Advancing CESM's Coupling Infrastructure with ESMF/NUOPC

Rocky Dunlap

NCAR / CSEG

ESMF Project Manager

June 19, 2018

CESM Software Engineering

Working Group Session

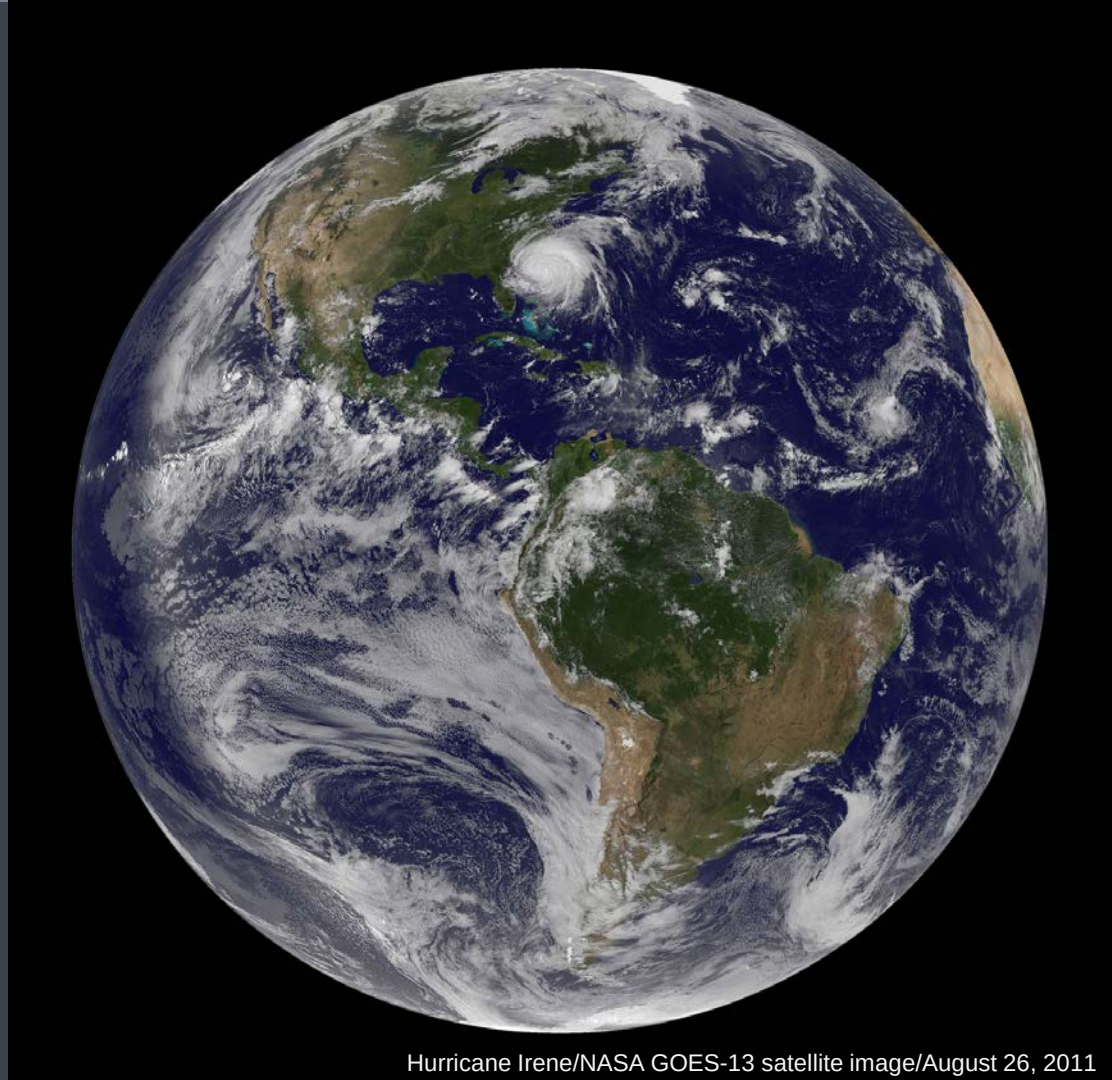

Hurricane Irene/NASA GOES-13 satellite image/August 26, 2011

Outline

- ESMF Overview and Applications Using It
- CESM adoption of ESMF/NUOPC and Implications
- Community Mediator
- ESMF Release Update including Optimization and Tools
- Upcoming Tutorials

Evolving CESM's Coupling Infrastructure

New coupling options and community mediator

CESM's coupling infrastructure keeps components separate while allowing them to communicate using a hub-and-spoke architecture with centralized coupler.

CESM's coupling infrastructure is being evolved to use **ESMF/NUOPC as the primary coupling framework**, adding **new coupling options** and capabilities and aligning CESM with coupling interfaces used in **community model components** and in **coupled systems** at major US modeling centers within NOAA, NASA, and the Navy.

This work includes development of a **new community mediator with a shared codebase** designed to be used at NCAR, NOAA/EMC, and NOAA/GFDL.

What are ESMF and NUOPC?

