The SciDAC2 CCSM Consortium Project

John B. Drake, Phil Jones

Kickoff Meeting: October 12, 2006, Boulder

Who are we?
(And what is SciDAC?)

- **Participating Institutions/Senior Personnel**
- ***Lead PI:* John B. Drake, Oak Ridge National Laboratory**
- ***Co-Lead PI:* Phil Jones, Los Alamos National Laboratory**
- Argonne National Laboratory (ANL) Robert Jacob
- Brookhaven National Laboratory (BNL) Robert McGraw
- Lawrence Berkeley National Laboratory (LBNL) Inez Fung*, Michael Wehner
- Lawrence Livermore National Laboratory (LLNL) Phillip Cameron-Smith, Arthur Mirin
- Los Alamos National Laboratory (LANL) Scott Elliot, Philip Jones, William Lipscomb, Mat Maltrud
- National Center for Atmospheric Research (NCAR) Peter Gent, William Collins, Tony Craig, Jean-Francois Lamarque, Mariana Vertenstein, Warren Washington
- Oak Ridge National Laboratory (ORNL) John B. Drake, David Erickson, W. M. Post*, Patrick Worley
- Pacific Northwest National Laboratory (PNNL) Steven Ghan
- Sandia National Laboratories (SNL) Mark Taylor

- **Scientific Application Partnerships**
- Brookhaven National Laboratory Robert McGraw
- Oak Ridge National Laboratory Patrick Worley
- Argonne National Laboratory Kotamarthi Rao *(contact Jay Larson)

- **Centers for Enabling Technology Collaborations**
- ESG - Dean Williams
- PERC – Pat Worley
- VIZ – Wes Bethel
- TOPS – David Keyes
- PRISMs – Dana Knoll

The Earth Climate System

The Grand Challenge problem is to predict future climates based on scenarios of anthropogenic emissions and changes resulting from options in energy policy

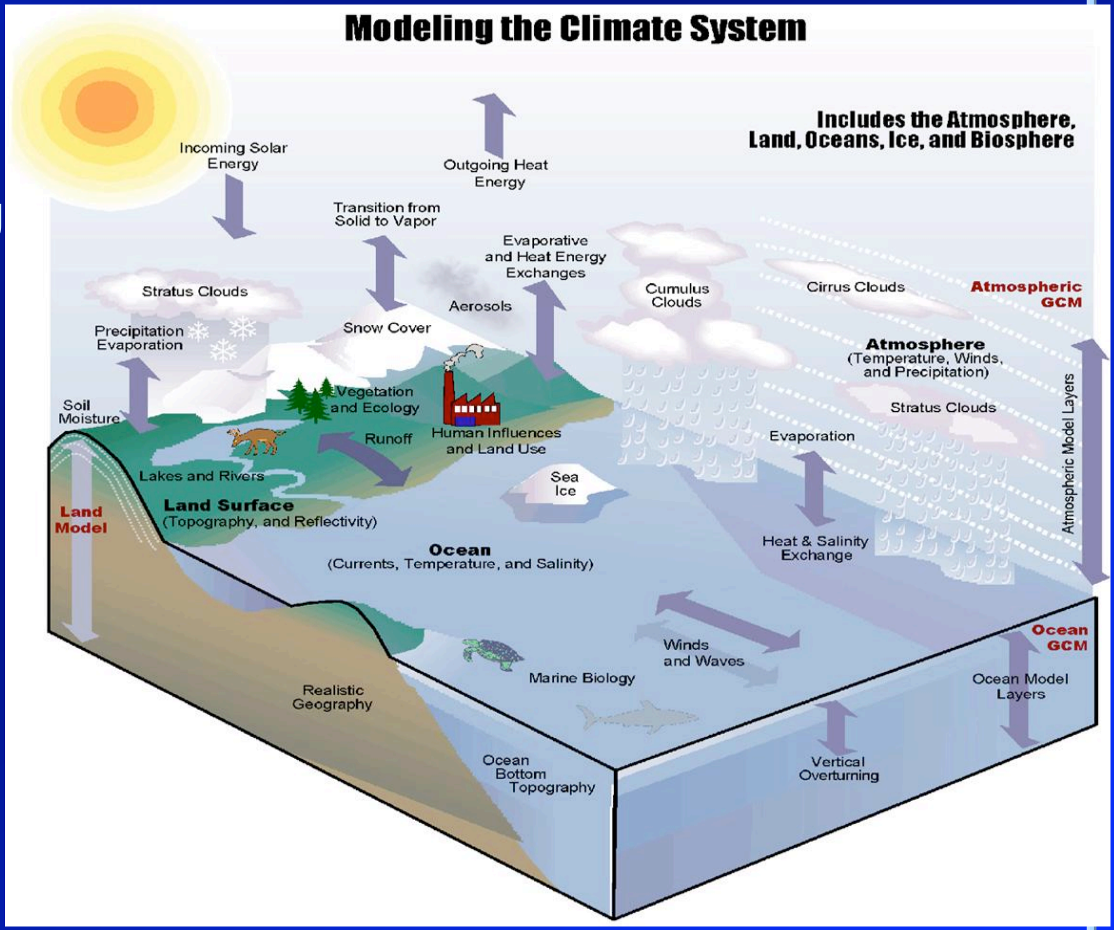

Finding: Natural climate variation does not explain the recent rise in global temperatures

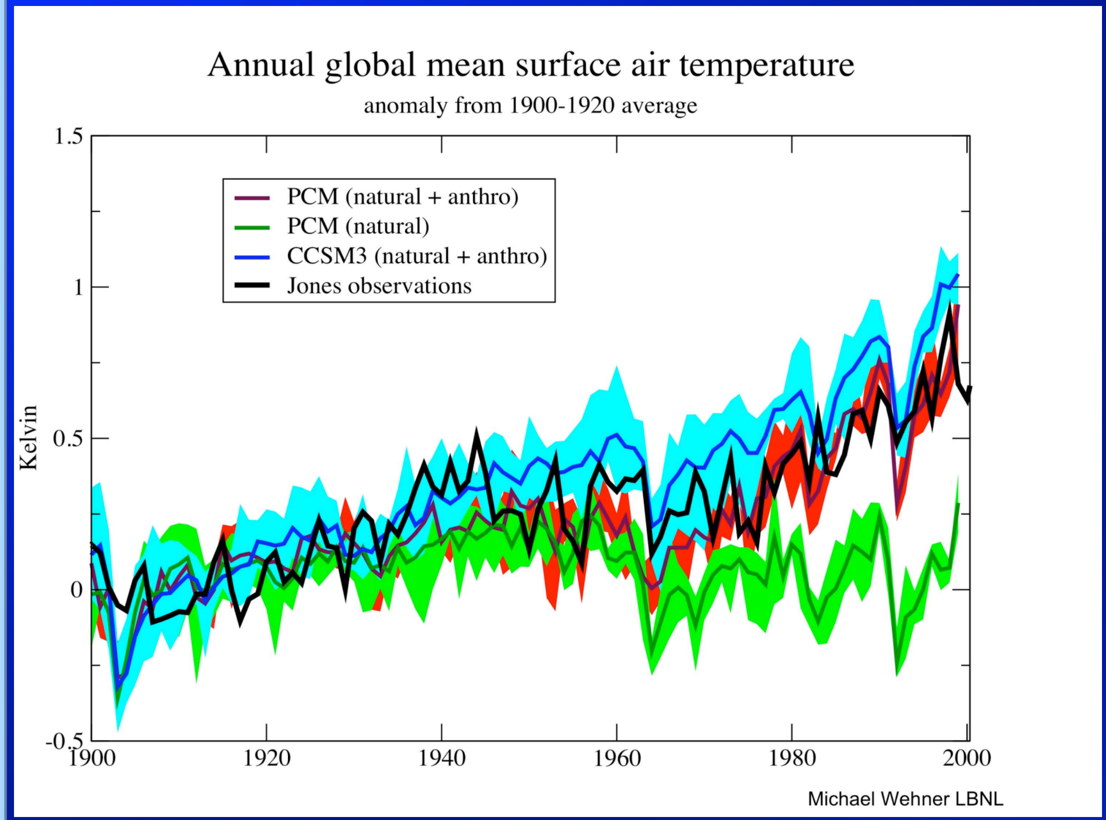

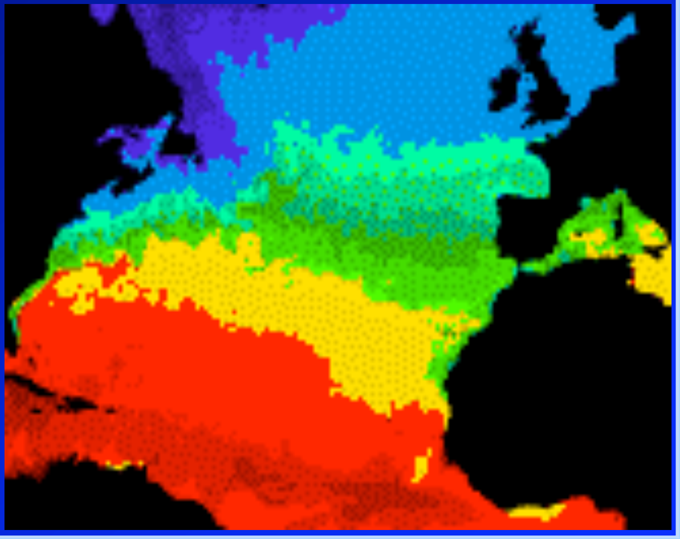

Why is it Important?

To the science/engineering community

- Discoveries of feedbacks between ecosystems and climate
- Fundamental science of aerosols effect in the atmosphere
- Advances in modeling and simulation science for climate prediction

To the public

- US Energy policy
- Contribution to international assessment of climate change and its causes

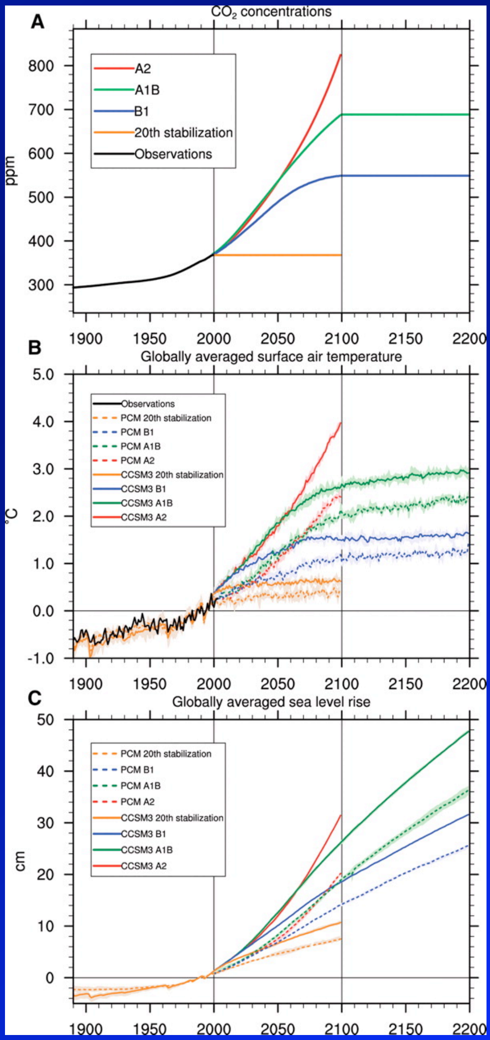

Arctic Thaw?

