CESM and ESMF

Advancing Earth System Modeling via New Infrastructure Capabilities

Mariana Vertenstein CESM Software Engineering Group Head



NSF

This material is based upon work supported by the National Center for Atmospheric Research, which is a major facility sponsor ed by the National Science Foundation under Cooperative Agreement No. 1852977

June 15, 2022

Outline

- New inter-component coupling capabilities via CMEPS
- New hierarchical modeling capabilities via CDEPS
- New ultra-high resolution simulations



New inter-component coupling capabilities CMEPS

Community Mediator for Earth Predictive Systems https://github.com/ESCOMP/CMEPS

Key part of NCAR/NOAA MOA





CMEPS architecture



Prognostic CMEPS compliant components

NUOPC is now the default coupling infrastructure in CESM and NOAA/UFS and will be in NORESM and CMCC-CM3+ Extensive validation done in multiple configurations of CESM, including multiple century long fully coupled simulations!



UCAR

CMEPS compliant data components (CDEPS)



Benefits of CMEPS-Introducing new Grids

- Easier to introduce new grids (1) no longer need offline mapping files
 - Before: all inter-component mapping files were created offline
 - 25 mapping files needed for a fully coupled pre -industrial control
 - **Now**: all non-custom mapping files!!! are generated at run time.
 - Only 4 mapping files are needed
- Easier to introduce new grids (2) no longer need offline land fraction files
 - Land and ocean fractions on atm/land grid is determined by mapping ocean mask conservatively to land grid
 - **Before**: each new component grid required generating new offline fraction files and updating CIME configuration files.
 - Now: land and ocean fractions are generated at runtime during model initializations!.



Benefits of ESMF-Introducing new Grids (cont)

- Easier to introduce new grids (3) land surface dataset generation is now parallel!
 - Before: needed to create 17 offline mapping files and use these as input to a surface dataset generation code that ran one processor . Took over 2 days to generate a surface dataset at 7.5 km MPAS grid.
 - Now: all mapping is done at run time and all I/O is parallel. Now takes 10 minutes to generate a surface dataset for a 7.5 km MPAS grid.
 - Now: ESMF and PIO2 enable mapping of 30 second (724M points) soil texture dataset.
- Creation of new surface dataset capabilities has leveraged ESMF features like dynamic masking for determining the standard deviation of surface elevation statistics



Benefits of CMEPS-New Land-Ice Capabilities

• Running both Antarctica and Greenland in one simulation (1)

- **Before**: proposed approach was to create a unified global grid. New global grids would have to be created for every combination. Results in combinatorial explosion!
- Now: create a 'nested state' where each ice sheet in CISM couples to a corresponding ice sheet in the mediator. Very extensible and user friendly approach new glaciers can be added easily.
- CMEPS has also been extended so that an arbitrary number of ice sheets can be coupled at run time.
- This has been now validated with the latest CISM updates
- Enabling Antarctic ocean <-> land-ice coupling (2)
 - Requires regridding ocn->cism fields at multiple levels. Each level has different mask due to different bathymetry.
 - **Before**: a different mapping file for ocn->cism mapping was required for each ocean level. Each ocean level field was passed separately.
 - *Now:* can do regridding leveraging ESMF **dynamic masking functionality** in the mediator. Only one field with multiple levels is passed.



Benefits of CMEPS-Greater Computational Efficiency

- Components sharing cores can now have different threading levels using ESMF-managed threading
- Before:
 - If component A is threaded 4 ways and component B is not threaded, if they are to share the same nodes, component B can only use ¼ of the cores in a node
 - This leads to idle cores and poor HPC resource utilization
- Now:
 - If component A is threaded 4 ways and component B is not threaded, if they are to share the same nodes, component B can use ALL of the cores in a node
 - This greatly increases the efficiency of the overall model
- Pre-industrial, fully coupled run (2° atm/1° ocn)
 - Now Model Cost: 2531 pe-hrs/simulated_year
 Now Model Throughput: 35 simulated_years/day
 - Before Model Cost:
 Before Model Throughput:
- 3140 pe-hrs/simulated_year
 - 31 simulated_years/day



Benefits CMEPS-New Exchange grid capability for calculating atm/ocn fluxes

- Exchange grid is the union of atm and ocn grids.
- Traditionally in CESM atm/ocn fluxes were computed on the ocean grid problems arise if the atm grid is much higher resolution



CMEPS now can compute atm/ocn flux calculation on either the ocean grid, the atm grid OR the exchange grid



Idealized Tropical Cyclone Experiments (Adam Herrington) "ne120pg3_g17" grid alias



Extensive simulations, including fully coupled 100 year runs, have been done to validate the exchange grid in CESM!!! This will be the default in CESM.







