

Recommendations for the Ocean Model Dynamical Core for CESM3

Prepared by the Ad Hoc Panel Appointed by the CESM Science Steering Committee

Members

Frank Bryan (NCAR) - Chair

Ping Chang (Texas A&M Univ.)

Steve Jayne (Woods Hole Oceanographic Inst.)

William Merryfield (Environment and Climate Change Canada)

Luanne Thompson (Univ. of Washington)

Ex-officio Members and Ocean Model Working Group Co-Chairs

Gokhan Danabasoglu (NCAR)

Baylor Fox-Kemper (Brown Univ.)

Executive Summary

The Advisory Panel was charged with recommending a new base model code (dynamical core) for future generations of the ocean component of CESM, beginning with CESM3. The new base model will replace POP, which has been the base code for the ocean component beginning with CCSM2 and all subsequent versions through the forthcoming CESM2. Community input on requirements for the future ocean component was solicited through a CESM-wide e-mail survey in February 2016, and during discussions at the winter 2016 Ocean Model Working Group meeting. A Town Hall organized by NSF and ONR at the 2016 Ocean Sciences Meeting in New Orleans provided additional background information on the current landscape of community-based ocean modeling. With this background information, a Request for Information (RFI – Appendix 1) was sent to the groups developing the following ocean models:

- Hybrid Coordinate Ocean Model (HYCOM)
- MIT General Circulation Model (MITgcm)
- Modular Ocean Model version 6 (MOM6)
- Ocean Model for Prediction Across Scales (MPAS-O)
- Nucleus for European Modelling of the Ocean (NEMO)
- Regional Ocean Modeling System (ROMS)

Information on the potential of each model as a CESM ocean component including technical characteristics and organizational and governance practices in the associated developer communities was solicited. Written responses were received from all groups. Representatives from each of the groups attended the 2016 Breckenridge CESM workshop and presented overviews of the models and development practices at a well-attended cross-working group session. The presentations were followed by a question and answer session and open discussion of concerns of all impacted CESM users.

Based on an evaluation of this background information and input, the panel makes the following recommendations:

1. CESM should adopt MOM6 as the provisional ocean component of CESM3 and begin working with GFDL and the CESM developer and user communities as soon as possible to design and develop a configuration(s) that meets the requirements of the full scientific scope of CESM.
2. CESM should work with LANL and DOE management to find a mutually acceptable organizational structure to facilitate continued collaboration on ocean model development and permit the use of MPAS-O within the CESM framework for focused research projects.
3. CESM should invest in developing, documenting, and supporting an “open architecture” for interfacing other ocean dynamical cores (e.g. HYCOM, MITgcm, MPAS-O, NEMO, ROMS) with the CESM framework and provide technical assistance and coordination to university based research groups interested in using alternative ocean components in CESM.
4. CESM should engage the regional ocean modeling community to identify requirements for downscaling CESM global simulations, and develop protocols and requisite infrastructure to support the growth of this research area within the CESM community.

Background

The CCSM Ocean Model Working Group (OMWG) began collaborating with members of the Climate, Ocean, and Sea Ice Modeling group (COSIM) at Los Alamos National Laboratory (LANL) around the year 2000 to replace the original CSM ocean component, a heavily modified configuration of the Modular Ocean Model (MOM), with the Parallel Ocean Program (POP). This switch was motivated by a combination of scientific, technical, and organizational considerations. Perhaps the most pressing among these was the shift occurring at that time from moderately parallel vector architectures to massively parallel microprocessor-based architectures in the NCAR computing arsenal. Along with the POP code being more suited to the available computing resources, it brought along a number of advances in model structure and algorithms such as a general orthogonal horizontal grid and free surface dynamics, while retaining many elements of the familiar B-grid Bryan-Cox-Semtner formulation of the Boussinesq primitive equations. This effort culminated in the release of CCSM2 with POP as the ocean component model in 2002. Over subsequent releases of CCSM and CESM there have been extensive developments in parameterizations and ocean process representation in the CESM version of POP. Much of this effort was conducted in collaboration with LANL, as well as the university-based ocean modeling community and other national labs, and through coordinated projects such as Climate Process Teams.

While advances in representation of sub-grid scale processes in the CESM version of POP have kept the model at the forefront of climate modeling, advances in the representation of resolved scale dynamics, i.e., the dynamical core, have not kept pace with advances in other ocean models. Through much of the decade of the 2000s work was underway at LANL on a hybrid vertical coordinate version of POP, HyPOP, intended to address a number of these issues. However, a production climate model ready version was not achieved.

