Present and Future Computing Requirements

Case Study: Climate Change Simulations with the Community Earth System Model (CESM)

Thomas Bettge

Warren Washington

Climate Change Research Section

Climate and Global Dynamics Division

NCAR Earth System Laboratory

National Center for Atmospheric Research

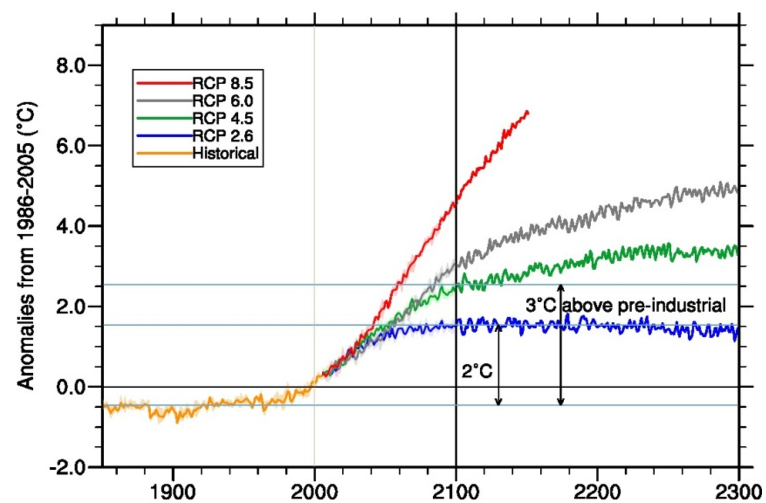

NERSC BER Requirements for 2017

September 11-12, 2012

Rockville, MD

Housekeeping

May 2009:   Trey White joined CCP group at NCAR.

May 2011: Trey completed his PhD in computer science " Algorithms for Advection on Hybrid Parallel Computers"

→ A GPU guy!!!

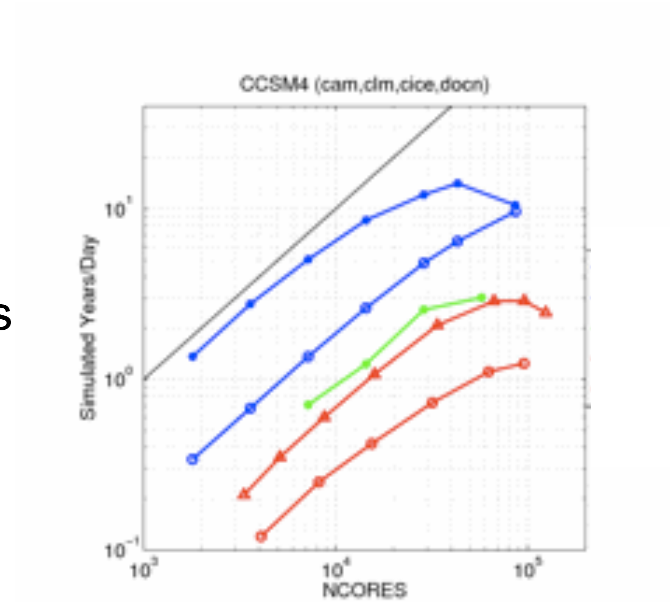

April 2012: Trey accepts job with DreamWorks Animation in Hollywood.

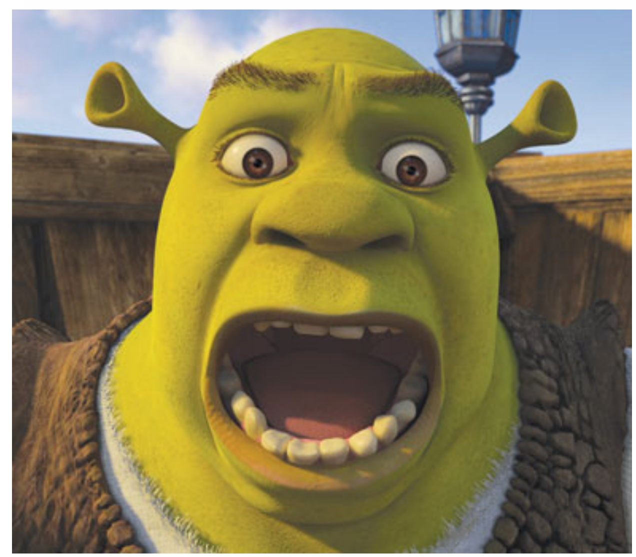

Preview

- ***Project Description***
  - Climate Change Prediction (CCP) and The Tool – CESM
- ***Computational Strategies***
  - CESM Two-Minute Tutorial
- ***Current HPC Usage***
  - Where is the parallelism?
- ***HPC Requirements for 2017***
  - How many hours did you say you needed?!
- ***Strategies for New Architectures***
  - Where is the speed-up?
- ***Summary***
  - New Science

Project Description

- ***The Climate Change Prediction (CCP) group at NCAR goals:***
  - To investigate mechanisms of climate variability and change
  - To detect and attribute past climate changes
  - To project and predict future climate changes
    - **"The future ain't what it used to be." –Yogi Berra**
- ***Present CCP focus (2012-2013)***
  - Simplified forcing experiments
  - Long pre-industrial control simulations (fully coupled model)
  - Future climate simulations using different emissions scenarios
  - Decadal hindcast and initialized prediction experiments <u>using highest resolution practical</u> to better quantify time-evolving regional climate changes
  - Evaluate natural variability (uncertainty) in all of the above <u>with large number of ensemble members</u>

