

Highlights from a workshop on the latest updates on coupling technologies and coupled applications in Earth System modelling



Sophie Valcke ^a

^a *CECI, Université de Toulouse, CNRS, CERFACS, France*

Corresponding author: Sophie Valcke, valcke@cerfacs.fr

The 5th Workshop on Coupling Technologies for Earth System Models

What: 152 leading researchers and practitioners in the field of coupling software for Earth System Models from all around the world registered to the 5th Workshop on Coupling Technologies for Earth System Models (CW2020), and about 80 people attended each on-line session. Five different sessions gathering 35 presentations and an additional 2-hour session were organised to discuss the latest updates on coupling technologies, coupled applications in Earth System modelling, computational performances of coupled models, links between data assimilation and coupling, and coupled model workflows.

When: September 21st-24th 2020 with one on-line session of 3.5 hours each day.

Where: The workshop was held virtually.

1. Introduction

The 5th Workshop on Coupling Technologies for Earth System Models (CW2020) brought together leading researchers and practitioners in the field of coupling infrastructure for Earth System Models. This workshop was the fifth in the series¹, started in 2010 in Toulouse and followed by the ones in Boulder USA (2013), Manchester UK (2015), and Princeton USA (2017).

¹ <https://portal.enes.org/community/community-workshops/>

152 people from all around the world registered, and about 80 people attended each session. The complete program is available on-line² as well as the presentations that were recorded³. A document was opened on-line for questions. Participants could answer some questions directly in the document during the presentation, leading to fruitful interactions between the workshop attendees, while the remaining questions were answered by the speaker at the end of the talk. At the end of the workshop, this document was 48 pages long! The different themes discussed and the conclusions and recommendations that can be drawn from the workshop are detailed here.

3. Overview of sessions and discussions

We provide here an overview of the coupling software, coupled applications and coupling issues presented and discussed during the workshop; this is not a chronological description of the sessions as the subjects addressed often overlapped. Web links are provided in the appendix for all tools and coupled applications mentioned here.

a. Coupling technologies and higher-level coupling layers

² See <https://portal.enes.org/community/community-workshops/fifth-workshop-on-coupling-technologies-for-earth-system-models>

³ Presentations are available on the IS-ENES3 Youtube channel <https://youtube.com/playlist?list=PLFvev1W5vG7O-h0H1AYrXHtkPQ3zAOZjt> .

Different coupling technologies were presented during the workshop. As already observed, coupling technologies can be broadly classified into two prominent categories, i.e. the “external coupler or coupling library”, and the “integrated coupling framework”. In the first approach, illustrated by presentations on OASIS3-MCT, YAC, C-Coupler2, and MOAB-TempestRemap, the component models remain separate executables and the original codes are modified as little as possible; the component codes are instrumented with calls to the coupling library Application Programming Interface (API), and the synchronization of the components is implicitly ensured by the coupling exchanges. The “integrated coupling framework” approach, for which ESMF/NUOPC is a quintessential representative (as are FMS at the Geophysical Fluid Dynamics Laboratory or CPL7 at the National Center for Atmospheric Research) involves splitting the original component codes into initialize, run and finalize units, adapting these units to the framework standard data structures and routine interfaces, and rebuilding a single integrated application based on these units, with a top-level driver controlling the component execution. Updates on the different couplers were presented, including details on their regridding/interpolation functionality, their community, and sample applications.

Different higher-level coupling layers were also discussed, such as ESM-interface, a coupling interface to both OASIS3-MCT and YAC, the MAPL ESMF-based coupling layer within the GEOS (Goddard Earth Observing System) data assimilation system, or CMEPS (Community Mediator for Earth Prediction System) that is the new NUOPC-compliant coupling architecture of the NCAR Community Earth System Model (CESM).

A specific issue, the compatibility or “coupleability” of the different coupling software, was discussed in more detail. Potential use cases are, for example, coupling one component interfaced with OASIS3-MCT with an ESMF/NUOPC-compliant component, or assembling an ESMF coupled system containing an OASIS3-MCT-coupled application. BMI (Basic

Model Interface), which is an attempt to define a standardized interface to geophysical model functions (e.g. control, time, grid), was presented as one approach favouring compatibility of models in general. It was however reported that most of the previous attempts to build generic coupling interfaces have been broken by additional specific requirements or use cases. A more pragmatic approach might be to simply accept the need to develop different interfaces, or “caps”, one for each specific coupling software. This may, however, be difficult to implement in practice across couplers that have fundamentally different design philosophies.