What does DOE want?

- Relation to Aerosol Science Program and Terrestrial Carbon Program
- Coordinated enterprise, relation to CETs and SAPs
- Reporting
  - Impacts and revised scope statement
  - 30 days: 4-6 slides for Dr. Orbach
  - 60 days: Management Plan, website, performance baseline
  - 6months: progress reports, highlights

DOE OBER Performance Targets
•2006 Deliver new measurements of clouds where observations are missing
•2007 Include realistic cloud simulations in a climate model
•2008 Measure ecosystem responses to climate change
•2010 Develop/validate new models predicting effect of aerosols on climate forcing
•2010 Provide climate model that links the Earth climate system with Earth's biological systems
•2013 Reduce differences between observed temperature & model simulations at sub continental scales using several decades of recent data
•2015 Deliver improved climate data & models for policy makers to determine safe levels of greenhouse gasses.

CCSM Development, the Climate End Station, and IPCC AR5: THE BIG PICTURE

- CES FY06 Allocation
  - 2 million CPU hrs on Phoenix Cray X1E
  - 3 million CPU hrs on Jaguar Cray XT3
- Need to coordinate 7 different CES subprojects

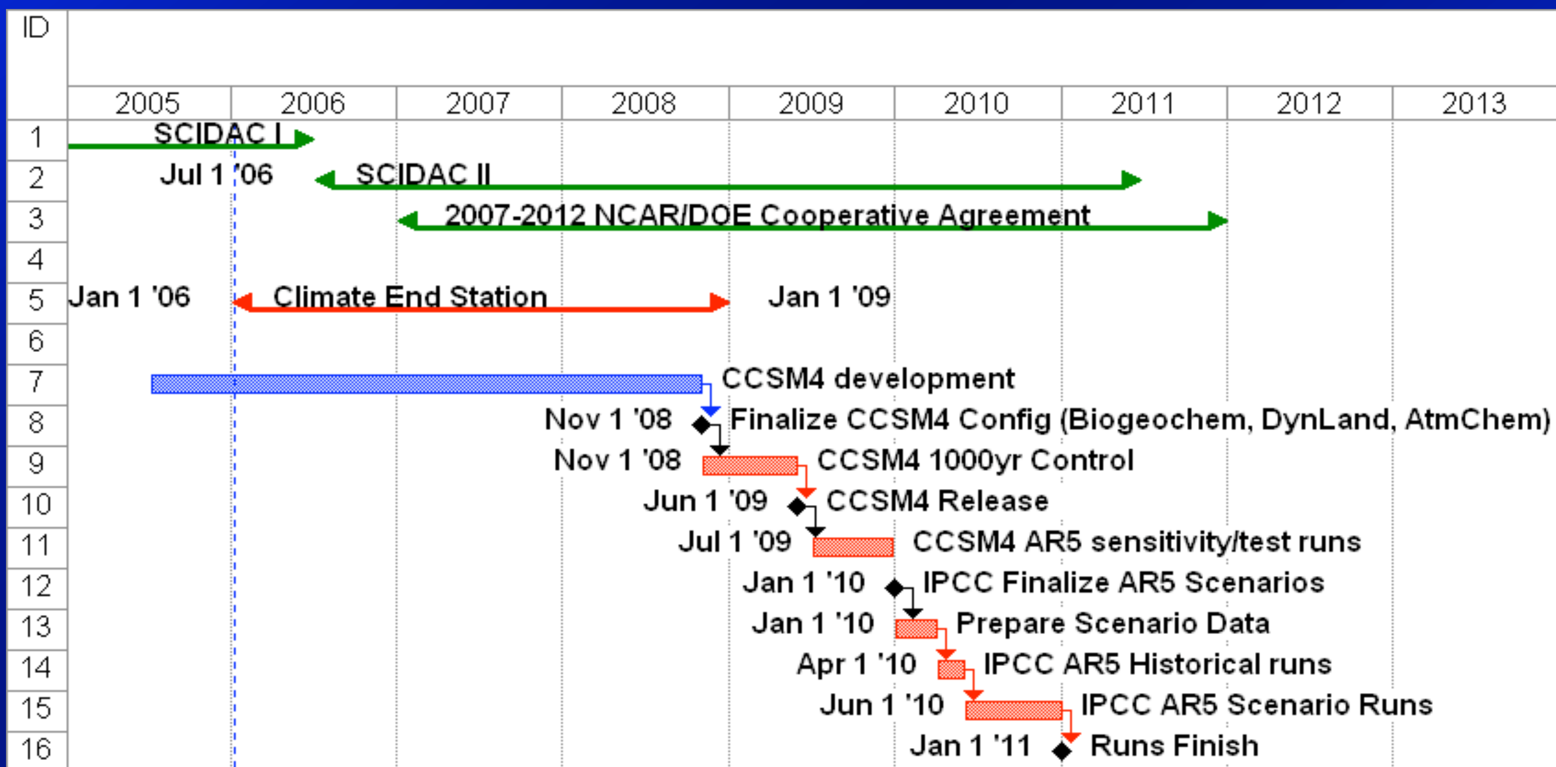

Proposal Targets

- Earth System Model
  - Terrestrial BGC and dynamic vegetation
  - Atm chemistry and aerosol dynamics
  - Ocn BGC
- Model Integration and Evaluation
  - Integration and unit testing
  - New cryosphere and ocean models
  - FV (cubed sphere), DG, others(icosahedral)
  - Frameworks for model evaluation
- Computational Performance
  - Scalablity, load balance, (fault recovery)

1st Generation Chemistry-Climate Model

- Components:
  - Processes for stratosphere through thermosphere
  - Reactive chemistry in the troposphere
  - Oceanic and terrestrial biogeochemistry
  - Isotopes of H_2O and CO_2
  - Prognostic natural and anthropogenic aerosols
  - Chemical transport modeling inside CCSM
- Prototype development:
  - SciDAC Milestone for 2005!
  - All pieces exist & run in CCSM3

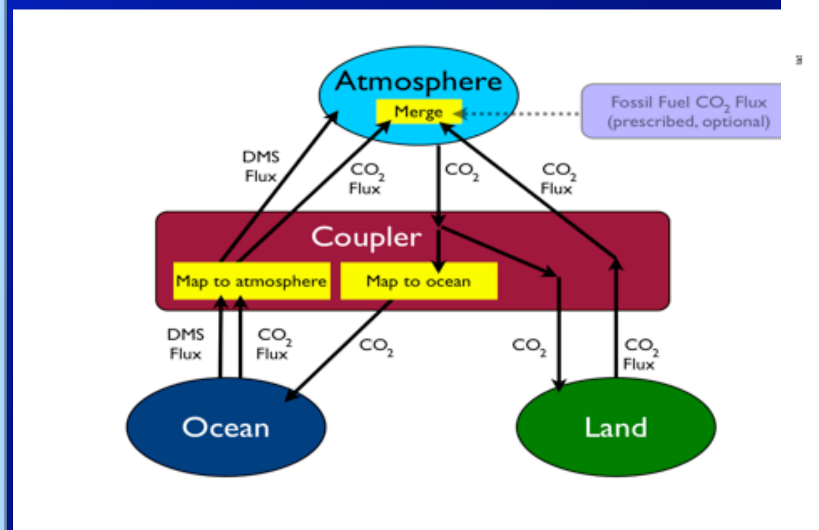

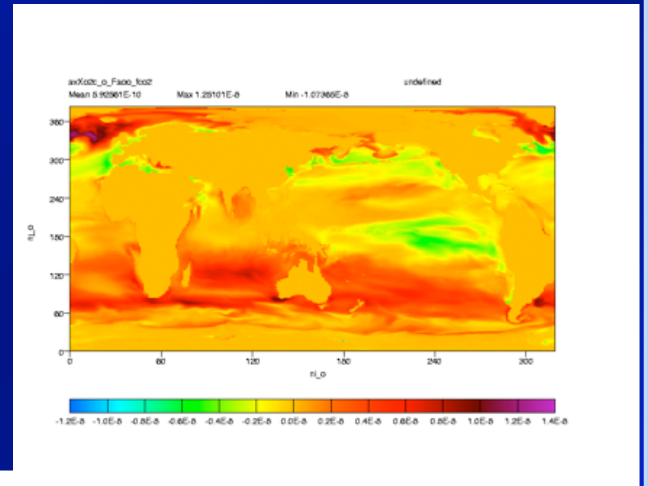

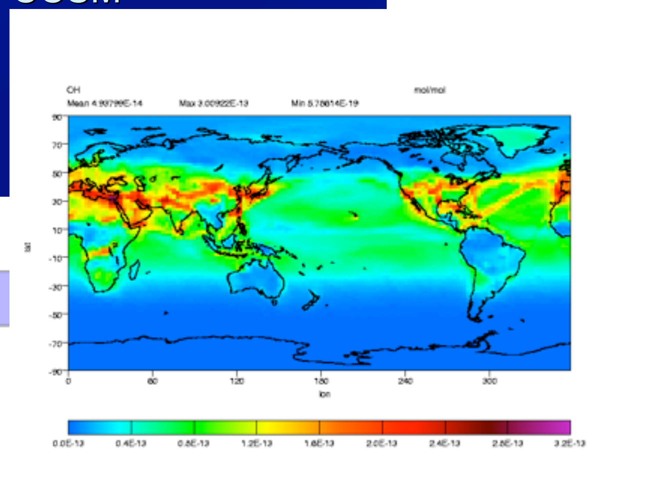

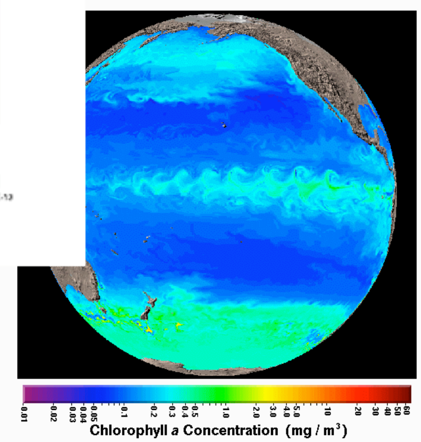

Maltrud, Shu, Elliot

Carbon Land Model Intercomparison (C-LAMP)