Project Description

CESM Project

Based on 20+ Years of Model development and application

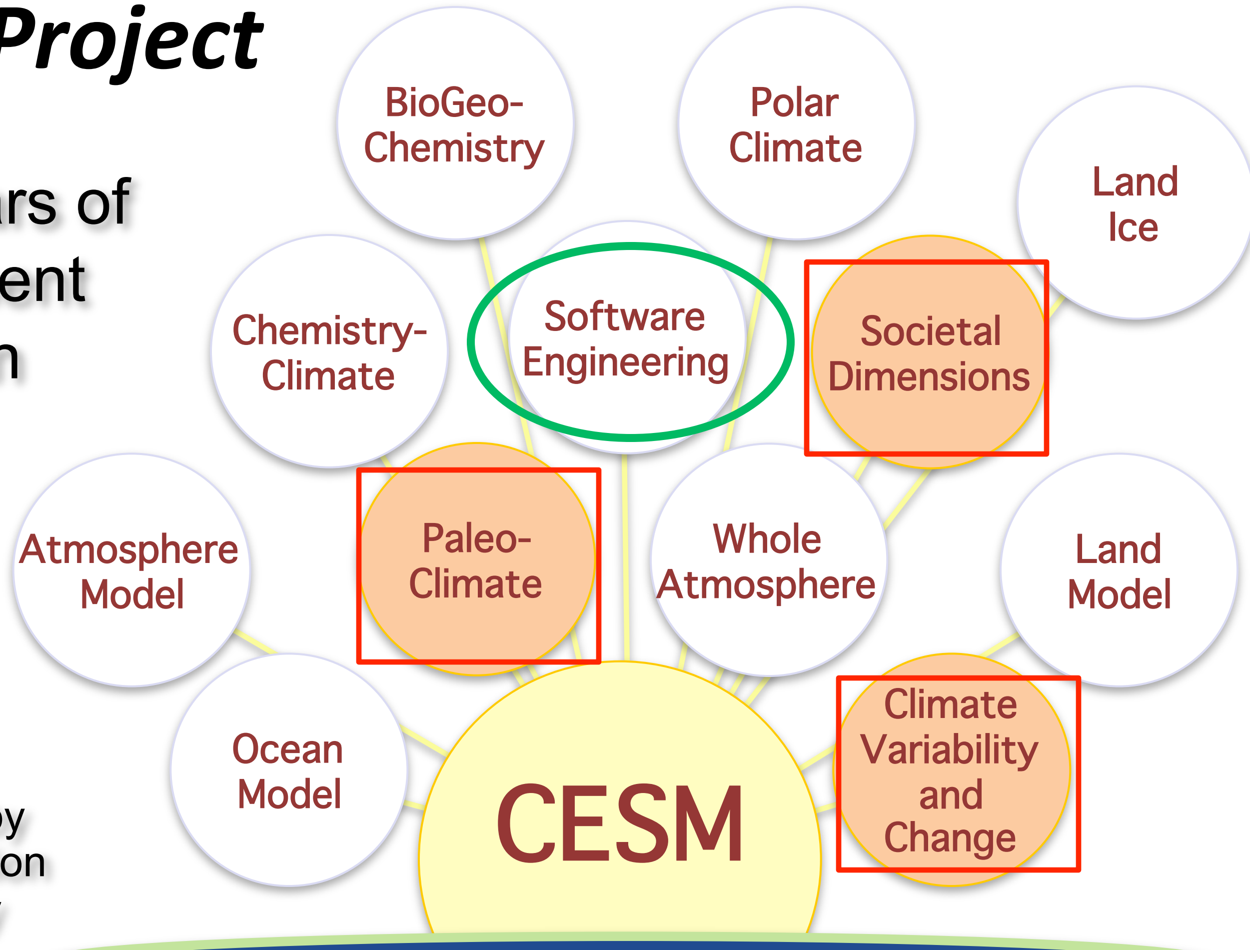

CESM is primarily sponsored by the National Science Foundation and the Department of Energy

http://www.cesm.ucar.edu/management

Project Description

- ***CESM Project can be classified by climate change research in three categories:***
  - CESM Development and Validation
    - R&D of scientific processes/methods and computational algorithms
    - Easy Accessibility to high-end computing platforms to test and validate
  - Production Climate Simulations using CESM (the focus of CCP)
    - Contribute to national and international missions and goals with scientifically and computationally validated CESM versions
    - Consistent Accessibility to stable high-end computing platforms for extended production
  - High Resolution Fully Coupled CESM Simulations
    - Test cutting-edge high resolution CESM configurations with extended integrations
    - Priority Accessibility to high-end computing platforms which will allow potential transformations in climate science in reasonable timeframes

Project Description

- ***Current CESM production climate change experiment***
  - Components: atmosphere, ocean, land, ice, coupler
  - 1^{o} horizontal resolution
  - Validated climate with and existing control (1850 forcing, 1000 years)
  - 2064 processors within a single executable
    - Tuned for optimal performance and optimal load balance
    - Components run both sequentially and concurrently
  - <u>10 simulated years per wallclock day</u> (5000 pe-hours per simulated year)
  - End-to-end experiment
    - Historical (156 years: 1850-2005)
    - Future (4 scenarios, 100 years each)
    - 5 ensemble members of each   (156 + 4*100)*5 = 2780 years
    - Total time required: <u>14M pe-hours</u>

Project Description

- ***Current CESM production … pragmatic limitations/challenges***
  - Need for certified model, validated control plus ~3000 years, and the attention span of the scientists ….. limits configuration to 1º (14M pe-hours!!!)
  - 1º resolution scalability limits processors count 2064 processors. Why? Cost and current model scaling with CAM Finite Volume dynamics:

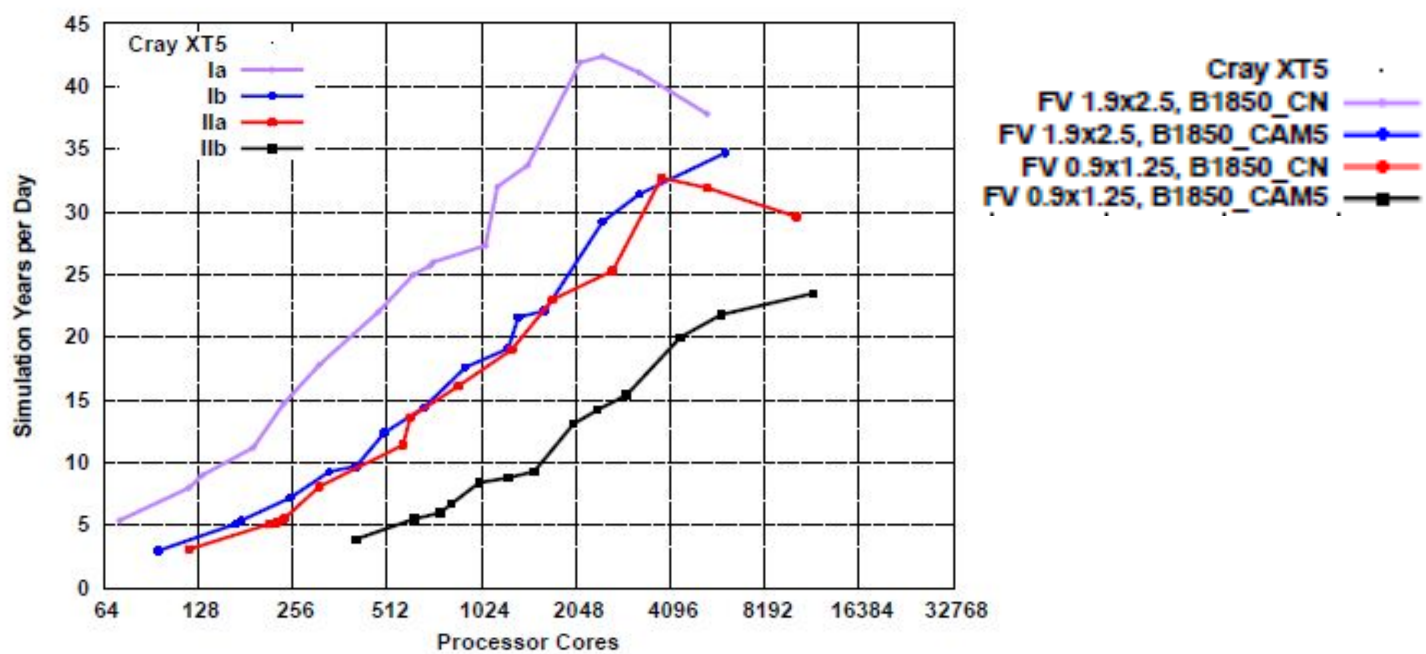

P. H. Worley, A. P. Craig, J. M. Dennis, A. A. Mirin, M. A. Taylor, and M. Vertenstein, *Performance of the Community Earth System Model,* in Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC11), Seattle, WA, November 14-18, 2011.

Project Description

- ***Current CESM production … pragmatic limitations/challenges***
  - CAM dynamics core – more scalable numerical method to higher resolutions
    - Finite Volume vs. Spectral Element at 0.25°

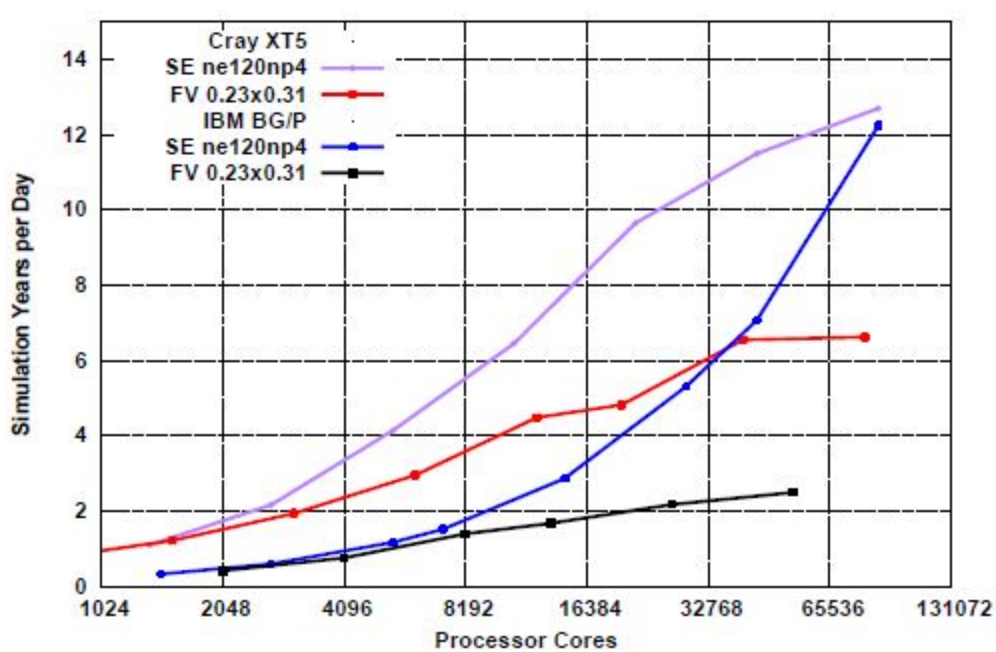

P. H. Worley, A. P. Craig, J. M. Dennis, A. A. Mirin, M. A. Taylor, and M. Vertenstein, *Performance of the Community Earth System Model,* in Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC11), Seattle, WA, November 14-18, 2011.

