Coprocessors, multi-scale modeling, fluid models and global warming.

Chris Hill, MIT

Outline

- Some motivation for high-resolution modeling of Earth ocean system.
  - the modeling challenge
- An approach
- Software tricks and techniques
  - Components
  - Data structures
- Status and plans

Why might we care about modeling global ocean at higher and higher resolution?

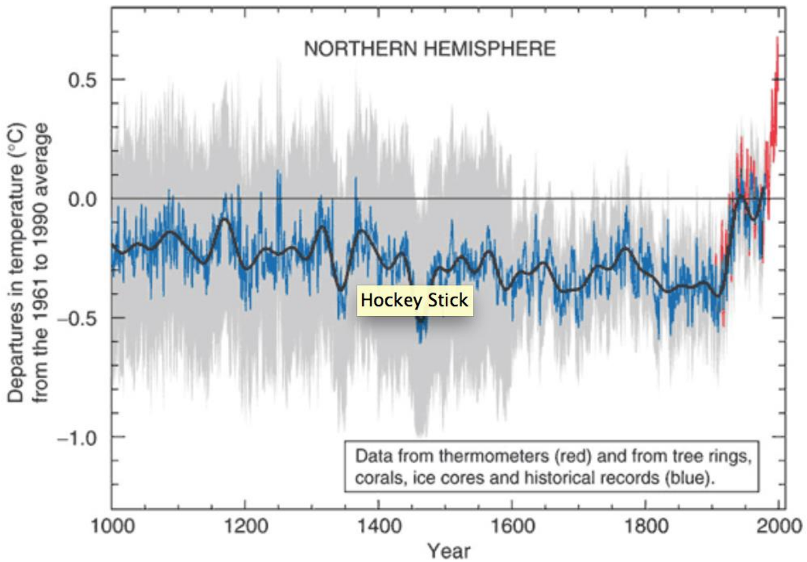

Northern Hemisphere Mean Temp Reconstruction

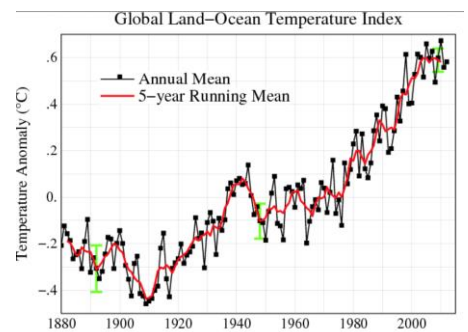

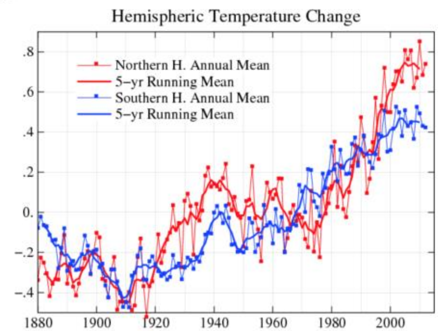

Human influence on the climate system is clear. This is evident from the increasing greenhouse gas concentrations in the atmosphere, positive radiative forcing, observed warming, and understanding of the climate system. {2–14}

IPCC, 2013

Rising CO2, 280 – 400PPM, net increase in “radiative forcing” due to net of changes in atmospheric chemical/particle composition.

Daily Mail, UK, 2012.

Global warming stopped 16 years ago, reveals Met Office report quietly released... and here is the chart to prove it

- **The figures reveal that from the beginning of 1997 until August 2012 there was no discernible rise in aggregate global temperatures**
- **This means that the ‘pause’ in global warming has now lasted for about the same time as the previous period when temperatures rose, 1980 to 1996**

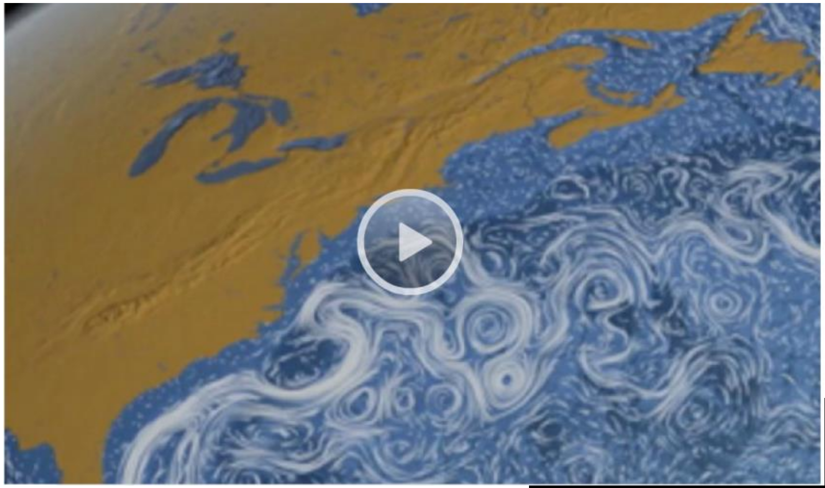

The Earth surface is 70% ocean

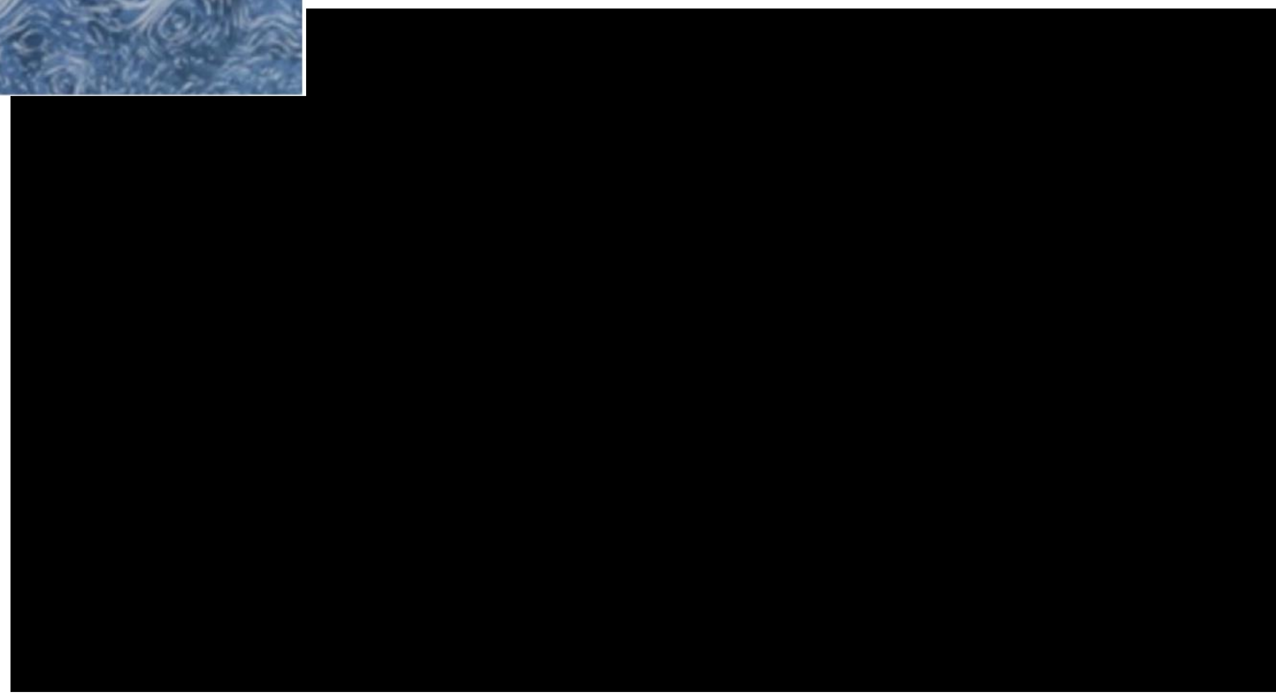

Ocean is fluid – in theory (and in practice) it can draw down and release heat on many time-scales, days, weeks, months, years, decades, centuries, millennia.
$C_v^{OCEAN}:C_v^{AIR}$ = 3000:1 (at sea-level, 300^{o}K).

