k\right) + F'\left(y_k\right)\left[y_{k+1} - y_k\right]$$

$$\Rightarrow F'\left(y_k\right)P^{-1}P\left[y_{k+1} - y_k\right] = -F\left(y_k\right)$$

P should be invertible and make the linear solve easier

- We have developed a block preconditioner based on the shallow water equation
- Blocks based on velocities and height, coupling between is relaxed
- FastMath's CAM-SE-Trilinos interface is used for this work
- Formulations should carry over to variable resolution grids

Optimization of implicit preconditioner

SciDAC Multiscale

Accomplishments:

Code optimizations for preconditioner

Interfaced Helmholtz Solver with Trilinos

Performed initial run on variable res grid

Tested preconditioner with Piro interface

Current Focus:

Implementing a multilevel solver for S solves

Building matrices for multilevel solvers in Trilinos

Testing on variable resolution grids

Extending preconditioner to 3D

Total Number of Unknowns	Original Implementation Total Time	New Implementation Total Time	Total Nonlinear Iterations	Total Linear Iterations	Total F Solve Iterations	Total S Solve Iterations
1,916,928	492	298	240	491	1870	14069
4,313,088	906	797	240	489	2237	20725
7,667,712	1906	1696	240	484	2235	28372

Jablonowski-Williamson Test case 5; 1 sim. day; Implicit method with Crank-Nicholson Time Stepping & 720s step size

Code optimizations produced some speedup in runtime

Main Cost: Schur Complement Solve -Developing Multilevel Solver

A. Lott, C. Woodward, K. Evans, A. Salinger, M. Taylor

Performance and GPU acceleration

Goal: Develop CAM-SE kernels providing GPU acceleration on next generation architecture

Accomplishments:

- Common build system that links Trilinos to CAM-SE **and** the GPU kernel development.
- Profile analysis to identify data and computation bottlenecks limiting performance of the CAM-SE/GPU implicit solve

Still to do:

- Provide algorithmic/code optimization to minimize data transfer through TITAN's hierarchy and computation for CAM-SE/GPU implicit solve

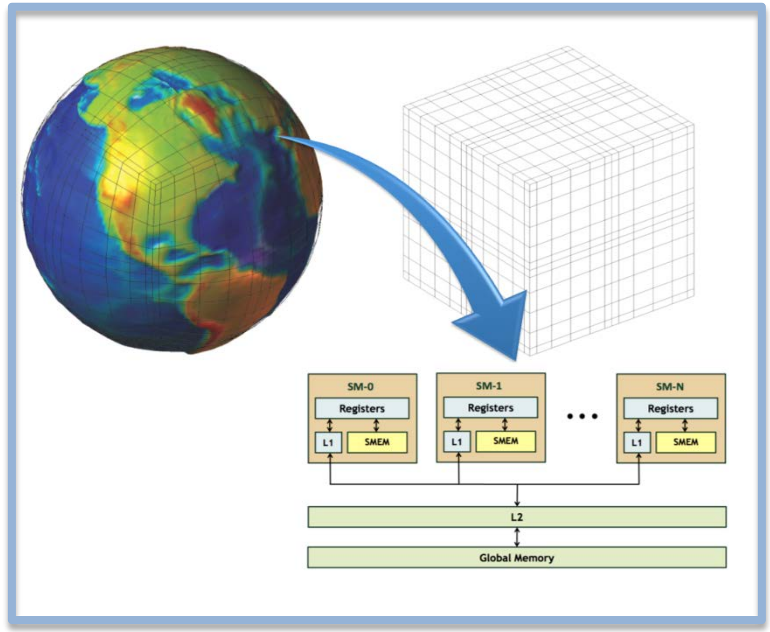

R. Archibald, M. Norman, K. Evans, C. Baker

Ocean: MPAS-O Optimization

Variable Mesh Dycores

Model for Prediction Across Scales (MPAS)

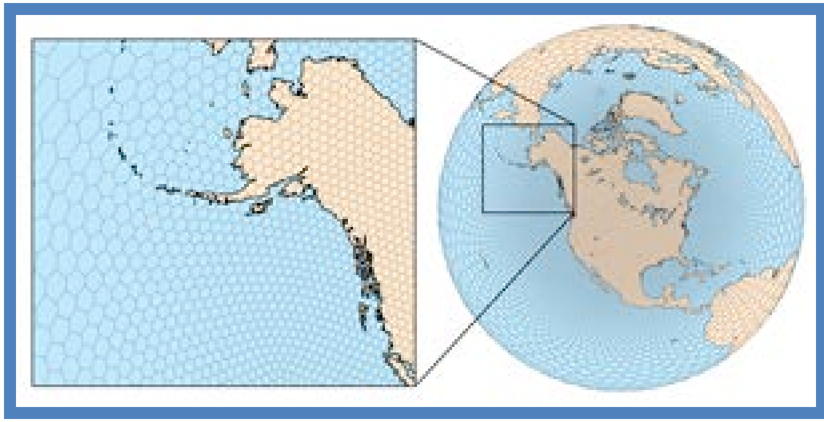

Spectral Element Dycore

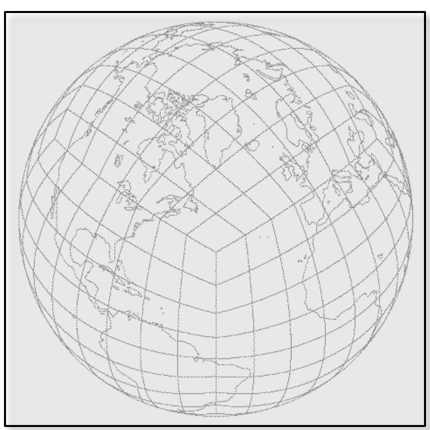

Physics-Dynamics Interface

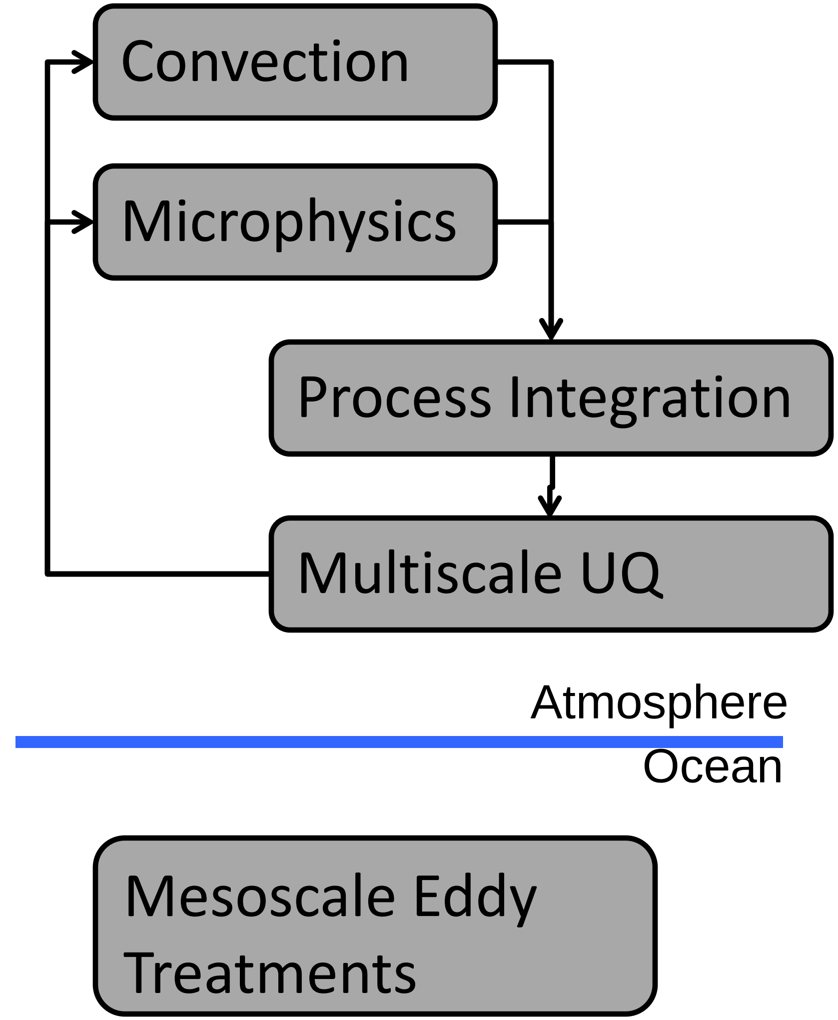

MPAS Ocean/Atmosphere Dynamical Core

1. MPAS is an unstructured-grid approach to climate system modeling.

2. MPAS supports both quasi-uniform and variable resolution meshing of the sphere using quadrilaterals, triangles or Voronoi tessellations.