Project Description

- ***By 2017……..*** **"You've got to be very careful if you don't know where you're going, because you might not get there." –Yogi Berra**
  - CESM (from CESM Planning Document) anticipates:
    - Increasing compute power to support increased model resolution
    - Improved understanding of processes represented in CESM
    - Improved understanding of process interactions
    - Improved numerical methods for simulation of geophysical flows
    - Improved observations of the atmosphere
  - CCP Goals
    - Refinement of simulations based upon 4x horizontal resolution . Production simulations will use O(30,000+) processors vs. O(2000) today.
    - Future emissions scenarios with refined source origin
    - Greater definition of regional climate changes
    - Explore ability to relate extreme weather events to climate change (i.e., tropical storms)
    - Production simulations for contribution to IPCC AR6 (2016)

Computational Strategies

- CESM includes atmosphere, ocean, land, and seaice components
- Conservative exchange of heat, water, momentum across components
- Changes in external forcing can be applied – solar input, GHGs, volcanic eruptions, etc.

CESM provides a virtual laboratory for experimentation.

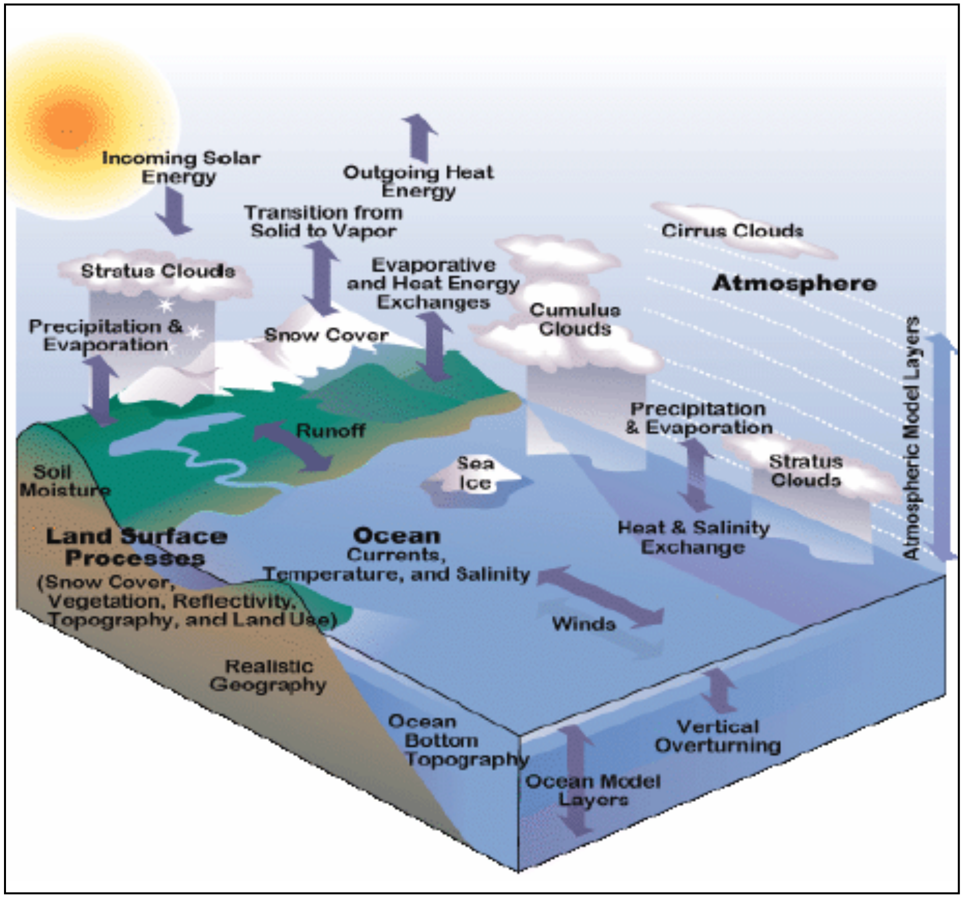

Computational Strategies

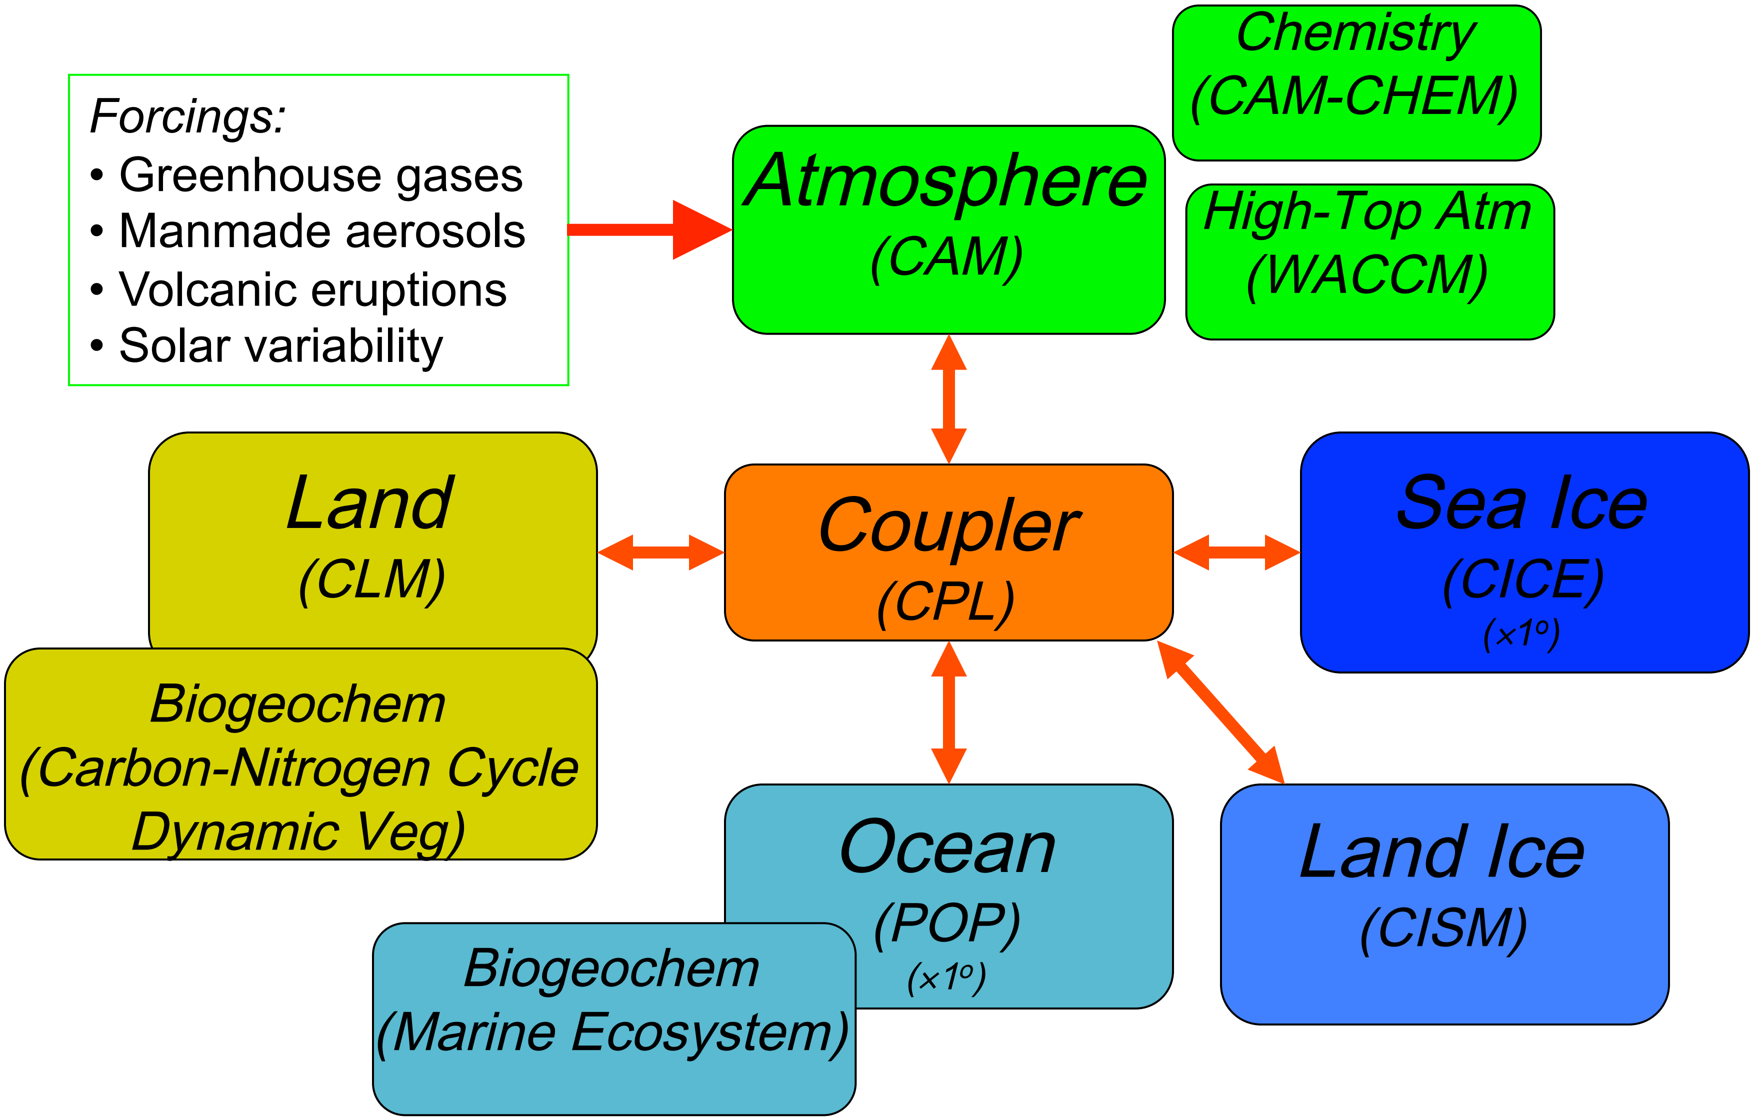

Computational Strategies

- ***Overall computational challenges involve CESM supporting varying:***
  - resolutions
  - component configurations
  - MPI and OpenMP parallelism
  - processor layouts
  - mixed sequential/concurrent execution
  - active and data driven components
  - ability to add new components
  - number of hardware platforms – **37** listed in:
    - www.cesm.ucar.edu/models/cesm1.0/cesm/cesm_doc_1_0_4/a3986.html
  - compilers, options, I/O infrastructure, performance libraries, etc.

