

MEETING SUMMARIES

SHARING EXPERIENCES AND OUTLOOK ON COUPLING TECHNOLOGIES FOR EARTH SYSTEM MODELS

BY SOPHIE VALCKE, ANTHONY CRAIG, ROCKY DUNLAP, AND GRAHAM D. RILEY

The Third Workshop on Coupling Technologies for Earth System Models [Coupling Workshop 2015 (CW2015)] continued the series started in Toulouse, France, in 2010 (Valcke and Dunlap 2011) and pursued in Boulder, Colorado, in 2013 (Dunlap et al. 2013).

A coupling technology for Earth system models (ESMs) is software that links together the model components of an ESM. The coupling technologies used in the geophysical community typically offer functions such as managing data transfers between components, interpolating data between the different component grids, coordinating the execution of the components, and providing other support utilities, for example, for logging and calendar management.

The three-day workshop included 22 talks distributed through four themes: “coupling technology”; “coupling systems”; “interoperability, related domains and metadata”; and “metrics, benchmarking and

THE THIRD WORKSHOP ON COUPLING TECHNOLOGIES FOR EARTH SYSTEM MODELS

WHAT: About 40 researchers and engineers from nine different countries in North America, Europe, and Asia came together to discuss the latest developments and applications of coupling technologies for Earth system models (ESMs). Topics included interoperability, recent efforts in characterizing the technologies and benchmarking their performance, and how to better share and govern coupling technologies within the ESM community.

WHEN: 20–22 April 2015

WHERE: Manchester, United Kingdom

performance.” Two working sessions gave the participants the opportunity to discuss issues linked to “Interoperability and Earth system framework description language” and “Future issues: Interoperability, sharing of models/infrastructure, governance, exascale.”

The workshop program, presentation slides, and abstracts are available on the CW2015 website (www.earthsystemcog.org/projects/cw2015). This current summary provides an overview of the presentations and of the discussions they generated, highlighting new and original aspects that came out of the workshop.

A significant part of the workshop was devoted to recent updates on coupling software capabilities in three existing technologies: the Earth System Modeling Framework (ESMF), the Chinese

AFFILIATIONS: VALCKE—CERFACS/CNRS, URA1875, Toulouse, France; CRAIG—Software Engineer Contractor, Seattle, Washington; DUNLAP—NOAA/ESRL, and CIRES, Boulder, Colorado; RILEY—University of Manchester, Manchester, United Kingdom

CORRESPONDING AUTHOR: Sophie Valcke, CERFACS, 42 Avenue G. Coriolis, 31057 Toulouse CEDEX 01, France
E-mail: sophie.valcke@cerfacs.fr

DOI:10.1175/BAMS-D-15-00239.1

In final form 9 September 2015

©2016 American Meteorological Society

Community Coupler (C-Coupler), and the Ocean Atmosphere Sea Ice Soil - Model Coupling Toolkit version 3.0 (OASIS3-MCT_3.0). New ESMF re-gridding features described include offline weight generation and their application, support for sets of disconnected points (e.g., a set of observations), and normalization options for conservative remapping. An update was also provided on the development of the emerging C-Coupler tool, with an evaluation of the importance of bitwise reproducibility of ESM output. The new version of the OASIS coupler was presented, OASIS3-MCT_3.0, which supports coupling exchanges between parallel components deployed in more diverse configurations than prior versions, such as within the same executable between components running sequentially on overlapping sets of tasks (or processes). This evolution of the OASIS coupler generated interest because it breaks down some barriers that were previously assumed with regard to models coupled via specialized communication calls and without a top-level driver.

The National Unified Operational Prediction Capability (NUOPC) is a software layer on top of ESMF that is designed to ensure a greater degree of interoperability of model components. This is accomplished through the use of four generic components: driver, model, mediator, and connector with well-defined interactions. The Cupid Integrated Development Environment is a training and development tool that facilitates adding NUOPC interfaces into existing models using automated code generation. Generated code may be customized for particular application needs.

A majority of the presentations detailed ESM coupled applications, most using one of the coupling technologies discussed above. Coupled systems using OASIS3-MCT were presented, including the climate model from l'Institut Pierre-Simon Laplace (IPSL), a regional ESM, COSMO-CLM based on the Consortium for Small-Scale Modeling (COSMO) atmosphere model developed by the Climate Limited-area Modelling (CLM) community, the Met Office high-resolution coupled model (~25 km in the atmosphere, ~10 km in the ocean) and the 3D coupling between the atmosphere and chemistry components, and the EC-Earth Consortium (EC-Earth) developed by a consortium of European national weather services and universities. ESMF applications were also discussed, in particular a regional ESM of the Anatolian Peninsula (Turkey).

Other coupled systems were presented, including the Community Earth System Model (CESM)

developed at NCAR, the Geophysical Fluid Dynamics Laboratory (GFDL) ESM, and the new coupled ocean-ice-wave-atmosphere model developed at the European Centre for Medium-Range Weather Forecasts (ECMWF). Ongoing developments of two coupled systems targeting exascale computing were also detailed, including the U.K. Lewis Fry Richardson (LFRic) system, which is intended to deliver a replacement for the Unified Model toward the end of this decade, and the Accelerated Climate Modeling for Energy (ACME) project, which is a newly launched project sponsored by the U.S. Department of Energy (DOE) that will apply leading-edge climate and Earth system models to climate change research imperatives.

One conclusion from these presentations on ESM coupled systems is that the ESM technical infrastructure must constantly evolve to address the increase in both model complexity and resolution. Improving parallel performance within components and increasing concurrency across components were both identified as important goals for developers of coupling technologies. It was stressed many times during the workshop that more parallelism, increased concurrency, larger ensemble runs, and multimodel ESM runs are essential to prepare for the exascale computing era.