Issues related to regridding/interpolation were heavily discussed, such as the interesting concept of the interpolation stack (user-controlled layering of methods) implemented in YAC, a new interpolation capacity in the ATLAS library developed at ECMWF, or the standardization of interpolation weight file format for flexible reuse. OBLIMAP, a specific regridding package for coupling general circulation models and ice sheet models was also presented.

Details were also shared on the implementation of the regridding in the different couplers, in particular for the conservative remapping. It was interesting to realize that the 2nd order conservative remapping in XIOS, YAC and ESMF is based on the same algorithm described in Kritsikis et al 2017. MOAB-TempestRemap includes a special scheme based on an advancing-front intersection algorithm and a resulting “supermesh” based on the source and target grids.

It was also proposed to develop a benchmark for comparing accuracy and performance of the different regridding implementation. Metrics defined in the CANGA-ROO project to evaluate the sensitivity, conservation, consistency, monotonicity, and performance of the tools

were discussed. Many coupling software developers expressed interest in participating in that regridding benchmark.

b. Coupled applications

Several talks on different applications coupling the components of the Earth System (i.e. ocean, atmosphere, sea-ice, land, as well as ice sheet, biosphere or coastal ecosystem constituents) were presented including :

- the US Navy Earth System Prediction Capability (Navy-ESPC), a global coupled atmosphere-ocean-sea ice prediction system ;
- the Unified Forecast System (UFS), the NOAA community-based operational weather and environmental modeling system ;
- the atmosphere-biosphere-land coupled system developed at the National Institute of Technology in India ;
- the ocean-wave-atmosphere coupled model developed at the "Laboratoire de l'Atmosphère et des Cyclones" (LACy) in La Réunion for the prediction of tropical cyclones ;
- EC-Earth4, a European community Earth System Model ;
- the ocean-atmosphere model ICONGETM, used at the Leibniz Institute for tropospheric research to study local processes over the Baltic sea ;
- TerrSysMP, the Terrestrial System Modeling Platform, coupling the ICON atmosphere model and the Community Land Model at the Juelich Research Center ;
- MOSSCO, the coastal ecosystem coupling physical, biogeochemical, ecological and geological processes at the Helmholtz-Zentrum Geesthacht center ;
- the ocean-sea-ice model assembling NEMO and neXtSIM, a sea-ice model using a Lagrangian moving mesh to better simulate the discontinuities in sea-ice properties;

- UKESM1, the Earth System Model developed in the UK, including in particular an interactive ice sheet model.

Details on how to conservatively couple runoff in global climate models were also introduced. It was noted that this type of coupling has a reverse rationale compared to other couplings, in the sense that the aim is to make sure that all source grid points discharge their runoff in at least one target grid point, whereas the usual rationale is to make sure that each target point gets one meaningful value. Different conservative methods with horizontal and/or vertical spreading of the different fluxes were evoked.

The talks on neXtSIM, the sea-ice model using a Lagrangian moving mesh, and on the interactive ice sheet model in UKESM1 generated interesting discussions related to dynamic coupling. These coupled models can suffer from crude and sub optimal coupling implementations. In some cases, the evolution of the coupling characteristics (e.g. the ocean solid boundary and 3D grid due to the advancing/retreating ice sheet) can be handled only by restarting the coupled model. This is linked to the need to recalculate the regridding weights dynamically, but also, for example, to the evolution of the ocean topography. Coupling infrastructures need to handle these dynamics grids more efficiently in the future for them to remain relevant.

c. Computational performances of coupled models

A few talks, reaching a high level of expertise, discussed the importance of running applications with proper process and thread affinity across diverse hardware architectures and memory hierarchies to optimize model performance. It was noted that the question is even more complex for climate coupled applications where the different components can have different ratios between MPI processes and (possibly nested) OpenMP threads. Furthermore, this ratio

may have to evolve at runtime given the different tasks to chain. For example, the same processes may be initially calculating the regridding weights and then later running a component model. It was observed that the question of how to optimally map the coupled application on available resources is a delicate and often platform-dependent issue. Tools such as HIPPO and Quo that dynamically rebind processes and threads to match the different phases of an application with calls from within the application code were introduced.

Tools (e.g. FLECSI by the Los Alamos Lab) that help manage the complexity of coupled models on different architectures were presented. Methods to define an optimal use of available resources and optimal load balance of the different components, based on performance metrics calculated by the coupler during a coupled run, were also discussed and compared with the possibility of using more complex profiling tools such as Extrae or Paraver.