- The Earth System Modeling Framework (**ESMF**) is community-developed, community-governed software for building and coupling model components in high-performance settings. ESMF is sponsored by multiple agencies.

- The National Unified Operational Prediction Capability (**NUOPC Layer**) is a set of extensions to ESMF that increases component interoperability and adds architectural options.

- **ESMF version 7.1.0r** was released on March 8, 2018.
- Download from: https://www.earthsystemcog.org/projects/esmf/download_710r

ESMF Metrics:

7000 downloads

150 components in use

7500 subscribers to info mailing list

40 platform/compilers regression tested nightly

8000 regression tests

Example Applications using ESMF/NUOPC

NOAA's Unified Forecast System (UFS), the new community modeling system intended for operations at EMC.

- NUOPC used in the coupling infrastructure layer, called the **NOAA Environmental Modeling System** (NEMS).
- **Coupled FV3GFS, MOM6 and CICE system** for S2S prediction in progress; previously delivered fully coupled GSM, MOM5, and CICE components, comparable performance with CFSv2.
- **FV3GFS asynchronous I/O** components built using ESMF.
- **Space weather, wave,** and **external land** coupled systems

Navy global coupled system with components: NAVGEM, HYCOM, CICE, LIS land and WAVEWATCHIII

Navy regional coupled system COAMPS including new coupling with LIS and WRF-Hydro; multi-nest coupling.

NASA's NU-WRF coupled with LIS and WRF-Hydro; exploring use of NUOPC in **coupled DA**.

Interoperability of **NASA's ModelE** and **GEOS5** via NUOPC.

CESM's Next Generation Coupling Infrastructure

Advanced coupling options through NUOPC software layer

New coupling options and capabilities:

- Options to **generate remapping weights online** or offline between a wide range of grids and meshes; exchange grid option
- Automatic **transfer of grids/meshes** from components to the mediator/coupler; mediator uses mesh internally
- **Multi-tile, nested,** and **regional grids**
- **Run sequence** is data driven and can be changed without recompiling; no more stub components
- **Optimization options** including reference sharing; built-in tracing/profiling

Component architecture CESM using ESMF/NUOPC coupling infrastructure

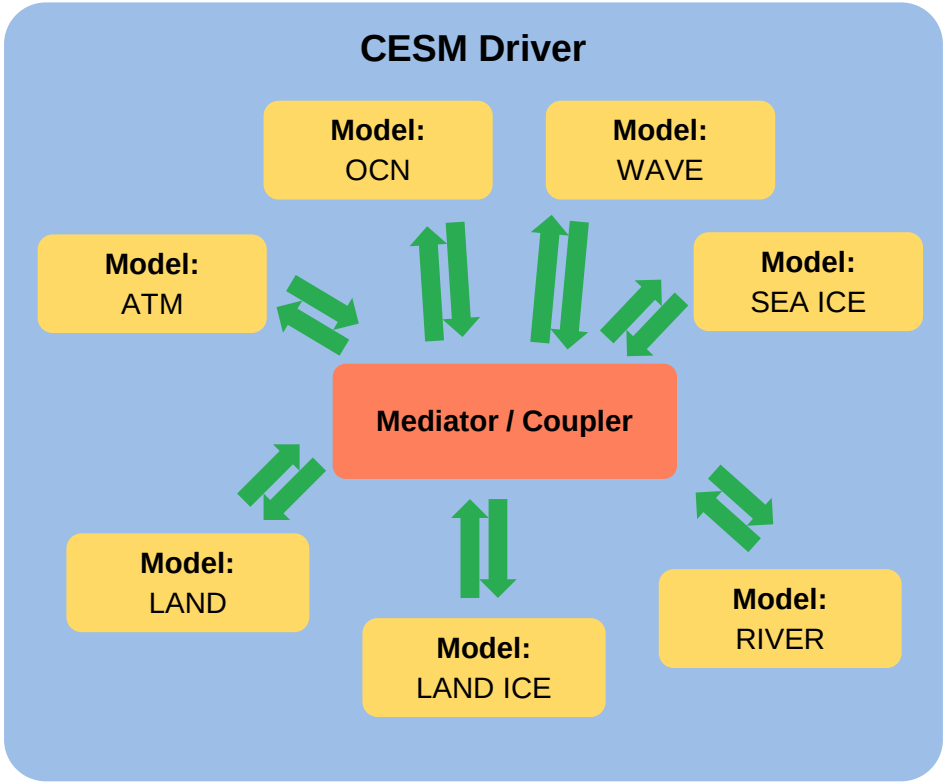

Community Mediator for Earth Prediction Systems - CMEPS

Goal: Deliver a flexible mediator (coupler) that can support CESM, NOAA Unified Forecast System (UFS) and NOAA/GFDL scientific coupling strategies and allow controlled experimentation using different coupling techniques.

There is a science of coupling and decisions for how to set up the coupling should be evidence-based. CMEPS allows for collaborative exploration of important questions:

- Should we couple components explicitly or implicitly?
- Should we use an exchange grid?
- Where and how should we compute fluxes (e.g., in mediator or components)?
- Which interpolation methods and options should we use?
- Which processes need to able to run concurrently or on different grids? How should components be sequenced with the coupler?

