Frameworks for Climate Modeling

V. Balaji
Princeton University

MODEST-7c Workshop
Drexel University
15 September 2006

GFDL and Princeton University

”Milestones in Scientific Computing”, from Nature (23 March 2006)

Among the milestones listed are:

- 1946 ”ENIAC, ... the first electronic digital computer”
- 1972 ”.. the first hand-held scientific calculator”
- 1989 ”Tim Berners-Lee ... develops the World Wide Web”
- ...
- 1969 Results of the first coupled ocean-atmosphere general circulation model are published by Syukuro Manabe and Kirk Bryan, paving the way for later climate simulations that become a powerful tool in research on global warming.

http://www.nature.com/nature/journal/v440/n7083/full/440399a.html

Weather and climate

“Climate is what you expect, weather is what you get.”

Climate is the study of long-term time averages.

The climate, however, has variations on all time scales: monsoons, ice ages.

More recently, the chemical and radiative properties of the atmosphere has been abruptly altered, with consequences that need to be explored.

Components of the Earth system

Atmosphere atmospheric fluid dynamics and thermodynamics, moist processes, radiative transfer, transport and chemistry of trace constituents.

Ocean World ocean circulation, ocean biogeochemistry.

Land surface Surface processes, ecosystems, hydrology.

Ocean surface Sea ice, wave processes.

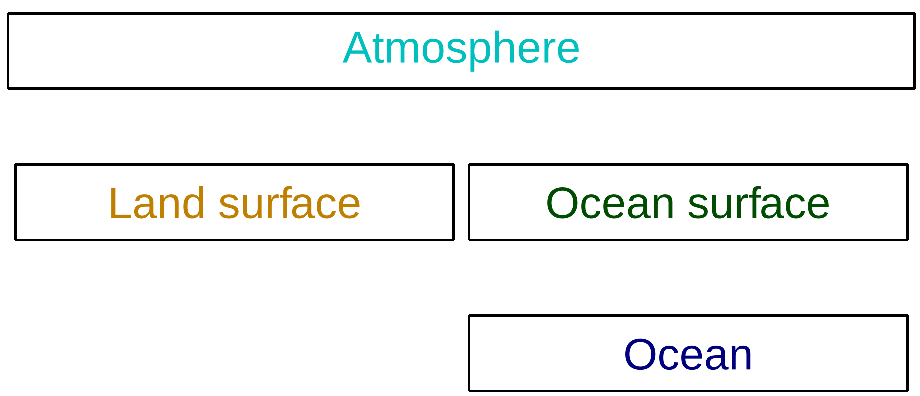

Complexity of climate simulations

Models have grown increasingly complex with time.

70s	80s	early 90s	late 90s	today	early 00s	late 00s
Atm	Atm	Atm	Atm	Atm	Atm	Atm
	Land	Land	Land	Land	Land	Land
		Ocn, SeaIce	Ocn, SeaIce	Ocn, SeaIce	Ocn, SeaIce	Ocn, SeaIce
			Aerosols	Aerosols	Aerosols	Aerosols
					C Cycle	C Cycle
		Aerosols			Ecosystems	Ecosystems
		Land C			Chemistry	Chemistry
	Ocn, SeaIce	Ocn Carbon			Ocn Eddies	Ocn Eddies
	Clouds	Chemistry	C Cycle	Ocn Eddies		Clouds

Components are developed “offline” (bottom left) and then are integrated into comprehensive coupled models.

Timescales and space scales

Aerosols μsecs to minutes; μm – cm.

Clouds minutes to hours; 10 m – 100 km.

Ocean eddies hours to days; 10 km.

Weather patterns days; continent scales.

Ocean currents days; ocean basin scales.

Short-term climate seasonal-interannual; El Niño; monsoons.

Climate change Decadal-centennial; ocean overturning and deep water formation.

Paleoclimate Millennia to deep time; ice ages; orbital changes.

Discretization in space and time

Each process has its own intrinsic time and space scales.

Older models did not allow subcomponents to be on independent grids and timesteps.

- Old way: sharing of data through arrays in common blocks.
- New way: independent model grids connected by a coupler.

Each model ***process*** component becomes an independent ***code*** component that can be separately instantiated, initialized, stepped forward, and terminated.

Gridded and ungridded components

Components may be associated with a grid (gridded component), associated with more than one component (coupler), or none at all (local computations, e.g chemistry).

Each component has its own data dependencies, decomposition strategies and coupling requirements. It must be possible to take full advantage of these features in a parallel code.

“Physics”

Processes that are at space and timescales too small to resolve are treated in parameterized equations: gridscale average response to gridscale average inputs.

Examples:

- Turbulence
- Moist convection
- Radiative transfer
- Boundary layer

These generally can be written as local (pointwise) or have a vertical column orientation: no horizontal dependencies.

This approach has limits: when the physical process becomes non-local on the gridscale. This limitation is most keenly felt now with the treatment of convective clouds in the atmosphere, and turbulent eddies in the ocean. Cloud-resolving AGCMs and eddy-resolving OGCMs to study climate are still about a decade away.

Types of grids

- Spectral grids.
- Lat-lon grids.
- Logically rectangular grids in generalized curvilinear coordinates:
  - Tripolar grids
  - Cubed sphere
- Reduced grids.
- Icosahedral and geodesic grids.
- Unstructured grids; catchment grids.

Any grids currently used in weather and climate modeling may be represented as a **mosaic** of ***logically rectangular grids (LRGs)*** or ***unstructured polygonal grids (UPGs)***.

Data assimilation

Data can appear at unpredictable locations in time and space (radiosondes, buoys, satellites) and have no clear radius of influence ("location streams").

Models will "slosh" if incompatible with data.

Data assimilation involves bringing models and data into acquiescence. Assimilation algorithms can be treated as gridded components.

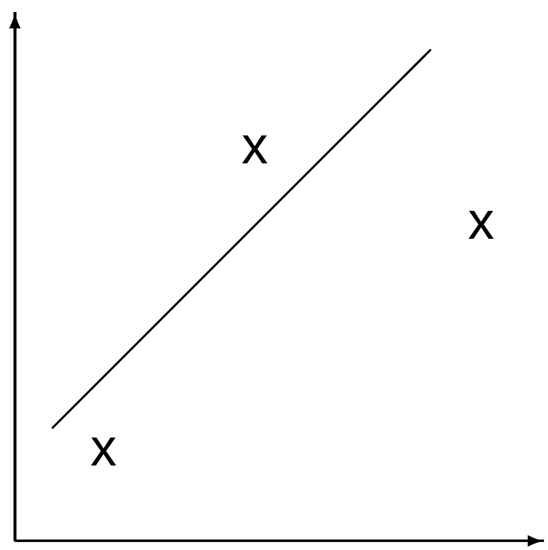

Model sizes and times

Climate Typically measured in simYears/wallDay. Resolution is adjusted till an experiment runs within a time useful for research. Current high-end research: $\sim 10^2$ variables integrated for $\sim 10^6$ timesteps on $\sim 10^6$ grid-points. (Earth Simulator considerably raises this bar, but there is currently no relevant experience to compare against).

Weather Operational forecasting is measured in simHours/wallHour. Resolutions about $4{\times}4$ higher; $\sim 10^3$ timesteps.

Data exchange

Data exchanges between components take the form of fluxes of heat, momentum and tracers.

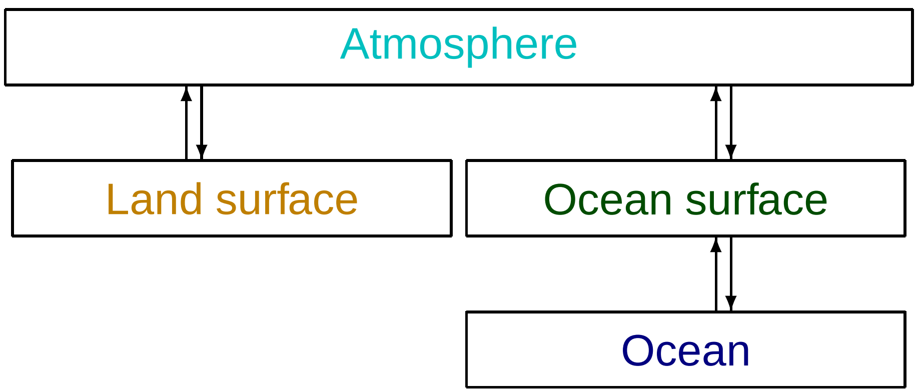

Conservation is very important for climate models; less so for weather.

Memory models

Shared memory signal parallel and critical regions, private and shared variables. Canonical architecture: UMA, limited scalability.

Distributed memory domain decomposition, local caches of remote data (“halos”), copy data to/from remote memory (“message passing”). Canonical architecture: NUMA, scalable at cost of code complexity.

Distributed shared memory or ccNUMA message-passing, shared memory or remote memory access (RMA) semantics. Processor-to-memory distance varies across address space, must be taken into account in coding for performance. Canonical architecture: cluster of SMPs. Scalable at large cost in code complexity.