Variations in ocean heat uptake can alter global mean surface temperature.

Ocean heat/buoyancy budget/exchanges.

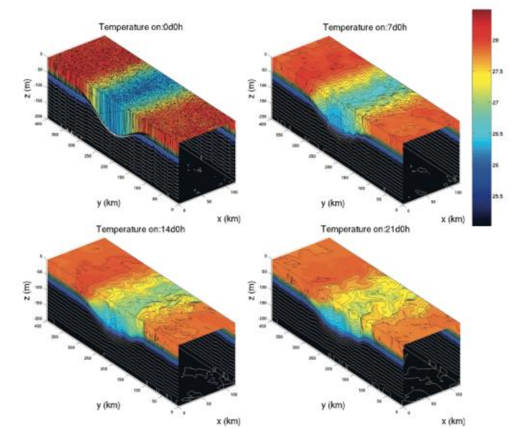

Figure 1. Restratification of an idealized hurricane wake, from initial conditions of a smooth wake plus noise with one snapshot per week. This simulation's parameters (buoyancy gradient, wake depth, etc.) are based loosely on Hurricane Frances.

Haney et al, JMR, 2012

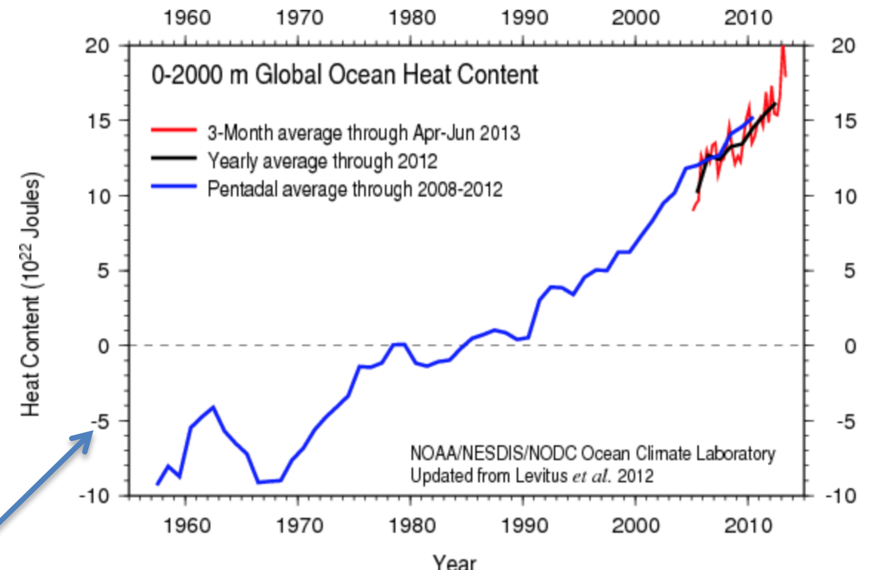

Best estimate, from somewhat sparse obs.

Heat/buoyancy loss/uptake processes have complex small-scale components.

Scales between two pictures are O(10m) process ➔ O(10,000km).
Default approach, in large-scale models, is to model small-scale as diffusive K grad(phi) term.
This presents a challenge to how much models can be used "ab-initio" to reason about ideas, test hypotheses about dominant processes etc…

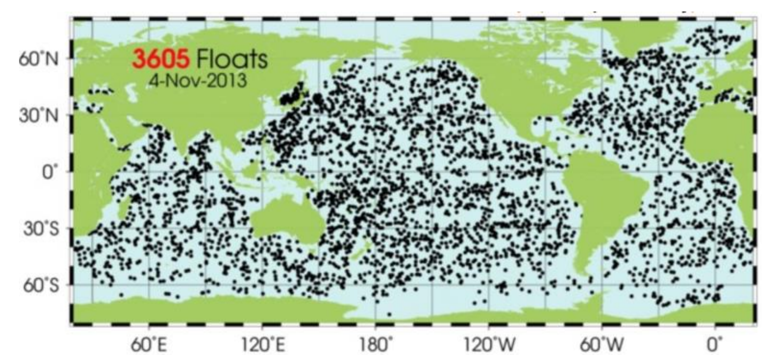

Outline

- Some motivation for high-resolution modeling of Earth ocean system.
  - the modeling challenge
- An approach
- Software tricks and techniques
  - Components
  - Data structures
- Status and plans

Multi-scale algorithm cartoon (2-level)

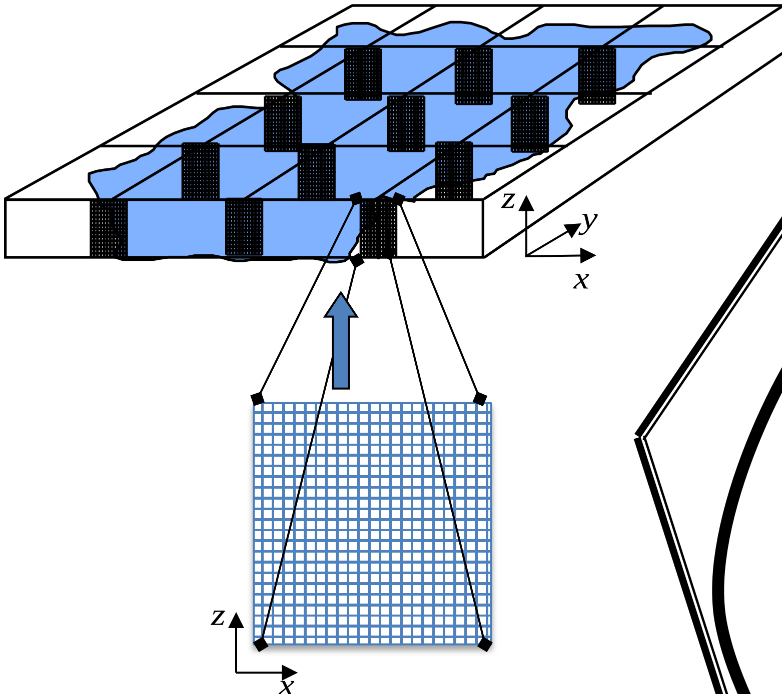

f – finer mesh models

c – coarser mesh model

$$\vec{v}_f = \vec{v}_c + \vec{v}' \qquad T_f = T_c + T'$$

Fine mesh fields ≡ coarse **plus** sub-grid scale variations

$$[T_f]_c = T_c, [\vec{v}_f]_c = \vec{v}_c$$

Ensure fine and coarse fields are consistent

$$\dot{T}_f = -\nabla \cdot (\vec{v}_f T_f) + \Im_f$$

Fine mesh tendencies

$$\dot{T}_c = -\nabla \cdot (\vec{v}_c T_c) + \Im_c + [\dot{T}_f]_c$$

Coarse mesh tendencies

core steps

- Embed local process models in large scale model
- processes more explicit, different physics, grid and timestep to large scale model
- Analogous to atmospheric superparam (e.g. *Grabowski, 2001*)
- Surprisingly analogous to “coupled” modeling in general.