The goal is *not* to find one final solution, but to flexibly support multiple options.

Community Mediator for Earth Prediction Systems - CMEPS

CMEPS Architecture: Represent different options through a set of *mediator phases.* Some phases are shared across coupled systems and some are specific to one or two systems.

CMEPS provides generic abstractions for merging multiple fields in the mediator and applying interpolation weights to sets of fields.

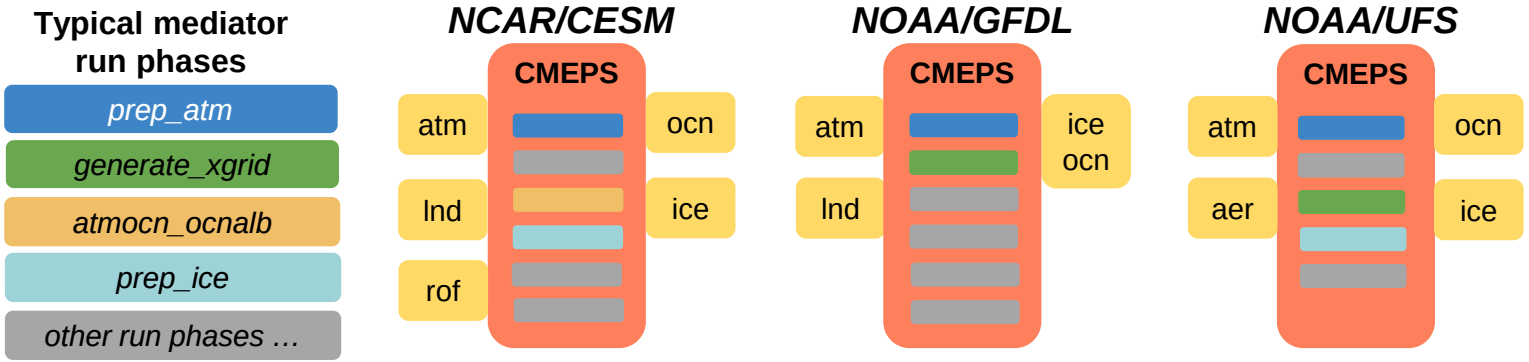

Diagram is illustrative only – components are not complete/correct.

CMEPS Approach and Milestones

Approach: Start with EMC's NOAA Environmental Modeling System (NEMS) Mediator, replicate CESM and GFDL coupling strategies, demonstrate how CMEPS can reproduce coupler behavior in these systems. Work on GitHub to promote collaborative development activities.

Milestones:

- ✓ **Q4FY17:** Couple all CIME prescribed data components with the community mediator.
- ✓ **Q2FY18:** Demonstrate running the community mediator within the NEMS framework with at least one active component. → *CMEPS v0.1 Milestone Release*
- **Q3FY18:** Run the community mediator with all active CESM components.
- **Q3FY18:** Develop and document a governance strategy for the community mediator.
- **Q4FY18:** Demonstrate that the ESMF/NUOPC Layer can replicate key GFDL coupling functions.
- **Q4FY18:** Demonstrate that the community mediator can replicate all NEMS coupling functions, and replace the NEMS mediator with the community mediator.

Working Configurations of CESM with NUOPC/CMEPS

I Compset:

Data atmosphere with active land model (CLMv5), Global Soil Wetness Project Phase 3 forcings. Global 4x5 degree grid.

Validation:

bit-for-bit reproduces CLM history of MCT version after 5 timesteps, which was previously validated

Test report / instructions:

https://github.com/ESCOMP/UFSCOMP/wiki/Test-Report:-I-Compset-(June-14,-2018)

F Compset:

Active atmosphere (CAM version 4.0 physics), active land model (CLMv5), prescribed sea ice (CICE), and a data ocean. Global 4x5 degree grid.

Validation:

roundoff level differences CAM/CLM history of MCT version after 5 timesteps

Test report / instructions:

https://github.com/ESCOMP/UFSCOMP/wiki/Test-Report:-F-Compset-(June-14,-2018)

G Compset:

Active ocean (MOM6) with data atmosphere, data ice, and data runoff. MOM6 uses gx1 global 1 degree grid.

Validation:

visual inspection of MOM6 forcings and history output after 5 day run shows physically realistic fields

Milestone documentation:

https://github.com/ESCOMP/UFSCOMP/wiki/Milestone:-CMEPS-0.1

ESMF Update

ESMF 7.1.0r Public Release

7.1.0r is the latest public release, March 8, 2018

New in ESMF:

- **2nd order conservative grid remapping** method
- New support for multi-tile grids, including API shortcuts for **cubed sphere grid** creation and **regridding to/from multi-tile grids**
- **Dynamic masking** during the application of interpolation weights
- **Extrapolation** of points that lie outside the source grid during grid remapping
- New capability to **read/write regridding weights** from/to file
- Fields/Arrays use **64-bit indices** when needed for large element counts
- New built-in **tracing** capability for application profiling

New in NUOPC:

- New support for **component hierarchies**
- **Run sequences** can be read in from configuration file or parsed from text strings
- **Share Field memory references** through Connectors for components running sequentially