A 2D example

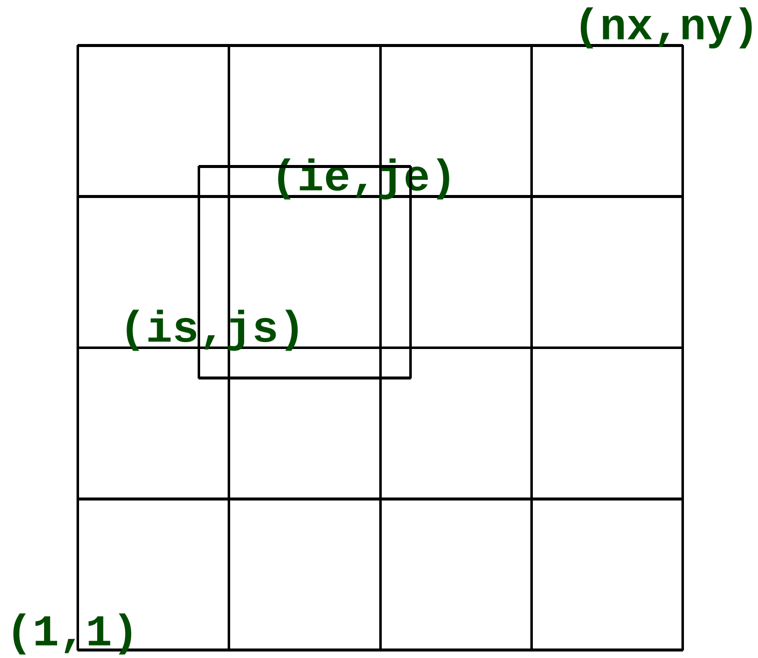

Consider a platform consisting of 16 PEs consisting of 4 mNodes of 4 PEs each. We also assume that the the entire 16-PE platform is a DSM or ccNUMA aNode. We can then illustrate 3 ways to implement a **`DistributedArray`**. One PET is scheduled on each PE.

Distributed memory

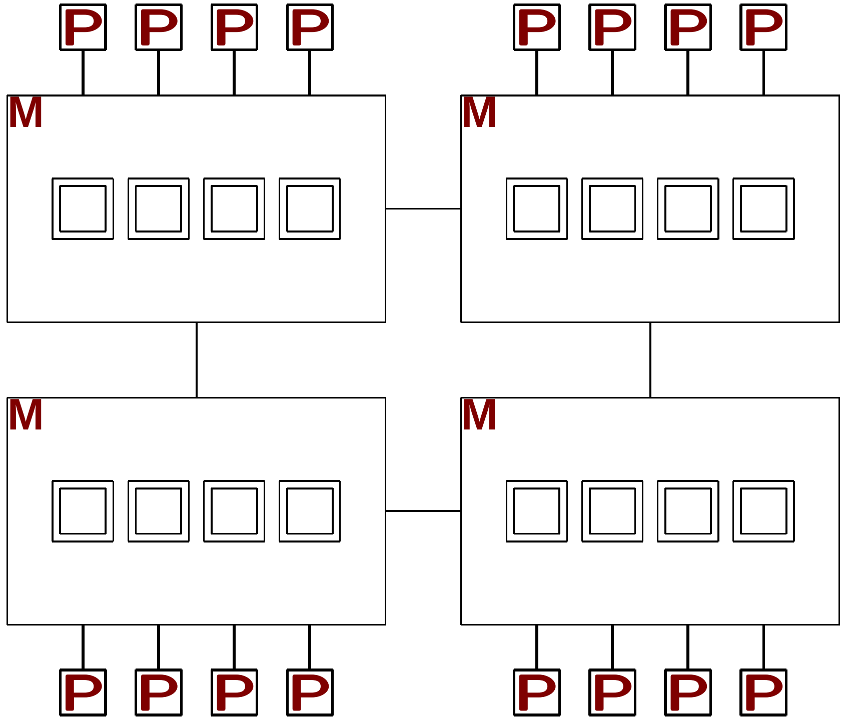

- each domain allocated as a separate array with halo, even within the same mNode.
- Performance issues: the message-passing call stack underlying MPI or another library may actually serialize when applied within an mNode.

Hybrid memory model

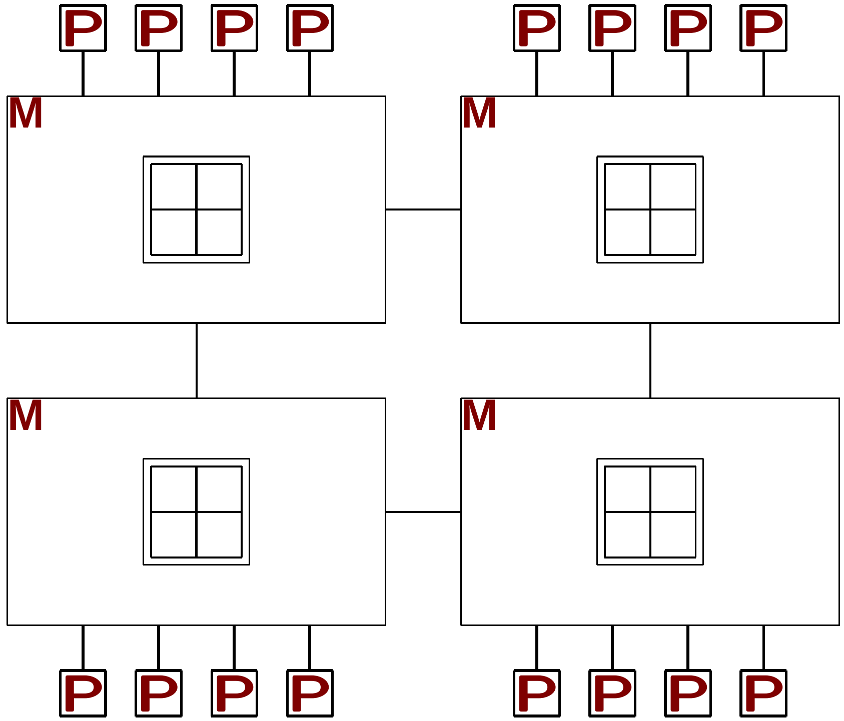

- shared across an mNode, distributed among mNodes.
- fewer and larger messages than distributed memory, may be less latency-bound.

Pure shared memory

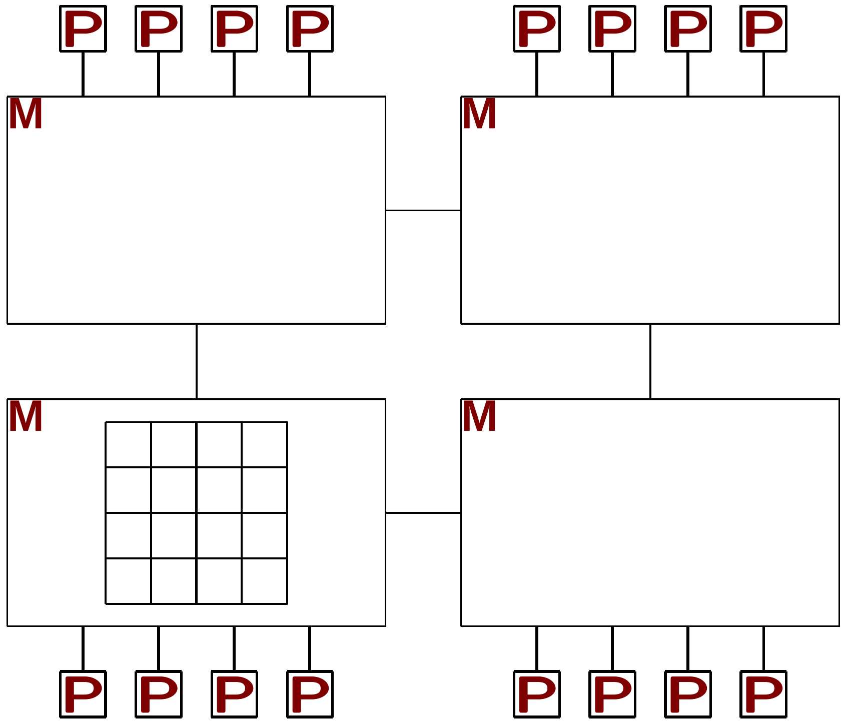

Array is local to one mNode: other mNodes requires remote loads and stores. OK on platforms that are well-balanced in bandwidth and latency for local and remote accesses. ccNUMA ensures cache coherence across the aNode.

Intelligent memory allocation on DSM

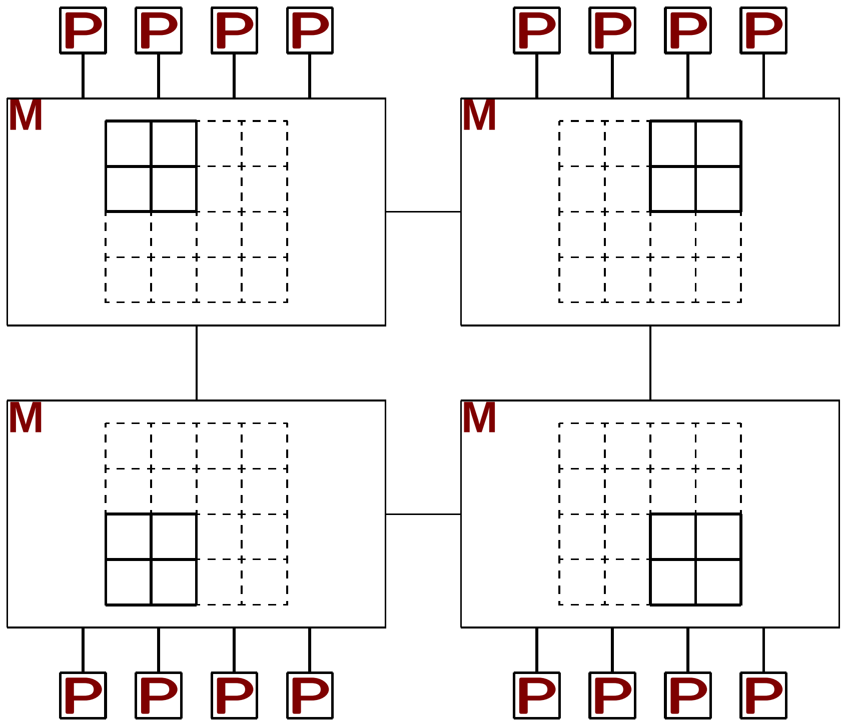

Better memory locality: allocate each block of 4 domains on a separate page, and assign pages to different mNodes, based on processor-memory affinity.