Parallel Formulation

- Sub-models can be expensive – one has to solve a complex sub-model at each large-scale model grid cell.
- However, sub-models are independent computations. So there is a lot of parallelism

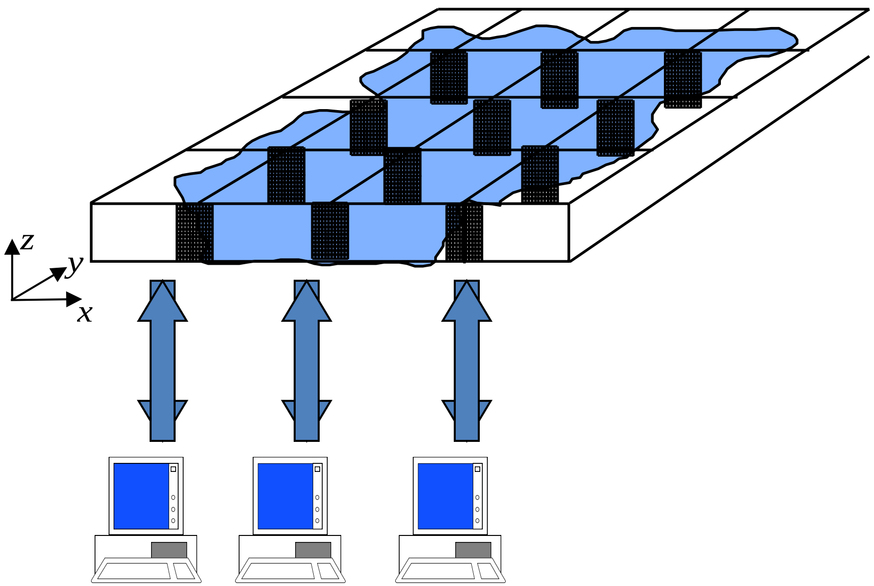

coarse model

fine scale models

$$\dot{T}_c = -\nabla\cdot(\vec{v}_c T_c) + \Im_c + [\dot{T}_f]_c$$

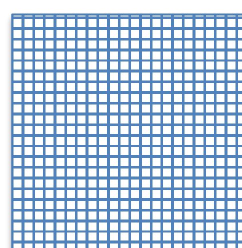

$$[T_f]_c = T_c, [\vec{v}_f]_c = \vec{v}_c$$

$$\dot{T}_f = -\nabla\cdot(\vec{v}_f T_f) + \Im_f$$

Parallel Formulation

- Sub-models can be expensive – one has to solve a complex sub-model at each large-scale model grid cell.
- However, sub-models are independent computations. So there is a lot of parallelism

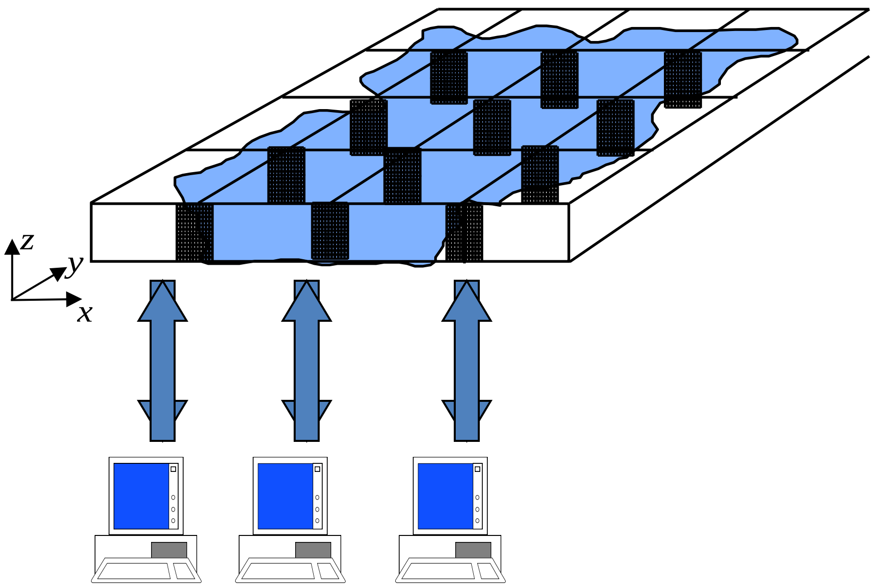

coarse model

fine scale models

$$\dot{T}_c = -\nabla\cdot(\vec{v}_c T_c) + \Im_c + [\dot{T}_f]_c$$

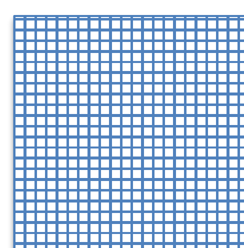

$$[T_f]_c = T_c, [\vec{v}_f]_c = \vec{v}_c$$

$$\dot{T}_f = -\nabla\cdot(\vec{v}_f T_f) + \Im_f$$

Proof of concept application to idealized, simplified ocean buoyancy loss

Super-parameterization in ocean modeling: Application to deep convection

Jean-Michel Campin *, Chris Hill, Helen Jones, John Marshall

Department of Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

2011

$Q=800Wm^{-2}$ $Q=800Wm^{-2}$

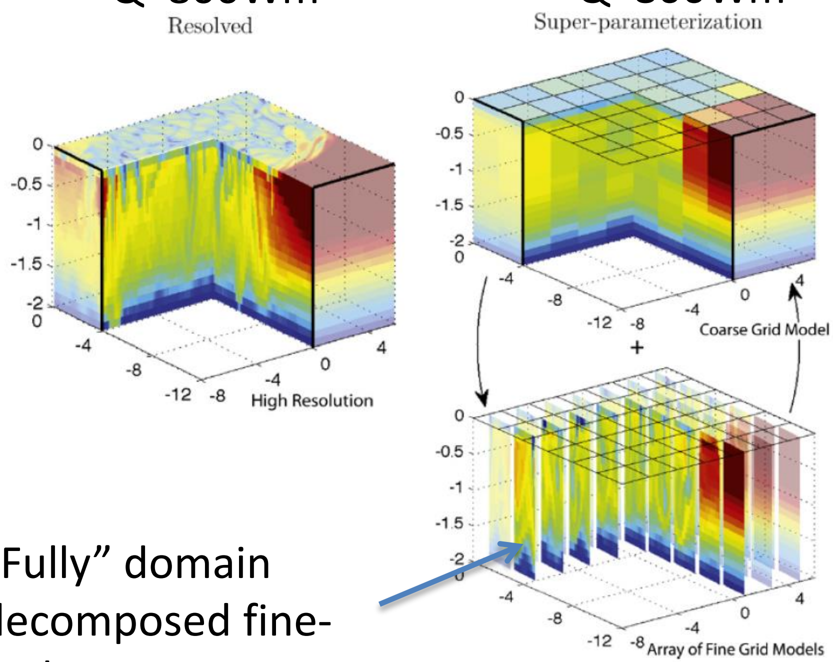

"Fully" domain decomposed fine-scale equations.