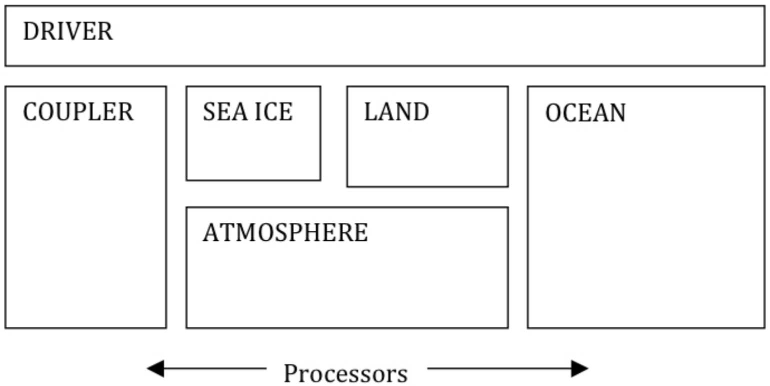

Computational Strategies

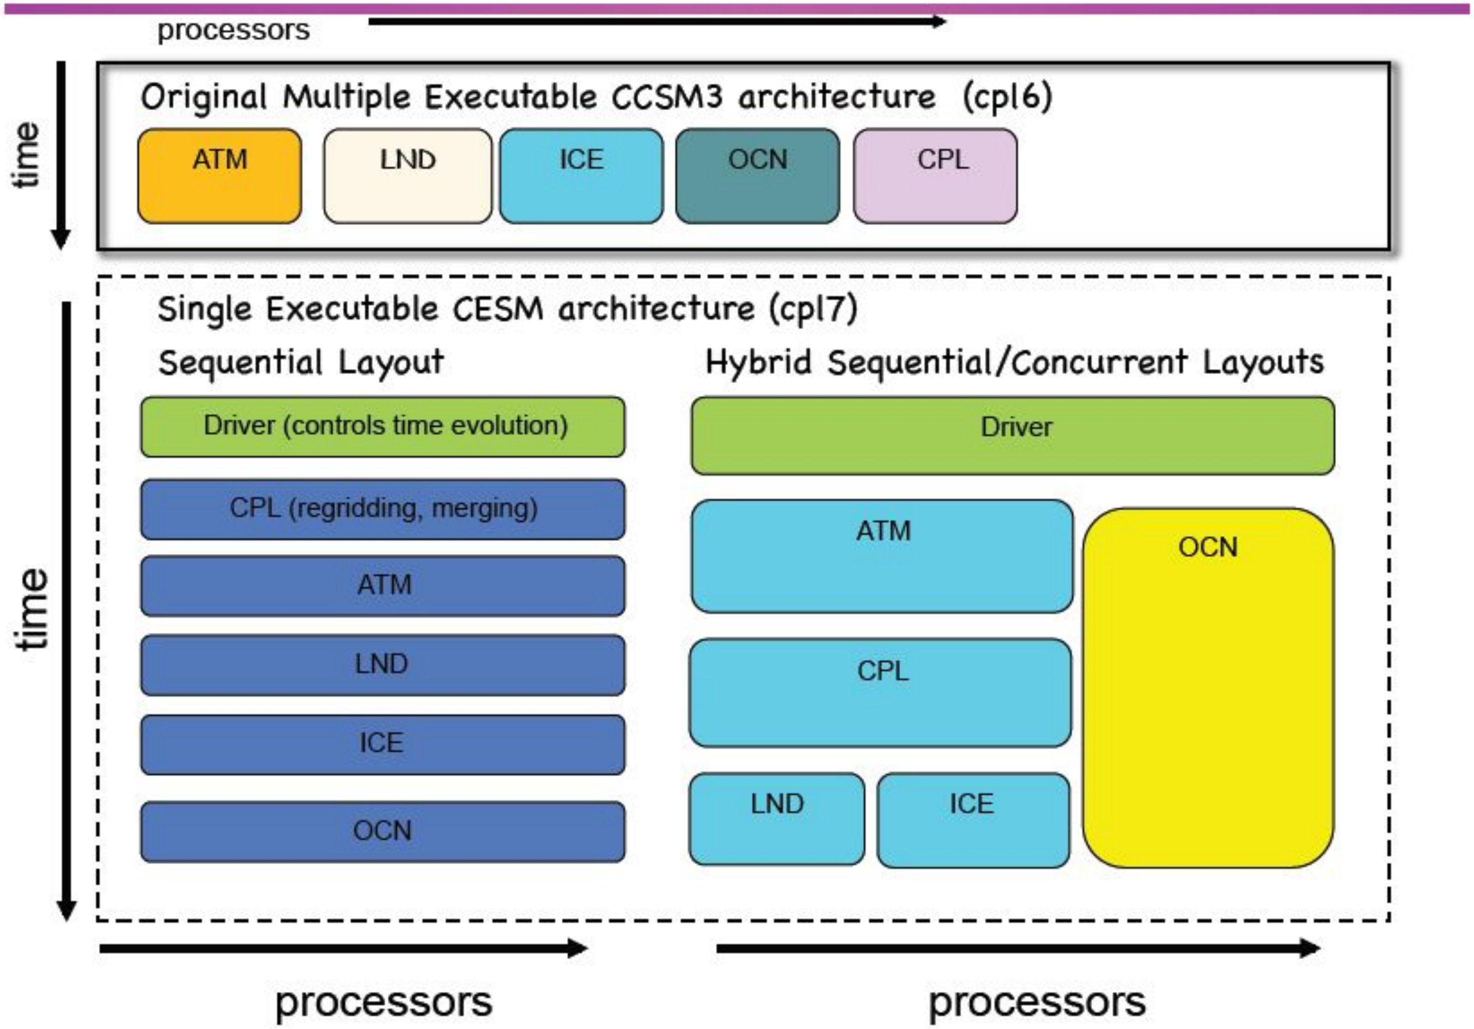

Current HPC Usage

		Mpe-hours*			
Site	*Machine*	*2011 Usage All CESM*	*2011 Usage CCP Only*	*2012 Usage All CESM*	*2012 Usage CCP Only*
NERSC	Cray XE6	32.5	22.1	23.8	15.3
OLCF	Cray XT5	48.0	8.8	37.5	8.7
ALCF	IBM BG/P	10.8	0	4.0	0
NCAR (CSL)	IBM p6	30.2	6.3	17.6	3.7

*Normalized to equivalent Cray XE6 pe-hours.

- CESM usage is aggregate of research, development, testing, climate change experiments, and high resolution cutting edge simulations.
- CCP-only is usage attributed to climate change simulations which we have used to contribute to national and international programs for climate model intercomparisons, climate change detection/attribution, and future climate change projections.

.

Current HPC Usage

- Achieve coarse parallelism by combining ensemble members into one job submission (to reach reg_med discount threshold)
- Run time: **36 hours** (queue limit)
- Data read: **<50 GB**
- Data written: **1 TB**
- Memory: **< 2 GB/pe**