3. MPAS is a software framework for the rapid prototyping of single-components of climate system models (atmosphere, ocean, land ice, etc.).

4. MPAS offers the potential to explore regional-scale climate change within the context of global climate system modeling. Multiple high-resolution regions are permitted.

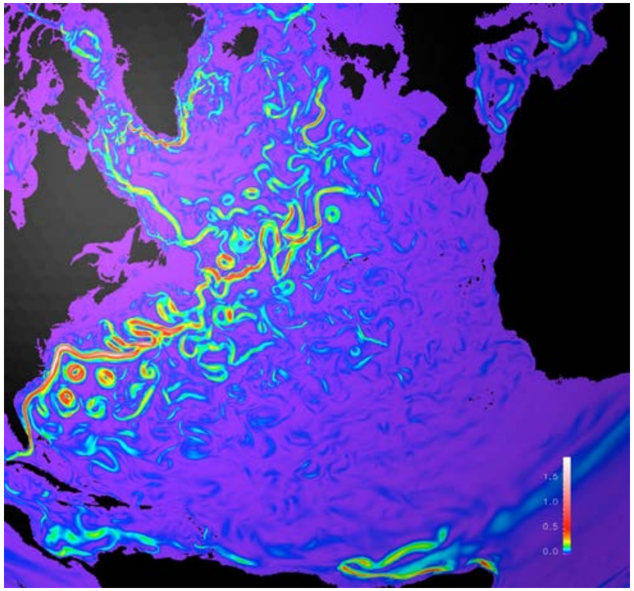

snapshot of kinetic energy from a global ocean simulation with 7.5 km resolution in the North Atlantic.

Ringler, 2012

MPAS
Model for Prediction Across Scales

The variable resolution mesh in MPAS

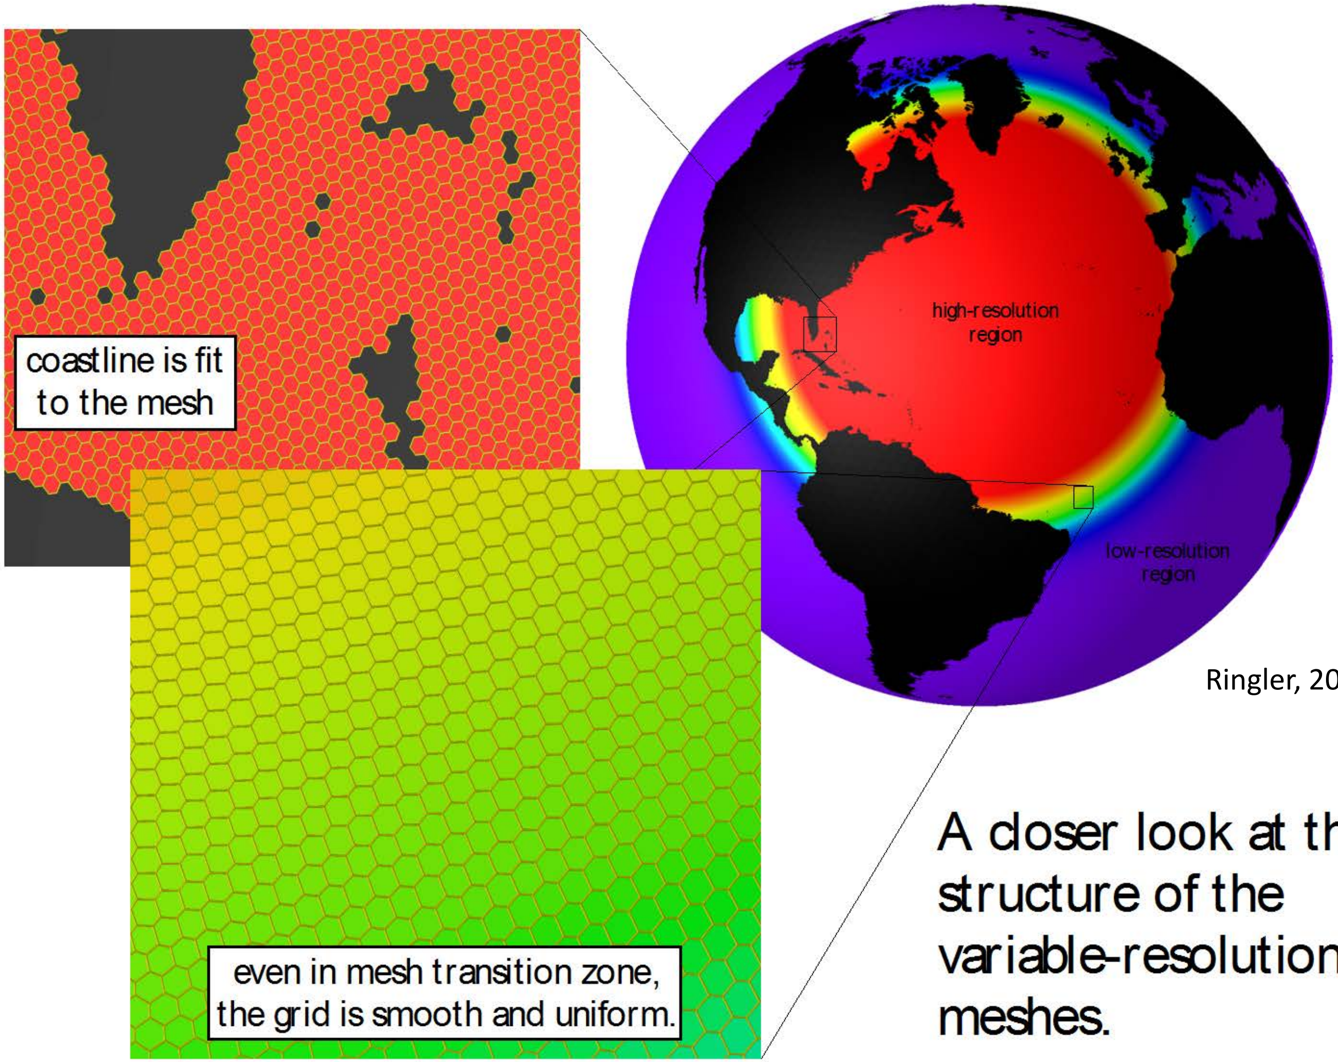

Ringler, 2012

A closer look at the structure of the variable-resolution meshes.