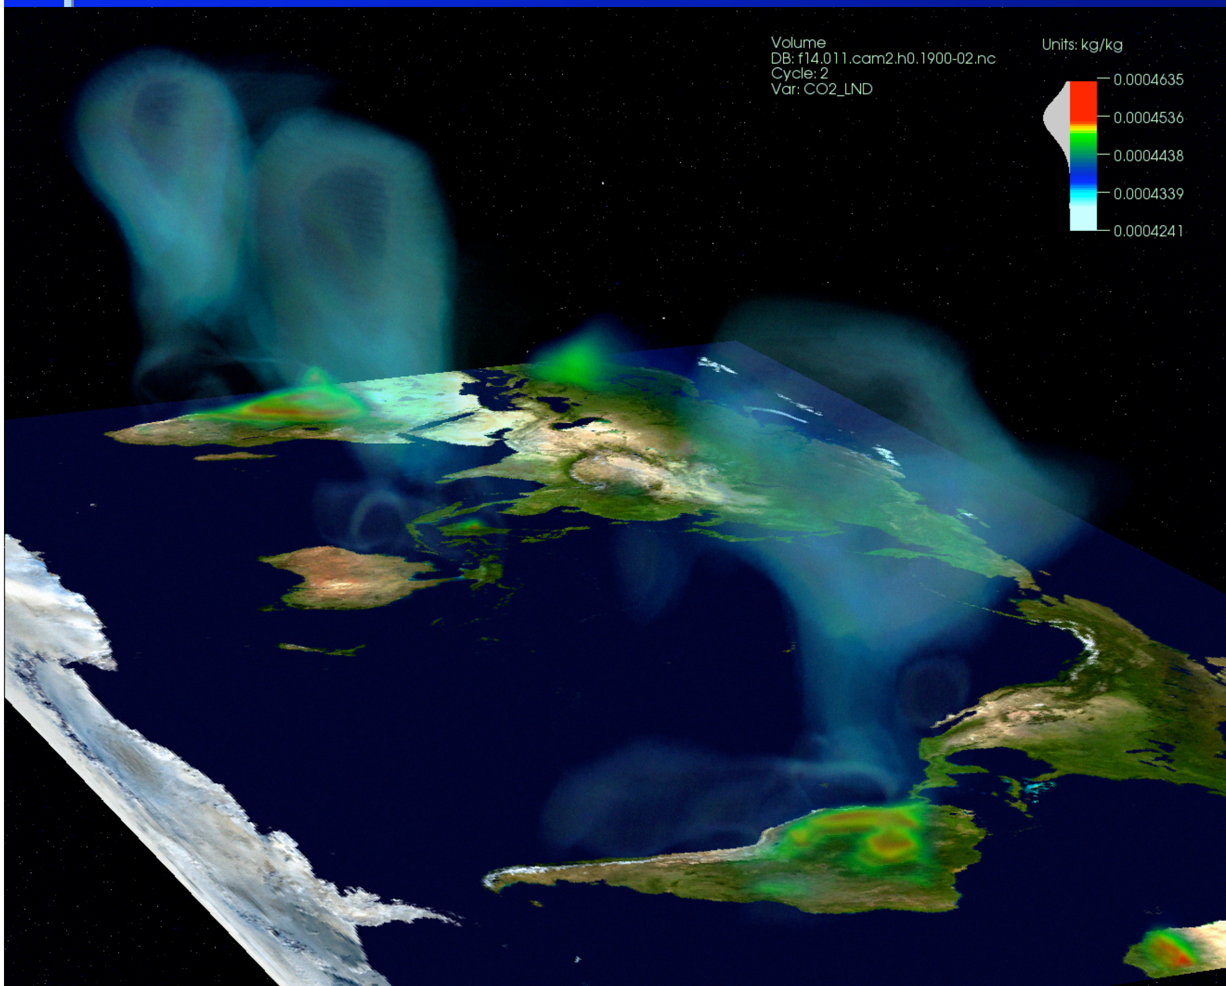

(courtesy J. Daniel)

- What are the relevant processes for carbon in the next version of the CCSM?
- Comparison of CASA', CN, and IBIS