Uniform interface to memory models: DEs and Layouts

Decomposition Element A DE is the logical atom of a distributed object, which may or may not map onto a single PE.

Layout A layout describes how DEs are organized along some logical dimensions. We may choose to map these dimensions on to shared and distributed memory.

Virtual Machine A VM describes the hardware and applicable memory semantics: a description of the architecture in terms of **aNode**s and **mNode**s.

Uniform interface to the grid index space: the `distGrid`

A **distGrid** describes how to distribute the index space of a component across a DE layout, and most importantly, the ***data dependencies*** between DEs.

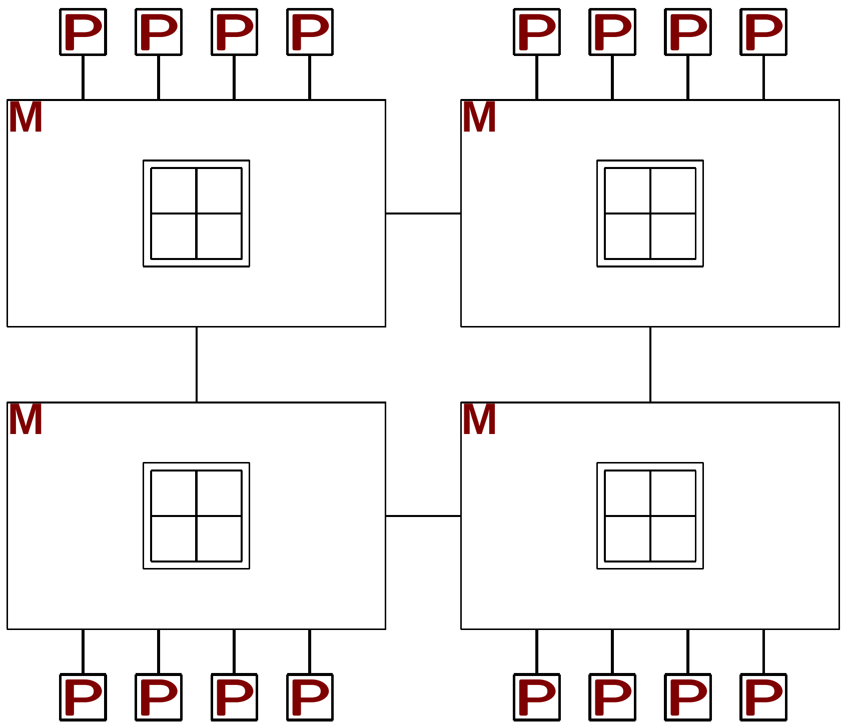

Halo updates and redistributions are typical operations performed using **`distGrid`**s.

Overlay physical information on the `distGrid`: `physGrid`

All the model numerics require knowledge of the physical locations, orientations, and interrelations of array indices.

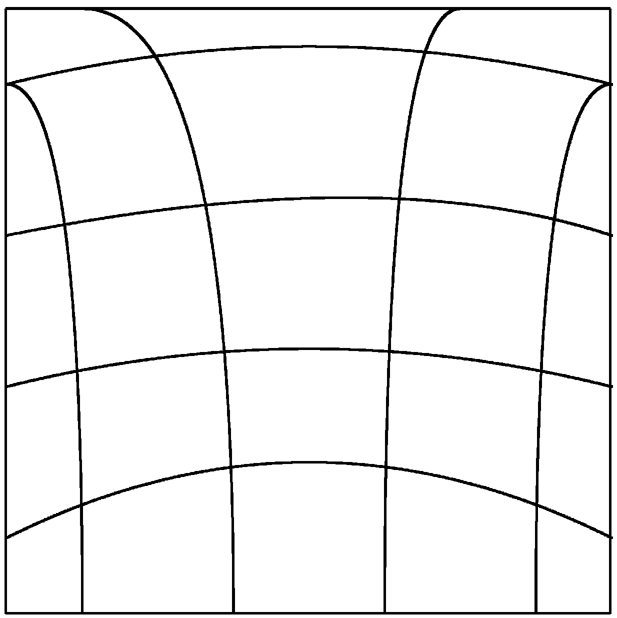

Differential operators (gradient, divergence) are typical operations requiring **`physGrid`**s.

Interpolate data between grids: reGrid

When components on independent **physGrid**s must exchange data, they require a **reGrid**.

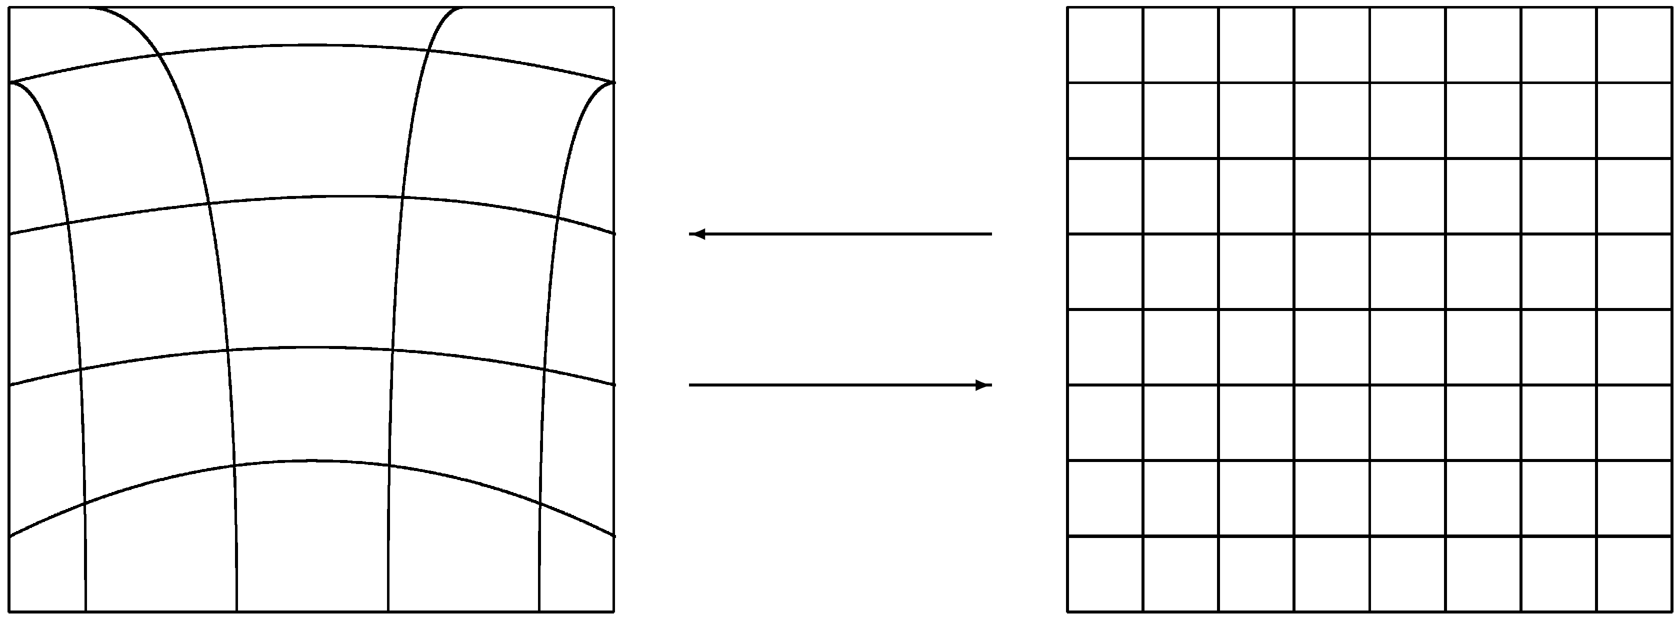

On parallel systems, we require all calls to be local; the **reGrid** object contains the optimized **route**s for communication.

Fields and bundles

Alongside the grid information, each model variable is present as a distributed array. It is useful, indeed essential, to associate **metadata** that describe the physical content (name, units, range of valid values) of the array. All this information together comprises a model **Field**.

Fields that share a grid may be treated collectively for certain grid operations. A collection of fields sharing a grid is called a **bundle**.

Model composition

We develop a hierarchy of components and assign them to PEs as appropriate.

A **coupler** mediates between pairs of components needing to exchange data, and matches fields in their **import** and **export States**.