Numerical control of refinement artifacts

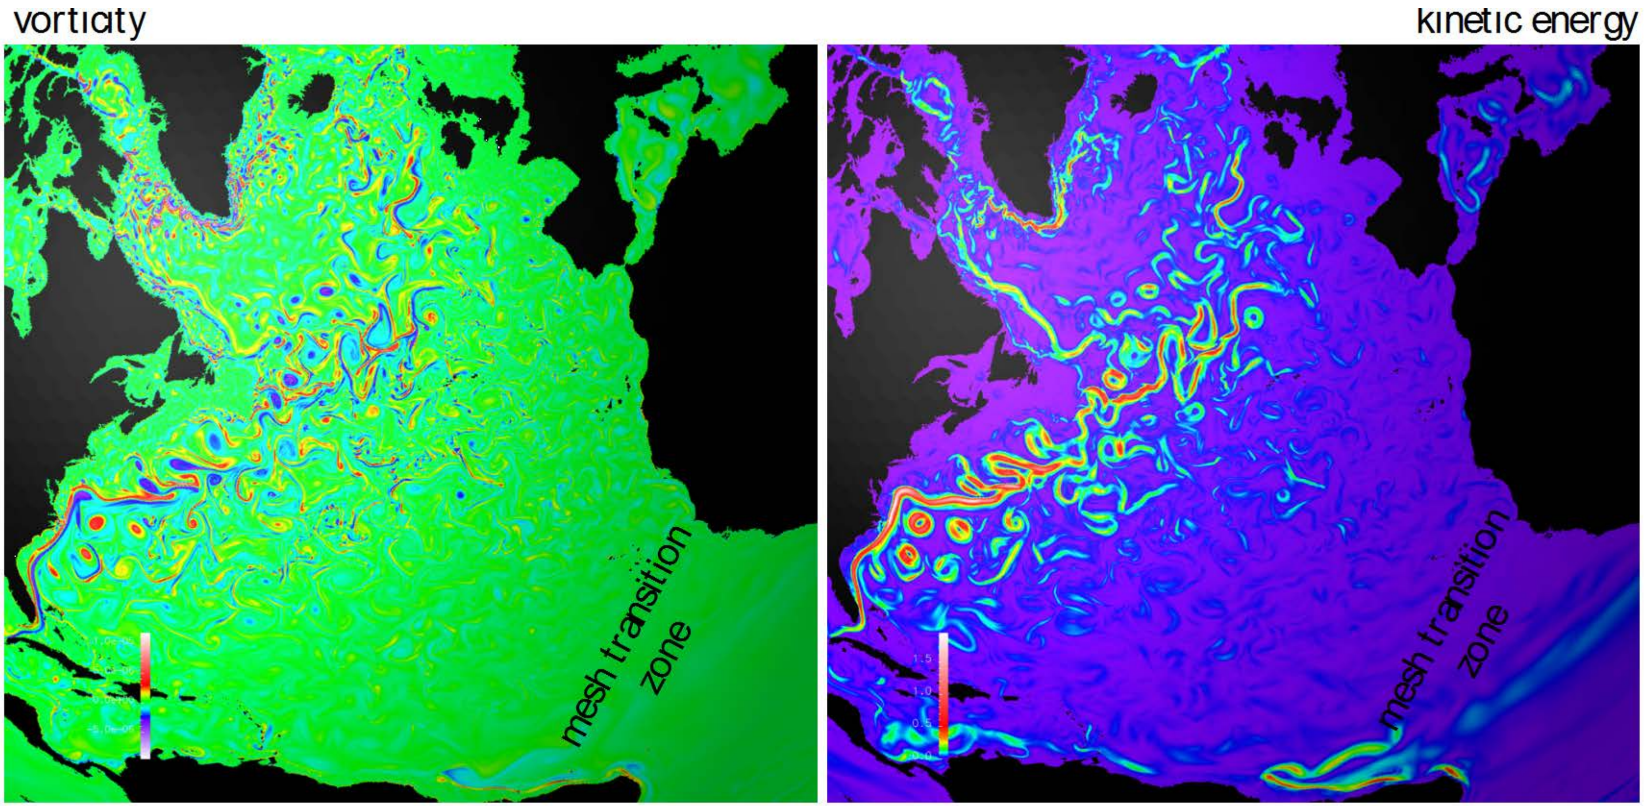

Ringler, 2012

We have sufficient evidence in shallow-water, 3D atmosphere and 3D ocean systems to declare this a non-issue as far as the dynamics are concerned.

Multiscale / SUPER Collaboration

- MPAS-Ocean has scaling limits using MPI only
- Adding OpenMP threading to efficiently leverage multicore systems
- TAU (SciDAC-3 SUPER) was used to evaluate MPI+OpenMP approach, suggest improvements
- Tests performed by Doug Jacobsen (LANL), Kevin Huck and Sameer Shende (U. Oregon)
- Observations, optimizations found:
  1. MPI block decomposition + OpenMP element decomposition reduces total instructions in computational regions (~10% faster) when compared to MPI block decomposition alone
  2. New weighted block decomposition using vertical elements (depth) balances work across processes (~5% faster)
  3. Guided schedule balances work across threads (~6% faster)
  4. Overlapping communication and computation will reduce synchronization delays when exchanging halo regions (underway)
- Evaluation of OpenMP implementation is ongoing coupled with new TAU development. Next steps include porting to MIC platform.