Atmospheric Chemistry for A2 Scenario

J-F Lamarque, S. Walters

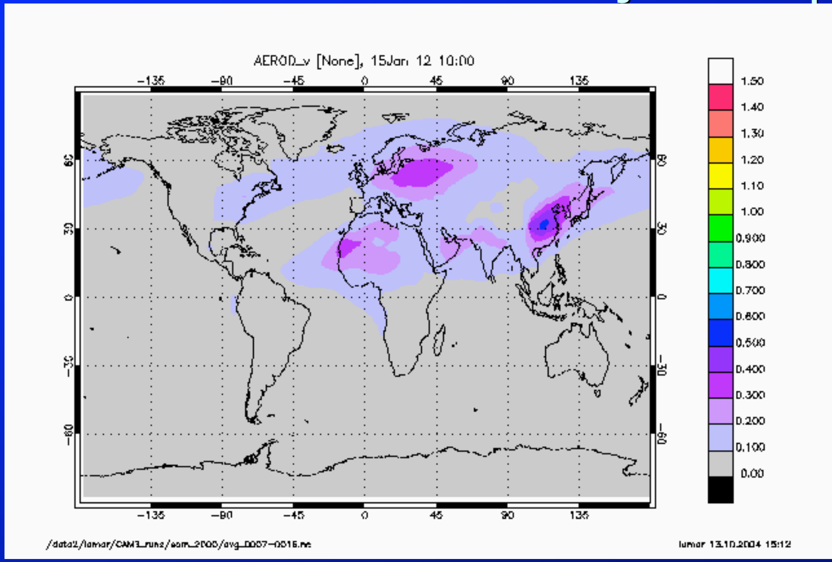

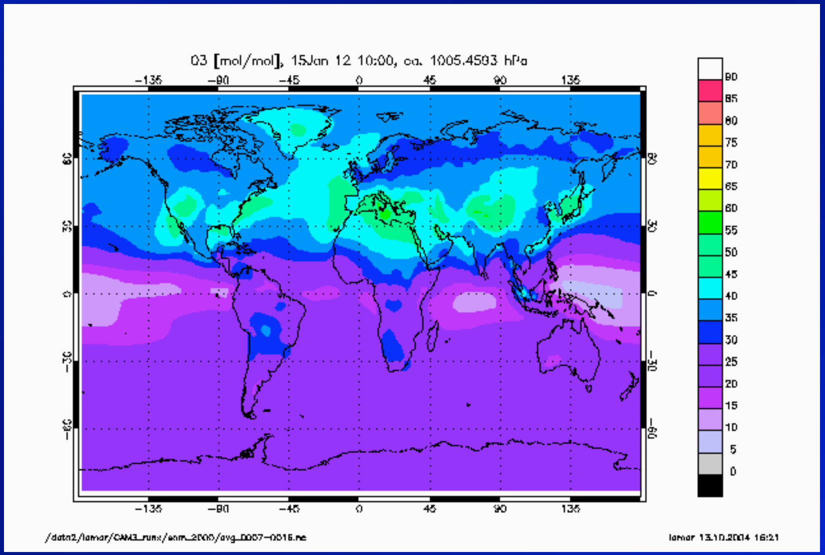

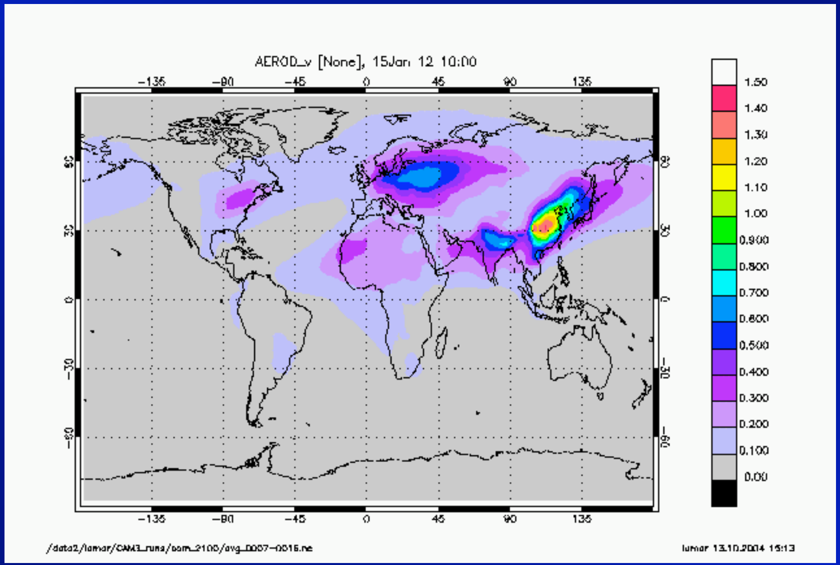

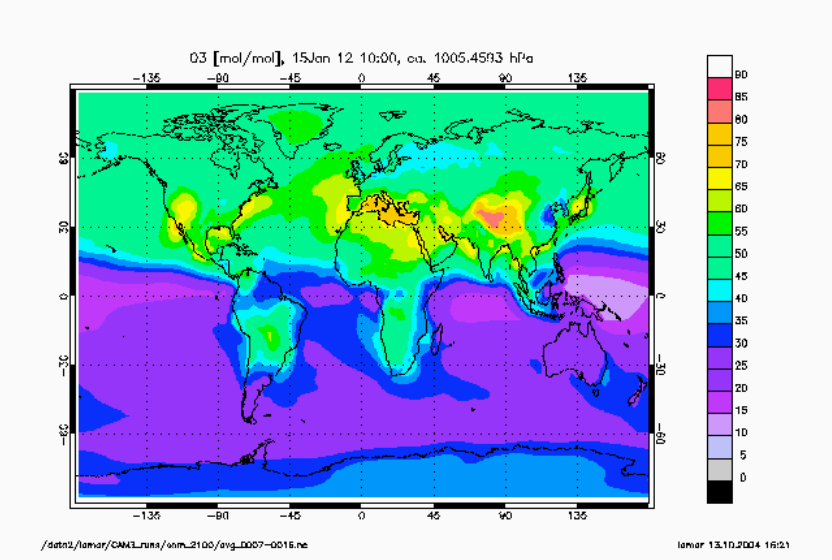

Multi-Century Coupled Carbon/Climate Simulations

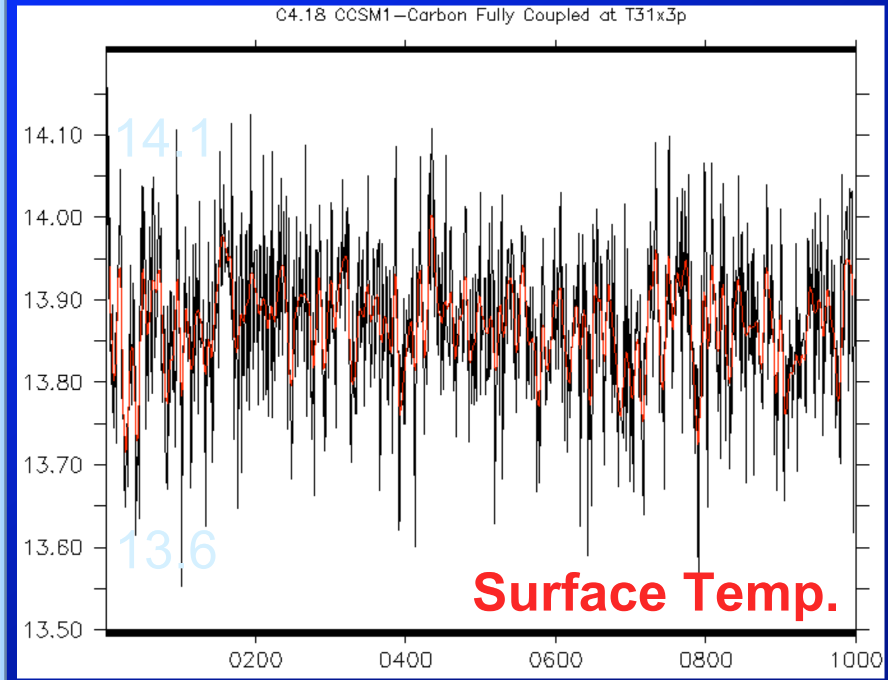

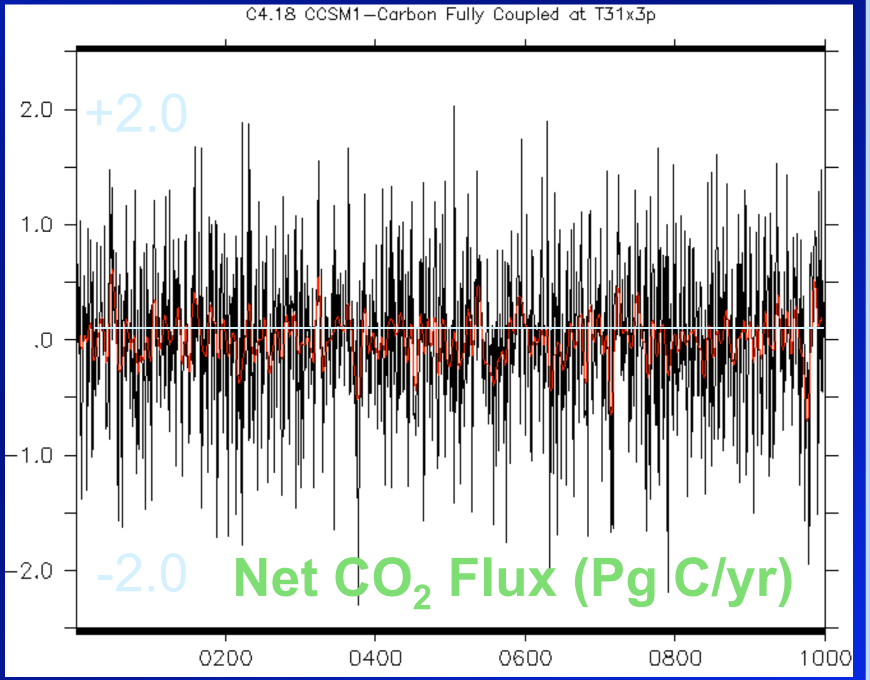

- Fully prognostic land/ocn BGC and carbon/radiation
- Atm-Land: 70 PgC/yr ↑↓; Atm-Ocean: 90 PgC/yr ↑↓
- Net Land+ocean: 0 ± 1 PgC/yr
- "Stable" carbon cycle and climate over 1000y
- Projection of climate change on natural modes
- Detection & attribution
- Future climate projections/fossil fuel perturbations

Doney and Fung

Constraints from Observations

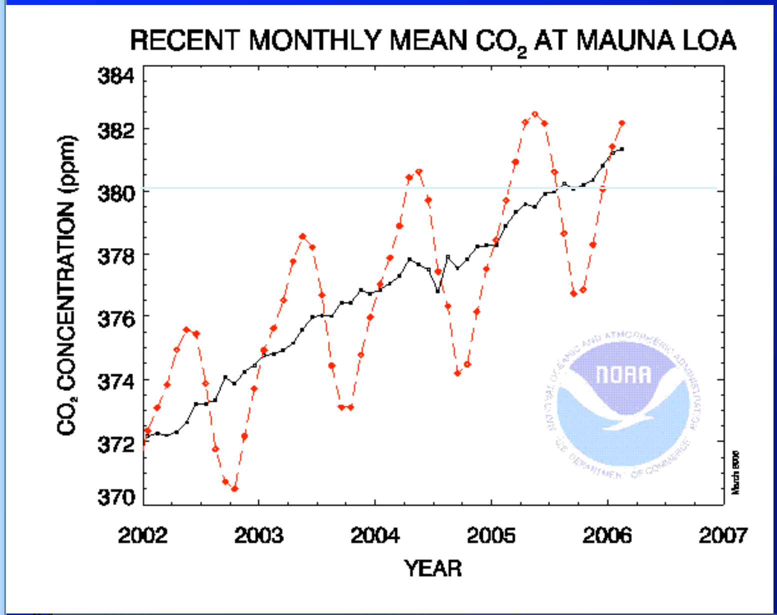

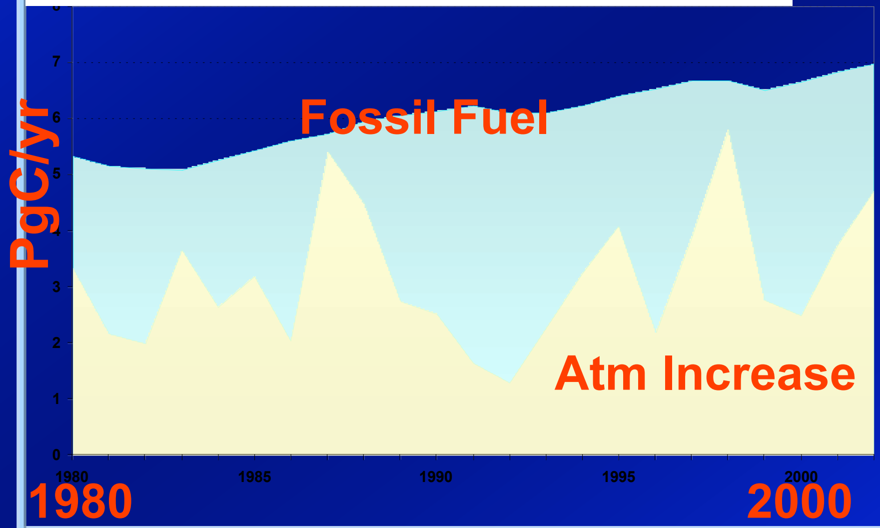

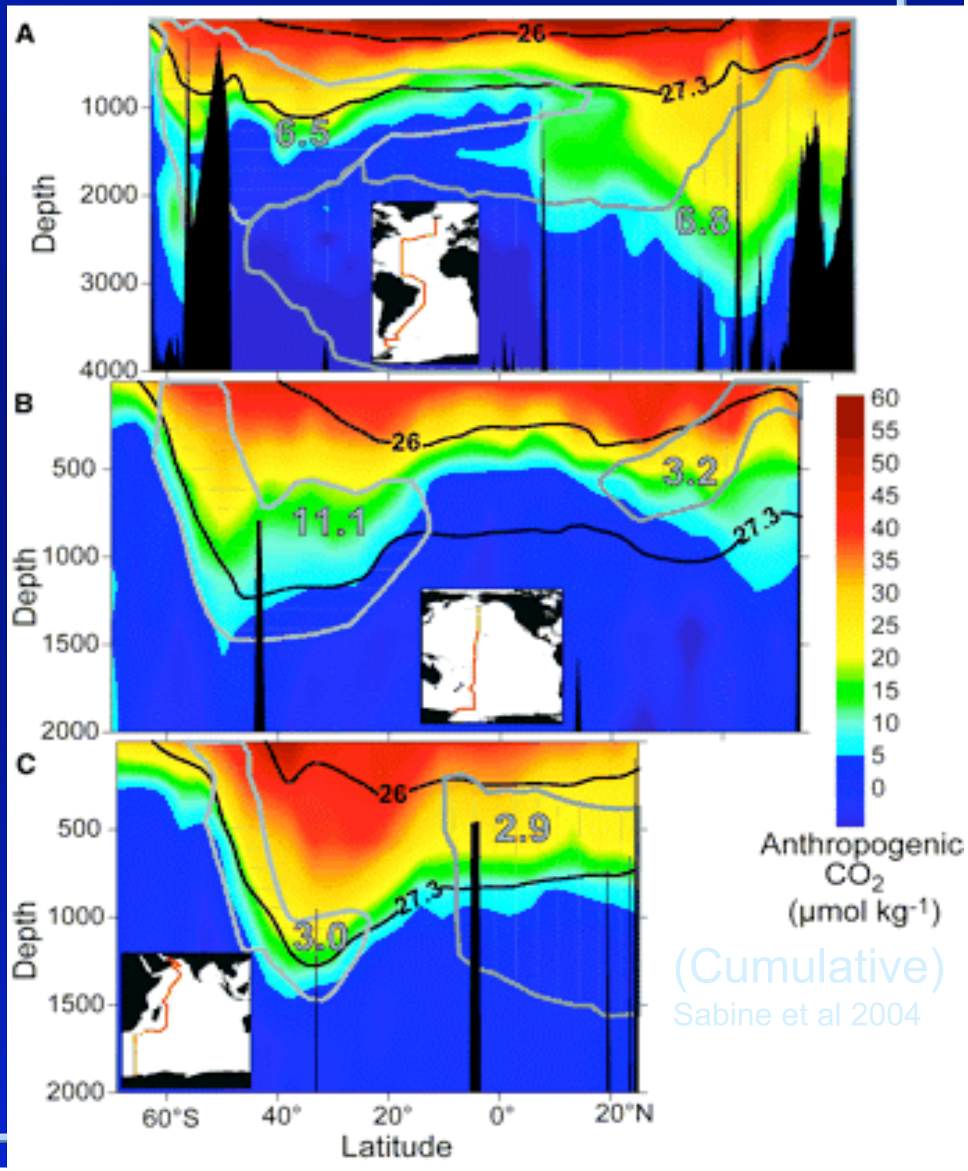

Proposal Targets

- Earth System Model
  - Terrestrial BGC and dynamic vegetation
  - Atm chemistry and aerosol dynamics
  - Ocn BGC
- Model Integration and Evaluation
  - Integration and unit testing
  - New cryosphere and ocean models
  - FV (cubed sphere), DG, others(icosahedral)
  - Frameworks for model evaluation
- Computational Performance
  - Scalablity, load balance, (fault recovery)

Integration and Evaluation of New Components in a Coupled Earth System Model

- Confidence in modeling the physical climate system does not extend to modeling the biogeochemical coupling
- Using observational data to validate and constrain the process models for terrestrial carbon cycle and atmospheric aerosols
- Atmospheric aerosol effects
  - Direct
  - Indirect
- Dimethel Sulfide from Ocean ecosystem
- Chemical coupling for Biogeochemistry
- Extending cryosphere to include ice sheets.
- New dynamical formulations and algorithms
- Carbon and climate coupling