- Resources: HPSS, project directories
- Software:
  - Parallel I/O PIO written by CESM staff

HPC Requirements for 2017

- ***The Standard Case End-to-End : Control-Historical-FutureScenarios w/ensembles -> 2780 simulated years***
  - Atmosphere resolution of 0.25^o
  - Validated dynamical core (likely CAM-SE)
  - New parameterizations and processes
  - Additional coupling across components with interactions
  - Software improvements

Resolution	Simulated Years Required	Cost per Simulated Year (pe-hours)	Data Storage per Simulated Year (GB)	Total Cost (pe-hours)	Total Storage Required (TB)
1^o (2012)	2780	5000	30	14M	84
0.25^o (2017)	2780	60,000	500	167M	1400

HPC Requirements for 2017

- ***Number of Compute Cores – Goal: <u>10 years/day</u>***

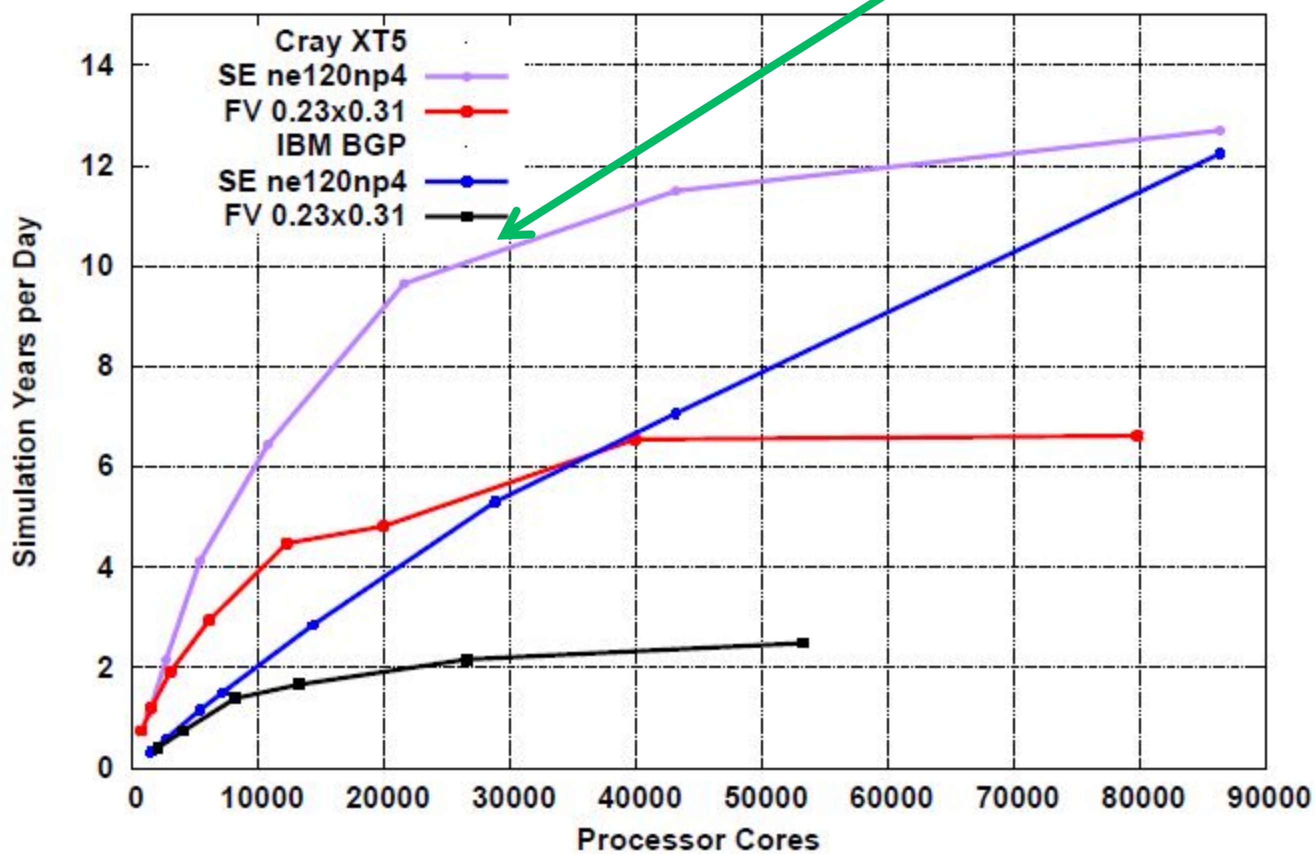

HPC Requirements for 2017

- *Storage (HPSS)*

year	files	volume (TB)
2000	34,354	2
2001	33,690	2
2002	198,950	9
2003	89,454	7
2004	123,043	17
2005	52,256	9
2006	98,932	10
2007	71,272	17
2008	493,446	98
2009	144,137	58
2010	248,066	156
2011	808,572	456
2012		
2013		
2014		
2015		
2016		
2017		