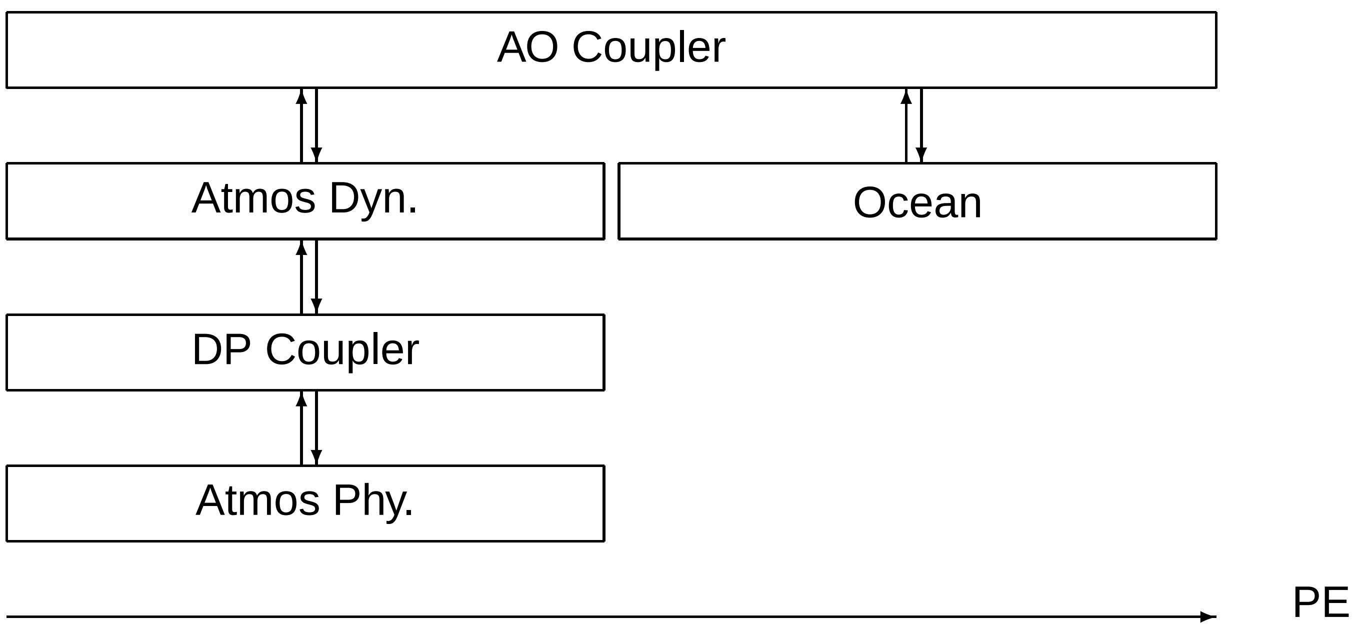

The coupler runs on the union of PEs of its components.

Architecture of an Earth System Modeling Framework: the sandwich

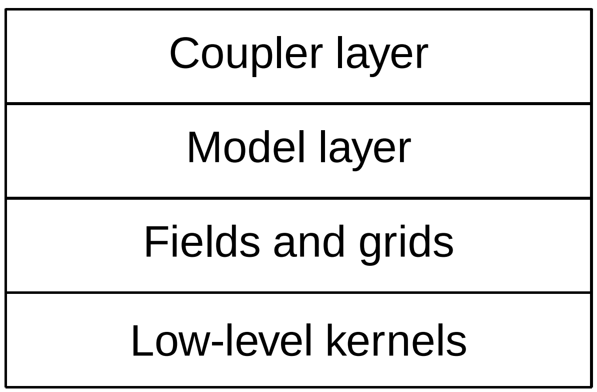

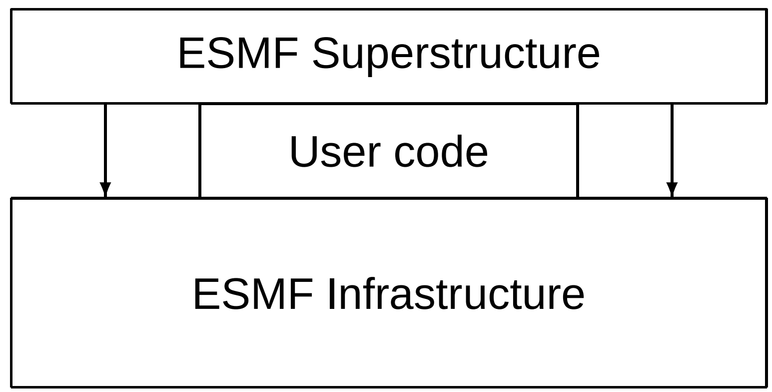

ESMF

The Earth System Modeling Framework is an end-to-end solution for the problem outlined here: supporting distributed development of models with many interacting components, with independent space and time discretization, running on complex modern scalable architectures. A capsule history of ESMF:

- The need to unify and extend current frameworks achieves wide currency (c. 1998).
- NASA offers to underwrite the development of an open community framework (1999).
- A broad cross-section of the community meets and agrees to develop a concerted response to NASA uniting coupled models, weather, and data assimilation in a common framework (August 1999). Participants include NASA/GMAO, NOAA/GFDL, NOAA/NCEP, NCAR, DOE, and universities with major models.
- Funding began February 2002: $10 M over 3 years.
- Current: ESMF superstructure deployed in several multi-component assemblies; ESMF infrastructure in active development for a uniform interface to distributed grids and arrays. Inter-agency Working Group established to coordinate funding; Joint Specification Team and Change Review Board established.