New in ESMPy:

- Create and regrid with **cubed sphere grids** on 6 PETs
- Expose **2nd order conservative** and **extrapolation** regrid options
- **Read/write regridding weights** from/to file

High Performance Grid Remapping

ESMF supports multiple conservative methods

Second-order conservative:

- Destination cell value is the combination of values of intersecting source cells modified to take into account the **source cell gradient.**
- Requires a **wider stencil** and more computation, so more expensive in terms of memory and time than first-order
- Preserves integral of field across interpolation, but gives **smoother results than first-order** (especially when going from coarse to fine grids)

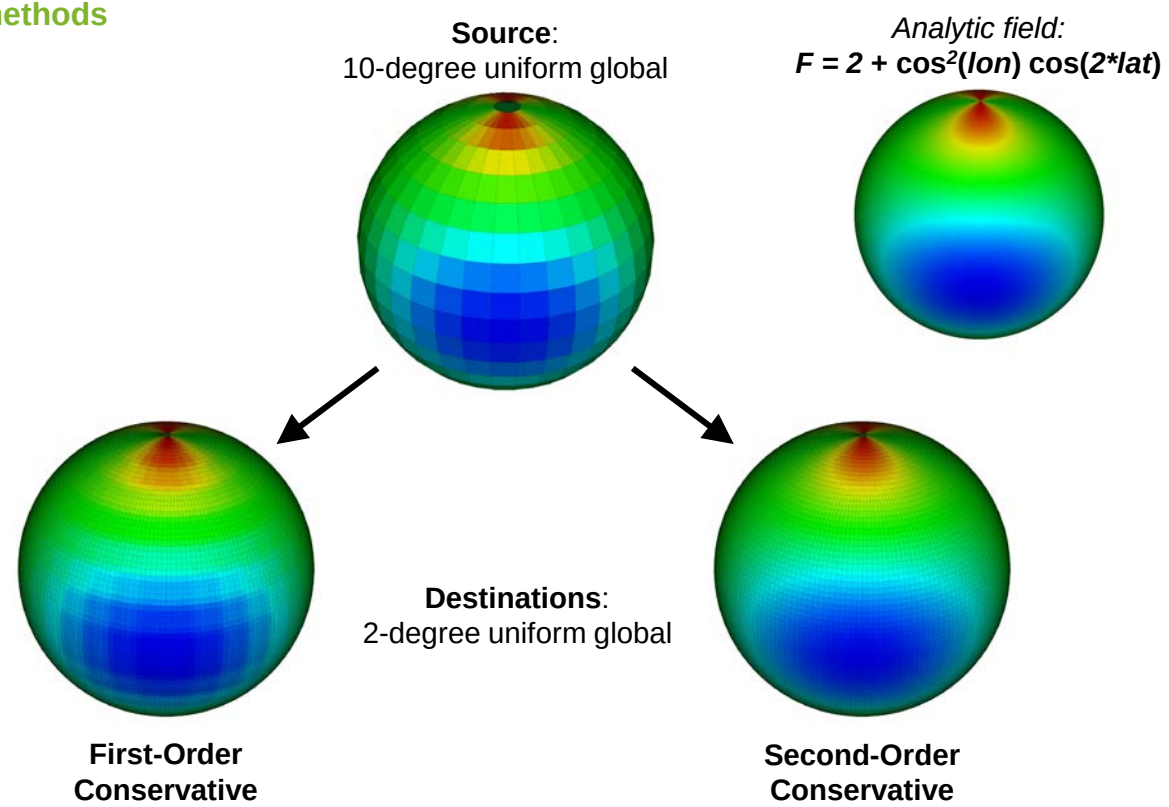

Cubed Sphere Support

Options for representing cubed spheres in ESMF

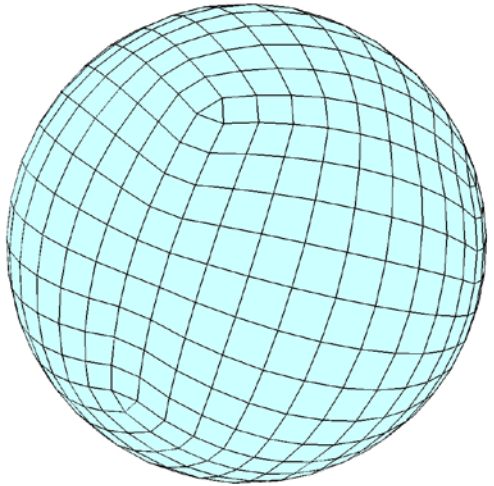

There are two ways cubed spheres are represented in ESMF:

1. **Unstructured Mesh**
   - data fields are 1D
   - more efficient for calculating regridding weights

1. **Multi-tile Grid**
   - data fields are 2D which more naturally matches shape of tiles

Both representations can be regridded to other ESMF geometry types (i.e. Grids, Meshes, and Location Streams).