In 2009, at the winter CESM OMWG meeting, the LANL COSIM group announced that they were discontinuing development of POP and HyPOP and would be focusing their efforts on development of an innovative, unstructured mesh ocean model with the intent that it would be proposed as a replacement for POP in CESM when it had reached an appropriate level of maturity. The model, designated MPAS-O, was based on the same underlying infrastructure being used within the NCAR Mesoscale and Microscale Meteorology Division in the development of the next generation mesoscale atmospheric model – the Model for the Prediction Across Scales (MPAS-A). The proposed ocean model included many promising and desirable attributes that addressed shortcomings in POP. Through the early part of the current decade the LANL-COSIM group documented systematic progress toward achieving a production climate model ready version, and regularly reported on their development progress and plans at CESM meetings. While there has been steady progress toward a skillful global ocean version of MPAS-O, plans to incorporate it as a CESM ocean component were interrupted in 2013 by the establishment of a distinct climate modeling effort, Accelerated Climate Model for Energy (ACME), within DOE. Timely access to code and restrictions imposed by DOE

on collaborative development brought the viability of this path toward a CESM ocean component model into question.

Simultaneously with the development of MPAS-O at LANL, a number of efforts to bring aspects of the POP dynamical core more up to date for CESM2 were attempted, e.g., a new time stepping scheme and a modified vertical coordinate. In part due to the dilution of human resources across the ACME and CESM efforts, and in part due to the significant technical gap between the late-1990s era algorithms underlying the POP code and the desired features, progress on these efforts faltered.

The growing obsolescence of the POP dynamical core is inhibiting progress in several areas of high priority CESM development including representation of near surface processes, sea level rise, biogeochemistry, and interfaces to the cryosphere and terrestrial hydrologic cycle. In response, over the past year the CGD Advisory Panel (CAP) and the CESM Advisory Board (CAB) have recommended a reevaluation of options for the future CESM ocean component model code. The CESM Science Steering Committee, in turn, appointed an Ad Hoc Advisory Panel to carry out an investigation of options for the ocean dynamical core for CESM3 and beyond. This report is the outcome of that investigation.

Investigative Process

Prior to formation of the panel, information relevant to its deliberations had been collected by the OMWG. This included the CGD Advisory Panel report from September 2015, input from the community on requirements for the future ocean component obtained through a CESM-wide e-mail survey conducted in February 2016 and extensive formal and informal discussions at the winter 2016 OMWG meeting. While not directly focused on the topic of a future ocean component model for CESM, a Town Hall on community ocean modeling organized by NSF and ONR at the 2016 Ocean Sciences Meeting in New Orleans provided an opportunity to assess the current state of ocean model development activity and the perspectives of most of the development groups considered in our investigation. An important outcome of this initial round of discussions was the decision that developing a new ocean model from scratch within CESM would be cost prohibitive and result in unacceptable delays in advancing the ocean component capabilities to meet the scientific priorities of the project. Instead, it was decided to seek a partnership with an existing ocean model development activity. At the same time, the CAP and CAB strongly advised against taking an existing model “off of the shelf”. A collaborative arrangement, in which both CESM and the partner group have an intellectual stake in the continuing development of the dynamical core should be sought.

Based on this background information, the OMWG co-chairs and the CESM Chief Scientist prepared a Request for Information (RFI) that was sent to six major ocean model development efforts:

- Hybrid Coordinate Ocean Model (HYCOM)
- MIT General Circulation Model (MITgcm)
- Modular Ocean Model version 6 (MOM6)

- Ocean Model for Prediction Across Scales (MPAS-O)
- Nucleus for European Modelling of the Ocean (NEMO)
- Regional Ocean Modeling System (ROMS)

Written responses to the RFI were received from all groups and shared with panel members as they arrived. Notably, responses were obtained from multiple sub-groups within the MITgcm and ROMS communities. All responses were posted on the public Ocean Model Working Group Web Page (Appendix 2) prior to the Breckenridge workshop to give the CESM community a chance to discuss them.

The panel met via video-conference prior to the 2016 Breckenridge workshop to discuss the written responses and conduct an initial assessment of the pros and cons of each option based on the material available at that time. A briefing on this initial assessment was provided to the joint meeting of the CESM Science Steering Committee and Working Group Co-chairs at the Breckenridge meeting. The panel met again in person (with members not able to attend joining by video-conference) in Breckenridge to discuss additional input received from the development groups, identify points of concern or gaps in the available information.