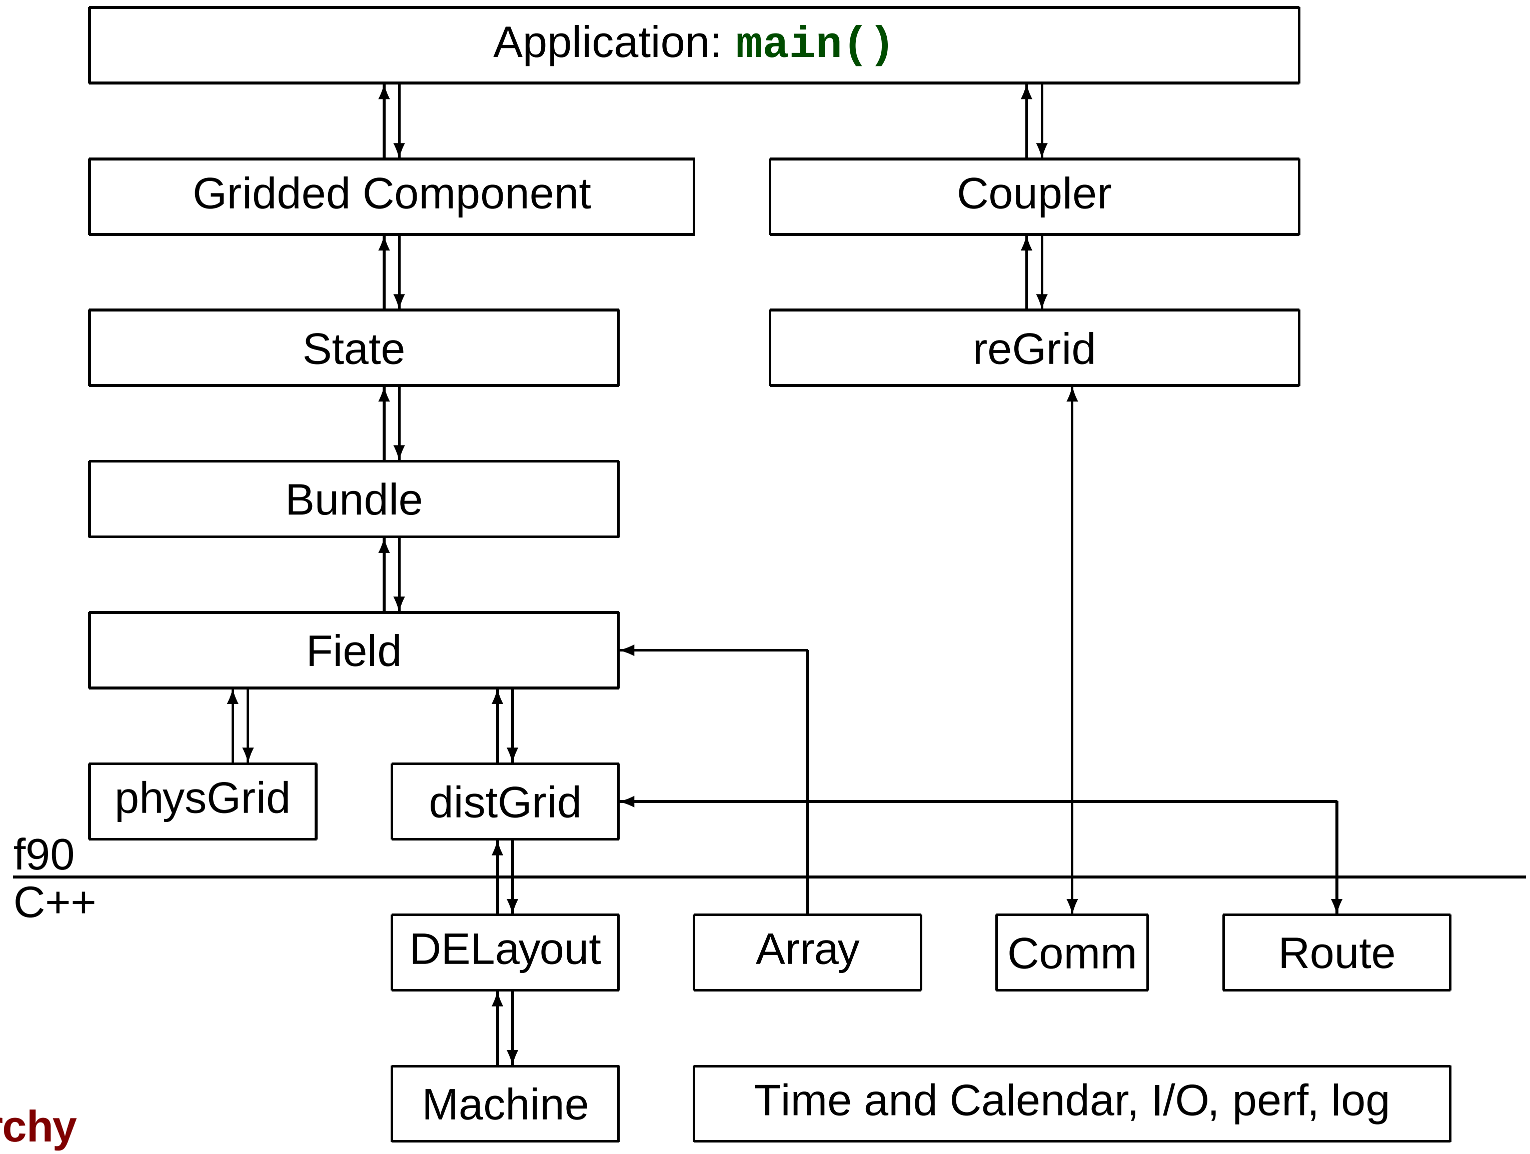

ESMF Class Hierarchy

Making an ESMF component

- Register the **`init`**/**`run`**/**`exit`** methods.

```
type(ESMF_GridComp) ::  comp
...
call ESMF_GridCompSetEntryPoint(comp,ESMF_SETRUN,update_ocean_model,...
```

(1)

- Create the import and export states from component internal datatypes. This can be done entirely with pointers: no data movement is involved.

```
ocean%u_surf=>expFMSArray(1)%ptr
...
expESMF_Field(i)=ESMF_FieldCreate(ocnGrid,expFMSArray(i)%ptr,...)
call ESMF_StateAddField(expState,expESMF_Field,...)
```

(2)

- Run the component.

```
type(ESMF_GridComp) ::  compOcn
...
compOcn=ESMF_GridCompCreate(vm,"Ocean",rc=rc)
call ESMF_GridCompSetServices(compOcn,OceanRegister,rc)
...
call ESMF_GridCompRun(compOcn,impOcn,expOcn,topClock,rc=rc)
```

(3)

ESMF features

- ESMF is usable by models written in f90/C++.
- ESMF is usable by models requiring differentiability.
- ESMF is usable by models using shared or distributed memory parallelism semantics.
- ESMF supports serial and concurrent coupling.
- ESMF supports multiple I/O formats (including GRIB/BUFR, netCDF, HDF, native binary).
- ESMF has uniform syntax across platforms.
- ESMF runs on many platforms spanning desktops (laptops, even!) to supercomputers.

Language interoperability

Models increasingly are built to use high-level abstract language features (f90, C++) to facilitate development process across large teams. ESMF is written to be usable by models written in both languages.

Incompatibilities:

- Languages use entirely different memory semantics. ESMF distinguishes between “deep” and “shallow” objects: depending on whether memory is managed by callee or caller.
- High-level data structures do not cross language boundaries. ESMF implements cross-language data structures by implementing an intermediate layer in a low-level language (f77, C) where intrinsic datatypes can be cross-matched. This includes a unique implementation allowing f90 array pointers to be manipulated in either language.

The ESMF **Implementation Report** presents details of its cross-language technology.

Recent developments in methodology

- Future projections of climate are performed at many sites, and a key goal of current research is to reduce the uncertainty of these projections by understanding the differences in the output from different models. This ***comparative study of climate simulations*** (e.g IPCC, ENSEMBLES, APE) across many models has spawned efforts to build uniform access to output datasets from major climate models, as well as modeling frameworks that will promote uniform access to the models themselves.

- As hardware and software complexity increase, we seek to encapsulate scalable data-sharing layers within an ***infrastructure***. Components of the physical climate system are now also code components, with coupling embedded in a standardized ***superstructure***. This has led to the emergence of Earth system modeling ***frameworks***, of which ESMF and PRISM are leading examples.

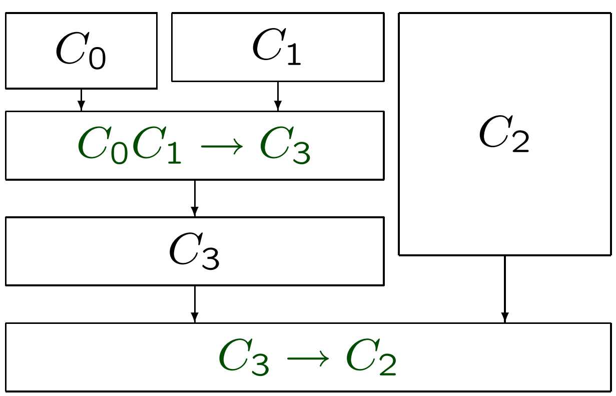

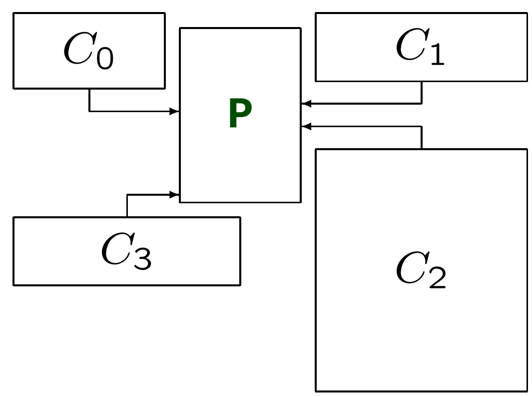

The IPCC AR4 archive at PCMDI

The IPCC data archive at PCMDI is a truly remarkable resource for the comparative study of models. Since it came online in early 2005, it has been a resource for $\sim$300 scientific papers aimed at providing consensus and uncertainty estimates of climate change, from $\sim$20 state-of-the-art climate models worldwide.

Model	Modeling Center
BCCR BCM2	Bjerknes Centre for Climate Research
CCCMA CGCM3	Canadian Centre for Climate Modeling & Analysis
CNRM CM3	Centre National de Recherches Meteorologiques
CSIRO MK3	CSIRO Atmospheric Research
GFDL CM2_0	Geophysical Fluid Dynamics Laboratory
GFDL CM2_1	Geophysical Fluid Dynamics Laboratory
GISS AOM	Goddard Institute for Space Studies
GISS EH	Goddard Institute for Space Studies
GISS ER	Goddard Institute for Space Studies
IAP FGOALS1	Institute for Atmospheric Physics
INM CM3	Institute for Numerical Mathematics
IPSL CM4	Institut Pierre Simon Laplace
MIROC HIRES	Center for Climate System Research
MIROC MEDRES	Center for Climate System Research
MIUB ECHO	Meteorological Institute University of Bonn
MPI ECHAM5	Max Planck Institute for Meteorology
MRI CGCM2	Meteorological Research Institute
NCAR CCSM3	National Center for Atmospheric Research
NCAR PCM1	National Center for Atmospheric Research
UKMO HADCM3	Hadley Centre for Climate Prediction

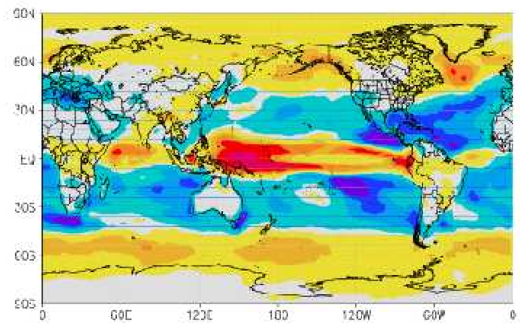

This figure, from Held and Soden (2005), is a composite analysis across the entire IPCC archive.

Computational load at GFDL:

- 5500 model years run.
- Occupied half of available compute cycles at GFDL for half a year (roughly equivalent to 1000 Altix processors).
- 200 Tb internal archive; 40 Tb archived at GFDL data portal; 4 Tb archived at PCMDI data portal.

I would argue that the IPCC experiment is ***already*** petascale!

The routine use of Earth System models in research and operations

Let's declare that 2000-2010 (the "noughties") is the decade of the coming-of-age of Earth system models.

Operational forecasting model-based seasonal and inter-annual forecasts delivered to the public;

Decision support models routinely run for decision support on climate policy by governments, for energy strategy by industry and government, as input to pricing models by the insurance industry, etc.

Fundamental research the use of models to develop a predictive understanding of the earth system and to provide a sound underpinning for all applications above.

This requires a radical shift in the way we do modeling: from the current dependence on a nucleus of very specialized researchers to make it a more accessible general purpose toolkit. This requires ***an infrastructure for moving the building, running and analysis of models and model output data from the "heroic" mode to the routine mode***.