Embedded multi-scale model significantly improves representation of buoyancy loss – capturing vertical and temporal characteristics well.

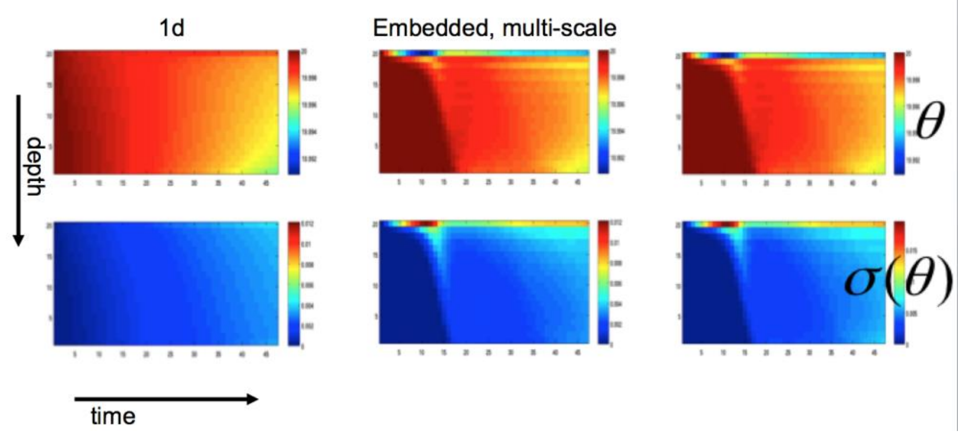

Some example differences

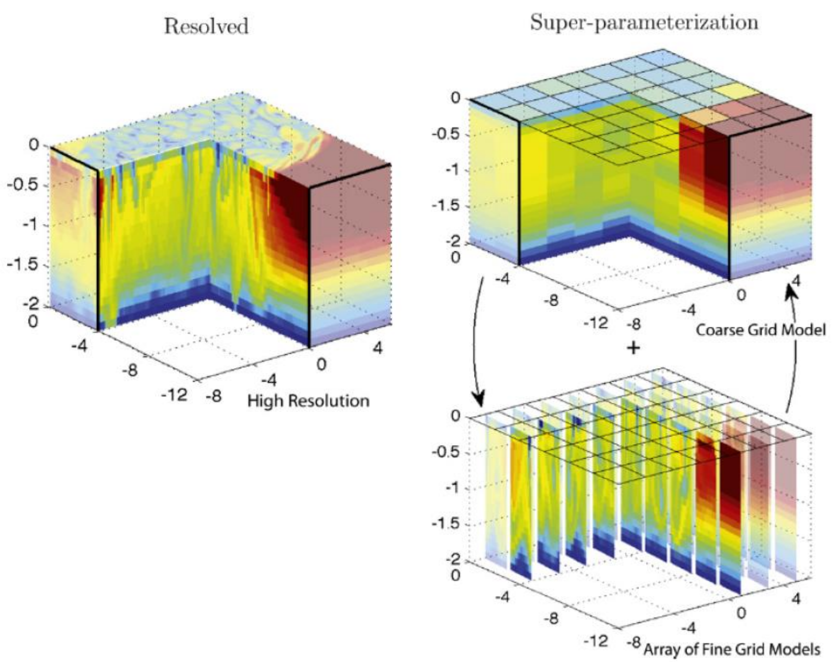

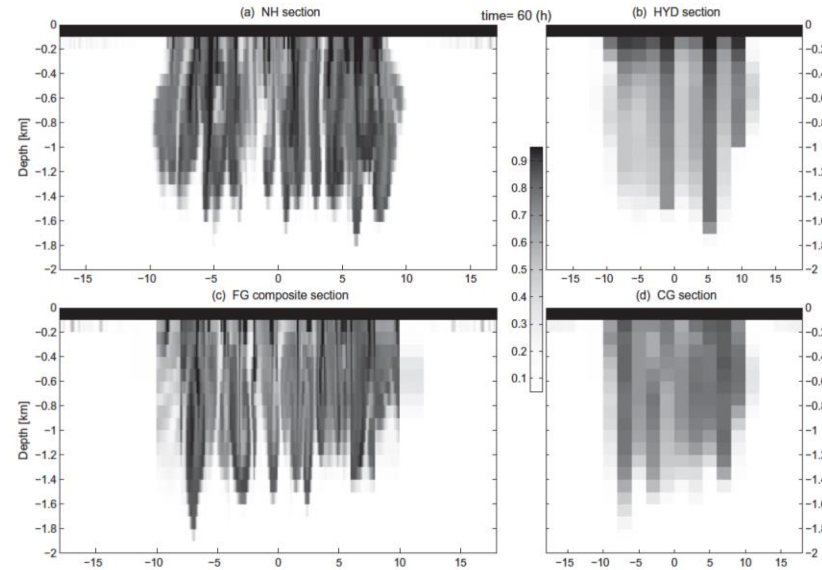

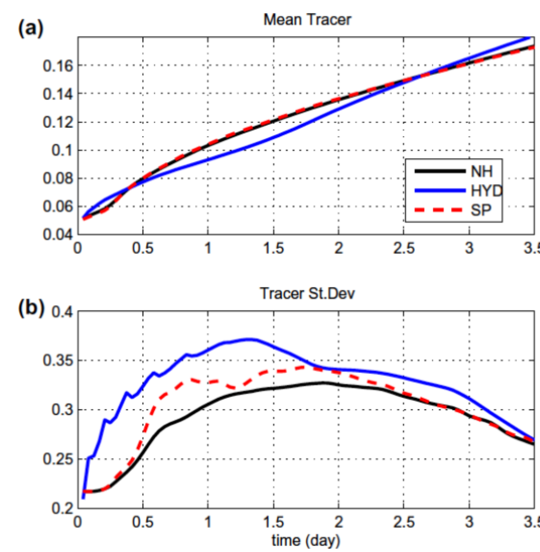

Computational operation count for multi-scale alg is 1/10 of fully resolved model and x10 of coarse alone model.
Domain decomp introduces extra parallelism (in theory).

Next step experiments and relation to emerging co-processor architectures.

Coarse and fine can also run concurrently.
Can we exploit this by adding extra co-pro hardware.

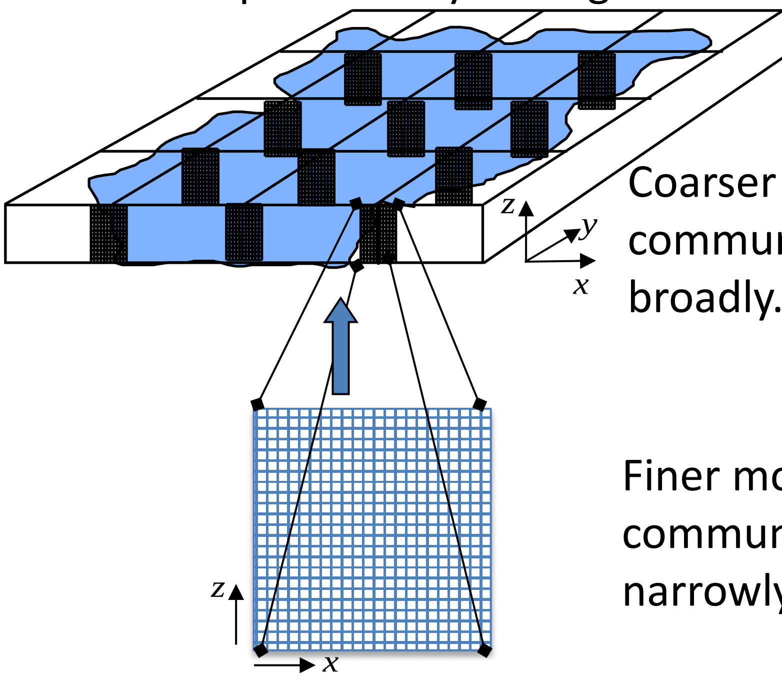

Coarser model – communicates more broadly.