MPAS_O Performance Improvements

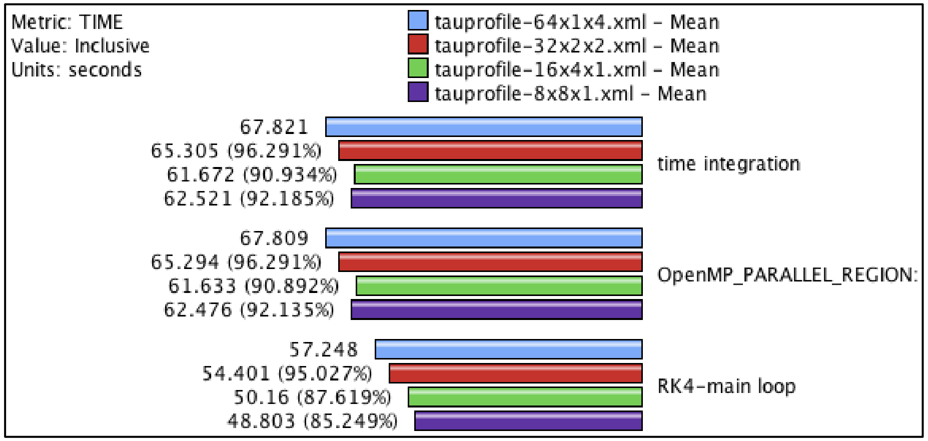

1. Increasing thread concurrency with constant resources (64 cores used in all cases).

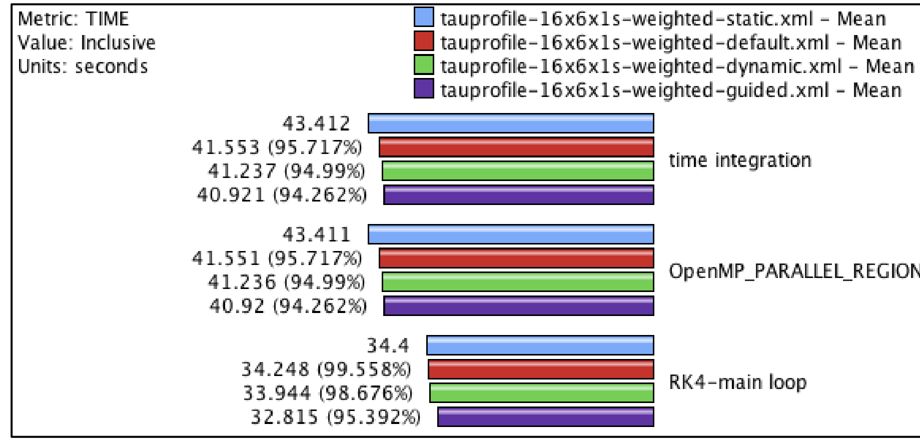

3. Adjusting thread schedule balances workload, reduces computation time

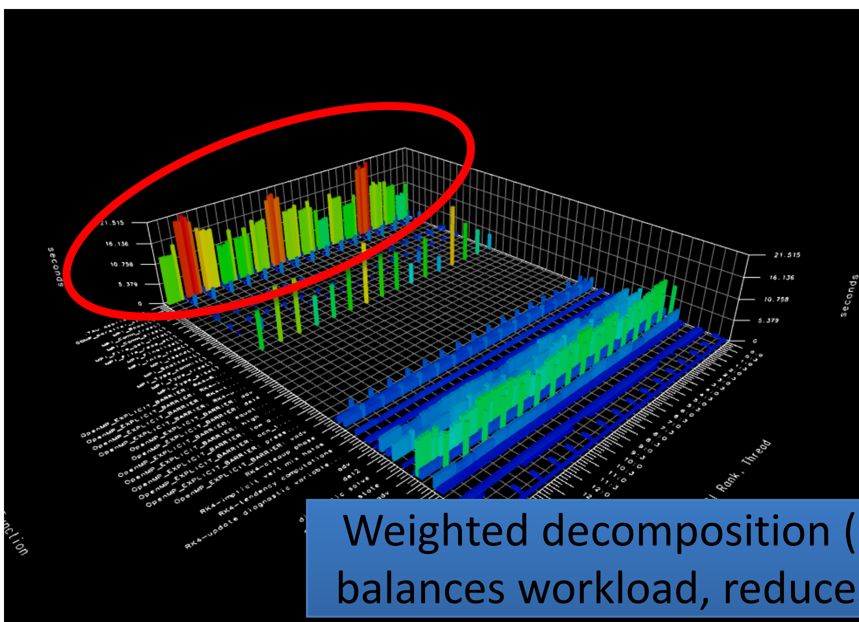

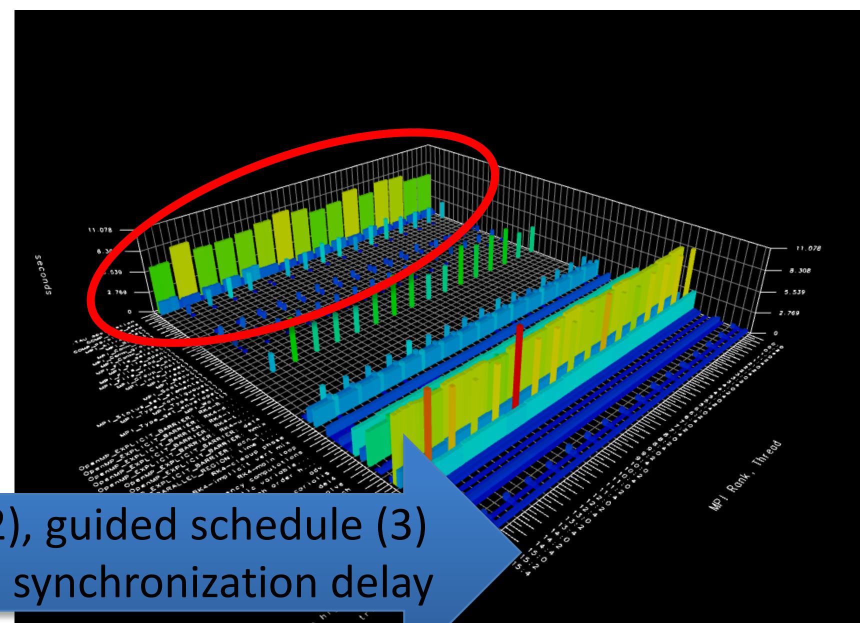

Weighted decomposition (2), guided schedule (3) balances workload, reduces synchronization delay

Process Integration: Physics Convergence

Variable Mesh Dycores

Model for Prediction Across Scales (MPAS)

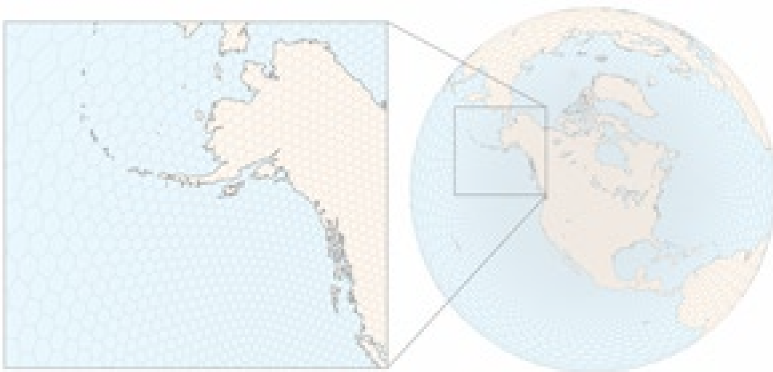

Spectral Element Dycore

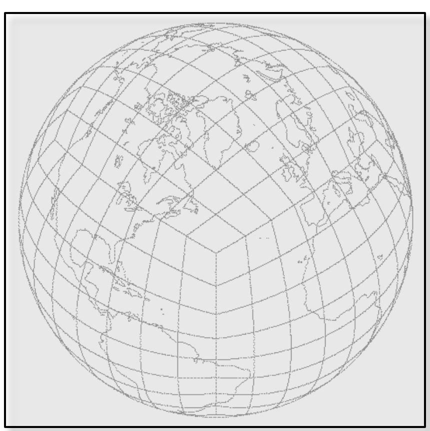

Physics-Dynamics Interface

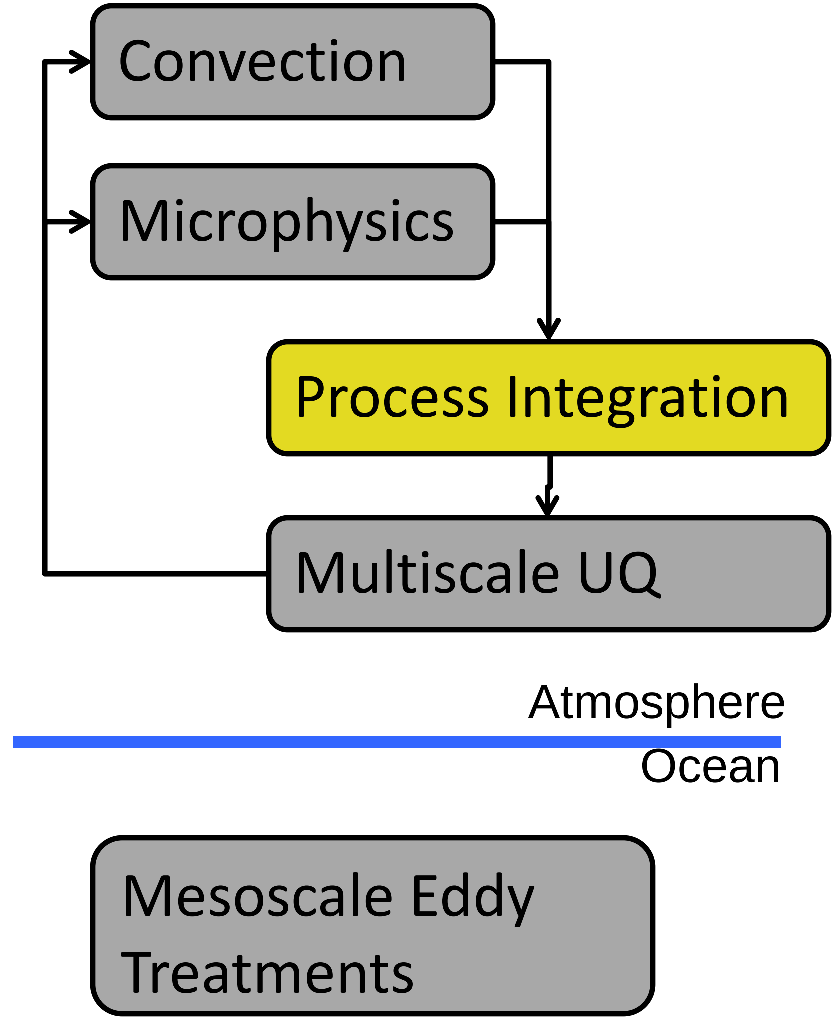

Time Step Sensitivity of CAM

- Global simulations showed strong time step sensitivity, with liquid water path and cloud cover increasing as timestep decreased
- As a proxy for the full 3D model, we are compiling 10 single column case studies spanning a variety of climate regimes to better explore time step sensitivity and convergence in CAM
- Initial results with timestep varying from 1 to 20 min show sensitivity similar to global model results