Another question specifically addressed was on the readiness of couplers to deal with heterogeneous hardware, e.g. CPU-GPU systems. Although some early prototyping has occurred, this issue has not been addressed properly yet in most couplers. Some doubts were raised about an efficient use of GPUs by couplers given their relatively small compute load. At the same time, the potential role of GPUs for offline regridding weights generation seems more viable. In all cases, coupling infrastructure working with top-level drivers should be aware of GPU resources so as to appropriately allocate resources to the different components. In particular, an abstraction layer was developed for ESMF to add the ability to map different components onto the available resources, including GPUs if available.

d. Data Assimilation (DA) and Numerical Weather Prediction (NWP)

At least four different prediction systems implementing weakly coupled data assimilation (i.e. data assimilation performed in each component separately) were presented during the workshop:

- at the UK Met Office, the system based on the Unified Model atmosphere, the JULES land surface, the NEMO ocean and the CICE sea ice models that should be in use for operational NWP from autumn 2021 onwards;
- ECMWF weakly coupled data assimilation including waves, land, ocean and sea ice components;
- Environment and Climate Change Canada system, built on four independent atmospheric, ocean, SST and sea-ice data assimilation components;
- the Navy Earth System Prediction Capability (Navy-ESPC) global coupled model based on the Navy's Global Environmental Model (NAVGEM) for the atmosphere, Navy's HYbrid Coordinate Ocean Model (HYCOM) for the ocean, and the Los Alamos National Laboratory's sea ice (CICE) for the sea ice.

A general discussion addressed the relationship between data assimilation tools (like OOPS, JEDI, PDAF, OpenPALM) and couplers and, more generally, between data assimilation systems and coupled applications. Even if some coupling frameworks, like FMS, include data assimilation functionality, it was observed that in general, the infrastructure used for coupling and for DA have been developed independently but that it may be useful to start converging these two tools to minimize duplication of similar developments like interpolation. A specific workshop could be organized on this subject.

e. Coupled model workflow

Many different solutions for managing coupled model simulations were discussed in different sessions during the workshop, which indicates that model coupling and workflow are two tightly linked issues:

- the eWaterCycle platform offers an integrated environment for hydrological modelling;
- ESM-Tools is a collection of scripts to download, compile, configure, and run different Earth system simulations;
- OMUSE is a Python environment for numerical experiments in oceanography and other climate sciences including, in particular, model setup and runs, run control, online data analysis, ensemble simulations and model comparison;
- EC-Earth4 workflow environment relies on different tools (CPLNG, ScriptEngine, Autosubmit, ocp-tool) to manage the different steps in the workflow (development, build, configure, run, monitor, postprocess);
- The workflow at Institut Pierre Simon Laplace (IPSL) uses the XIOS IO server together with dr2xml, a package for translating a Coupled Model Intercomparison Project (CMIP) Data Request into XIOS configuration files, to directly produce CMIP6 publication-ready data files;
- CESM and E3SM (Energy Exascale Earth System Model) are using the python-based CIME Case Control System.

An interesting discussion dealt with the difficulty of coherently configuring the coupled system, the coupler, and the component models. One challenge is that coupling information, that needs to be consistent across the application, can sometimes be duplicated in multiple components.

4. Summary and Recommendations

The 5th Workshop on Coupling Technologies for Earth System Models provided excellent interactions and discussions on the latest updates on coupling technologies, coupled applications in Earth System modelling, computational performance of coupled models, links between data assimilation and coupling, and coupled model workflows.

In general, the community would like to see a reduction in the number of couplers available by merging the ones that offer the same functions and follow the same design philosophy. In practice, this may be difficult to achieve, but coupler developers should continue to collaborate towards using and merging common features. Various approaches that address either coupling interface standardization or implementation of multiple interfaces were discussed during the workshop.

The implementation of regridding in different couplers, in particular the conservative remapping, raised a lot of interest. It was proposed to develop a benchmark for comparing the quality and the performance of the regridding algorithms in the different couplers and this effort should be taken forward by the coupler developers.

The questions of dynamic coupling and coupled model performance also generated interesting discussions. Methods using coupling information to define an optimal load balance of the different components in the coupled application were presented. It was also stressed that couplers need to ensure (or at least be compatible with) a proper execution of the component models with respect to process and thread affinity, given the specific architecture and memory hierarchy of the computing platform. The discussion also showed that more reflection is needed on how couplers could exploit new heterogeneous hardware, e.g. CPU-GPU systems.