Finer model – communicates more narrowly.

Interaction between coarse and fine occurs through $[T_f]_c$
This can provide a favorable compute to co-processor communicate balance ➔ a better way forward to improve model for a given "transistor"/"power"/"dollar"/"space" budget.

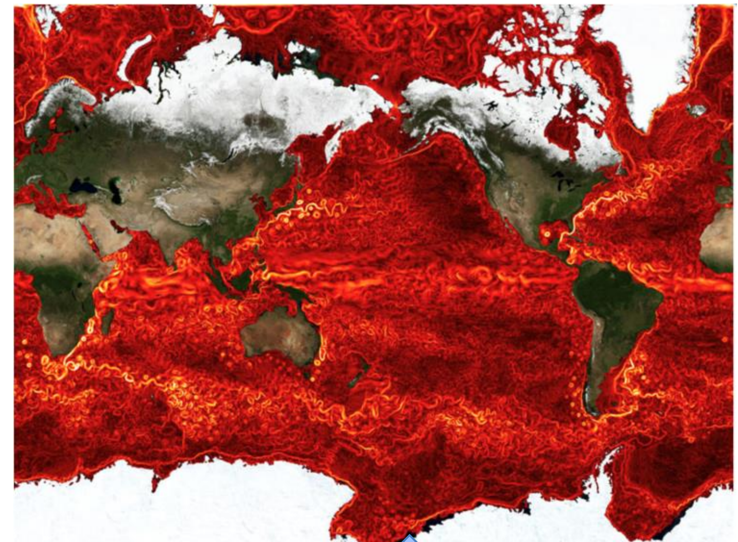

(~1000 CPU cores)

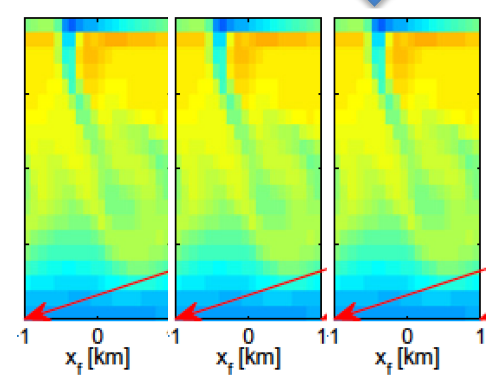

x ~16million

(~100,000 co-pro cores)

Outline

- Some motivation for high-resolution modeling of Earth ocean system.
  - the modeling challenge
- An approach
- Software tricks and techniques
  - Components
  - Data structures
- Status and plans

Component Perspective

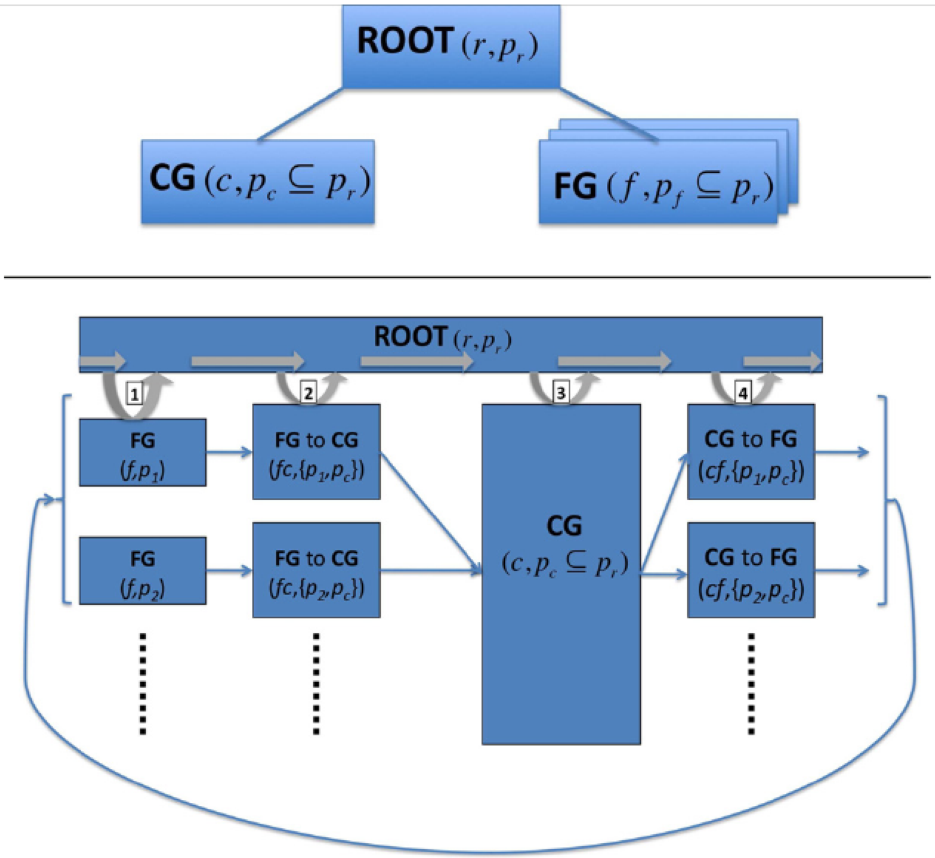

We use component based system, ESMF (Earth System Modeling Framework), to orchestrate distributed, concurrent execution.

Provides object like high-level environment for abstracting mathematical building blocks.

Overlaid on standard MPI parallelism for CPUs.

Fig. 4. Top panel: The computational elements of the SP scheme involving **CG** and **FG** models organized as components under a parent component **ROOT** using the Earth System Modeling Framework. The **ROOT** acts to schedule **CG** and **FG**s pairing in a set of computational threads/processes p_r that includes all the **CG** threads/processes p_c and all the **FG** threads/processes p_f. Bottom panel: The sequence of events in a model timestep of **SP** coordinated by **ROOT**. The numbers 1–4 correspond to the algorithm steps 1–4 described in Section 3.1. In step (1) the **FG** models are integrated forward. There are $n_x \times n_y$ independent **FG** models, one for each grid cell. These models are spread over threads/processes p_n, where $n = 1, m$ within the set p_f. In step (2) updated tendencies from the **FG** models are mapped to the **CG** model grid and transferred to the **CG** model threads/processes p_c. In step (3) the **CG** model executes on a set of threads/processes p_c. Finally, in step (4), **CG** model state information is distributed to the **FG** models. Control returns to the **ROOT** component in between each step, allowing it to coordinate **SP** model integration.

Hill, C., DeLuca, C., Balaji, V., Suarez, M., Da Silva, A., 2004. The architecture of the Earth System Modeling Framework. Computing in Science and Engineering 6 (1), 18–28.