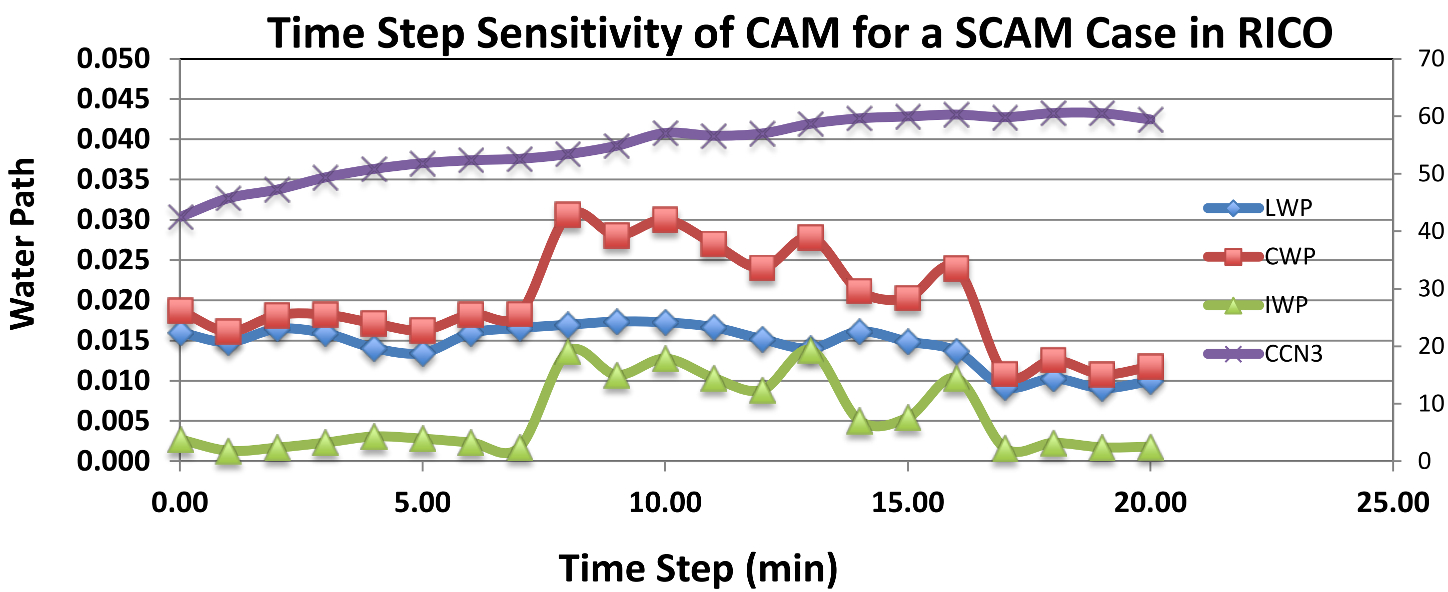

Process Integration

Importance:

Time-integration issues corrupt physics code and cause Δt sensitivity

Progress:

1. Discovered time-stepping issue in PBL scheme (see fig)
   a. developing idealized model to isolate the essence of the problem
2. Used a suite of single-column simulations to clarify time-convergence issues in important climate regimes
   a. confirmed q_l increase w/ decreased Δt found in global simulations

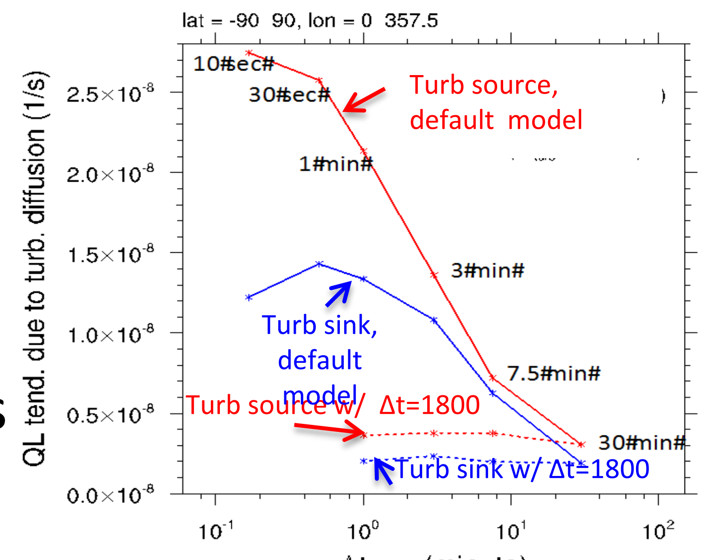

Fig. Global-ave turbulent liquid water tendency as a function of Δt in CAM5 (solid) and in a variant with turbulent Δt=1800 s regardless of model Δt (dashed)

3. Identified a new variant of operator splitting (sequential-tendency splitting) which maintains better process coupling yet allows for long model timesteps

UQ: Toward Multi-fidelity Calibration

Variable Mesh Dycores

Model for Prediction Across Scales (MPAS)