From heroic to routine in other fields

The ***polymerase chain reaction*** was awarded a Nobel prize not long ago. Later, you could get a PhD for developing PCR in different contexts. Now you order online and receive samples through the mail...
Transgenic implants in different organisms are another example... below, you see a service provided by a lab at Princeton University which will develop and store transgenic mice and other organisms.

What will the transition from heroic to routine look like in our field?

The IPCC data pipeline at GFDL

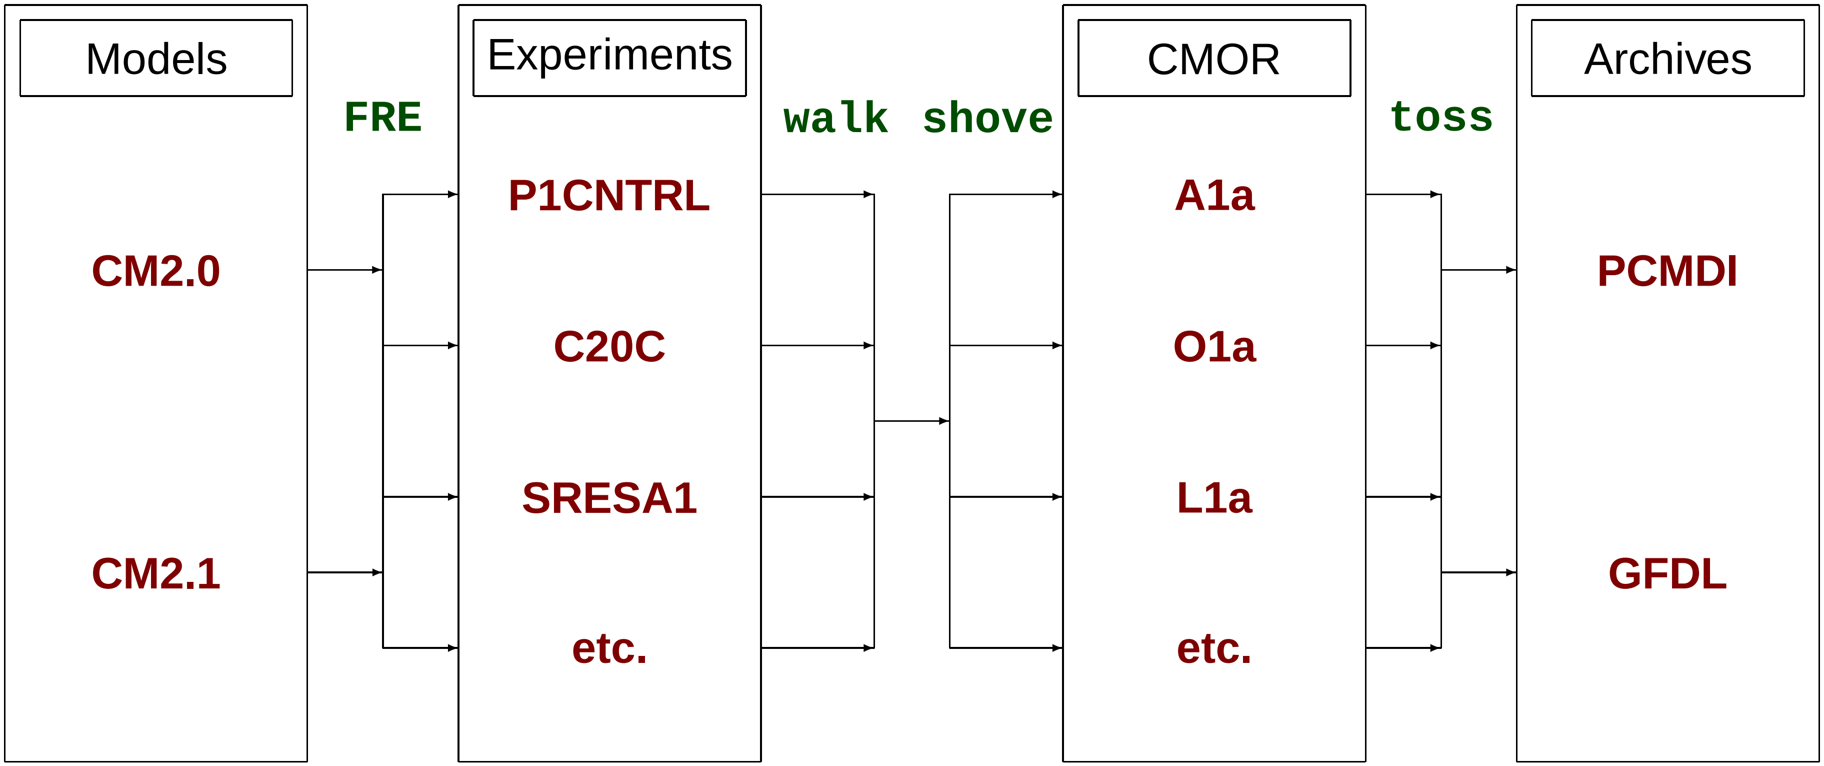

The process was time- and data-intensive, with multiple access episodes for the same datasets. Clearly it would be ideal if FRE already produced compliant data.

Current problems with CMOR-compliant data

- A principal difficulty is CMOR's restricted view of model grids: only simple latitude-longitude grids are permitted. This is because the current crop of visualization and analysis tools cannot easily translate data among different grids. Shown at right are the ***tripolar grid*** (Murray 1996, Griffies et al 2004) used by MOM4 for GFDL's current IPCC model CM2. Below is the ***cubed sphere*** (Rancic and Purser 1990) planned for the Finite-Volume atmosphere dynamical core for the next-generation GFDL models AM3 and CM3. If there were a ***grid metadata standard***, regridding operations could potentially be applied by the end-user using standard-compliant tools.
- The model descriptions demanded by CMOR do not contain enough information about the models, and are added after the fact. If there were a ***model metadata standard*** such as NMM in force, comprehensive model descriptions could be automatically produced. The end-user could better diagnose specific differences between different models in an archive.

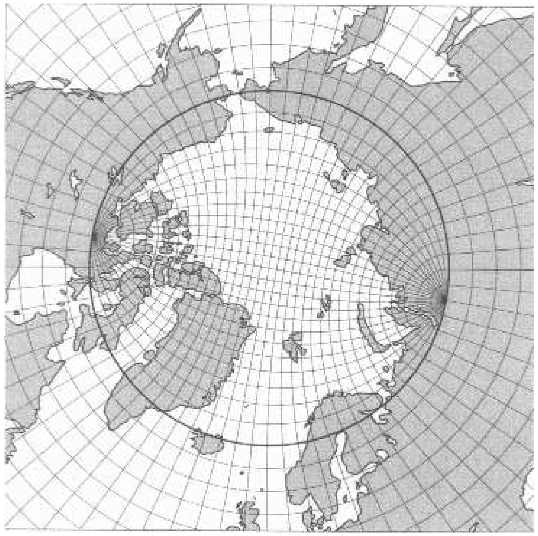

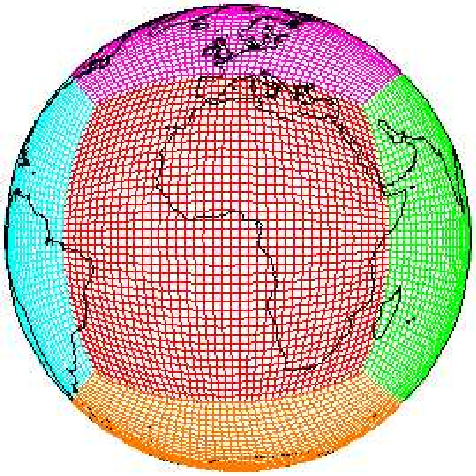

Can an experiment like IPCC be run at higher resolution?

Possible key challenges for the next IPCC:

- Robust estimates of regional climate change.
- Interactive carbon dynamics: inclusion of land-use change, ocean carbon uptake, marine and terrestrial biospheric response to climate change.
- Increased resolution in the atmosphere (even before we get to cloud-resolving scales) will lead to better characterization of storm track changes and hurricane intensity projections in a changed climate. Target: 1°or 0.5°model for IPCC AR5.
- Increased resolution in the ocean is even more critical: key mechanisms of ocean mass and energy transport are currently unresolved. Targets: 0.25°("eddy-permitting") models next time around, 0.0825°("eddy-resolving") still out of reach.

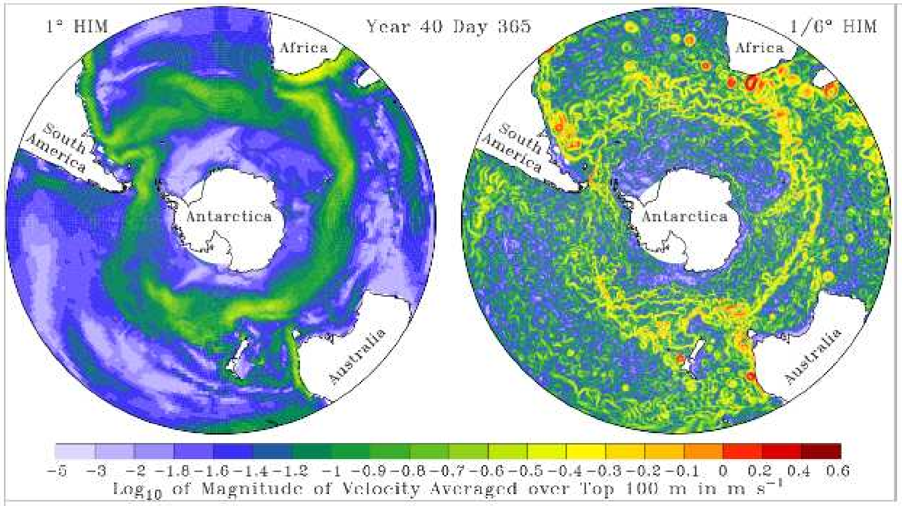

Shaded is the speed of the surface currents in the ocean model at two resolutions.