Three **new APIs** to allow easier creation of cubed spheres in ESMF:

- **`ESMF_MeshCreateCubedSphere(tileSize, … )`** - create mesh from parameters
- **`ESMF_GridCreateCubedSphere(tileSize, … )`** - create multi-tile grid from parameters
- **`ESMF_GridCreateMosaic(filename, … )`** - create from GFDL Gridspec format mosaic file

Integration of MOAB into ESMF

Underneath ESMF_Mesh and ESMF regridding code is a custom built 3D finite element engine

We are replacing that code with the externally developed **MOAB - Mesh Oriented dAtaBase**:

- Get new capabilities:
  - Higher-order elements (elements with data on more than just corners and centers)
  - Data on edges
  - More flexible internal fields
- Follow development advances

Status:

- Internal code separated to allow creation of original ESMF mesh or MOAB mesh
- Can compute ESMF conservative regridding weights using these Meshes
- Bilinear regridding working for spherical and Cartesian grids/meshes
- Update to 4.9.2 version of MOAB library, which had refactored code organization

Smart Resource Mapping on Emerging Platforms

Coupled model optimization

Increasing complexity of coupled models and computing architectures reduces potential for manual optimization

- Coupled models: more processes modeled, more components
- Hardware: CPUs + GPGPUs + MICs, threading and SIMD parallelism on CPU and device cores, multiple programming models

Approach: Framework handles smart mapping of components to hardware resources

- Goal is efficient mapping of multi-component, coupled systems to heterogeneous resources
- New capability added to ESMF to recognize accelerator devices
- Optimization includes components which may use different programming models and optimization strategies: OpenMP, OpenCL, OpenACC, Intel MIC
- Approach designed to preserve kernel-level optimizations and adapt to multiple computing platforms

Smart Resource Mapping on Emerging Platforms

Coupled model optimization

Smart mapping sequence (funded under ONR NOPP)

- Capture information about the current performance of component run phases, and the requirements and capabilities of each component
  - e.g., *can*, *cannot*, or *must* utilize accelerator device
- Generate dependency graphs that capture relationships among components
  - e.g., whether one component phase must run after another
- Analyze the dependency graphs to generate a set of potential scenarios that satisfy the dependencies
  - e.g., concurrent vs. sequential processor layout
- Assess the potential scenarios using a minimization algorithm and make a selection
- Test the new configuration
  - Ideally new configurations tested without a restart, but restarts may be required depending on the constituent models
- Iterate over this process to build up a knowledge base

Cupid Tool: Visualize Call Stack

NUOPC Call Stack View assists with post-run debugging and performance analysis

- Shows entry/exits of NUOPC execution phases including timing information
- Stack process traces to see concurrency (or lack thereof)
- Helps to optimize across platforms: The tracing capability is built into ESMF (no additional setup) and supported on all platforms supported by ESMF