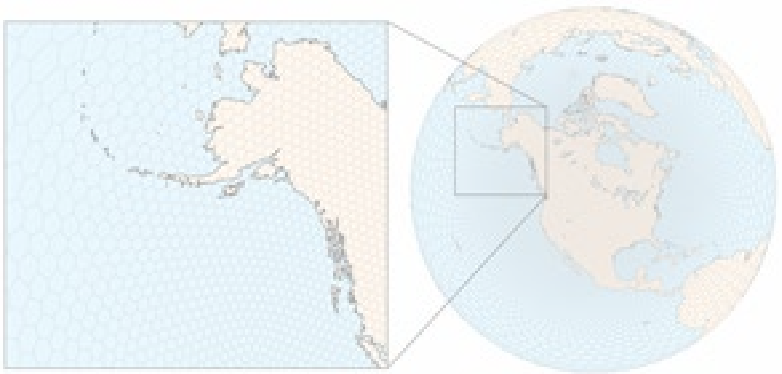

Spectral Element Dycore

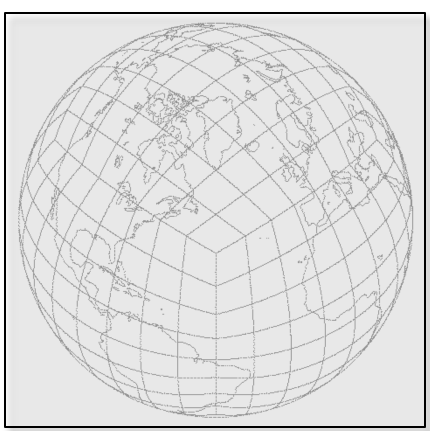

Physics-Dynamics Interface

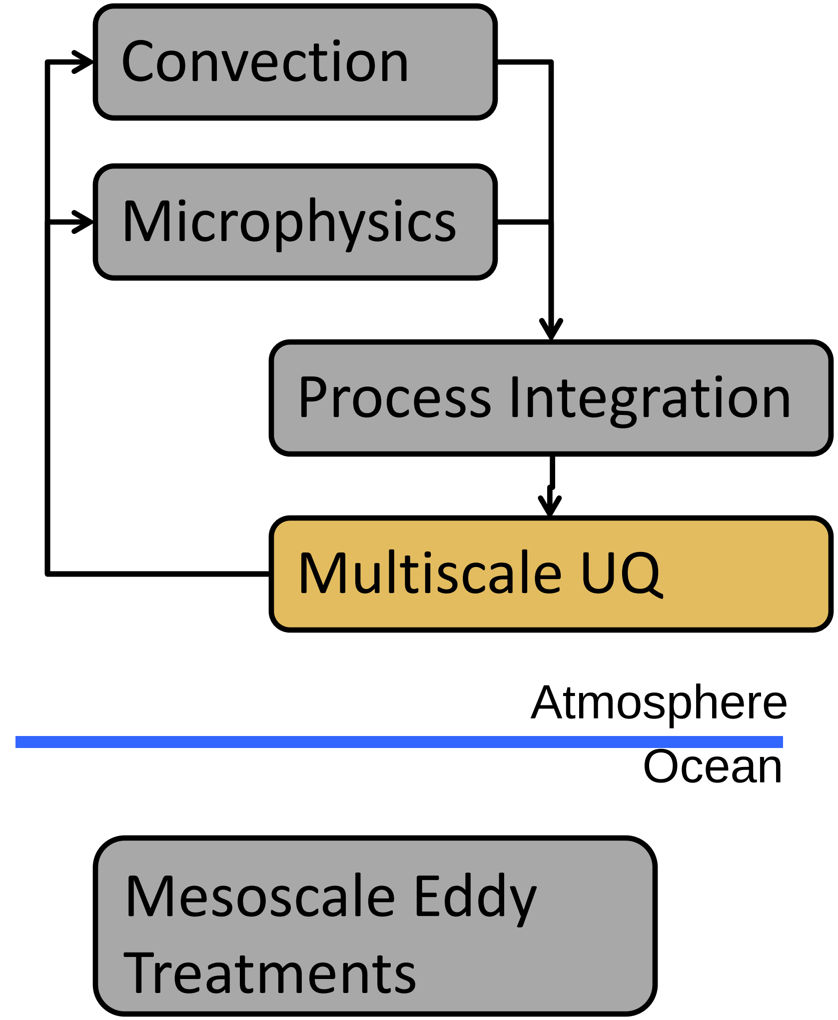

Sensitivity to sub-cycle physics parameter variations

Goal: determine the sensitivity of CAM5 to changes in 6 sub-cycle parameters

Challenge: there are 22,500 unique configurations among these 6 parameters

Approach: ran 272 Latin hypercube ensemble simulations and used random forests to calculate sensitivities

Random forests (RFs) are popular in supervised machine learning (Breiman 2001)

RFs use randomized decision trees to partition the input space (see example on right)

RF predictions are based on an ensemble of trees

RFs assess sensitivities via feature selection

Example of a 2D decision tree

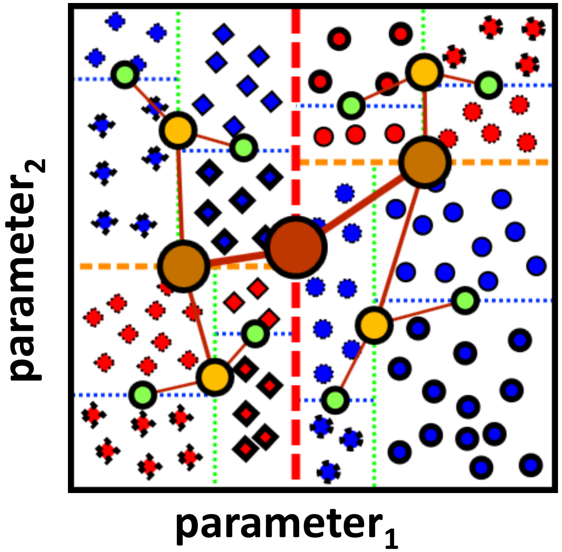

UQ: sub-cycle parameter sensitivities

Year 1: define metrics to quantify model development changes, identify scale-dependent physics parameters, and begin sensitivity studies at coarse model resolutions

Sensitivity to discrete sub-cycle physics parameters

Used Latin hypercube to sample discrete values of 6 sub-cycle physics parameters.
Calculated model sensitivity to changes in sub-cycling parameters.
Demonstrated statistical method to model changes from discrete factors.

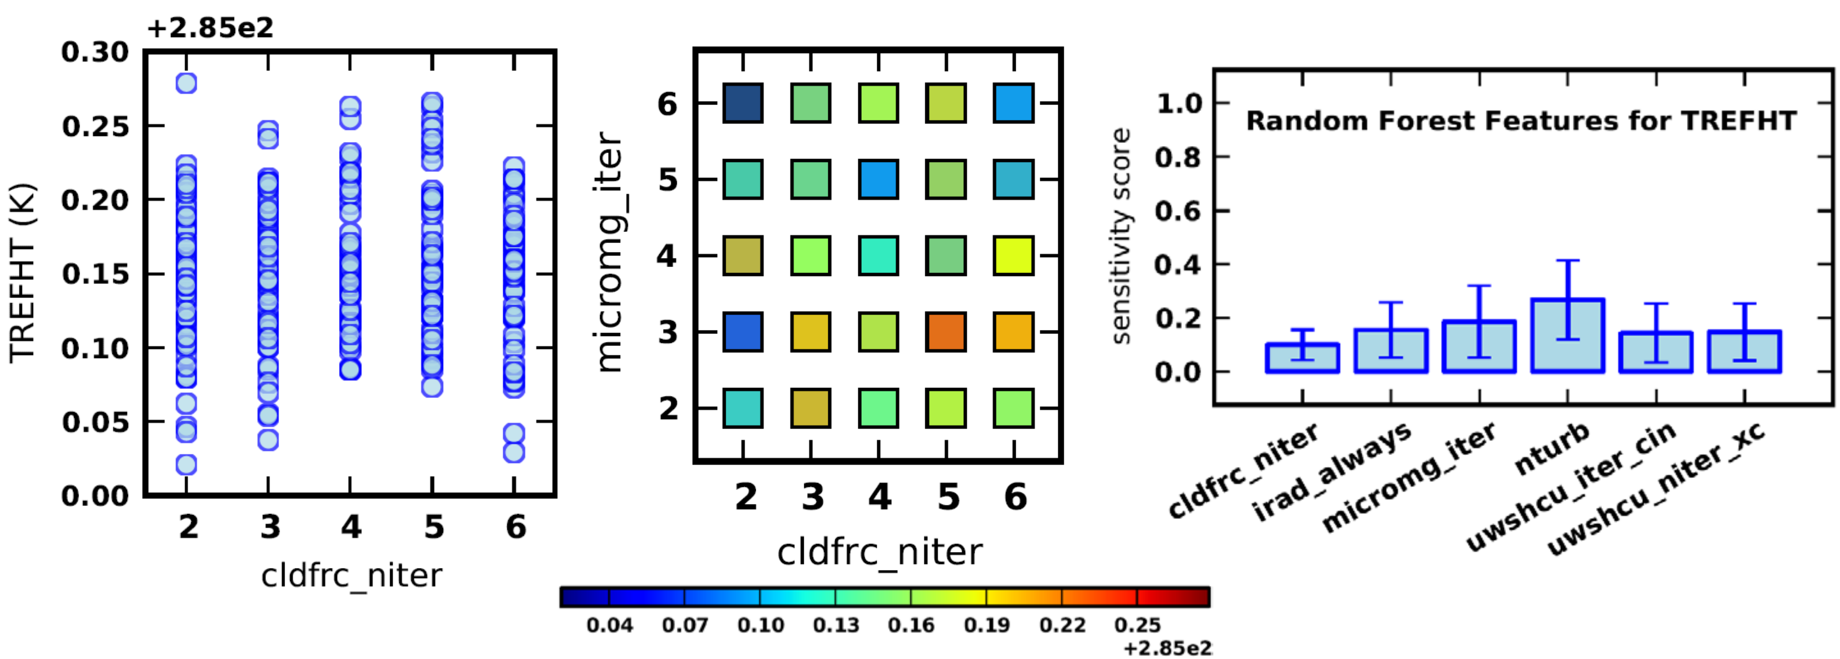

UQ: toward multi-fidelity calibration

Year 1: define metrics to quantify model development changes, identify scale-dependent physics parameters, and begin sensitivity studies at coarse model resolutions