Suarez, M., Trayanov, A., Hill, C., Schopf, P., Vikhliaev, Y., 2007. MAPL: a high-level programming paradigm to support more rapid and robust encoding of hierarchical trees of interacting high-performance components. In: CompFrame'07: Proceedings of the 2007 symposium on Component and Framework Technology in High-Performance and Scientific Computing. ACM, New York, NY, USA, pp. 11–20.

Implementation

- Generally applicable implementation not tied to a specific problem
  - cast computational formulation as a coupled modeling scheme and borrow ESMF technology from the coupled modeling community e.g.

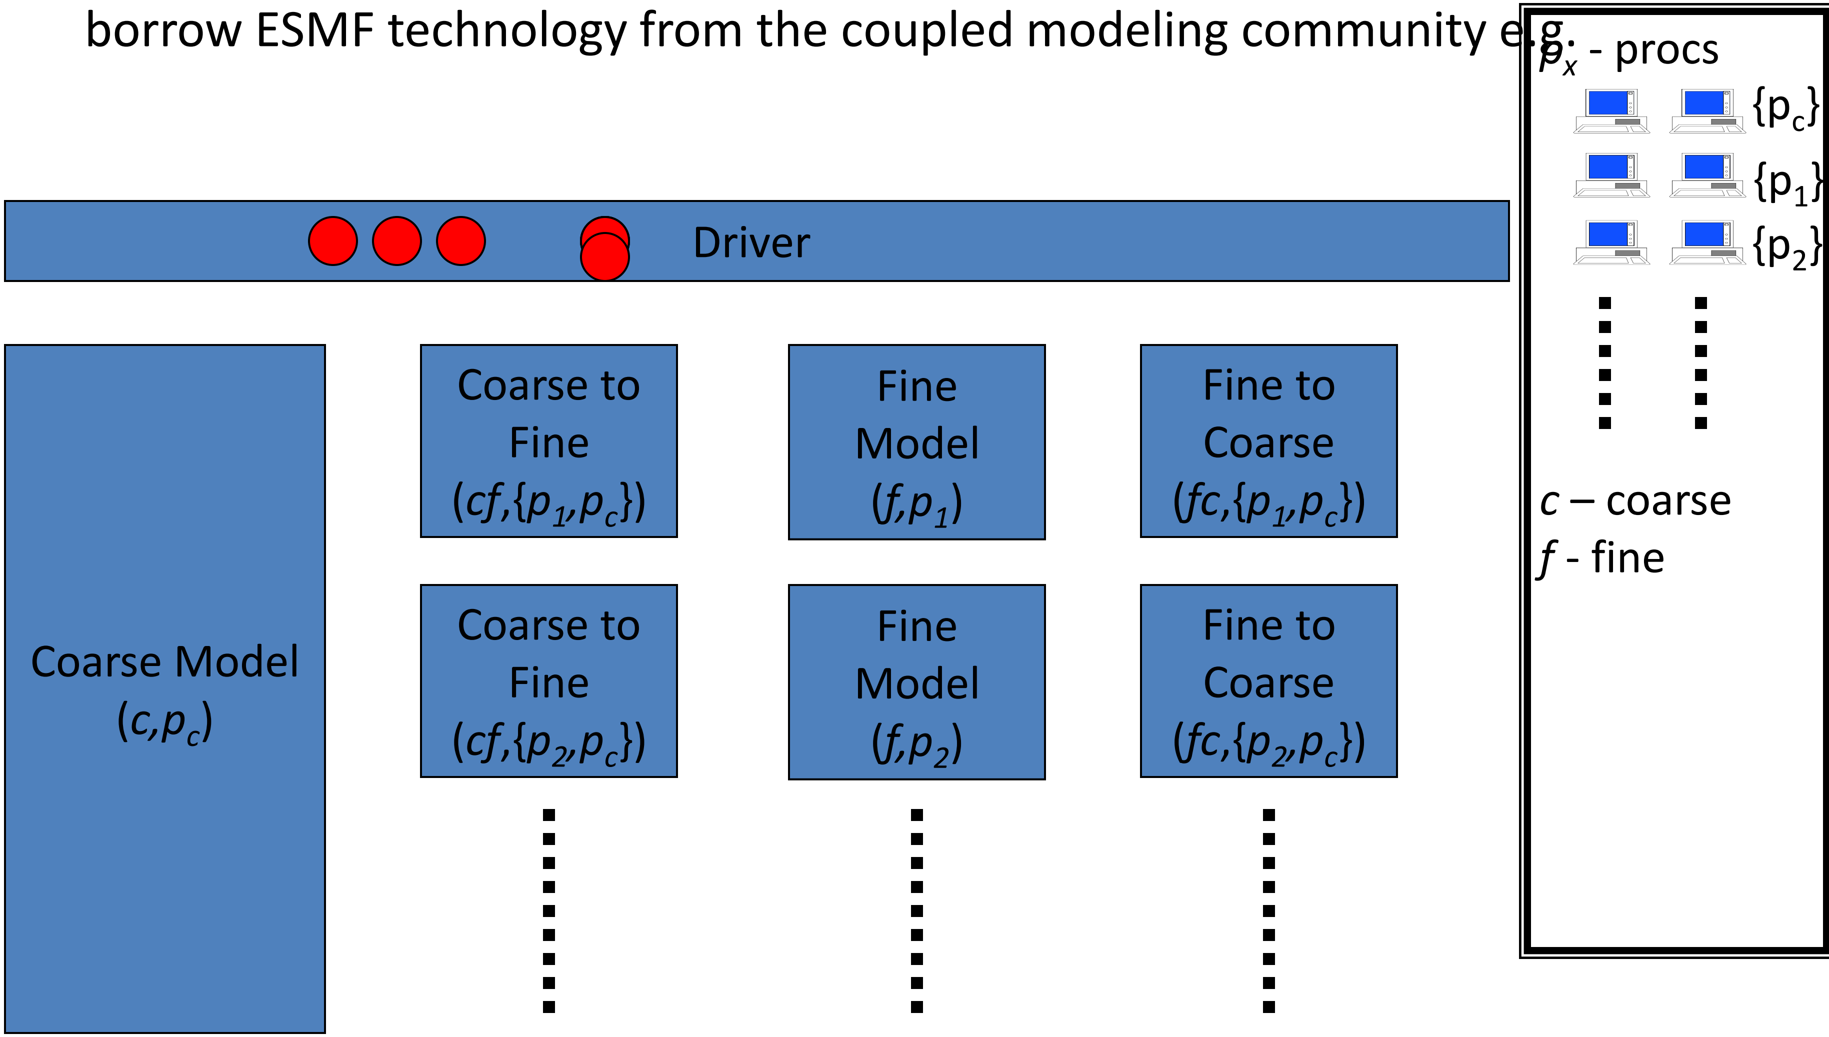

Implementation

- Generally applicable implementation not tied to a specific problem
  - cast computational formulation as a coupled modeling scheme and borrow ESMF technology from the coupled modeling community e.g.