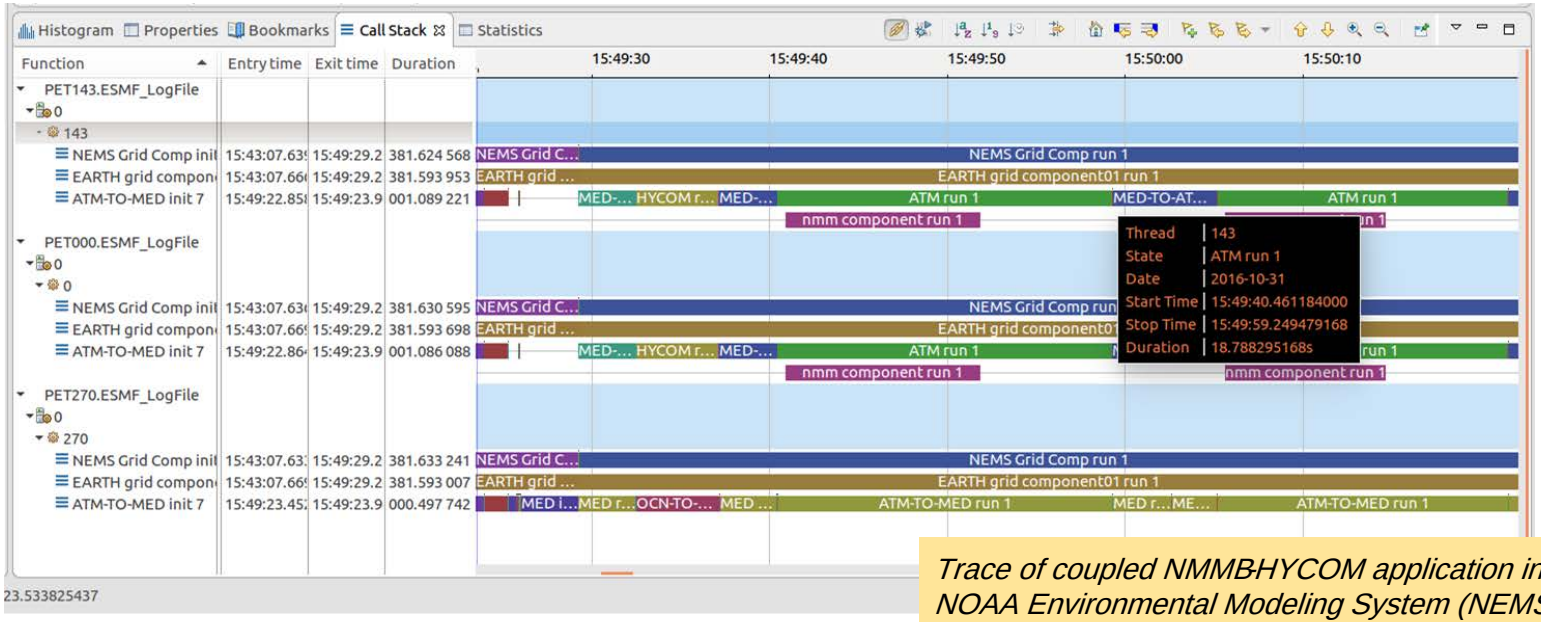

Trace of coupled NMMB-HYCOM application in the NOAA Environmental Modeling System (NEMS)

Cupid Tool: Tracing and Profiling Tool

Built -in and user-defined timer regions

- ESMF can automatically instrument component phases with profiling timers.
- User-defined timer regions supported.
- Code region timing statistics available on a per-process basis.
- Cross PET comparison helps determine load imbalance.

Histogram | Properties | Call Stack | Statistics | Timing Tree (Global) | Timing Tree (Per PET) | State System Explorer

Level	Total time	Self time	Count	Mean
Global (includes all traced PETs)				
[COAMPS] RunPhase1	59788.056 s	96.523 s	256	233.547 s
Per-PET Statistics				
[ATM] RunPhase1	40506.050 s	40506.050 s	20480	1.978 s
[LND] RunPhase1	18512.040 s	6025.686 s	20480	903.908 ms
Per-PET Statistics				
nest:3	9222.175 s	42.758 s	20480	450.302 ms
nest:2	2503.961 s	16.544 s	20480	122.264 ms
nest:1	760.218 s	12.153 s	20480	37.120 ms
[LND-TO-ATM] RunPhase1	577.826 s	577.826 s	20480	28.214 ms
[ATM-TO-LND] RunPhase1	95.618 s	95.618 s	20480	4.669 ms
Per-PET Statistics				
[COAMPS] Init #1	16922.906 s	14.975 s	256	66.105 s
[COAMPS] FinalizePhase1	29.353 s	4.244 s	256	114.659 ms
Per-PET Statistics				
[ATM] FinalizePhase1	24.076 s	24.076 s	256	94.048 ms
[ATM-TO-LND] FinalizePhase1	506.557 ms	506.557 ms	256	1.979 ms
[LND-TO-ATM] FinalizePhase1	480.263 ms	480.263 ms	256	1.876 ms
[LND] FinalizePhase1	45.423 ms	45.423 ms	256	177.432 µs

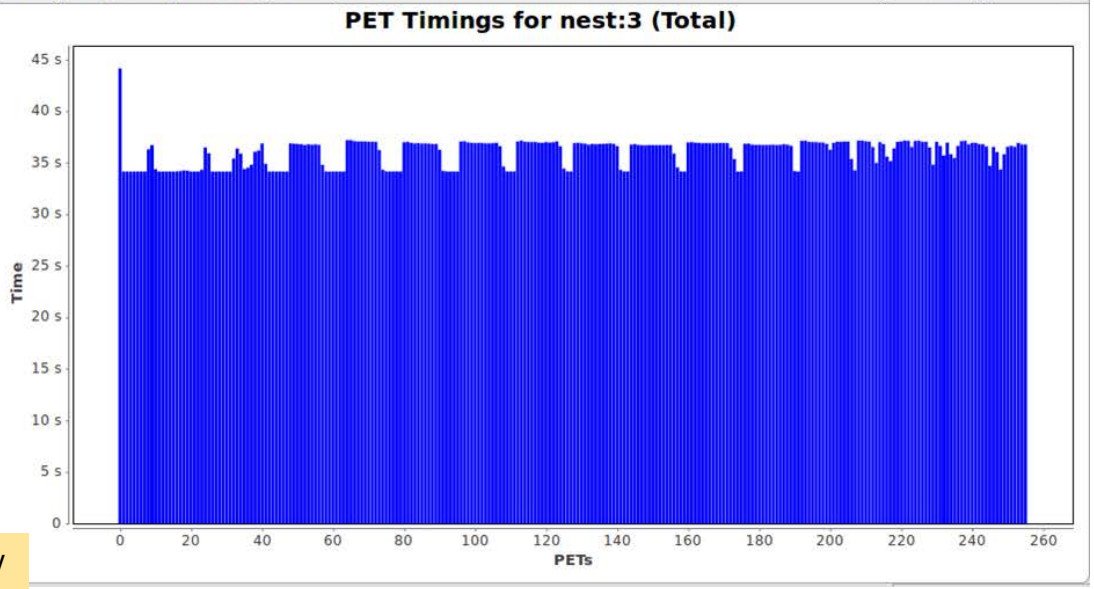

Trace of Navy's COAMPS with coupled atmosphere-land-hydrology.