The exi Mediatc cells, flu stability The ap; couplir

Status:

inge Grid Implementation in the idiator

OLFSR2C

tpability is being added to the CMEPS sgy for robustly handling fractional surface sit atm-surface coupling, and increased model is (desirable for S2S regimes).

ised models to support exchange grid ad in GFDL and NASA modes.

n xgrid online from separate atm and ocn 96/.25deg).

- Atmosphere-ocean fluxes can be computed in the Mediator. Runtime options (agrid, ogrid, xgrid) allows the user to choose the desired grid for surface flux computations.
- Seamless model development from atmosphere-only to fully-coupled. CCPP surface schemes can be called from CMEPS.
- Initial runs with the exchange grid option used for a-o fluxes show close correspondence between "prototype 8" reference simulations.

teams have started joint development calls to coordinate hange grid implementation in UFS ESMF and efforts on t

NWS/OSTI & OAR GGPS - HFIP - W

GGPS - HFIP - Week 3&4 - JTTI - EPIC



UFS *cpld_control_c36_p8* configuration (C96 atm/1⁺ ocn-ice). Top left: Reference simulation SST. *Top right* a-o fluxes computed on atm grid in CMEPS. Bottom left: a-o fluxes computed on ocn grid in CMEPS. Bottom right' a-o fluxes computed on exchange grid in CMEPS.



Time series of SST in Nino3.4 region comparing reference simulation to runs with a-o fluxes computed on the atm grid, ocn grid, and xgrid. Images courtesy Utuk TuruncogluMCAR/ESME.

https://vlab.noaa.gov/web/ufs-r2o



Component Update: CICE6 and WW3

- CICE6:
 - now default ice component in cesm2_3_beta08 shared cap with NOAA/EMC
- WW3:
 - new shared cap with NOAA/EMC provides access and collaboration to latest WW3 development code base (including new OpenMP capability)
 - default in upcoming CESM3 development tags
 - will facilitate exploration of multi-grids and unstructured grids
- new WW3/CICE6 coupling
 - wave fields break up the sea ice into smaller floes and the sea ice concentration and thickness feeds back onto the wave energy field.
 - sending new spectral data (25 fields) from WW3 to CICE6 with CMEPS can pack this into 1 field (with undistributed dimensions)



New Data Assimilation with JEDI through the NUOPC Driver

JEDI is a model agnostic data assimilation framework that implements various DA algorithms in a generic way.

The JEDI **Object Oriented Prediction System** (OOPS) layer provides a generic interface to forecast models for DA

The NUOPC Driver provides full range of access to the complete forecast model

- Access to individual component states
- Access to the combined coupled model state
- ESMF regridding methods to interpolate between model grid and observation location
- ESMF redistribution methods to move data from model memory layout to JEDI memory layout
- ESMF reference sharing for efficient high volume data access
- The mediator provides a place for temporal interpolation or aggregation to support different DA windows (e.g. fast ATM and slow OCN)
- Mechanism for sending modified states back to the individual components and/or the combined coupled state to the mediator





Hierarchical Model Development Capability

CDEPS

Community Data Models for Earth Prediction Systems https://github.com/ESCOMP/CDEPS





Hierarchical Model Development: A simple-to-more-complex comprehensive approach to identify systematic biases and improve models.





CDEPS

- CDEPS contains ESMFNUOPC compliant data components that are modular and flexible: Can be used in any ESMF/NUOPC compliant modeling system
- CDEPS handles the ability to ingest multiple data sources with different spatial and temporal resolutions. Also provides ability to customize the ingested data (e.g. unit conversions)
- All data is read with parallel IO (PIO2)- can easily ingest 2d or 3d fields!
- Automated regridding capability: 1) online regridding of 2D/3D fields, 2) support for different regridding types such as conservative, patch, 3) extrapolation and 4) various time interpolations (coszen, bilinear, etc)
- Inline data models: CDEPS share code provides an interface that can be called directly from prognostic components and is used throughout CESM (future targeting of aerosol ingestion, nudging)



CDEPS (cont)

CDEPS provides data component support in both NOAA/UFS and CESM and supports many different forcing scenarios





Ultra-high resolution cloud resolving model configuration



Earthworks and ESMF

EarthWorks is a five-year NSF funded university based project to develop a global coupled model, based on CESM and the CESM coupling framework (CMEPS) that uses **a single uniform MPAS 3.75km global grid** for the atmosphere, ocean, sea-ice and land surface.