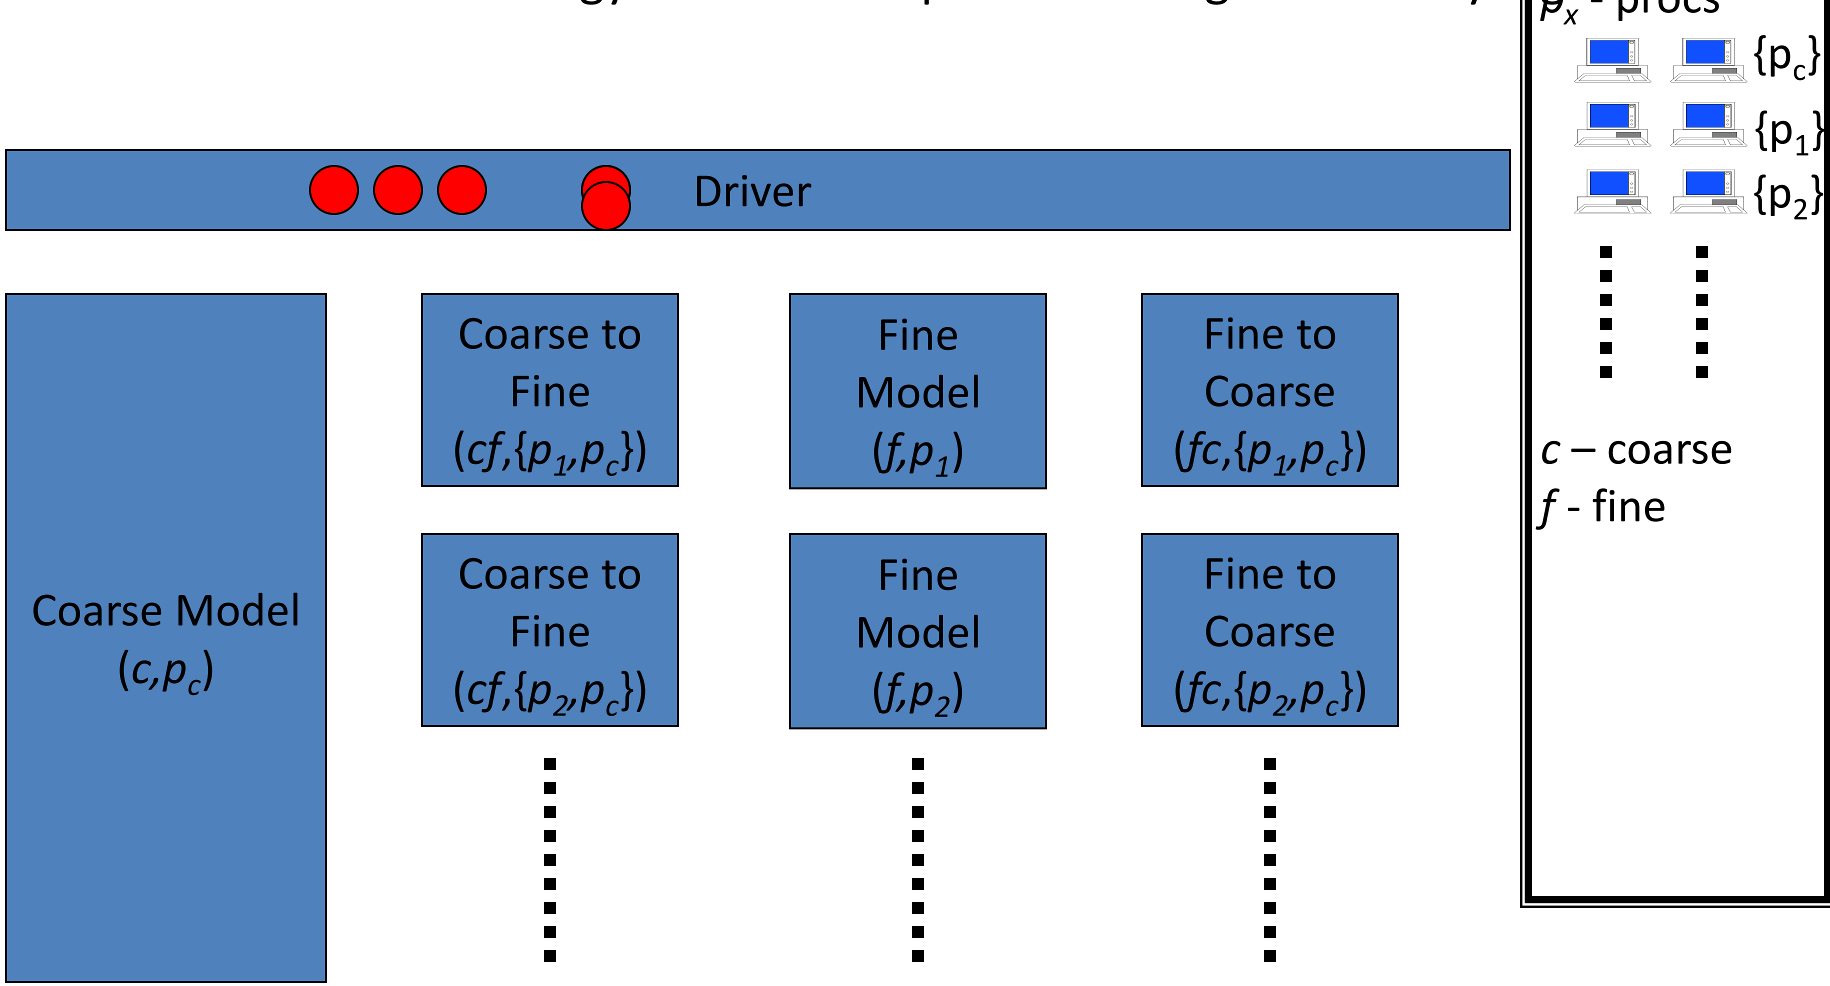

Implementation

- Generally applicable implementation not tied to a specific problem
  - cast computational formulation as a coupled modeling scheme and borrow ESMF technology from the coupled modeling community e.g.