Petascale methodologies

As much emphasis must be placed on methodologies to facilitate scientific analysis of multi-model ensembles on distributed data archives, as on the computational technology itself.

Some current efforts:

ESC Earth System Curator, funded by NSF. Partners GFDL, NCAR, PCMDI, Georgia Tech. Will be used to promote the existence of a model and grid metadata standard, and build a prototype relational database housing these metadata. Will build tools for model configuration and compatibility checking based on automatic harvesting of metadata from code.

MAPS Modeling, Analysis and Prediction System? funded by NASA, partners NASA/GSFC, GFDL, MIT. Proposes to build a configuration layer for a subset of coupled models based on PRISM config files, and conformant with grid and metadata standards. Will attempt to promote a "standard coupling architecture" and develop a standard for exchange grids for ESMF.

GO-ESSP and CF should be the medium of exchange for standard-building. CF is seeking funding and WGCM backing to become a mandated activity. GO-ESSP is the ideal medium for the actual technical work of standard-building.

IPCC! PCMDI and other data centres should be core participants.

With a complete metadata hierarchy defined, one can envisage the convergence of modeling and data frameworks into a single environment: a model ***curator***.

Scenario 1: dynamically generated data catalogues

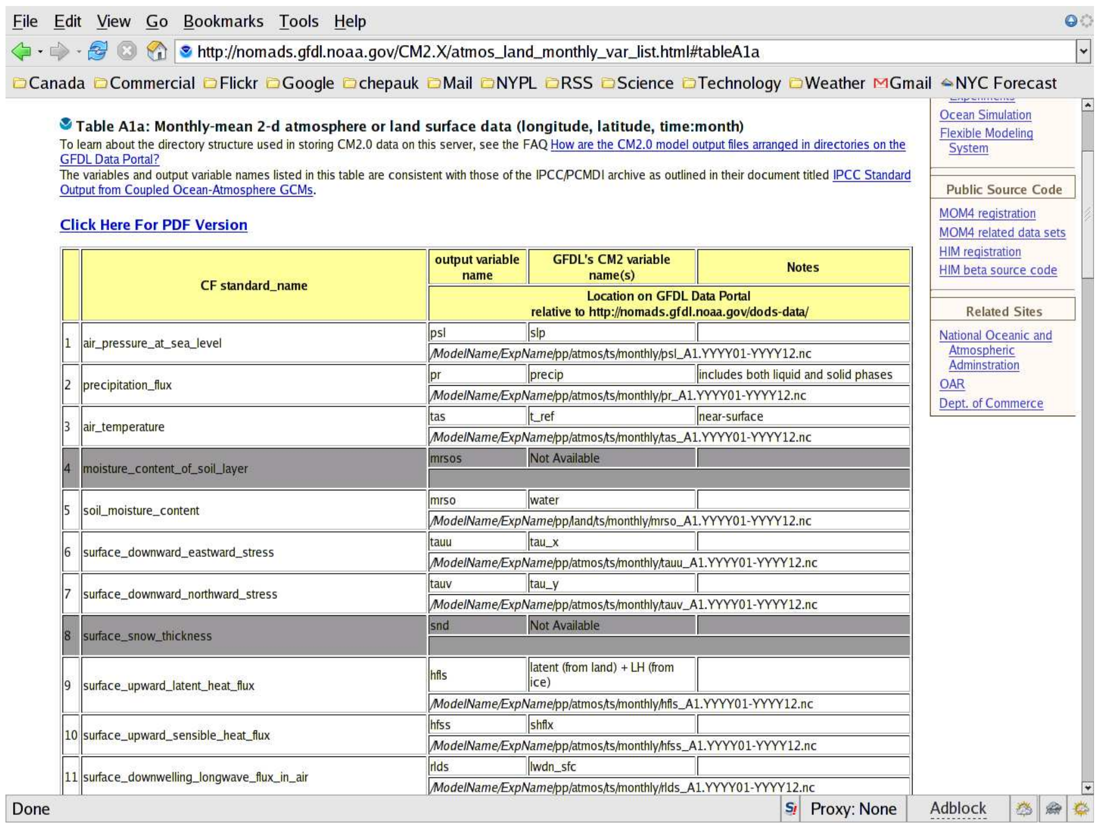
File Edit View Go Bookmarks Tools Help

http://nomads.gfdl.noaa.gov/CM2.X/atmos_land_monthly_var_list.html#tableA1a

Canada Commercial Flickr Google chepauk Mail NYPL RSS Science Technology Weather Gmail NYC Forecast

Table A1a: Monthly-mean 2-d atmosphere or land surface data (longitude, latitude, time:month)

To learn about the directory structure used in storing CM2.0 data on this server, see the FAQ How are the CM2.0 model output files arranged in directories on the GFDL Data Portal?

The variables and output variable names listed in this table are consistent with those of the IPCC/PCMDI archive as outlined in their document titled IPCC Standard Output from Coupled Ocean-Atmosphere GCMs.

Click Here For PDF Version

	CF standard_name	output variable name	GFDL's CM2 variable name(s)	Notes	Location on GFDL Data Portal relative to http://nomads.gfdl.noaa.gov/dods-data/
1	air_pressure_at_sea_level	psl	slp		/ModelName/ExpName/pp/atmos/ts/monthly/psl_A1.YYYY01-YYYY12.nc
2	precipitation_flux	pr	precip	includes both liquid and solid phases	/ModelName/ExpName/pp/atmos/ts/monthly/pr_A1.YYYY01-YYYY12.nc
3	air_temperature	tas	t_ref	near-surface	/ModelName/ExpName/pp/atmos/ts/monthly/tas_A1.YYYY01-YYYY12.nc
4	moisture_content_of_soil_layer	mrsos	Not Available		
5	soil_moisture_content	mrso	water		/ModelName/ExpName/pp/land/ts/monthly/mrso_A1.YYYY01-YYYY12.nc
6	surface_downward_eastward_stress	tauu	tau_x		/ModelName/ExpName/pp/atmos/ts/monthly/tauu_A1.YYYY01-YYYY12.nc
7	surface_downward_northward_stress	tauv	tau_y		/ModelName/ExpName/pp/atmos/ts/monthly/tauv_A1.YYYY01-YYYY12.nc
8	surface_snow_thickness	snd	Not Available		
9	surface_upward_latent_heat_flux	hfls	latent (from land) + LH (from ice)		/ModelName/ExpName/pp/atmos/ts/monthly/hfls_A1.YYYY01-YYYY12.nc
10	surface_upward_sensible_heat_flux	hfss	shflx		/ModelName/ExpName/pp/atmos/ts/monthly/hfss_A1.YYYY01-YYYY12.nc
11	surface_downwelling_longwave_flux_in_air	rlds	lwdn_sfc		/ModelName/ExpName/pp/atmos/ts/monthly/rlds_A1.YYYY01-YYYY12.nc

Ocean Simulation

Flexible Modeling System

Public Source Code

MOM4 registration

MOM4 related data sets

HIM registration

HIM beta source code

Related Sites

National Oceanic and Atmospheric Adminstration

OAR

Dept. of Commerce

Done Proxy: None Adblock

Already in use at PCMDI, DDC, GFDL Curator, elsewhere: metadata requires extension.

Scenario 2: statistical downscaling of climate change projections

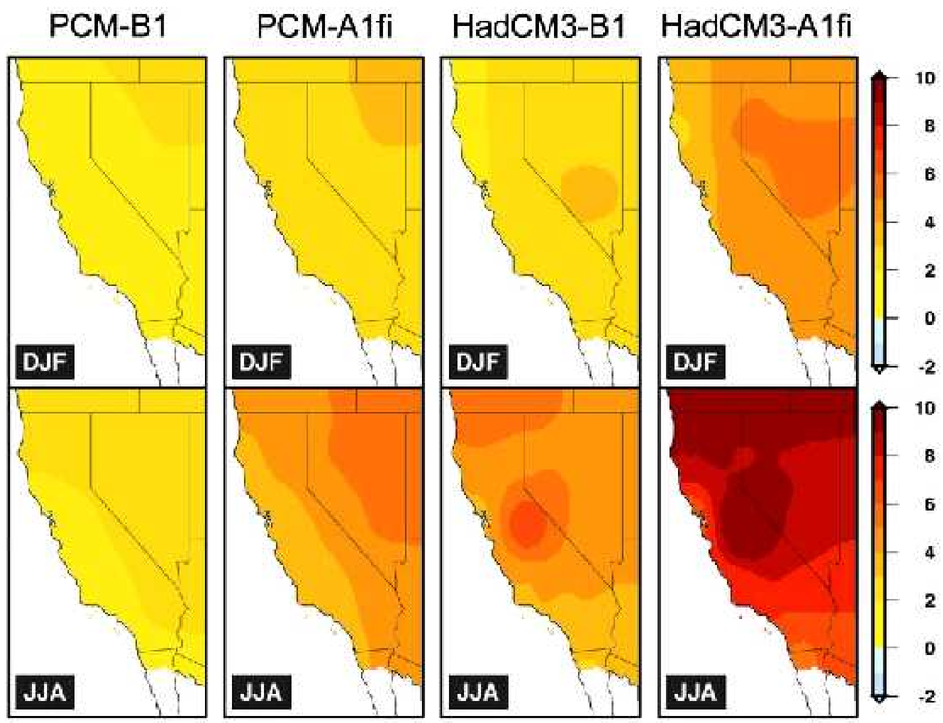

Hayhoe et al, ***PNAS***, 2004: ***Emissions pathways, climate change, and impacts on California.***

Uses daily data for “heat degree days” and other derived quantities. Requires data beyond that provided by IPCC AR4 SOPs (1960-2000).