- MPAS atmosphere non-hydrostatic dynamical core
- A resolved stratosphere
- High-resolution CAM-ish physics
- The Community Physics Framework (CPF, a.k.a. CCPP)
- The MPAS ocean model developed at Los Alamos and used by E3SM
- The MPAS sea ice model, which is based on CISE and designed to work with the MPAS ocean model
- The Community Land Model (CLM)
- The Community Mediator for Earth Prediction Systems (CMEPS)





Earthworks and ESMF (cont)

- Targeting new ultra high resolution configurations in CESM exposed memory bottlenecks in ESMF mesh ingestion and IO
- Working closely with the CSEG group at NCAR, preliminary scalability issues have been resolved and are part of the new ESMF release
- Continued work on EarthWorks configurations will enable us to push ESMF and CESM to new high resolution configurations that will benefit all users
 - memory reductions
 - surpassing 32-bit integer limit (mesh size)
 - smart data movement with model data on GPUs







New ESMF 8.3 Release on June 8

The main new features that will be included in the release are:

- Scalable mesh creation from file. Substantial improvements were made to the ESMF_MeshCreate() function that reads in unstructured meshes from file. The new function removes a memory limitation that prevented reading in high resolution meshes. The function has been successfully tested on a global ~3km resolution mesh.
- **Major upgrade to internal ParallellO (PIO) library.** PIO is a parallel IO package developed at NCAR that is used by ESMF internally to read and write ESMF data structures to/from file, such as ESMF Meshes and Fields. The very old version of PIO v1.x has been upgraded to PIO v2.5. Users can also now link to an externally built PIO library. The new library will provide more flexible parallel output configurations.
- API updates to support asynchronous output of multiple static and moving nested domains. A new API was created to redistribute Grid coordinates from a source decomposition to a different destination decomposition. This is used in conjunction with regridding to support asynchronous regridding and file output in contexts with dynamically changing grids.
- Exchange grid improvements. A new
 ESMF_FieldGetArea() API allows the user to retrieve cell areas of fields built on an ESMF_XGrid. Fixed the XGrid regridding weight calculation to re-use weights generated during XGrid construction.
- Internal MOAB mesh updates. Updated the internal MOAB mesh library to 5.3.1 and introduced a large number of test cases to ensure consistency between the native mesh and new MOAB mesh. This is part of a long term effort to replace the native mesh implementation with the more performant and flexible MOAB library.



Thank you!

Questions?



Benefits of CMEPS-Interagency collaboration

• New code sharing

- NOAA-NCAR MoA coupling collaboration using CMEPS required NCAR and NOAA to "speak the same language" (i.e., ESMF) and this greatly accelerated the ability to share code and components (currently MOM6, CICE6 and WW3)
- How is CMEPS being developed?
 - Collaboratively between NCAR and NOAA
 - Openly on GitHub to allow community code contributions and encourage collaboration and innovation
 - All CMEPS PRs are reviewed by both NOAA and NCAR software engineers and tested with both NOAA/UFS and CESM use cases
- Exercising CMEPS with different components and coupling strategies has resulted in a much more robust and more widely tested system



Benefits of CMEPS- much easier debugging

@1800 < MED med phases prep ocn accum avg MED -> OCN :remapMethod=redist OCN @900 MED med phases prep atm MED med phases prep ice MED -> ATM :remapMethod=redist MED -> ICE :remapMethod=redist ATM ICF ATM -> MED :remapMethod=redist ICE -> MED :remapMethod=redist MED med fraction set MED med phases prep ocn map MED med phases aofluxes run MED med phases prep ocn merge MED med phases prep ocn accum fast MED med phases history write @ OCN -> MED :remapMethod=redist MED med phases restart write @

- Simple syntax for driver looping structure - component coupling frequency and order of component execution – replaces thousands of lines of previous driver code
- Generic connectors transfer data between mediator and components and are generated automatically
- Components can have multiple named phases
- Run sequence can be changed without recompiling
- Sequential and concurrent execution in separate runtime configuration
- Run sequence generated automatically by CIME for a given case