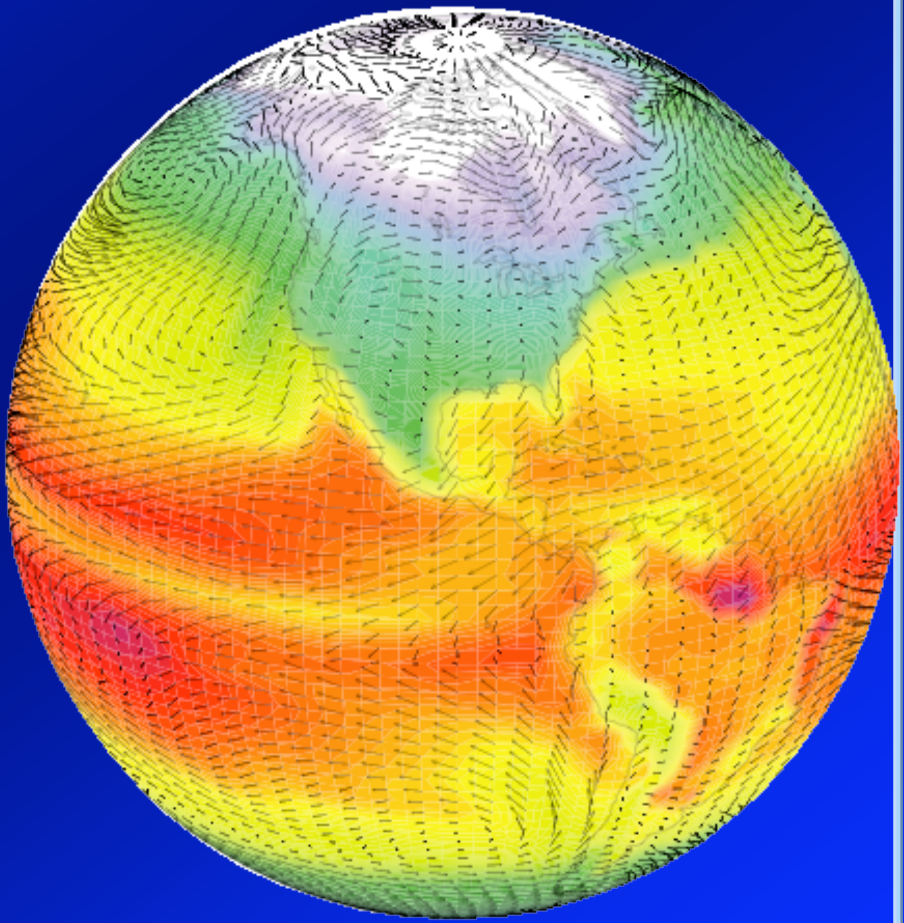

"Increasing Resolution vs. Actual Thinking"