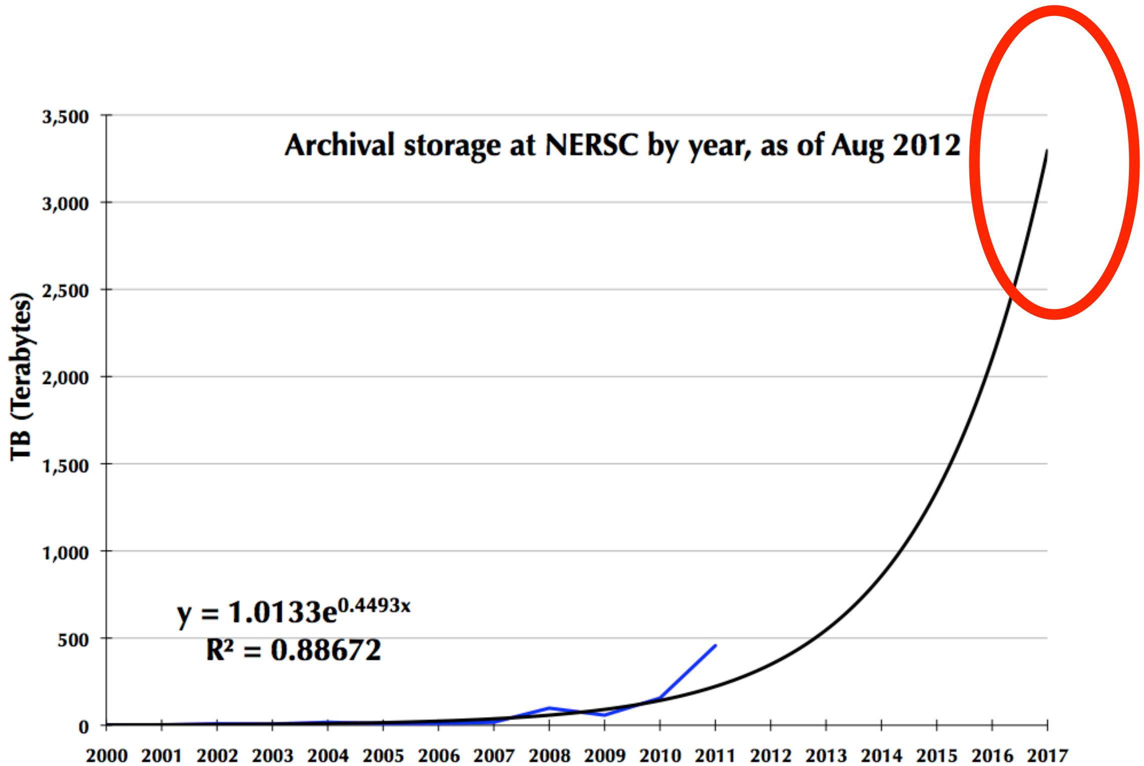

Strategies for New Architectures

- ***Stay abreast of developments in many-core architectures***
  - NOAA Earth System Research Laboratory
  - OLCF : AMIP version of CESM using CAM-SE has been designated a benchmark for acceptance of the machine
  - CESM Software Engineering Working Group
- ***Preparing***
  - OLCF team of application readiness software engineers has redesigned the SE dynamics for efficient use of the GPU hardware, and has achieved an overall 2.6x speedup of the CAM-SE dynamics[1]
- ***Opportunities***
  - CAM physics have not yet been evaluated for GPU use
    - Cloud resolving superparameterizations – fertile ground?
  - CESM software engineers are beginning to look at designing a GPU enabled version of POP

[1]Norman, M., J. Larkin, R. Archibald, I. Carpenter, V. Anantharaj, P. Micikevicius, and K. Evans: *Porting the CAM – Spectral Element Code to Utilize GPU Accelerators.* Cray User Group, Stuttgart, Germany, April 29 – May 3, 2012.

Summary

CVCWG/CCP Research 2013-2017

2013
- Decadal prediction
- Single-forcing runs
- High resolution CAM5/SE
- Special DOE runs

Decadal prediction runs using hindcast or DART assimilation for initial states: CESM1 at 1°
Long term runs: CESM1, complete CMIP5 simulations at 1° resolution, concentration driven, CESM1/BGC concentration driven RCP runs, 0.25° pre-industrial control
CCSM4 at 0.5° CMIP5 RCP future climate change runs
Special DOE scenario runs with CESM1, including short-lived GHG and black carbon
CESM1 runs with specified ice-sheet meltwater
High-resolution time-slice runs with CAM5/SE at 0.25° and 0.125° resolution
Perturbed physics ensemble with CCSM4 at 1° resolution
Single-forcing sensitivity experiments with CCSM4 and CESM1/CAM5

2014
- BGC RCP runs
- CESM1 at 0.25°
- Perturbed physics
- CESM1/IAM at 1°

Decadal prediction runs using DART assimilation for initial states: CESM1 at 1°
Long term runs: CESM1 at 0.25° resolution, historical, start RCP future climate change scenario runs
Additional CESM1/BGC emission-driven RCP runs at 1° resolution
Special DOE scenario runs with CCSM4 or CESM1
CESM1/GLIMMER-CISM runs at 1° including Greenland and Antarctic ice sheets
High-resolution time-slice runs with CAM5/SE at 0.125° resolution and test runs at 0.0625° resolution
Perturbed physics ensemble with CESM1 at 1° resolution
Coupled CESM1/IAM runs at 1° resolution
Single-forcing sensitivity experiments with CESM1

2015
- RCPs at 0.25°
- Improved BGC
- Greenland and Antarctic Ice Sheets

Decadal prediction: DART assimilation and CESM1 at 0.25°
Long term: CESM1 at 0.25° RCP future climate change emission scenario runs
CESM1/BGC with improved BGC physics, RCP emission-driven future climate change scenario runs
CESM1/GLIMMER-CISM with improved ice sheet physics and Greenland and Antarctic ice sheets
High-resolution time-slice experiments with CAM5/SE at 0.0625° resolution
Perturbed physics ensemble with CESM1 and 0.25° resolution
Coupled CESM1/IAM runs
Single-forcing sensitivity experiments at 0.25° resolution

2016-2017
- CMIP6 for IPCC AR6
- CAM5/SE at 0.0125°
- Decadal prediction with CESM2

Decadal prediction: DART assimilation and CESM1 at 0.25°
Long term: CESM1 at 0.25° RCP future climate change emission scenario runs
Additional runs with CESM1/BGC , RCP emission-driven future climate change scenario runs
CESM1/GLIMMER-CISM with improved ice sheet physics and Greenland and Antarctic ice sheets
High-resolution time-slice experiments with CAM5/SE at 0.0625° and 0.0125° resolutions
Coupled CESM1/IAM runs
CMIP6 runs for IPCC AR6
New CESM2 at 0.25° using concentrations
CESM2/BGC at 0.25° using emissions
Decadal-prediction experiments with CESM2 at 0.25°
High-resolution time-slice experiments with CAM6/SE at 0.0625° and 0.0125° resolutions