Three talks presented coupling challenges under the topic of related domains. The first one, on environmental modeling, highlighted difficulties when modeling phenomena covering wide ranges of scales and processes. For example, a study on how the mercury left by gold mining in the Sacramento River resurfaces and spreads to the ocean through the river delta involves heterogeneous models covering the urban, channel, estuary, coastal, sea, and global scales and domains. The Modular System for Shelves and Coasts (MOSSCO) is an ESMF application developed in Germany that integrates models and data from different disciplines related to coastal research. These two coupled applications implement the Basic Model Interface (BMI; described in more detail below) in their components. Finally, the development of a metadata standard describing environmental numerical models and their interfaces to facilitate their discovery and evaluation of suitability for a particular application was presented, with a first application in a hydrometeorological model chain.

As an introduction to the session on benchmarking and performance, the first talk reported on efforts currently ongoing in the Infrastructure for the European Network for Earth System Modelling

phase 2 (IS-ENES2) project to benchmark coupling technologies, which was initiated during the previous Coupling Workshop in Boulder in 2013. Mind maps gathering the characteristics of the coupling technologies were produced, and the coupling characteristics were analyzed to determine the benchmark priorities. A benchmark suite has been defined consisting of a number of precoded stand-alone components running on different grids, a number of coupling configuration specifications, and a list of test cases. The objective is to provide reference implementations using common technologies. The benchmarking work presently ongoing at the Met Office was also described, and the presentation highlighted the difficulties in providing comparable and meaningful results when applying the same measurements to different technologies. Finally, a talk detailed how the load balancing and optimal deployment on available computing resources is particularly complex for multicomponent coupled applications. For this purpose, standard but complex tools [e.g., VTune, Paraver, Scalasca, Tuning and Analysis Utilities (TAU)] are available and tools included in the coupling software are also being developed.

The workshop also noted that the interoperability of coupling technologies or, more generally, the interoperability of modeling frameworks, is currently a major challenge. This is the main driver of the Earth System Bridge, a project funded by the U.S. National Science Foundation to develop innovative approaches to allow different Earth system modeling frameworks to work together. Again, the importance of more standardization for interoperability was reaffirmed. A framework-agnostic BMI with universal adapters to go from a BMI to a modeling framework's native component interfaces was proposed, and ongoing efforts to establish a "crosswalk" between the Climate and Forecasting (CF) Standard Names and the Community Surface Dynamics Modeling System (CSDMS) Standard Names were described. Interoperability was also the main theme of a working session that helped the participants understand the Earth System Framework Description Language (ES-FDL). The ES-FDL is a language developed by the Earth System Bridge team that provides standard descriptions of the design and implementation of modeling infrastructure in the Earth sciences domain.

With regard to benchmarking, a process to characterize coupled ESMs, which started during the 2013 Coupling Workshop via a breakout session, produced a list of the scientific and technical

requirements needed to build a geophysical coupled system from independent models. Based on this list, the Earth System Bridge project proposed an initial series of mind maps highlighting the primary functions and characteristics of coupled systems. These were later extended and complemented in the framework of the IS-ENES2 project. The top-level mind map "CouplingSystem" references five sub-mind maps: "CouplingTechnology," "Components," "Metadata," "Composition," and "Deployment," which together describe all the characteristics of the coupling technology and the characteristics required of the components. The full set of mind maps is available online (<https://github.com/IWCCT/es-fdl>). The development of a questionnaire to populate the mind maps was then presented and discussed with the participants. A discussion of benchmarks highlighted the importance of defining both meaningful questions and answers to allow objective characterization of both functional and nonfunctional aspects of coupling technologies.

The second working session led to an animated discussion on community aspects of coupling technology development and how to fund, govern, and work together, especially in the context of the exascale computing systems expected before the end of the decade. A vast majority of the participants agreed that coupling software developers should continue to share experiences and code at the working level on a personal basis and through future Coupling Workshop meetings. People were also encouraged to contribute to a special issue of *Geoscientific Model Development* (2015, Vol. 8, No. 42) on model infrastructure integration and interoperability (MI3) recently launched in conjunction with this Coupling Workshop (CW2015). More interaction with the computer science community, and with other communities addressing multiphysics and multiscale coupled problems, was proposed and encouraged.

Moving forward, it was suggested to increase the activity of the International Working Committee on Coupling Technology (IWCCT), established during the previous workshop to facilitate communication among teams developing modeling infrastructure. Opinions were divided as to whether the scope of the committee should be extended to other elements of the geoscience modeling environment software, such as workflow management, testing and validation, and input/output (I/O). It was suggested that future Coupling Workshops could be coordinated with other technical workshops focusing on similar themes to facilitate discussion on overlapping issues. There was

also interest in framing an international modeling infrastructure project, given sufficient resources. Finally, there was clear enthusiasm for continuing the series of Coupling Workshops and the next one is scheduled for 2017 in the United States.

ACKNOWLEDGMENTS. We thank all the workshop participants who actively took part in the discussions and the IS-ENES2 project, funded under the European FP7-INFRASTRUCTURES-2012-1 call (Grant Agreement 312979), for financial support.

REFERENCES

- Dunlap, R., M. Vertenstein, S. Valcke, and T. Craig, 2013: Second Workshop on Coupling Technologies for Earth System Models. *Bull. Amer. Meteor. Soc.*, **95**, ES34–ES38, doi:10.1175/BAMS-D-13-00122.1.
- Valcke, S., and R. Dunlap, 2011: Report from the workshop “Coupling Technologies for Earth System Modelling: Today and Tomorrow.” *CLIVAR Exchanges*, No. 57, International CLIVAR Project Office, Southampton, United Kingdom, 38–39.