Regarding data assimilation, it was proposed to organize a specific workshop on data assimilation and coupling software infrastructures, in order to minimize duplication of developments and ensure compatibility.

Finally, the workshop showed that many different solutions exist for managing the coupled model simulation workflow. Again, the community should aim at more software sharing on this aspect of workflow management, in order to reduce the total development effort and minimise duplications.

It was stated that the next Coupling Workshop could be held before the end of the IS-ENES3 project, i.e. before December 2022, and the possibility of organizing it back-to-back with the next Physics Dynamics Coupling in weather and climate models workshop (see <https://pdc-workshop-abstracts.com/event/2/>) was very well received.

Acknowledgments.

The 5th Workshop on Coupling Technologies for Earth System Models was organized in the framework of the IS-ENES3 project that has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824084.

Data Availability Statement.

The workshop program and abstracts are freely available on the workshop web page at <https://portal.enes.org/community/community-workshops/fifth-workshop-on-coupling-technologies-for-earth-system-models> . The presentations were recorded and are available on the IS-ENES3 Youtube channel (<https://youtube.com/playlist?list=PLFvev1W5vG7O-h0H1AYrXHtkPQ3zAOZjt>).

APPENDIX - WEB LINKS FOR TOOLS AND COUPLED APPLICATIONS

ATLAS: <https://confluence.ecmwf.int/display/ATLAS/Atlas>

BMI: <https://bmi-spec.readthedocs.io/>

C-Coupler2: <https://gmd.copernicus.org/articles/11/3557/2018/>

CANGA-ROO : <https://github.com/CANGA/Remapping-Intercomparison>

CESM: <http://www.cesm.ucar.edu>

CIME: <https://e3sm.org/resources/tools/other-tools/cime/>

CMEPS: <https://github.com/ESCOMP/CMEPS>

CPL7: <https://www.cesm.ucar.edu/models/cesm1.0/cpl7/>

E3SM: <https://e3sm.org>

EC-Earth : <http://www.ec-earth.org>

ESM-interface: <https://www.esm-project.net>

ESM-Tools: <https://www.esm-tools.net/get-esm-tools/>

ESMF: <http://earthsystemmodeling.org>

eWaterCycle: <https://www.ewatercycle.org>

Extrac: <https://tools.bsc.es/extrac>

FLeCSI: <https://laristra.github.io/flecsi/>

FMS: <https://www.gfdl.noaa.gov/fms/>

GEOS: <https://gmao.gsfc.nasa.gov/GEOS/>

HIPPO: https://cerfacs.fr/wp-content/uploads/2019/06/Globc-WN-Maisonnavé-oasis_hippo-2019.pdf

JEDI: <https://www.jcsda.org/jcsda-project-jedi>

MAPL: <https://dl.acm.org/doi/abs/10.1145/1297385.1297388>

MCT: <https://www.mcs.anl.gov/research/projects/mct/>

MOAB-TempestRemap: <https://gmd.copernicus.org/articles/13/2355/2020/>

MOAB: <https://press3.mcs.anl.gov/sigma/moab-library/>

MOSSCO: <http://www.mossco.de>

Navy-ESPC: <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2020EA001199>

NUOPC: <http://earthsystemmodeling.org/nuopc/>

OASIS3-MCT: <https://portal.enes.org/oasis>

OBLIMAP: <https://github.com/oblimap/oblimap-2.0>

OMUSE: <https://gmd.copernicus.org/articles/10/3167/2017/>

OOPS: <http://www.data-assimilation.net/Events/Year3/OOPS.pdf>

OpenPALM : http://www.cerfacs.fr/globc/PALM_WEB/

Paraver: <https://tools.bsc.es/paraver>

PDAF : <http://pdaf.awi.de/>

QUO: <https://github.com/lanl/libquo#citing-quo>

TerrSysMP: <https://www.terrsysmp.org>

UFS: <https://vlab.ncep.noaa.gov/web/environmental-modeling-center/unified-forecast-system>

UKESM: <https://ukesm.ac.uk>

XIOS: <http://forge.ipsl.jussieu.fr/ioserver/wiki>

YAC: <https://dkrz-sw.gitlab-pages.dkrz.de/yac/>

REFERENCES

Kritsikis, E., M. Aechtner, Y. Meurdesoif, and T. Dubos, 2017: Conservative interpolation between general spherical meshes, *Geosci. Model Dev.*, **10**, 425–431, <https://doi.org/10.5194/gmd-10-425-2017>