Summary

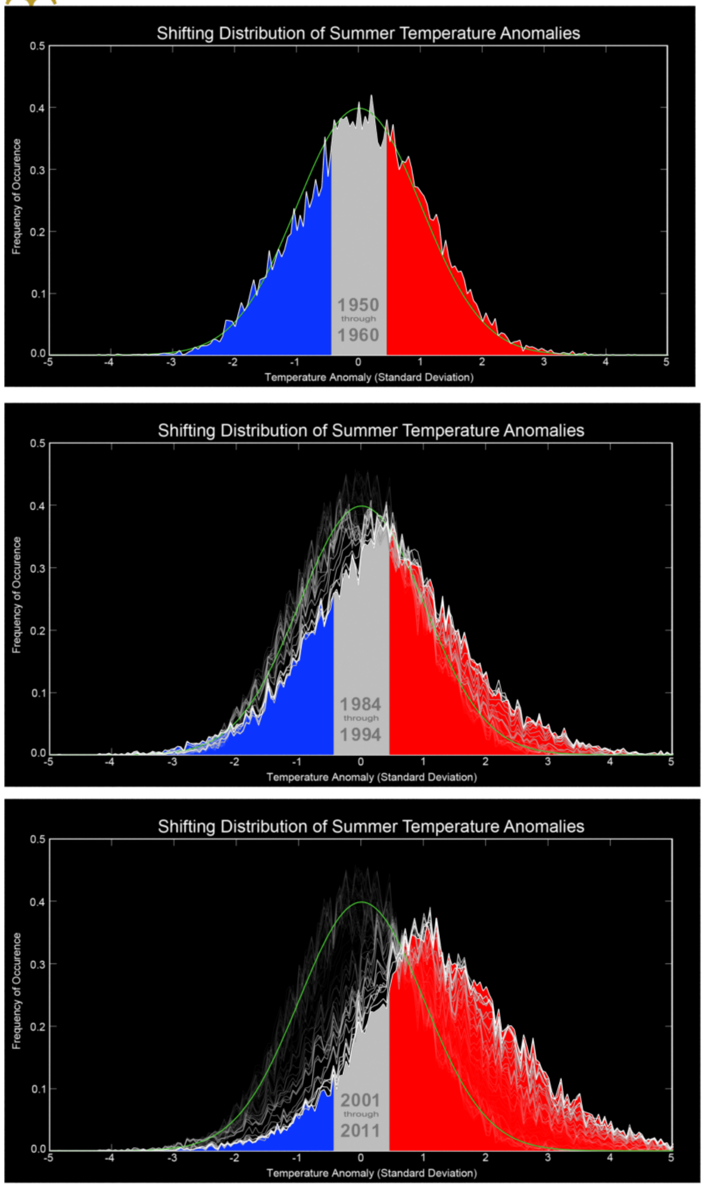

- ***Distribution* of Northern Hemisphere Summer Temperature Anomalies, 1951-2011***
  - **Climate Science - use the same analysis in simulated historical and/or future emission scenarios, <u>at high resolution</u>, to (1) validate model regional climates, and (2) project potential weather related extremes for regional climates with increased confidence.**

* http://svs.gsfc.nasa.gov/goto?3975

Summary

Climate Science – Abrupt Climate Change

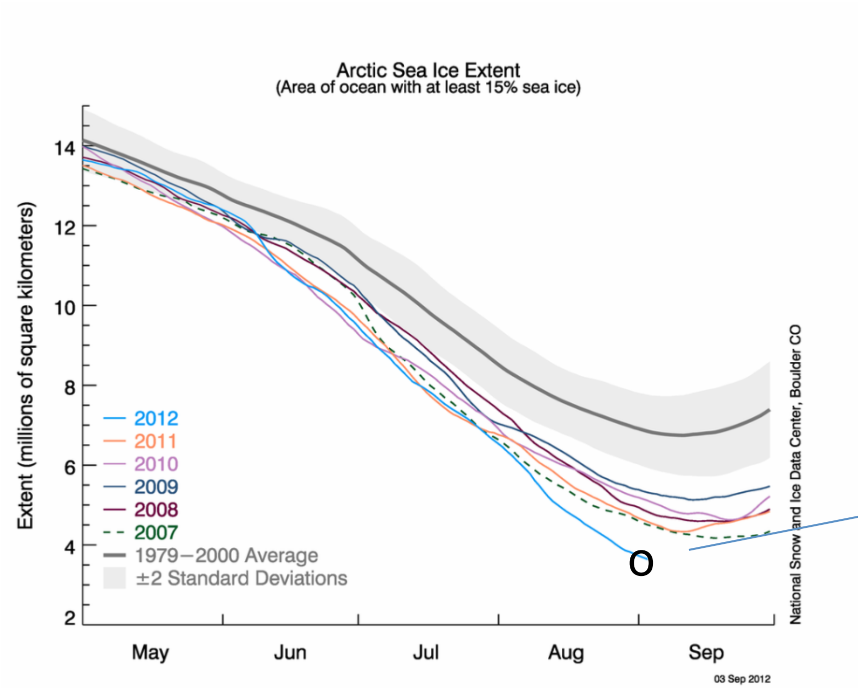

Summer Arctic Ice Extent

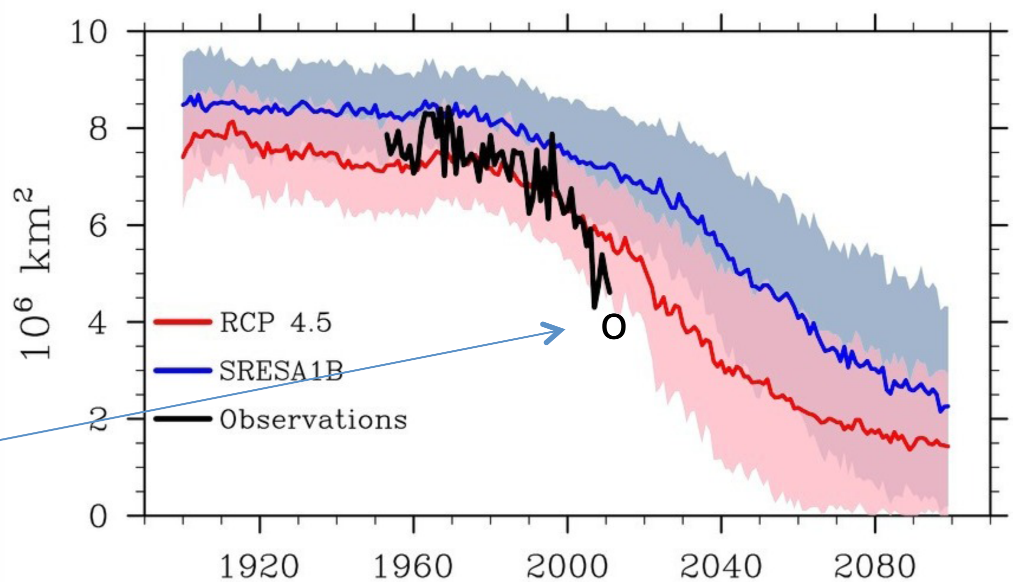

September Ice Extent

1. Project Description

List of Pi(s)/Institution(s)

- Summarize your project(s) and its scientific objectives through 2017
- Our present focus is …
- By 2017 we expect to …

2. Computational Strategies

- We approach this problem computationally at a high level by …
- The codes we use are …
- These codes are characterized by these algorithms: …
- Our biggest computational challenges are …
- Our parallel scaling is limited by …
- We expect our computational approach and/or codes to change (or not) by 2017 in this way …

<u>3. Current HPC Usage</u> (see slide notes)

- Machines currently using
- Hours used in 2012 (list different facilities)
- Typical parallel concurrency and run time, number of runs per year
- Data read/written per run
- Memory used per (node | core | globally)
- Necessary software, services or infrastructure
- Data resources used (HPSS, NERSC Global File System, etc.) and amount of data stored

4. HPC Requirements for 2017

(Key point is to directly link NERSC requirements to science goals)

- Compute hours needed (in units of Hopper hours)
- Changes to parallel concurrency, run time, number of runs per year
- Changes to data read/written
- Changes to memory needed per (core | node | globally)
- Changes to necessary software, services or infrastructure

5. Strategies for New Architectures

- Our strategy for running on new many-core architectures (GPUs or MIC) is …
- To date we have prepared for many core by …
- We are already planning to do …
- To be successful on many-core systems we will need help with …

5. Summary

- What new science results might be afforded by improvements in NERSC computing hardware, software and services?

- Recommendations on NERSC architecture, system configuration and the associated service requirements needed for your science

- NERSC generally refreshes systems to provide on average a 2X performance increase every year. What significant scientific progress could you achieve over the next 5 years with access to 32X your current NERSC allocation?

- What "expanded HPC resources" are important for your project?

