Software Infrastructure for a Unified Modeling System:

ESMF, NUOPC Layer, and NEMS

Cecelia DeLuca
NOAA Environmental Software Infrastructure and Interoperability (NESII)
NOAA ESRL/University of Colorado

Unified Weather-Climate Modeling: Challenges and Opportunities
September 28, 2016

Outline

- Introduction to ESMF, NUOPC Layer and NEMS
- Coupling infrastructure capabilities and flexibility needed to span weather to climate time scales
- Coupling in a NEMS unified modeling system
- Open issues and next steps

ESMF, NUOPC Layer, and NEMS

- The Earth System Modeling Framework (**ESMF**) is community-developed software for building and coupling model components.
- The National Unified Operational Prediction Capability (**NUOPC Layer**) is a set of extensions to ESMF that increases component interoperability
- The NOAA Environmental Modeling System (NEMS) is one of the U.S. modeling systems that are using ESMF and the NUOPC Layer

Metrics:
~7000 downloads

~150 components in use

~4000 individuals on info mailing list

~40 platform/compilers regression tested nightly

~6500 regression tests

~1M SLOC

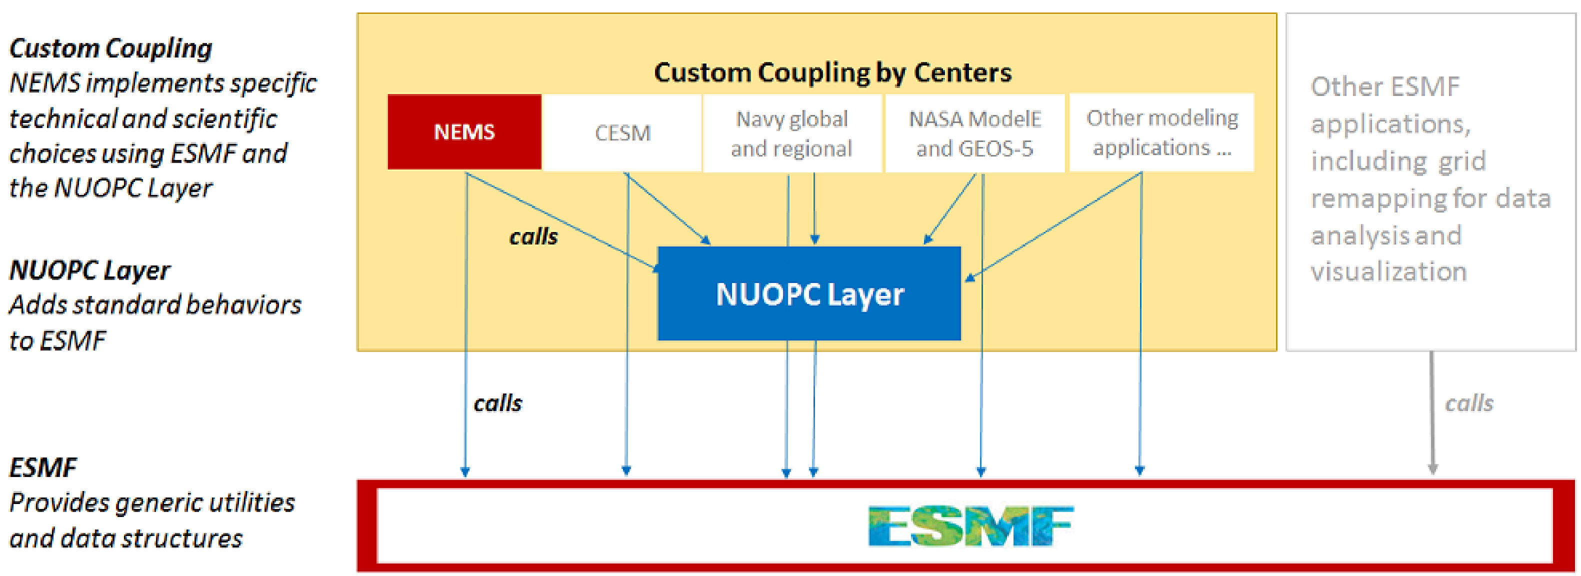

Coupling Infrastructure Across Timescales

- ESMF and NUOPC Layer software infrastructure are used in modeling systems that span a range of space and time scales.
- Examples:
  - Navy COAMPS coupled atmosphere-ocean supports a low-res global model with a high-res regional nest, with direct ocean - wave interaction, for weather and ocean prediction
  - GEOS-5 is a coupled modeling system, structured as a hierarchy with 50+ components, that is used for weather prediction through decadal time scales
  - NUOPC CESM supports high resolution ocean coupling at climate time scales, with separate atmosphere, ocean, sea ice, land and other components
  - ESMF/NUOPC applications are described in more detail in Theurich et al 2016
- Different predictive time scales can warrant different coupling approaches.