CCM3.6.6 T42 -> T239

AMIP 1 -> AMIP 2

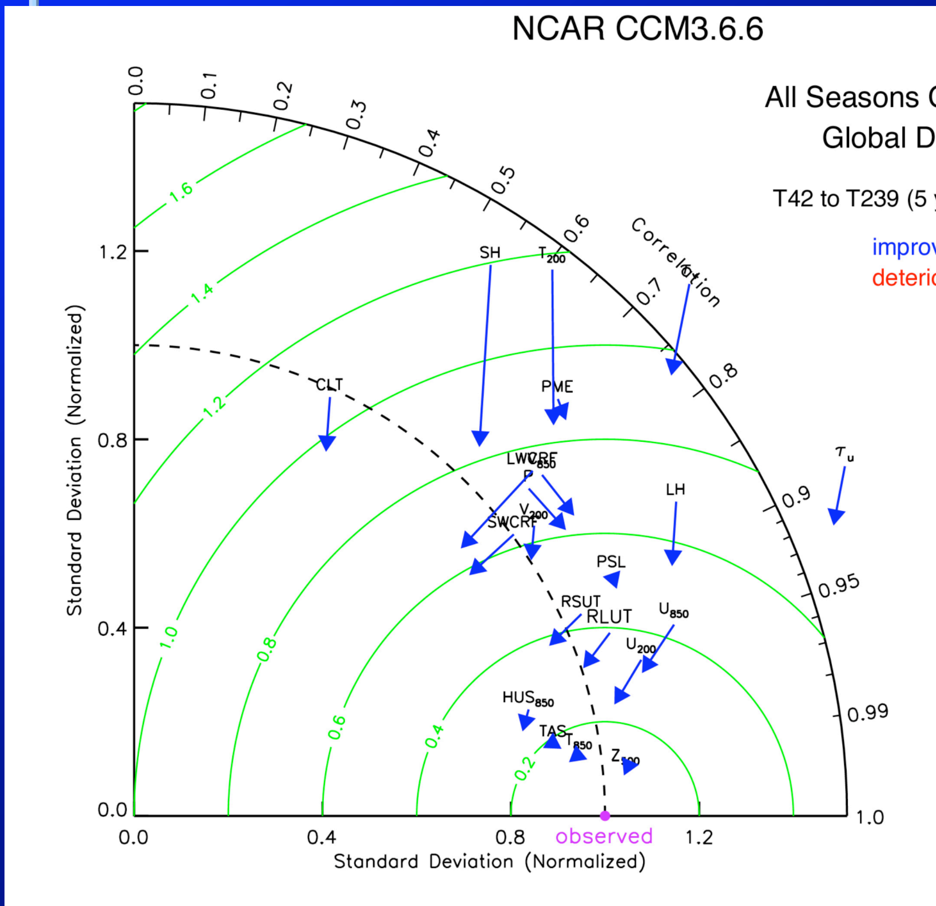

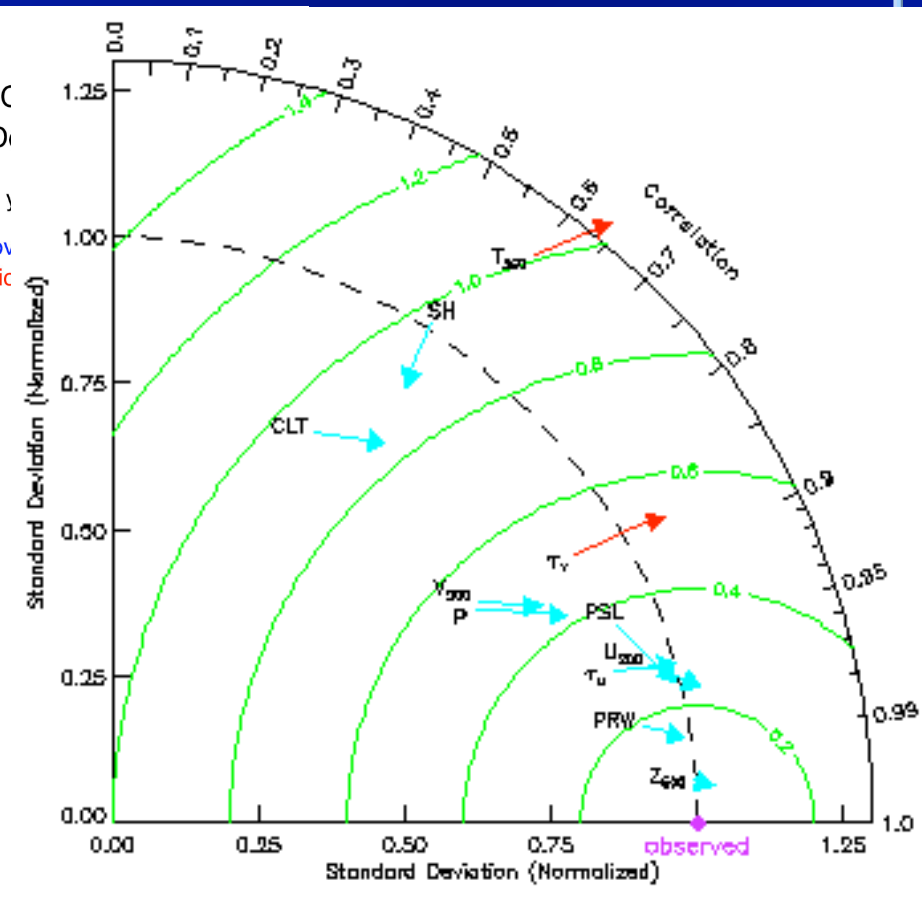

From P. Duffy presentation to CCSM Workshop June, 2003

Eddy-Resolving Ocean

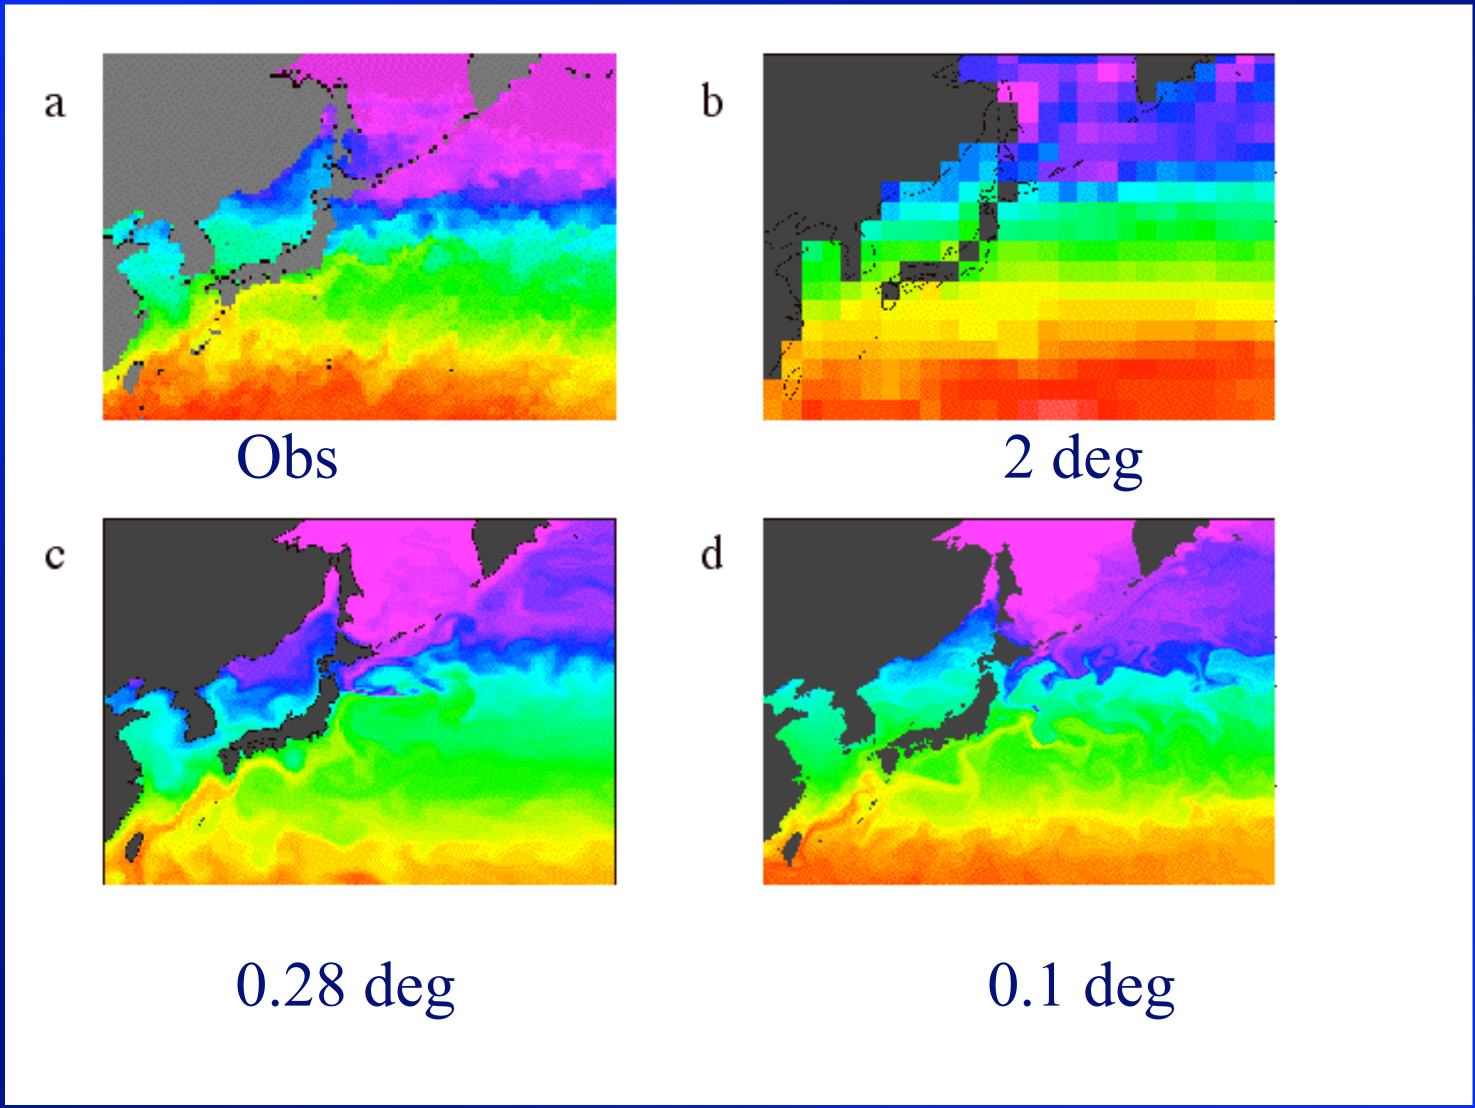

More Accurate Climate Models: Resolution Case Study

- FY06 Milestones
  - **High resolution ocean and sea ice , POP2 and CICE**
  - **High resolution atmosphere model bias studies,**
  - **Biogeochemical intercomparison simulations from C-LAMP**
  - **Climate Change scenarios stabilization with CCSM3.0 at T85**
- FY07 Milestones
  - **Bias studies with high resolution atmosphere/ocean coupling,**
  - **Dynamic ecosystem feedback simulation,**
  - **High res ocean THC and deep water formation,**
- FY08 Milestones
  - **Fully coupled physical climate at high resolution**
  - **Chemical coupling of climate and ecosystems**
  - **Climate sensitivity of high resolution coupled model.**

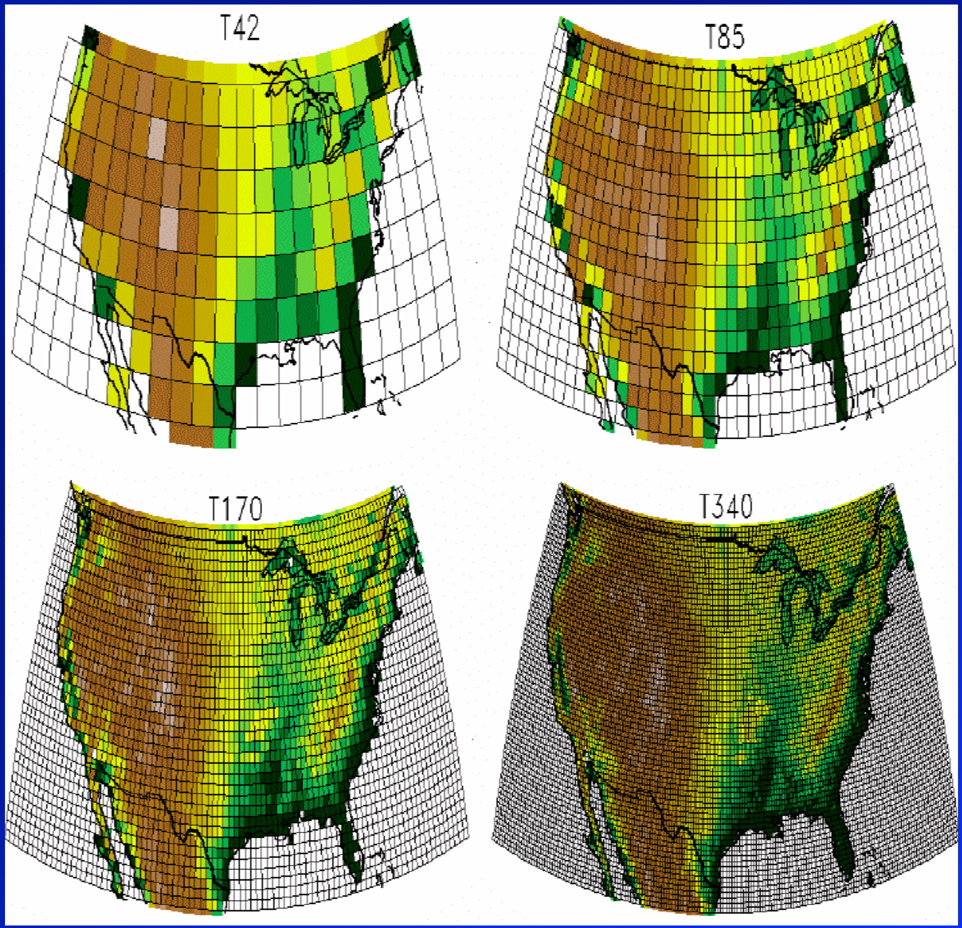

Proposal Targets

- Earth System Model
  - Terrestrial BGC and dynamic vegetation
  - Atm chemistry and aerosol dynamics
  - Ocn BGC
- Model Integration and Evaluation
  - Integration and unit testing
  - New cryosphere and ocean models
  - FV (cubed sphere), DG, others(icosahedral)
  - Frameworks for model evaluation
- Computational Performance
  - Scalability, load balance, (fault recovery)

Hardware and Software …

- Scalable and Extensible Earth System Model (SEESM)
  - Must meet the CCSM4 release schedule (tight NCAR coupling)
  - Look beyond to the 5 year horizon of the project

Community Climate System Model (CCSM3.1)

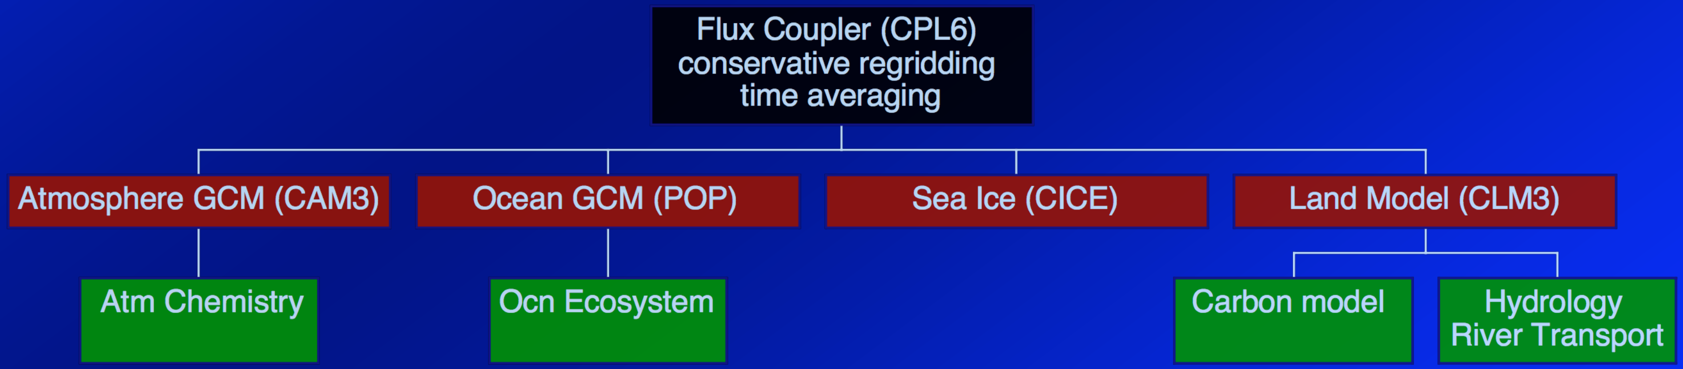

Scalability and Capability

The state of the code (prototype resolutions)

- scalability
  - Atm ~ 1000 procs without chemistry
  - Atm ~ 5000 procs with chemistry
  - Ocn ~ 4000 procs
  - Fully coupled low res production ~ 500procs
  - SciDAC2 will scale to 25K and 100K
- current capability
  - Coupled physical atm, ocn, land, sea ice
- planned capabilities
  - + atm chem, ocn ecosystem, land carbon, dynamic vegetation, ice sheets, aerosol indirect effect

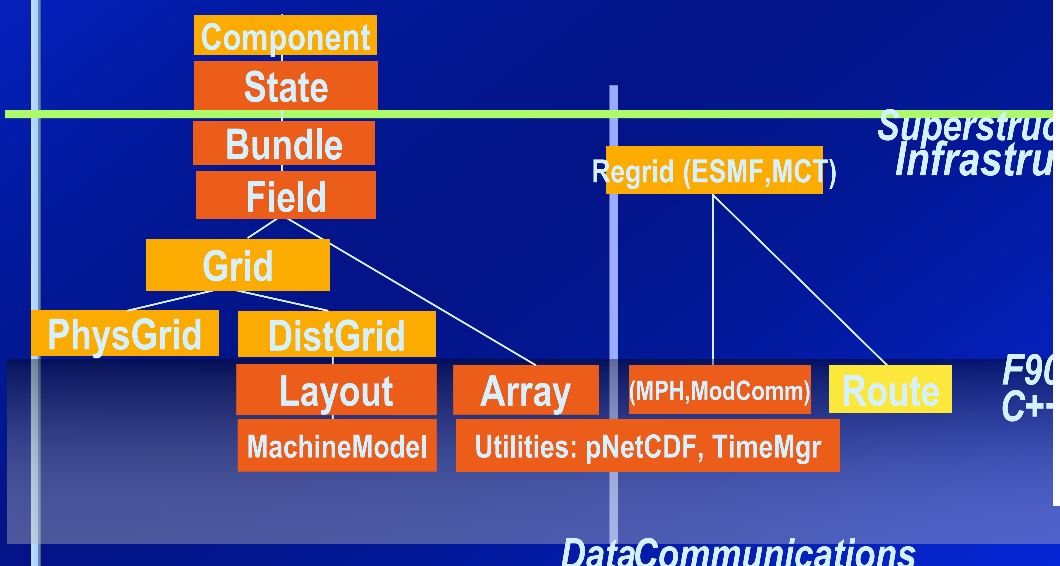

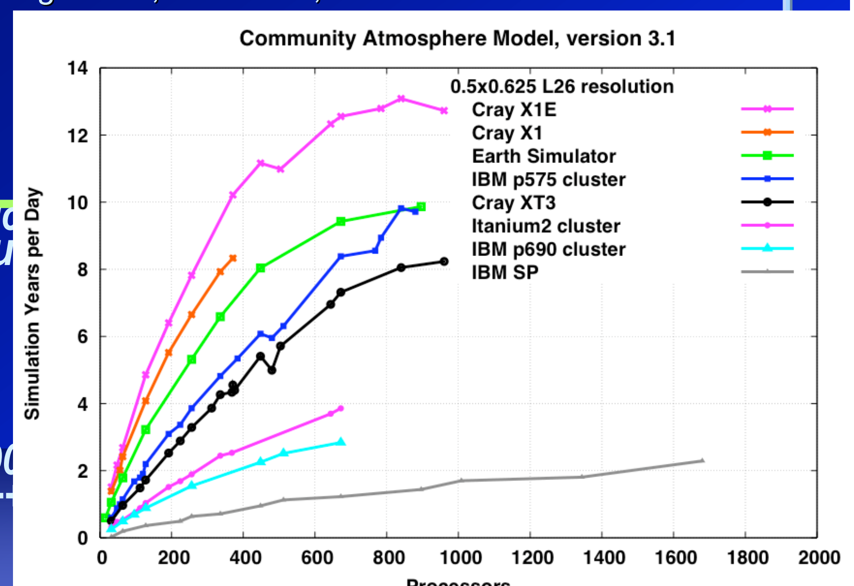

Single Source Performance Portability

- Load balance within atmospheric component
- Load balancing between components
- Tunable data structure size for cache and vector performance portability