- General discussion

Computational Strategies

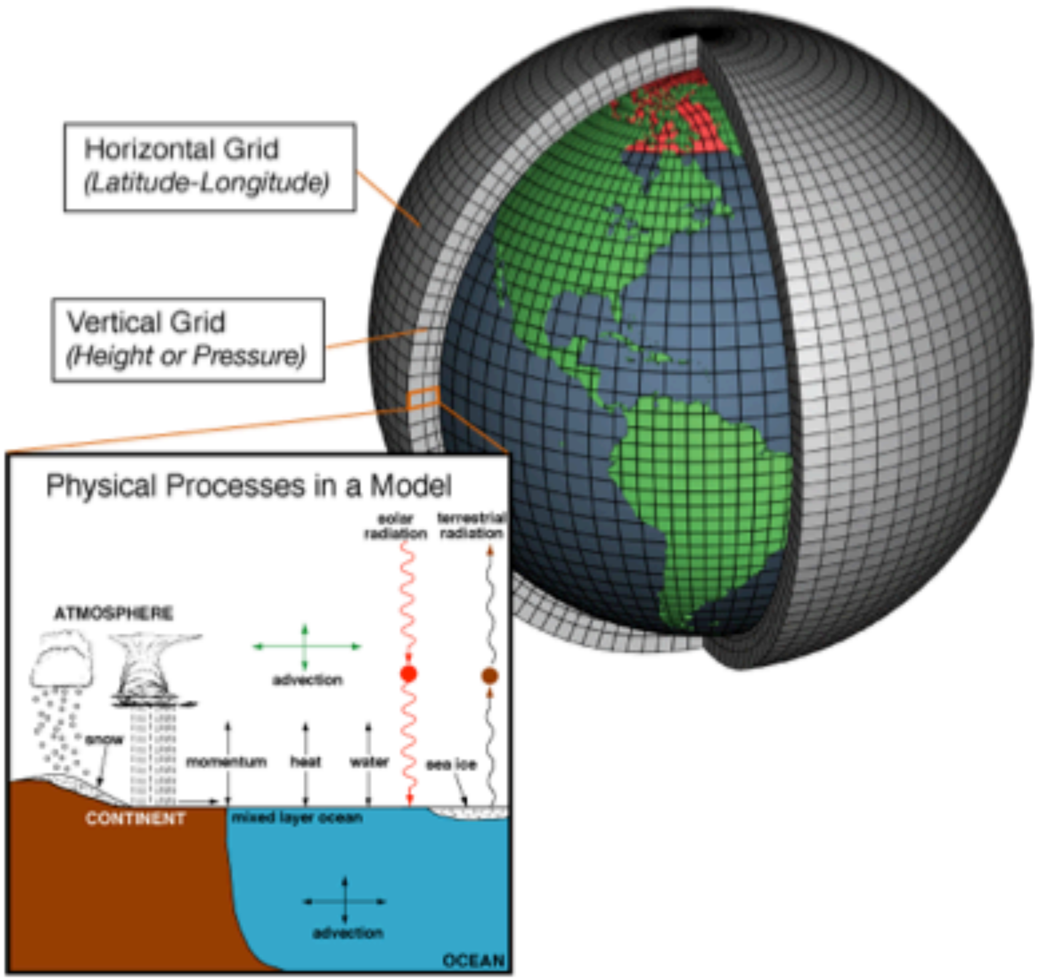

- Systems of differential equations that describe fluid motion, radiative transfer, etc.
- Planet divided into dimensional grid to solve equations
- Atmosphere and land traditionally on same horizontal grid
- Same for ocean and seaice
- Sub-gridscale processes are parameterized

Computational Strategies

- Hardware Platforms – **37** listed in:
  - **www.cesm.ucar.edu/models/cesm1.0/cesm/cesm_doc_1_0_4/a3986.html**
- Major computational challenges
  - hardware platforms
  - MPI and OpenMP parallelism
  - processor layouts
  - mixed sequential/concurrent execution
  - active and data driven components
- The coupler performs data exchange insuring conservation, remapping, diagnostics, timing, etc.

Computational Strategies

CESM 1º Finite Volume Climate Change Model may not want to show this

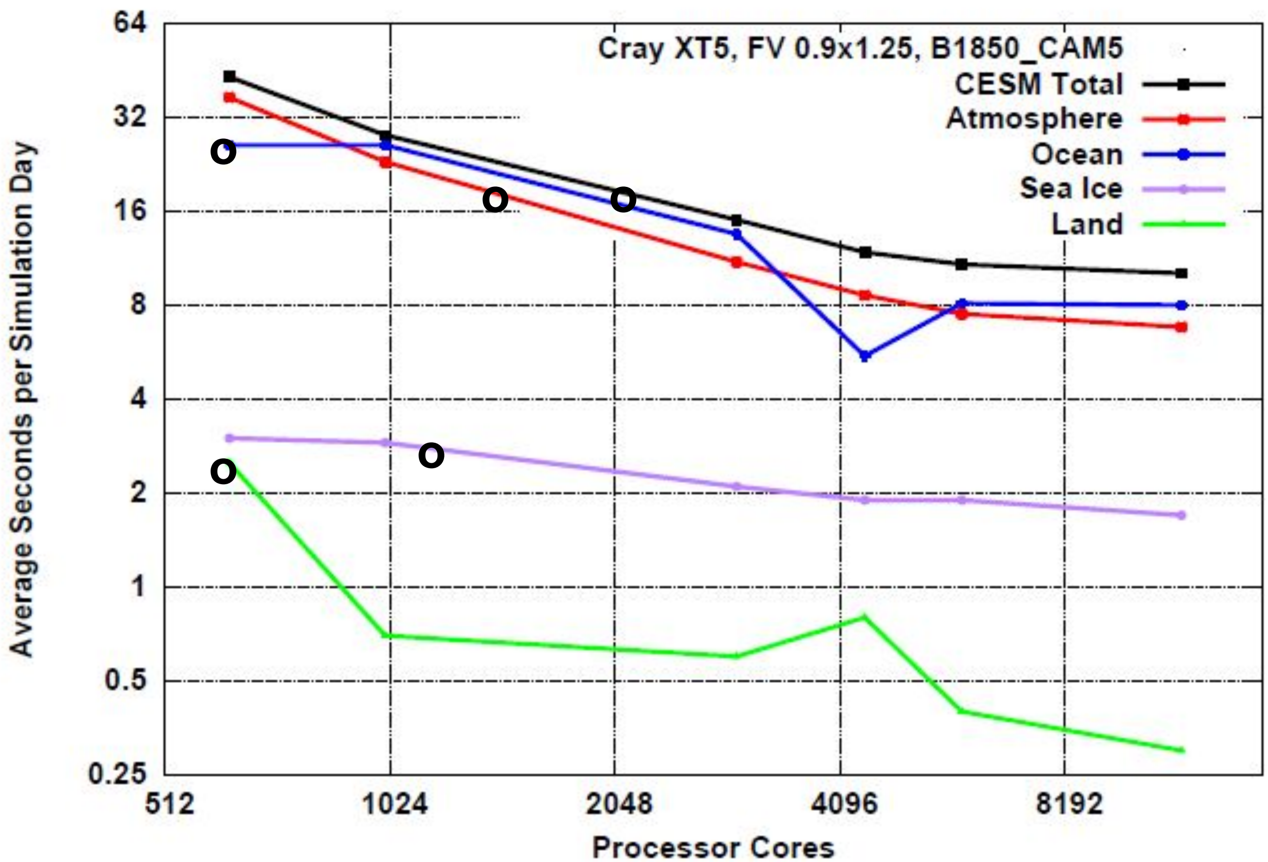

Computational Strategies

- Major CESM codes:
  - CAM .. Community Atmopshere Model
  - POP .. Parallel Ocean Program
  - CLM .. Community Land Model
  - CICE .. Community Ice Model
- Need to insert here what algorithms characterize these models
  - Time dependent forward stepping
  - CAM dynamics
    - Finite Volume (FV) decomposition
    - Spectral Element (SE) decomposition