An important event during the investigative process was the cross-working group session on the future ocean component held at the Breckenridge workshop. All members of the panel participated either in person or remotely. The session was very well attended with audience members representing a cross-section of CESM working groups. Representatives from each of the model development groups listed above gave presentations on the key characteristics of their models, the organization and culture of their development communities, and the potential for their respective models serving as a component of CESM. The presentations are available on the meeting website (Appendix 3). A lively discussion followed the presentations, with questions posed to individual model representatives as well as broader questions and feedback to the panel regarding the implications of the choice of model on CESM research in paleoclimate, biogeochemistry, sea-ice, and data assimilation, and on software engineering practices.

Following the Breckenridge workshop the panel completed its investigation through e-mail communication, and solicited additional information from select model development groups to clarify some remaining issues.

Summary of Options

Each of the models considered has features and capabilities that address shortcomings that are inhibiting progress with the POP model. There has been considerable cross-fertilization among the represented modeling groups so many algorithms and capabilities are shared. However, unsurprisingly, no one model meets every technical requirement listed in the RFI. The nature of governance and collaboration within the existing development communities and visions for collaborating with CESM vary widely. Here we summarize the key distinguishing features of each option.

HYCOM

The HYCOM model was the first ocean model to use the Arbitrary Lagrangian-Eulerian (ALE) vertical coordinate, and grew out of the MICOM model that pioneered the use of isopycnal coordinates. Its strongest user community is in the US operational ocean forecasting centers (Navy and NOAA) where it is typically run at much higher resolution than would be the case for standard CESM configurations. The operational global HYCOM system at NAVOCEANO uses a 3DVar assimilation methodology, and HYCOM can interface to any data assimilation system that uses direct or incremental insertion. HYCOM has been coupled to CESM and run in non-eddy-resolving configurations. It has a modest-sized community of university users with applications in regional simulations and process-studies. There has not been much recent development of the dynamical core aspects of the model. It would require considerable development work in the implementation of the eddy mixing parameterization, and in implementing natural boundary conditions at a minimum before beginning to build a CESM configuration. The model documentation is rather fragmented and disorganized. There is a very small, centralized team responsible for maintaining the released code.

MITgcm

The MITgcm was developed with a strong focus on applications in ocean state-estimation. A companion adjoint code is available. The MITgcm is the base code for the ECCO 4DVar assimilation system run in both the US and Germany, and it has been interfaced with the NCAR DART ensemble Kalman filter assimilation system. It has an extensive university-based user community with many users doing process-based and idealized geophysical fluid dynamics investigations. The availability of a non-hydrostatic dynamical core option is an important asset to the process modeling community. A number of the algorithms that have spread through the level-coordinate ocean modeling community over the last decade (partial bottom cells, alternative time-dependent vertical coordinates, non-Boussinesq capability, etc) originated in the MITgcm development community. In addition to global ocean reanalyses, it has been used as a forward model in global configurations at a range of resolutions. It has not been used extensively in coupled climate simulations. It has not been coupled to CESM at this point. The documentation is fairly comprehensive. At the present time, the development community is rather dispersed and self-governing, with modest organizational support at the home institution.

MOM6

While this model has a venerable name, the code is a very recent, nearly from scratch, effort. It builds on the experience gained with the legacy models of its development team (MOM, GOLD, MITgcm) with many innovations that address the requirements for an ocean component of future Earth System Models – wetting and drying, efficient advection and time-stepping for high tracer counts, strict attention to conservation, etc. It uses the CVMix vertical mixing module and shares a number of other parameterizations with CESM that have been developed within CPT projects. There will still be

considerable effort required to adapt specific formulations of sub-grid scale mixing for CESM configurations. It has not been coupled to CESM at this point, but coupled integrations with the GFDL system have been carried out. The FMS coupler infrastructure used at GFDL is ESMF compliant, so it is not expected that there would be any insurmountable challenges in bringing the model into the CESM framework. The development trunk code is publically available. The documentation is currently in progress, but that for the predecessor models is thorough and extensive – the best available. It is reasonable to expect the same for MOM6. The user base of MOM6 is limited at this time, but can be expected to grow considerably as users of the predecessor codes migrate to it. The host institution, from the scientists to the top management levels, has communicated very strong support for collaboration with CESM in model development.