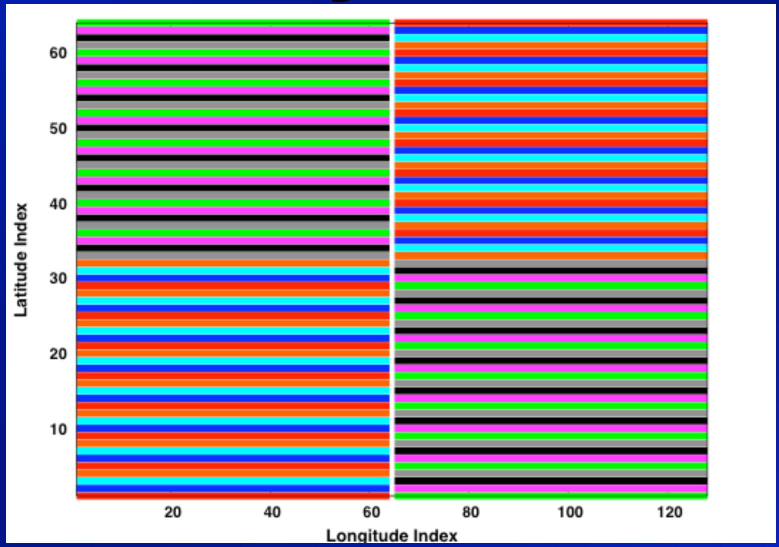

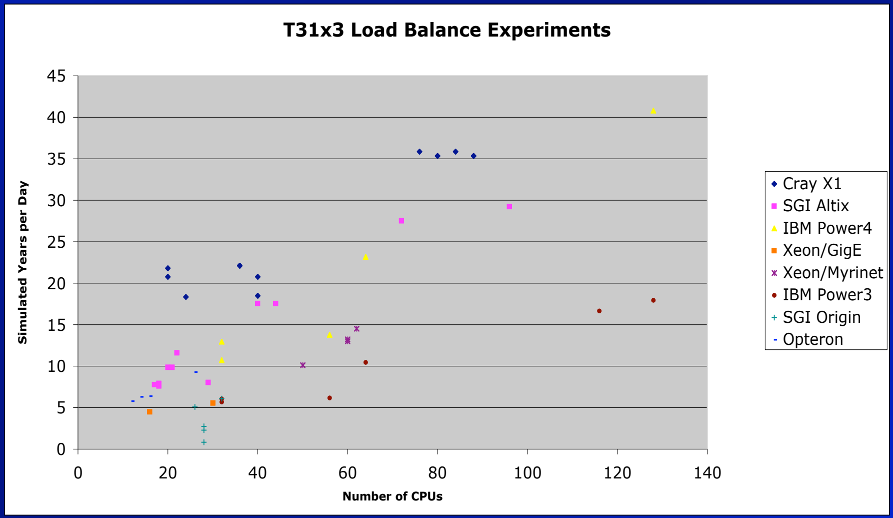

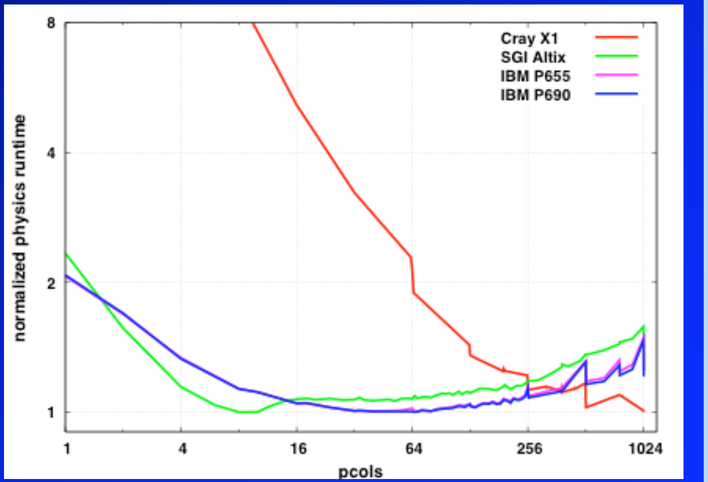

Present Parallel Algorithms

- Two dimensional domain decompositions
  - Independent atmospheric columns for radiation calculation
  - Patches for atm, sea ice and ocean dynamics and semi-implicit solvers
  - Independent particle tracking for semi-Lagrangian and incremental remapping transport algorithms
  - Clumps for land points and plant functional types
- Concurrent parallel components
- Parallel coupler component that remaps fluxes in space and time
  - Argonne Model Coupling Toolkit (MCT)
  - Berkeley Multi-Program Handshaking (MPH)
- Transpose based communicators: optimize components and localize communication between components

"The method providing access to polar data becomes an important consideration when actual programming is attempted. .. Using the broadcast register of the SOLOMON II system to provide a variable to the north row of PE's .."
-A.B. Carroll (1967) *Application of Parallel Processing to Numerical Weather Prediction*

Reviewers Comments

- Is the potential scalability of a many-tracer code compatible with an AMR code? These questions are not even contemplated in the proposal
- The proposal is so vague about what variables will be constrained by the assimilation that the reader is left guessing.
- This is a very comprehensive proposal
- Despite its strength and scope, in a first reading this proposal didn't seem very responsive to the High Performance Computing aspects of SciDAC.
- A couple of activities seemed a bit disconnected of other ongoing activities:1) There is a explicit task to extend the finite volume dynamical core to the cubed-sphere. On going development by S.-J. Lin at GFDL, in coordination with NASA/GSFC and NCAR scientists, is producing a version of the finite-volume dynamical on the cubed-sphere. No mention of such activity, or whether it meets or does not meet the CCSM requirements, is made in the proposal.2) It was not clear how introduction of aerosols in CCSM relates to the effort of Phil Rasch's group at NCAR. 3) There is mention in the proposal about development of coupling frameworks, with only mention in passing about using low level utilities from the ESMF.
- highly likely to produce a successful computational infrastructure for the next several generations of the CCSM.
- poor understanding may lead to great uncertainty in the earth system model, much more than the current climate model. The improvement of this uncertainty may pose the major challenge in the coming decades to our climate community.
- One perhaps needs to look back seriously at history of CCSM for some lessons: there has been little effort outside NCAR to test the full CCSM, simply because it is too complex and too computationally expensive.

Climate-Science
Computational End Station Allocation

PI: Warren Washington (NCAR), partners: CCSM, COSIM, PCMDI, SciDAC, NASA-GSFC, PNNL, CCRI(Universities)

- Extensible *community models* available for computational science
- Coordination of effort among agencies and institutions
- *Scalability* from 500 to 5,000 to 50K processors

Testing of Methods

(spherical geometry)

- Barotropic vorticity equation (Charney, Fjortoft, von Neumann (1950)) - single prognostic eqn with elliptic diagnostic equation
- Shallow water equation test set (Williamson, et al, 1992) - u,v, h equations
- Held-Suarez test for baroclinic models - u,v, p, p_{surf}, T , (w is diagnostic)
- Aqua-planet - full moist physics but no topography (Neale&Hoskins, 2001)
- AMIP
- CMIP
- C-LAMP, C4MIP

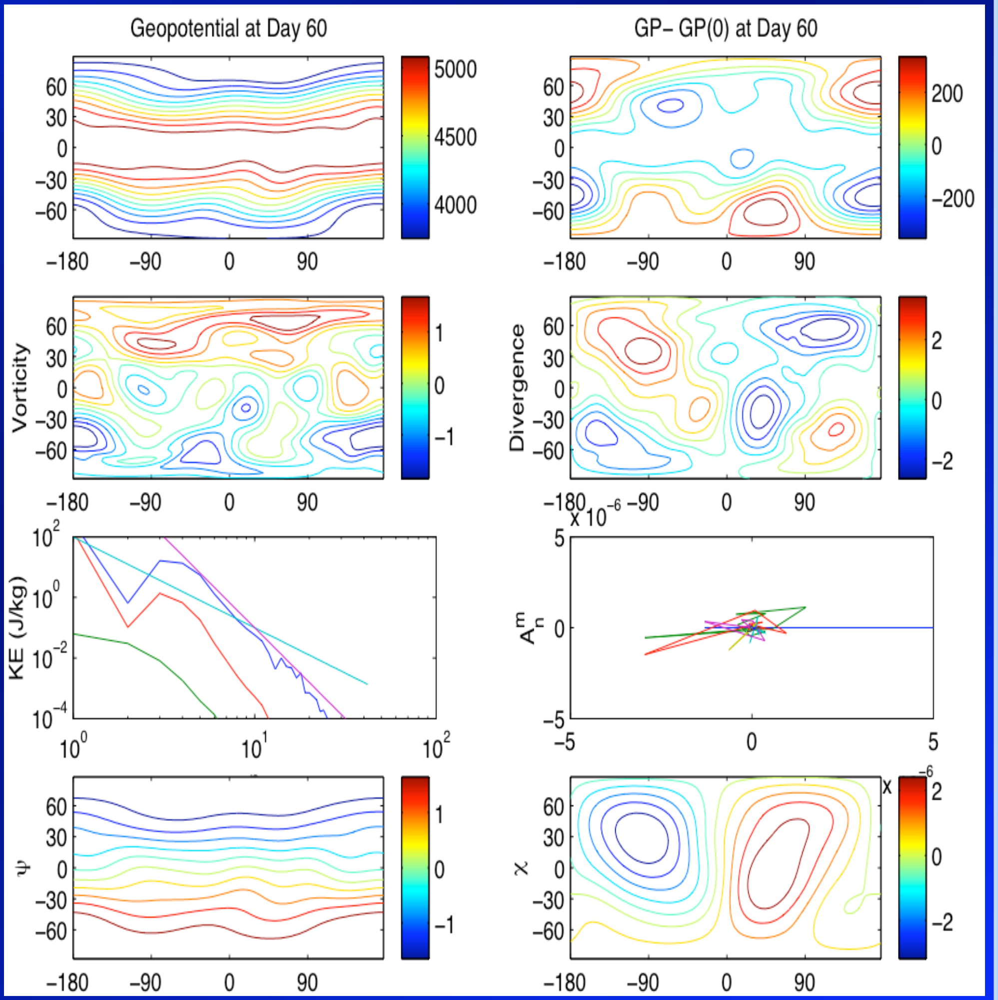

"In retrospect, the shock problem seems relatively easy." - J. Dukowicz(2000)