ESMF Training Program

Upcoming Webinars	
July 24, 2018 - 11 MT / 1 ET	Building coupled applications with NUOPC
September 25, 2018 - 11 MT / 1 ET	Overview of ESMF and NUOPC

Topic	Duration	Frequency	Format
ESMF/NUOPC overview tutorial	1.5 hours	quarterly	webinar
ESMF/NUOPC overview tutorial	1 day	as requested, at least yearly at NRL, EMC, and NASA	on-site
Advanced topics: Building a NUOPC cap	1.5 hours	as requested	webinar
Site-specific tutorials	4 hours	as requested	webinar
Cupid tutorial	1 day	as requested, at least yearly at NRL, EMC, and NASA	on-site

Visit earthsystemcog.org/projects/esmf/tutorials to register and for up-to-date training schedules.

ESMF Support List

General questions, technical support, and feature requests

esmf_support@list.woc.noaa.gov

Thank you!

ESMF Sponsors:

Extra slides

CMEPS v0.1: Active Component Milestone

Q2FY18: Demonstrate running the community mediator within the NEMS framework with at least one active component.

- Configuration is MOM6 running with prescribed data atmosphere, ice, and runoff, coupled using CMEPS
- Distribution includes the CIME workflow (i.e. case control system), model components, and mediator
- Started with the MOM6 NUOPC cap used in NEMS
- Runs on NOAA (Theia) or NCAR (Cheyenne) computers - portability important for collaborative development.
- Available on GitHub under the Earth System Community Modeling Portal (ESCOMP), which is managed by NCAR: https://escomp.github.io/

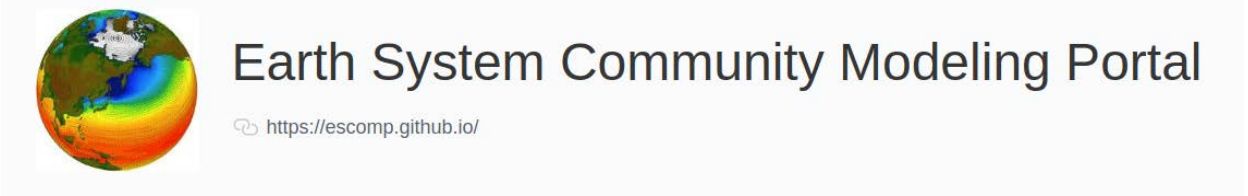

ESMF Projects

Current support

Title	Start	End	Sponsor
An Integrated, Observation-Driven Hydrological Modeling System Using LIS and WRF-Hydro Enabled by ESMF	12/1/2015	11/30/2019	NASA MAP
Extending Interoperability of ESMF-Based Models at NASA	8/1/2015	7/31/2018	NASA MAP
NOAA Environmental Software Infrastructure and Interoperability team	5/1/2015	4/30/2018	NWS
Coupling Between the Whole Atmosphere Model (WAM) and the Ionosphere-Plasmasphere Electrodynamics (IPE) Model	8/1/2016	7/31/2018	NOAA SWPC
Modeling and Data Infrastructure in Support of NOAA's Global Models	8/1/2015	7/31/2018	NOAA CPO
An Integration and Evaluation Framework for ESPC Coupled Models	8/1/2013	11/30/2018	ONR NOPP
An Integrated Hydrological Modeling System for High-Resolution Coastal Applications	4/1/2015	3/31/2018	NRL
Earth System Modeling Framework and NUOPC Layer Development	10/1/2015	9/30/2018	NRL
Development of Unified Forecast System Aerosol Component	9/1/2017	2/28/2018	NOAA GSD
Building a Distributed Earth System Model Development Community	6/1/2017	5/31/2018	ESPC

Examples of Other Application Activities

- Continued development of the Navy global coupled system including NAVGEM, HYCOM, CICE, LIS la
- Continued development of the Navy regional modeling system based on COAMPS, with LIS land and
- Coupling of LIS land and WRF-Hydro with NU-WRF at NASA, exploring representation of coupled data
- Work with the ModelE and GEOS-5 teams to merge/refactor/interface NASA's infrastructure with new
- Completion of CESM and high resolution HYCOM runs under Navy funding
- Collaboration with the Joint Effort for Data assimilation Integration (JEDI) community data assimilation
- Development of NUOPC-based asynchronous I/O component and implementation in HYCOM

NOAA Environmental Modeling System (NEMS) Application Status

Focus on bringing in FV3GFS as a new atmosphere model:

- Spring 2018 release of FV3GFS driven by NEMS, optimized using multiple asynchronous I/O components.
- Coupled FV3GFS, MOM6 and CICE system for S2S prediction in progress; previously delivered fully coupled GSM, MOM5, and CICE components, initialized for a cold start and optimized for comparable performance with CFSv2.
- Two-way space weather coupling in progress; previously completed development version of Whole Atmosphere Model (WAM) and the Ionosphere-Plasmasphere Electrodynamics (IPE) models validated under NEMS, with a one-way (WAM>IPE) 3D coupling exchange.
- WAVEWATCHIII forced by FV3GFS atmosphere completed; previously implemented WAVEWATCHIII coupled two-way to GSM, including nesting

Model Component Liaison (MCL) Committee

Chartered October 2017

Membership Representatives of Earth system model components that have implemented the ESMF/NUOPC interoperability standards

Importance **The MCL Committee embodies an Earth system model code base that U.S. federal agencies rely on**

Current Challenge *Usability* – making sure components can be accessed, understood and improved by partners, and verified for correct operation

Objectives Develop guidelines for model component repositories, documentation, and testing, and work with component developer teams to implement them

Approach Enable more experienced model component development teams to mentor and guide newer and smaller teams

Sample Activities

- Survey of teams to understand status relative to guidelines
- Learn about tools (Cupid, Component Explorer, …)
- Real examples - e.g. how can we coordinate use, evolution, and testing of MOM6 and its NUOPC cap (wrapper) across sites?