Progress toward multi-fidelity calibration

Began design of multi-fidelity UQ studies

Adapted UQ interface to 3-4 uniform resolutions (f45, f19, f09, f05)

Implemented scale-dependent UQ parameters in new release CESM v1.2

Leverage low fidelity simulations to calibrate high fidelity model

(2^o CAM5) (1^o CAM5) ($½^o$ CAM5)

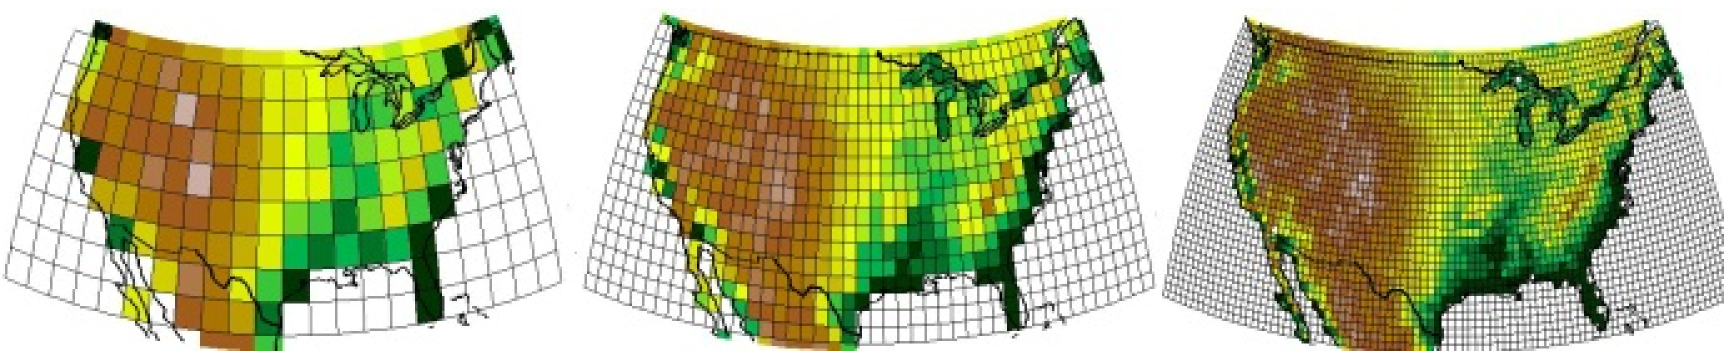

$$\underbrace{Y_{high}(p)}_{\text{high resolution feature}} = \overbrace{\tilde{Y}_{low}}^{\text{low resolution surrogate}} + \underbrace{\delta_{high\text{-}med} + \delta_{med\text{-}low}}_{\text{resolution discrepancies}} + \overbrace{\delta_{low}}^{\text{surrogate discrepancy}} + \underbrace{\delta_{high\text{-}obs} + \varepsilon_{obs}}_{\text{observation discrepancy and error}}$$

UQ: Sub-grid Physics Quadratures

Variable Mesh Dycores

Model for Prediction Across Scales (MPAS)

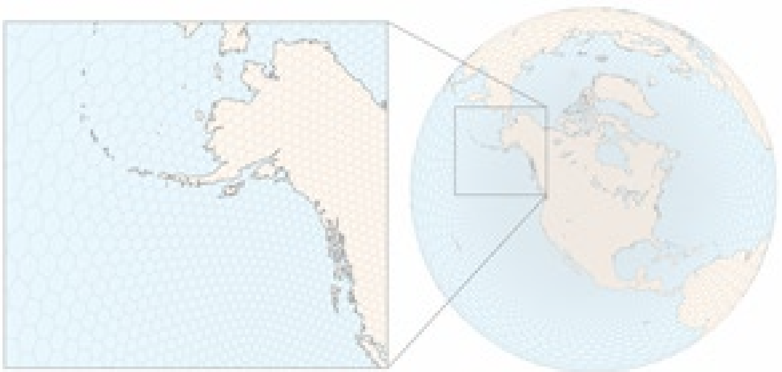

Spectral Element Dycore

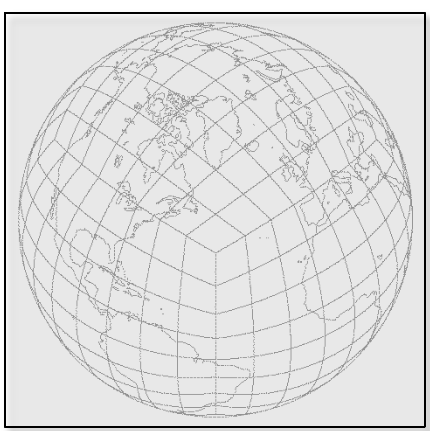

Physics-Dynamics Interface

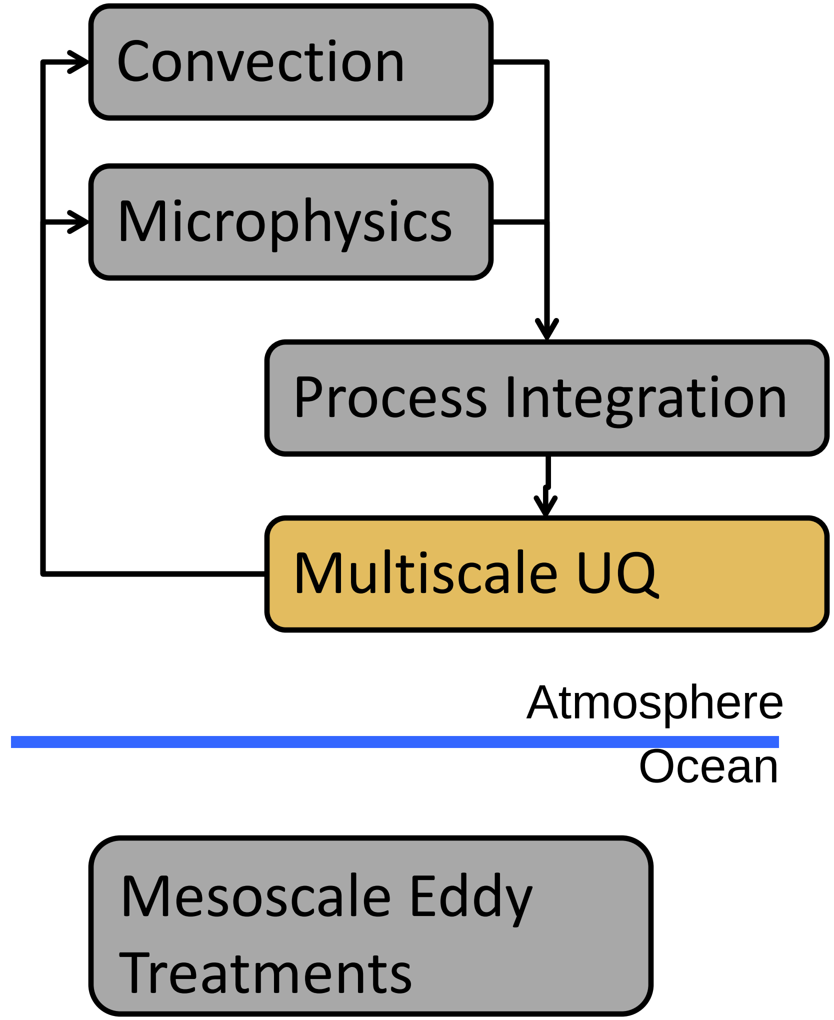

Better sub-grid physics PDF quadrature

Kenny Chowdhary (SNL), Bert Debusschere (SNL), Vince Larson (UWM)

Autoconversion – conversion of cloud droplets to rain droplets (measured as rate of mass transfer).