ESMF and NUOPC capabilities allow for customization of coupling techniques for specific problems.

Capabilities: Exchange Grid and Interpolation

There are multiple techniques for performing interpolations and computing fluxes.

Exchange grids were introduced at GFDL as a technique for accurate conservative grid remapping (Balaji et al 2006)

An exchange grid between two component model grids is the grid formed by the union of the bounding lines of the component model grids. At GFDL:

- Exchange grid cell areas are used as the weights for conservative interpolation
- The surface flux calculation (atmosphere <> ice/ocean/land) is performed on the exchange grid

Another approach is to merge and compute fluxes in a “hub” coupler on the grid of one of the components being coupled (the “spokes”). A variety of interpolation methods may be used (e.g. CESM, Craig et al 2012).

ESMF supports both exchange grid (Xgrid) and “hub and spoke” style coupling. It also supports connections directly between components through simple, generic “connectors”.

Capabilities: Implicit and Explicit Coupling

Components may be coupled using implicit or explicit techniques:

- Explicit coupling – unknowns at the next timestep are computed using known values from the timestep before it
- Implicit coupling – unknowns at the next timestep are defined by coupled sets of equations that include values at the next timestep, so the solution requires either solving a matrix or an iterative technique

Implicit methods allow for longer timesteps and tend to produce more stable solutions, since changes in values can be smoothed further in time and space.

Explicit methods may be able to achieve accurate results with less computational effort for smaller time steps.

ESMF and the NUOPC Layer support both implicit and explicit coupling exchanges.

Capabilities: FEM Engine for Grid Flexibility

ESMF uses an internal parallel finite element mesh (FEM) framework to implement grid remapping in a common 3D space (originally with a native FEM, now testing a DOE FEM code called MOAB).

This enables it to represent and remap virtually any grids.

This FEM approach (used in ESMF since ~2008) is now being implemented in the ACME project.

ESMF also has options for how the lines are drawn between gridpoints, including straight lines and great circle lines.

The infrastructure for supporting weather to climate scales needs to be able to accommodate grids from components representing a wide range of physical processes, and potentially both 2D and 3D coupling interactions. ESMF does this.

Capabilities: Interpolation Methods

ESMF includes a variety of interpolation methods:

- Bilinear
- Higher order patch recovery (Khoei et al 2007, Hung et al 2004)
- First order conservative (higher order conservative in next release)
- Nearest neighbor

There are tradeoffs in using these methods; for example, the conservative remapping method introduces larger interpolation errors relative to the patch method, but will do a better job of preserving the integral of the data as it moves to the destination grid.

The choice of method may therefore depend on whether short-term accuracy or long-term conservation is more important.

Flexibility of Resource Mapping

When there are a limited number of components in an application (e.g. just atmosphere and ocean, as in HWRF) – the performance may be acceptable if they are always run concurrently on mutually exclusive processor sets, as separate programs or executables.

When there are *more* components in an application (e.g. atmosphere, ocean, ice, wave, ionosphere, etc.) it is useful for performance optimization be able to overlap components on the same processors OR run them concurrently. If the components are in separate executables, overlapping components becomes more difficult. U.S. climate models that have many components (e.g. CESM, GFDL models, GEOS-5) tend to use this single executable approach.

ESMF promotes a single executable approach, though it can also support multiple executables.

Flexibility in Time vs Process Split Coupling

Implicit / explicit choices and choices about resource mapping are also tied to decisions about how to advance coupled components:

Time split coupling – two components are calculated sequentially, each based on the state used by the other.

Process split coupling – two components are based on the same state and their tendencies are added to produce the updated state.

(Williamson 2002)

Process split coupling is associated more with explicit solutions and concurrent execution of components, time split coupling with sequential execution.

ESMF and the NUOPC Layer do not prescribe which approach to use.