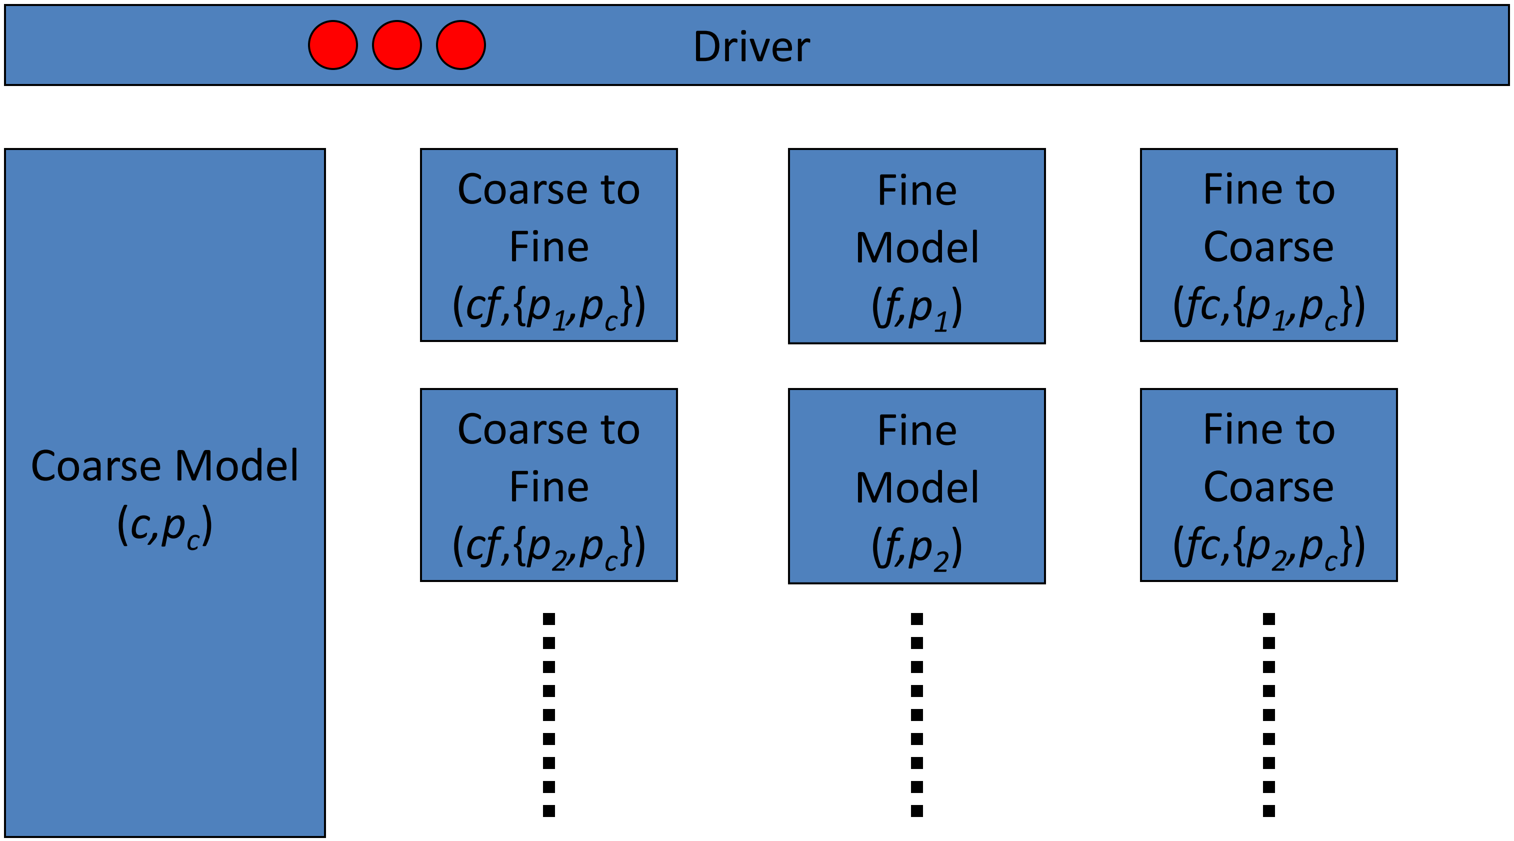

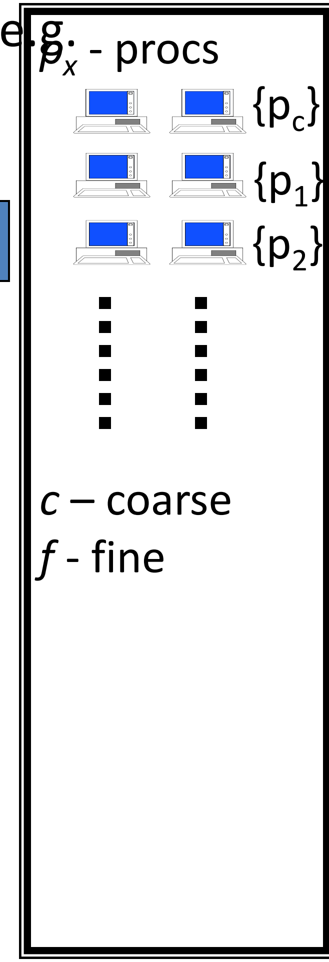

Co-processor abstraction

- Similar to high-level abstraction
  - create abstract data type handles.

    ```
    TYPE(ACCEL_OBJ_T) :: V1, V2, V3, AC
    ```

  - define overloaded functions for mapping from native CPU language (F90/95) to co-processor handles

    ```
    CALL ACCEL_TO(Va,V1,AC)
    ```

  - and for basic, kernel operations (vector, matrix, custom solve etc…)

    ```
    CALL ACCEL_VADD(A3,A1,A2,AC)
    ```

  - co-processor kernels are hand written (CUDA or C)

Outline

- Some motivation for high-resolution modeling of Earth ocean system.
  - the modeling challenge
- An approach
- Software tricks and techniques
  - Components
  - Data structures
- Status and plans

Testing and developing

- Targeting production environments at TACC, ORNL/NICS, Blue Waters, NASA Ames.
- Development leveraging local
  – compute resources at new collaborative BU, Harvard, MIT, Northeastern, UMass facility – MGHPCC.


  – intellectual resources that are also part of this collaboration.
    - these include a very broad set of interactions that are working to evolve the HPC ecosystem at “all” levels from hardware/systems/architectures ➔ algorithms/applications.

Note, by special request - a shameless plug for some of our regional activities!

We need ways to test and evolve ideas easily and dynamically.

The Massachusetts Open Cloud (MOC)

Submitted by: Boston University Hariri Institute for Computing and Computational Science & Engineering

Award Abstract #1246003

CC-NIE Integration: Multi-Wave - A Dedicated Data Transport Ring to Support 21st Century Computational Research

Award Abstract #1229059

**MRI Consortium: Acquisition of a Heterogeneous, Shared, C
Instrument to Enable Science and Computing Research by tl
Green High Performance Computing Consortium**

C3DDB Phase II Application

I. This is a proposal to build a Commonwealth Computational Cloud for Data Driv
(**C3DDB**) that will be a resource for boosting research that connects world l
science and emerging, innovative big-data analytics.

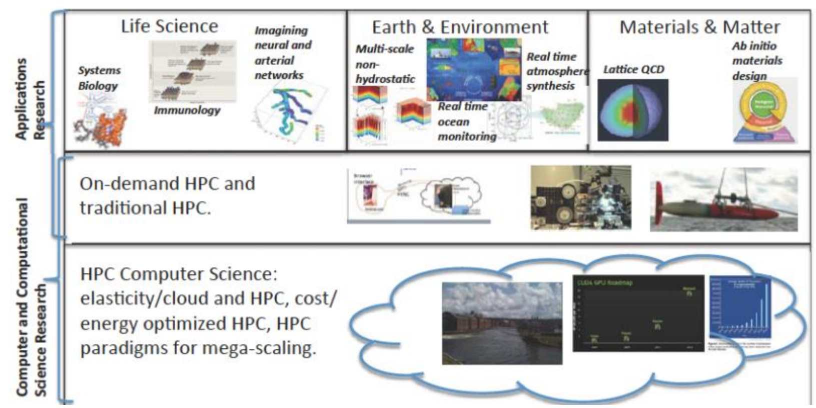

Through these, and more, we are strengthening our regional collaborative activities significantly.

Summary/discussion

- Domain-decomp and “overset” mesh in a multi-scale paradigm may be a way to move forward faster in planetary scale modeling, v. monolithic resolution.
  - Needs to be tested
- There are appealing potential synergies with co-processor developments around taking a hierarchical view.
  - Interesting software challenges around realizing in practice
  - Hierarchical algorithms have relevance for super-large HPC activities in general. Physical facility network layouts are increasingly hierarchical.
  - Multiple scales in physics of problems of interest – getting mathematics to cooperate remains challenging.