Scenario 3: disease vectors in a changing climate

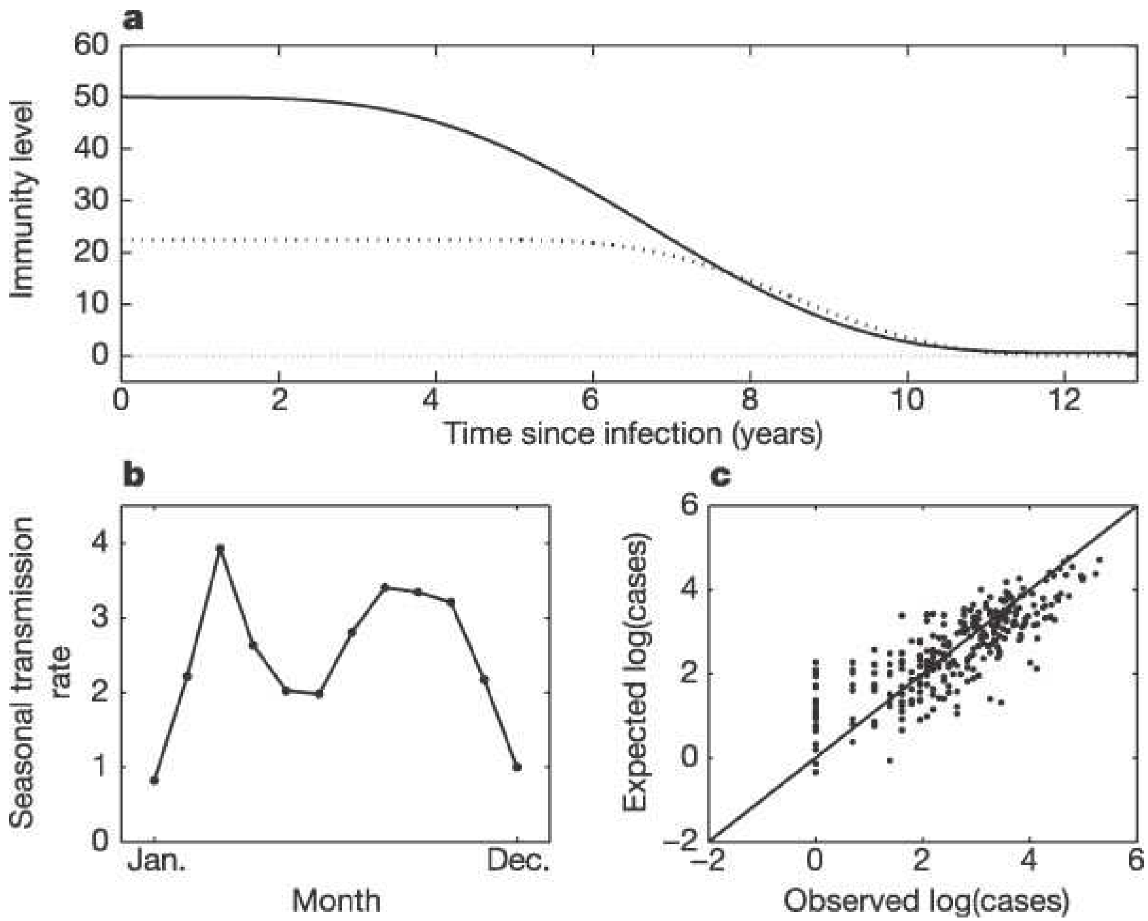

Koelle et al, ***Nature***, 2005: ***Refractory periods and climate forcing in cholera dynamics.***

Requires monthly forcing data, no feedback.

Scenario 4: alternate energy sources

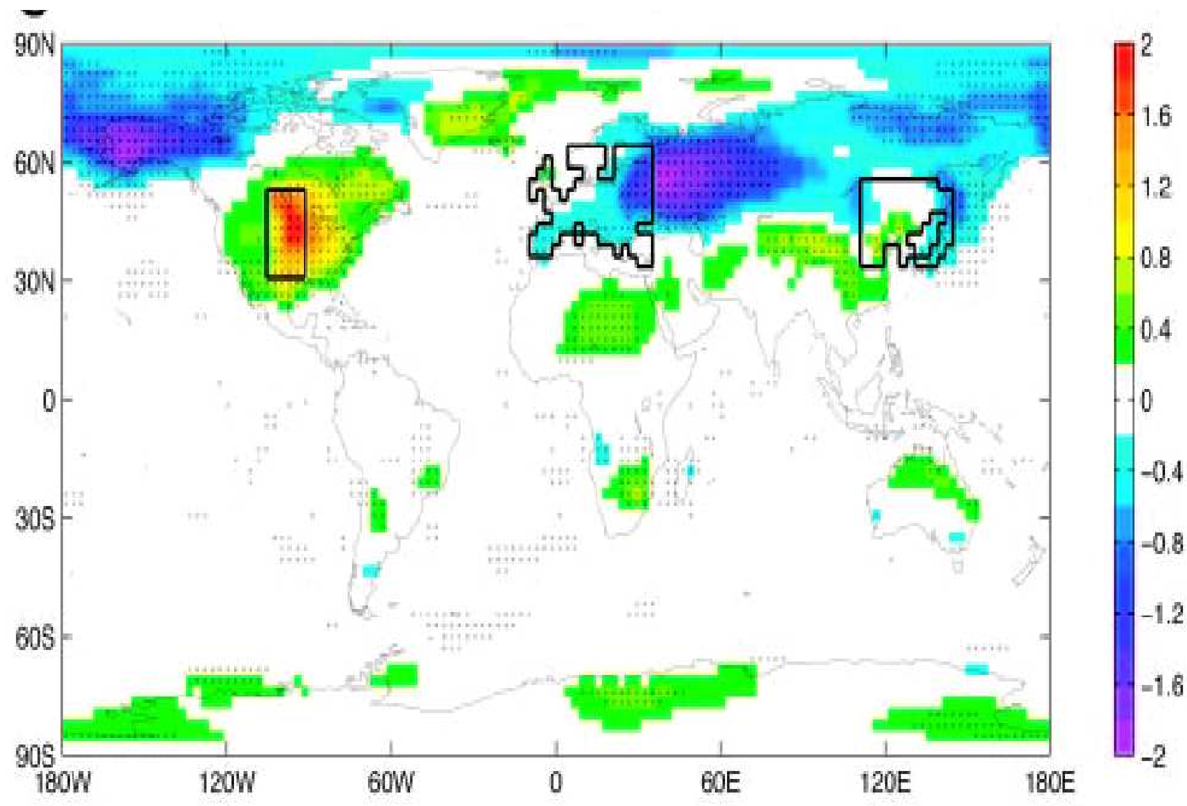

Keith et al, ***PNAS***, 2005: ***The influence of large-scale wind power on global climate.***

Feedback on atmospheric timescales: but does not require model to be retuned.

Taking stock halfway through the noughties

- Earth system models are evolving into powerful tools for advancing our understanding, and well on their way to being operational tools in support of policy and industrial strategy.
- The principal research path for consensus and uncertainty estimates of climate change is the comparative study of models.
- The building of appropriate standards has been identified as a key element in uniting modeling and data communities.
- This requires convergence and cross-fertilization between model and data frameworks: by developing a clear understanding of the architecture of Earth system models, PRISM and ESMF also point the way to a metadata hierarchy to be used in building curators.
- Leadership in standards will come from custodians of international multi-model data archives well connected to data consumers, and will be embedded in the modeling frameworks.
- Research is needed into hierarchical data storage, use of pattern recognition and feature detection for data reduction, remote data analysis and visualization.
- While building petascale systems, let's not neglect the low end... see e.g TGICA Data and Capacity Building Initiative for developing and transition economies (part of IPCC).

Selected web references

http://www.esmf.ucar.edu General website for ESMF: documentation, code, examples, contacts.

http://prism.enes.org General website for PRISM: PRogram for Integrated Earth System Modeling.

http://www.gfdl.noaa.gov/~fms The GFDL Flexible Modeling System. Also links production models and climate model simulation data.

http://gmao.gsfc.nasa.gov Focus on short-term climate variability.

http://mitgcm.org The MIT GCM. Considerable emphasis on data assimilation.

http://www.wrf-model.org The NCAR/NCEP Weather Research and Forecasting model.

http://www.ccsm.ucar.edu The NCAR Community Climate System Model.

http://www.ipcc.ch Intergovernmental Panel on Climate Change. Comprehensive scientific synthesis of current thinking on climate.