Flexibility in Run Sequence

ESMF and NUOPC support flexible, parameterized run sequences that are specified at run-time. This approach enables components and couplers (mediator or "med") to be reconfigured easily to support different applications and component calling sequences:

Hurricane sequence

```
runSeq::
  @7200.0
    OCN -> MED
    MED MedPhase_slow
    MED -> OCN
    OCN
    @3600.0
      MED MedPhase_fast_before
      MED -> ATM
      ATM
      ATM -> MED
      MED MedPhase_fast_after
    @
  @
::
```

```
runSeq::
  @1800.0
    MED MedPhase_slow
    MED -> OCN
    OCN
    OCN -> MED
    @600.0
      MED MedPhase_fast_before
      MED -> ATM
      MED -> ICE
      ATM
      ICE
      ATM -> MED
      ICE -> MED
      MED MedPhase_fast_after
    @
  @
::
```

Global seasonal sequence

Coupling Choices

When building infrastructure for weather to climate scale systems, there are many choices, which have scientific and performance impacts, for example:

- Which processes should be represented as components? Which processes need to able to run concurrently or on different grids? How should components be sequenced?
- Should we use an exchange grid?
- Should we couple some components implicitly?
- Where and how should we compute fluxes?
- Which interpolation methods and options should we use?

Most of these decisions are not “switches” but choices that require significant development investment. Empirical testing (implement all options, and test if they make the forecast better) is not likely to be feasible or efficient.

NOAA Environmental Modeling System

- NEMS is a rapidly evolving unified modeling system that will be the basis of critical operational modeling applications at NOAA, including *10 day, six week, and 9 month forecasts.*
- Global and regional modeling applications are anticipated.
- Coupled components will include:
  - Atmosphere
  - Ocean
  - Land
  - Wave
  - Sea ice
  - Coastal/storm surge
  - Hydraulics/hydrology
  - Ionosphere
  - Aerosol/chemistry

https://www.earthsystemcog.org/projects/couplednems/

A Unified Software Framework for Multiple Applications

Modeling Application	ATM	OCN	ICE	WAV	LND	AER	HYD	ION	CST
UGCS-Weather	o	o	o	o	o	o		o	o
UGCS-SubSeasonal	o	o	o	o	o	o		o	o
UGCS-Seasonal	o	o	o	o	o	o		o	o
WAM-IPE (Space Wx)	o							o	
HYCOM-Ice (RTOFS)		o	o						
Wave Prediction	o			o					
Regional	o	o	o		o		o		
Regional Nest	o	o							
CMAQ Air Quality	o					o			

- Above is a simplified version of the spreadsheet of NEMS applications: http://tinyurl.com/nems-apps

NEMS Status

- Delivered development version (0.2) of the Unified Global Coupled System – Seasonal (UGCS-Seasonal) running under NEMS, with fully coupled GSM, MOM5, and CICE components, initialized for a cold start and optimized for comparable performance with CFSv2.
- Delivered development version (0.1) of a single domain regional configuration running under NEMS, with two-way coupled NMMB and HYCOM components.
- Delivered development version (0.4) of WAM and the Ionosphere-Plasmasphere Electrodynamics (IPE) models validated under NEMS, with a one-way (WAM>IPE) 3D coupling exchange.
- Delivered development version (0.2) of WRF-Hydro and LIS/Noah land validated running with GSM, MOM5, and CICE under NEMS, showing correct technical exchanges implemented among components, not yet scientifically viable.
- Implemented WAVEWATCHIII coupled one-way to GSM (GSM>WWIII), including nesting, running under NEMS.

See: http://tinyurl.com/nems-apps

Open Issues and Next Steps

The current NEMS system uses a hub and spokes architecture with all explicit coupling, like CESM and CFSv2. Is this the right approach?

The evolution of infrastructure for NEMS requires answering these questions, and many more:

- Which processes should be represented as components? Which processes need to able to run concurrently or on different grids? How should components be sequenced?
- Should we use an exchange grid?
- Should we couple some components implicitly?
- Where and how should we compute fluxes?
- Which interpolation methods and options should we use?

To answer the questions, a system architecture team has been assembled by new EMC director Mike Farrar. In addition to EMC, the working group engages experts in the integrative science of building coupled applications, from GFDL, ACME, CESM, NASA, Navy, and other centers.

Thank you!

Any questions?