Supercomputers at ORNL

18 TF Cray X1 / X1E

- Will not expand
- 1024 processors
- Vector processing for sustained performance

50 TF Cray XT3

- Expandable to 100+TFLOPS
- Max 10,000+ processors
- MPP Compute system for large-scale sustained performance
- Based on Sandia "Red Storm" collaboration

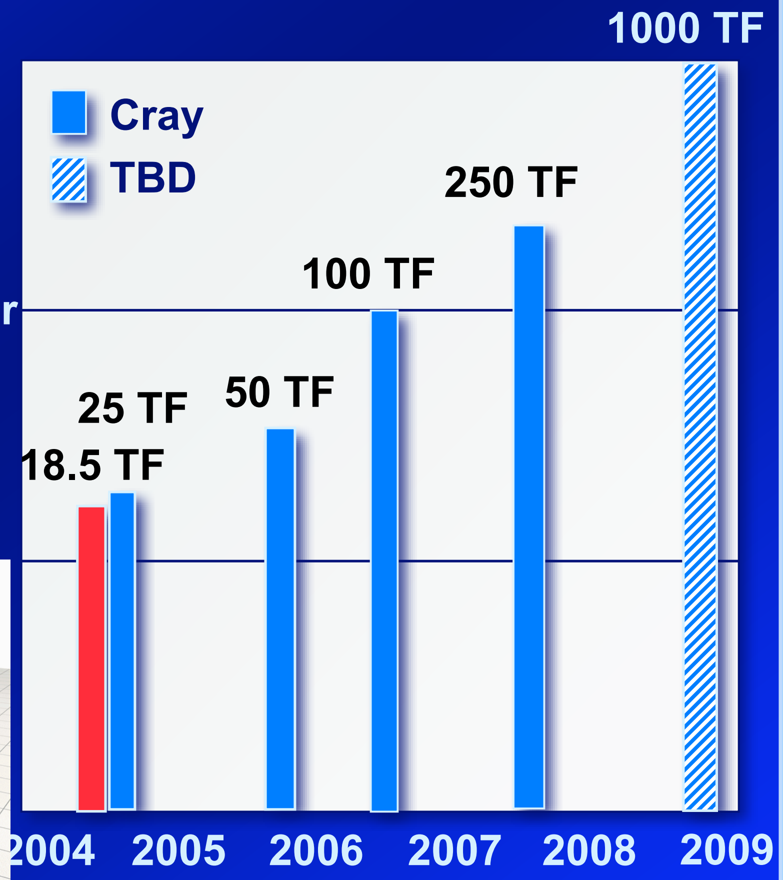

Leadership Computing Facility

(also BG/L at Argonne)

Computational Requirements

Issue	*Motivation*	*Compute Factor*
Spatial resolution	Provide regional details	10^3-10^5
Model completeness	Add "new" science	10^2
New parameterizations	Upgrade to "better" science	10^2
Run length	Long-term implications	10^2
Ensembles, scenarios	Range of model variability	10
Total Compute Factor		10^{10}-10^{12}

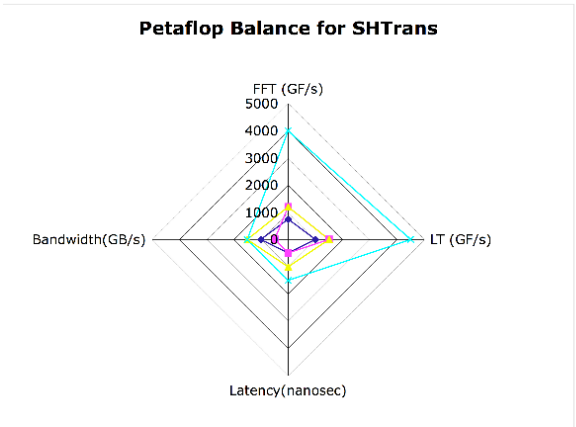

A Science Based Case for Large-Scale Simulation (SCaLeS), SIAM News, 36(7), 2003 - David Keyes

Establishing a PetaScale Collaboratory for the Geosciences
UCAR/JOSS, May 2005

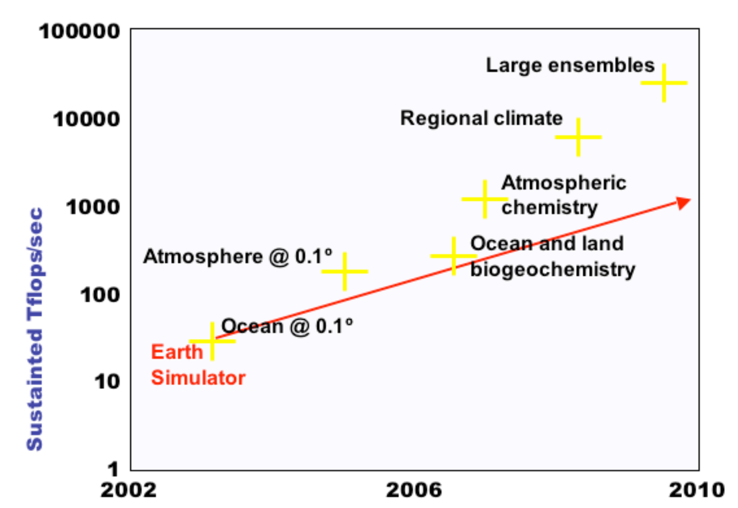

Will CCSM4 be ready by June 2008?

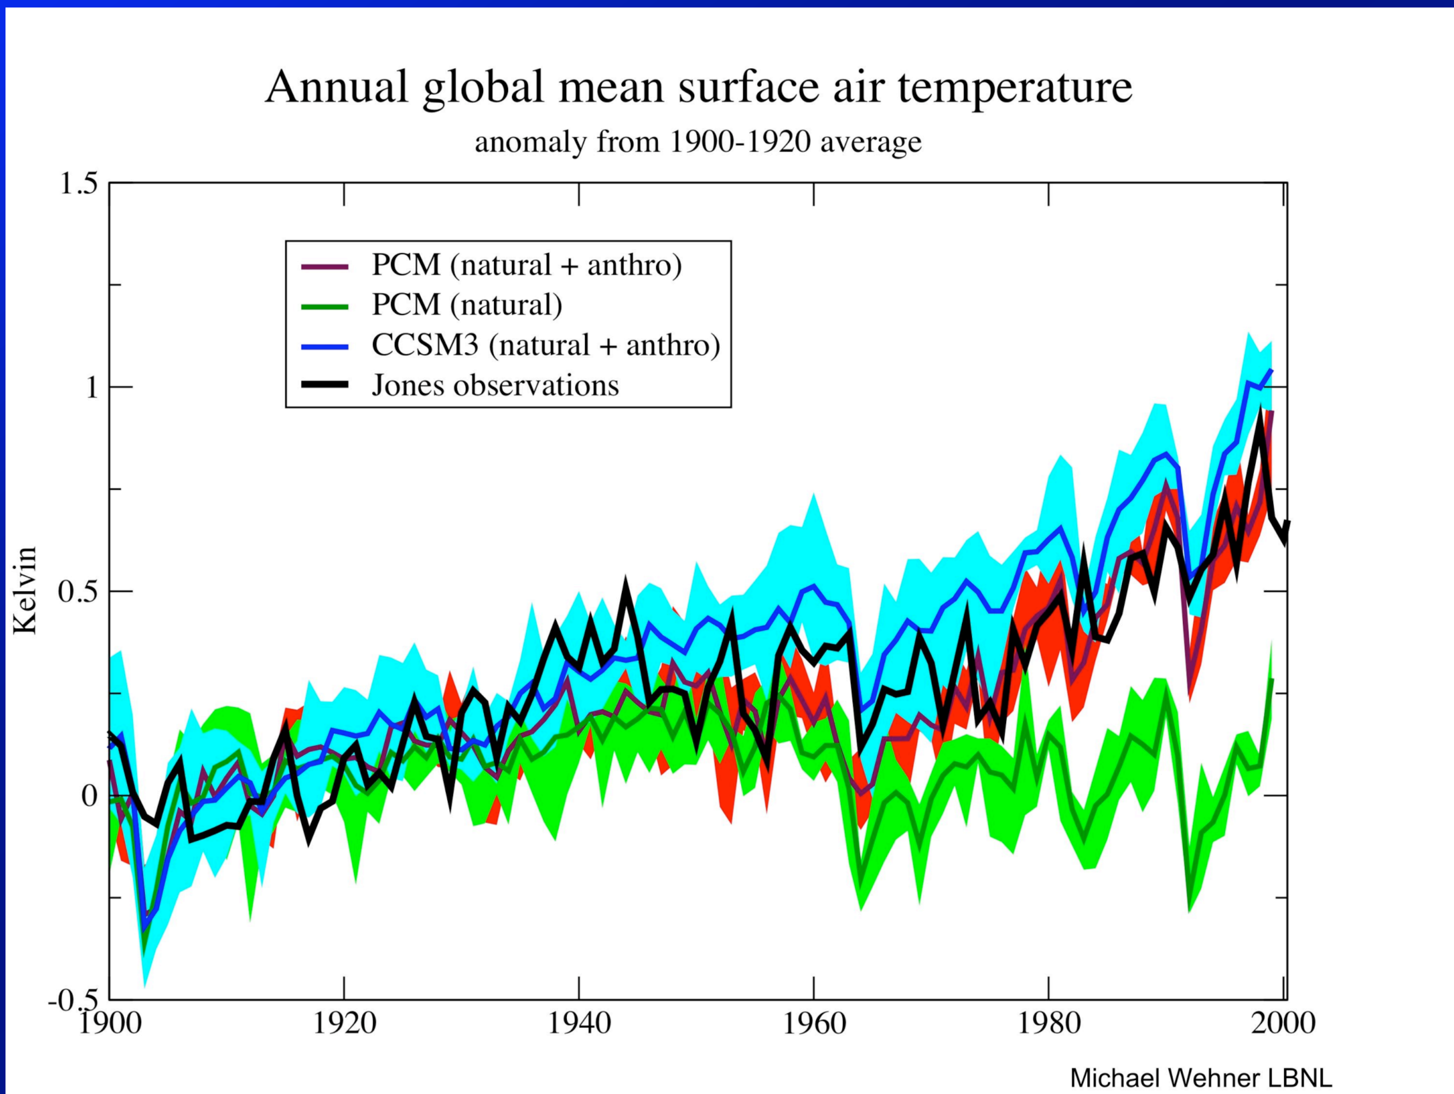

Summary

- SciDAC2 CCSM Consortium will collaborate with NSF and NASA projects to build the next generation Earth System Model
- The NLCF Climate End Station provides a significant portion of the development and climate change simulation resources
- Scalability and Extensibility are required for petascale science applications