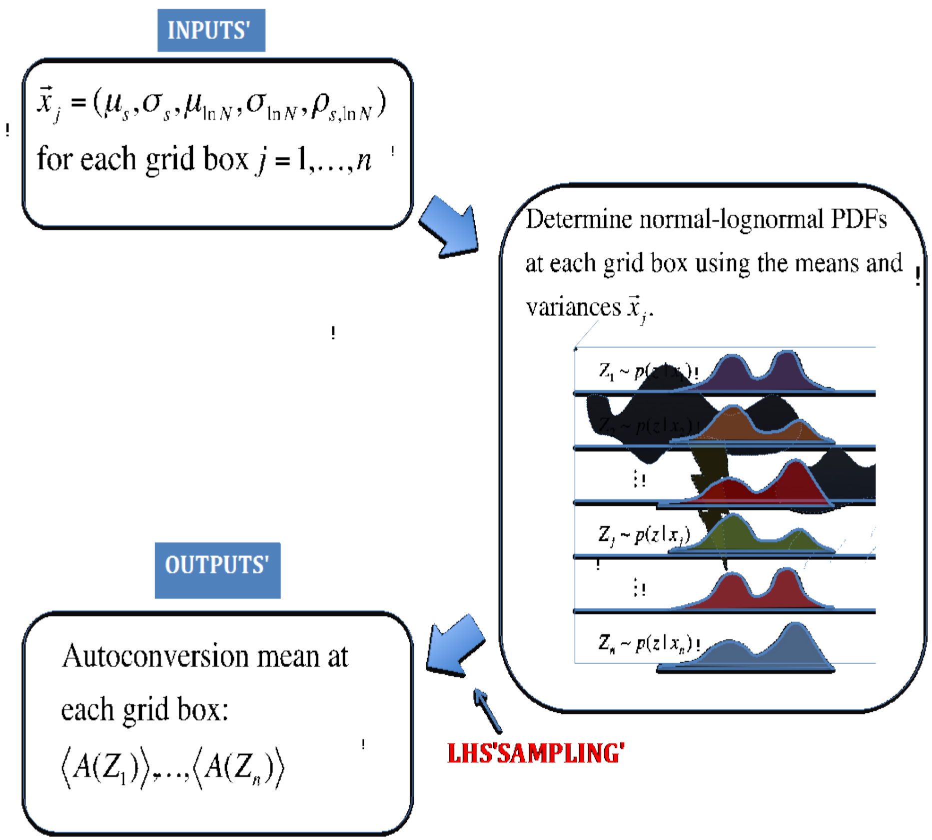

- Autoconversion depends on the cloud water mixing ratio, s, and the cloud droplet concentration N.
- In the CLUBB model, It is assumed that s and N take on a joint normal-lognormal probability density function (PDF), which depends on 5 parameters (the means and variances) in x_j.
- Currently, at each time step, and every grid box, the normal-lognormal PDF is sampled in each grid box using Latin Hypercube sampling.
- The samples are then used to compute the Autoconversion mean at each grid box and time step.

Better sub-grid physics PDF quadrature

- We can replace Latin Hypercube sampling with a quadrature based approach.

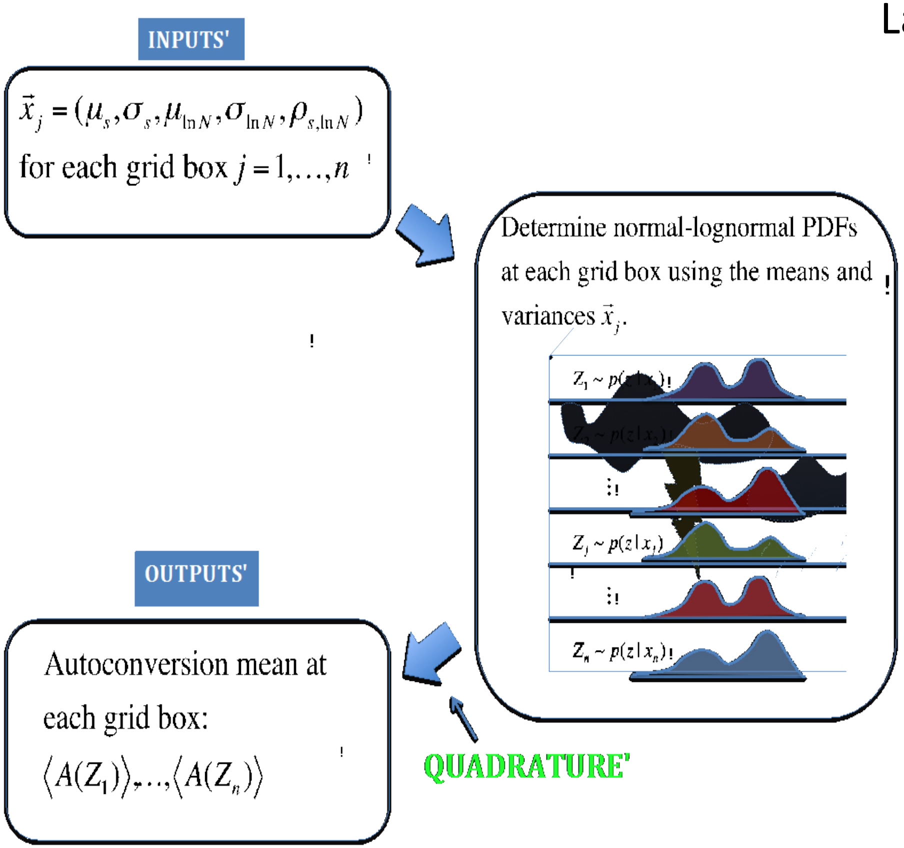

Latin Hypercube Sampling vs Quadrature

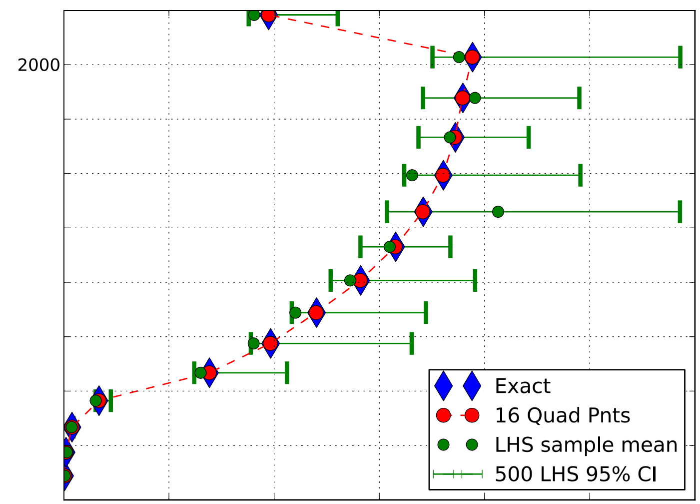

Using a quadrature approach, we can bypass the random sampling and calculate the Autoconversion mean using far fewer points, with even greater accuracy.

Selected interactions with SciDAC Institutes

- FastMATH: Implicit solver introduction and optimization for variable grids
- SUPER: Performance portability and acceleration of ocean dycore
- QUEST: Acceleration of quadratures for sub-grid integrals over PDFs