Atmosphere	NAVGEM, FV3GFS, CAM, COAMPS
Ice	CICE, SIS, KISS
Ocean	MOM, HYCOM
Land	LIS/Noah, CLM/MOSART
Space	Navy-HITIDES, IPE
Hydrology	WRF-Hydro
Wave	WAVEWATCHIII
Chemistry	GOCART/GSD, GEOS-Chem
Coupler	CMEPS

Members of the MCL committee represent NUOPC-compliant codes that are in the Earth System Prediction Suite (Theurich et al 2016)

Performance and Memory Optimizations Ongoing

- **Do-no-harm**:
  - Overhead of ESMF/NUOPC component interfaces is small (for ESMF, ~μs)
  - Set of NUOPC prototypes demonstrates preservation of accelerator and other component-specific optimizations
- **Scalable** - key methods (e.g. sparse mat mul) tested to ~ 16K processors by ESMF team, ~30K processors by customers (e.g. NASA)
- Component interfaces and sequential/concurrent modes support **increasing task parallelism and optimized mappings to hardware**
- Data communications between components can **preserve locality:**

Components with the same grid and decomposition:
- **direct reference sharing**
- **local memory-to-memory copy**

Components on disjoint processor sets:
- **redistribution**
- **parallel grid remapping**

Performance sudies available at: https://www.earthsystemcog.org/projects/esmf/performance/

Interpolation Weight Generation Performance

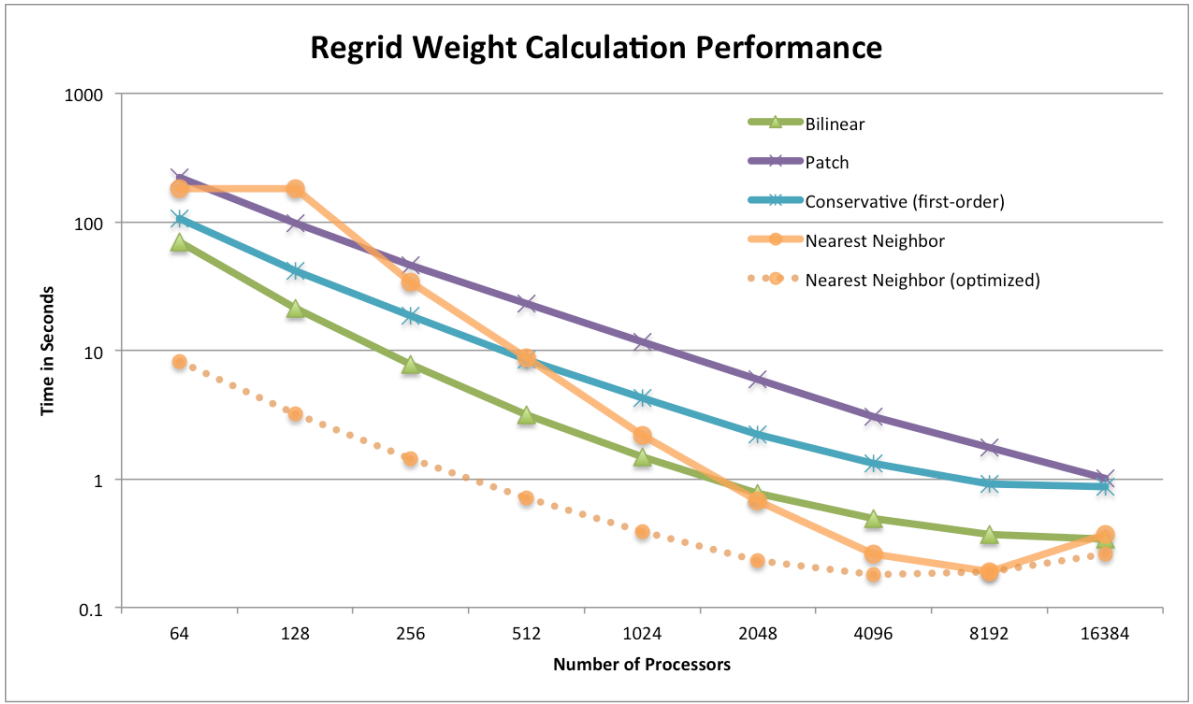

Source: regular cubed sphere grid (~25 million cells)
Destination: uniform latitude longitude grid (~17 million cells)
Platform: IBM iDataPlex cluster (Yellowstone at NCAR)
ESMF version: ESMF_7_1_0_beta_snapshot_24