MPAS-O

This model is one of a very small number of global ocean models using unstructured horizontal grids. Due to the complexity of the geometry of ocean boundaries and strong inhomogeneity of the ocean mean flow and variability (western boundary currents, coastal and equatorial upwelling) this feature is very attractive. At the same time it will require significant reworking of pre- and post-processing tools, and does incur a performance penalty relative to structured grid codes. It has been coupled to both CESM and ACME using versions of the CIME interfaces. Coupled integrations are underway within ACME. It uses the CVMix and MARBL modules for diabatic mixing and biogeochemistry respectively, providing a leg up in configuring versions suitable for CESM. Exploiting regional refinement fully will require considerable research in the area of scale-aware parameterizations, especially for mesoscale eddy stirring and mixing. MPAS-O has been interfaced with DART. The current user and developer base is very small, centered at LANL with a small number of external collaborators. The development group proposes to establish a consortium that would oversee the governance of the collaboration and rights and responsibilities with respect to code access. The exact terms of such an agreement would need to be negotiated with DOE management. The use of the unstructured mesh would require access to the accompanying unstructured mesh sea-ice model. Establishing terms of access to that code would also need to be negotiated through a separately organized consortium.

NEMO

The NEMO model is used by most of the major climate and ocean modeling centers in Europe with a very large user community. It has a long history of successful applications in both the eddy-parameterized and eddy-resolving regimes. A number of widely used innovations (tripole grids, two-way nesting, natural boundary condition formulations, fast barotropic solvers) were developed in this model. There is currently no option for an ALE type vertical coordinate, and the recently developed “z-tilda” coordinate is not fully supported at this time. NEMO includes linear-tangent and adjoint models for ocean dynamics that are used operationally at the Met Office and ECMWF. Other centers (MERCATOR and INGV) are using NEMO with an ensemble Kalman filter approach to

data assimilation. It has been coupled to CESM, replacing POP as the ocean component of the CMCC coupled climate system model. It has thorough documentation. A unique aspect of the development effort is an established formal consortium with members committing manpower resources to community development priorities. Though not all centers using NEMO are consortium members (notably ECMWF is not), CESM was encouraged to consider full membership in order to have a say in development directions.

ROMS

The ROMS model is the most widely used regional ocean model in the world with perhaps the largest user base of any of the models considered here. It is formulated with particular attention to accurate numerics and suitability to highly turbulent flows. It uses some of the same nesting infrastructure as NEMO. There is a strong data assimilation activity within the ROMS community including active development of an interface to DART. While it has been coupled to atmospheric models through the CESM infrastructure for regional applications, ROMS has not been used as a component of a global climate model. There would be a considerable amount of work in adapting the model to eddy-parameterized global configurations. The ROMS development community is the most fragmented of all the options considered here. They self-identify as a “federation” with distinct development communities that collaborate but maintain separate code bases. The level of public code access and documentation is variable across the federation.

Recommendations

While each of the options considered could serve as the basis of the provisional ocean component model for CESM3, the panel narrowed its discussion to a choice between the two most recently developed, and in our opinion, innovative of the models: MOM6 and MPAS-O. The more established global models (HYCOM, MITgcm, NEMO) are somewhat lower risk options, but each lags the more recent models in algorithmic sophistication to varying degrees. ROMS is a higher risk option as it is unproven in the global eddy-parameterized regime.

The MOM6 and MPAS-O models are the most forward looking in terms of addressing issues like time-varying boundaries as required for interfacing to ice-sheets, and in terms of flexibility of vertical coordinate choices. The relative immaturity of the models does present some risk in that there is little direct experience with their performance in standard climate simulations – neither have formally participated in the Coordinated Ocean Reference Experiments (CORE) or CMIP. However, both the MOM6 and MPAS-O developers provided the panel with example diagnostics from realistic forced and coupled simulations that provide confidence that high quality global climate simulations can be obtained with either.

It is the panels view that the unique capabilities of the MPAS-O unstructured mesh make it a very attractive option for focused regional downscaling and process studies, but that these advantages are outweighed by several risk factors. At a technical level, the vast

majority of CESM users who are most interested in global applications may not exploit the additional capabilities, but would pay the performance penalty of the unstructured mesh formulation even for uniform resolution configurations. A stronger concern is the uncertainty associated with terms, rights, and responsibilities of the consortium required as a prerequisite to a collaboration with DOE. This organizational hurdle stands in stark contrast to the openness and position taken by GFDL staff and management. Thus, while we cannot recommend MPAS-O as the best option for the provisional CESM3 ocean model at this time, we strongly encourage NCAR and DOE management to formulate a consortium or other mutually agreeable governance structure that will allow the two decades of collaboration between LANL-COSIM and CESM to thrive. The MOM6 model is world-leading in virtually every other algorithmic aspect and will provide CESM with an excellent base for improving the ocean component through many generations. The GFDL development team have emphasized that they do not expect, and in fact discourage, CESM from adopting their configuration “off of the shelf”. We concur that maintaining a diversity of formulations of parameterizations and even some dynamical core choices is important to the health of the US climate modeling enterprise. We therefore recommend:

- 1. CESM should adopt MOM6 as the provisional ocean component of CESM3 and begin working with GFDL and the CESM developer and user communities as soon as possible to design and develop a configuration(s) that meets the requirements of the full scientific scope of CESM.**
- 2. CESM should work with LANL and DOE management to find a mutually acceptable organizational structure to facilitate continued collaboration on ocean model development and permit the use of MPAS-O within the CESM framework for focused research projects.**

There is a strong desire in the climate and ocean modeling communities to carry out climate research with CESM using alternative ocean components. This is evident from the fact that four of the six models considered here have already been coupled to CESM or are using the CESM coupling infrastructure. In each case however, the work to couple a new ocean component was conducted independently, with little or no coordination with parallel efforts with other ocean models. Some of the efforts have struggled to keep current with the evolving CESM infrastructure. In the interest of broadening the utility of CESM across the climate and oceanographic research communities, in reducing duplication of effort, and providing a basis for probing simulation fidelity and sensitivity to model formulation, we encourage CESM to work with interested parties to standardize this process. We also note that this is not the first time, and may not be the last time, CESM will choose to replace or significantly revise its ocean dynamical core. Having access to a base of experimental evidence of how a variety of ocean models perform with CESM in fully coupled configurations will contribute to a scientifically based decision process when these choices are made in the future. We therefore recommend:

- 3. CESM should invest in developing, documenting, and supporting an “open architecture” for interfacing other ocean dynamical cores (e.g.**

HYCOM, MITgcm, MPAS-O, NEMO, ROMS) with the CESM framework and provide technical assistance and coordination to university based research groups interested in using alternative ocean components in CESM.

Throughout the process of gathering information on requirements for future CESM ocean models and in the process of evaluating these options, there has been consistently strong interest in having regional climate downscaling capabilities as part of CESM. This interest has come from both the global modeling community looking to downscale global simulations and the regional ocean modeling community interested in applying their local models in a variety of climate scenarios accessible through CESM. The MPAS-O approach of regional mesh refinement provides one avenue towards this objective, but the overwhelming majority of users interested in this class of application are already working with local models (most often ROMS) that they want to nest within CESM. The capability to do so has been developed in a limited way for specific projects, but has not been supported or made broadly available by CESM. Given the diversity of regional models and scale ranges of interest, it is unlikely CESM itself would or should directly support a particular regional configuration. However, it would be extremely useful to work with the regional modeling community to establish protocols and configurations for a standard interface on the global model side against which regional modelers could develop specific applications. We therefore recommend:

- 4. CESM should engage the regional ocean modeling community to identify requirements for downscaling CESM global simulations, and develop protocols and requisite infrastructure to support the growth of this research area within the CESM community.**

Acknowledgements

The panel would like to sincerely thank representatives from all of the modeling groups that contributed material, advice, and thoughtful discussion throughout this decision process. It is extremely gratifying to work with a community of scientists with such a strong spirit of collaboration.

Appendixes

1. RFI

https://www2.cesm.ucar.edu/sites/default/files/working/groups/omwg/files/CESM3_Ocean_RFI.pdf

2. RFI Responses

- HYCOM <https://www2.cesm.ucar.edu/sites/default/files/working-groups/omwg/files/HYCOM.pdf>
- MITgcm (1) https://www2.cesm.ucar.edu/sites/default/files/working-groups/omwg/files/MITgcm_1.pdf
- MITgcm (2) https://www2.cesm.ucar.edu/sites/default/files/working-groups/omwg/files/MITgcm_2.pdf
- MOM6 <https://www2.cesm.ucar.edu/sites/default/files/working-groups/omwg/files/MOM6.pdf>
- MPAS-O <https://www2.cesm.ucar.edu/sites/default/files/working-groups/omwg/files/MPAS-O.pdf>
- NEMO <https://www2.cesm.ucar.edu/sites/default/files/working-groups/omwg/files/NEMO.pdf>
- ROMS (1) <https://www2.cesm.ucar.edu/sites/default/files/working-groups/omwg/files/NEMO.pdf>
- ROMS (2a) https://www2.cesm.ucar.edu/sites/default/files/working-groups/omwg/files/MYROMS.org_RFI_response.pdf
- ROMS (2b) https://www2.cesm.ucar.edu/sites/default/files/working-groups/omwg/files/MYROMS.org_Technical_Summary.pdf

•

3. Breckenridge Workshop Cross-Working Group Presentations

- Frank Bryan (Introduction)
<http://www.cesm.ucar.edu/events/ws.2016/presentations/cross/bryan.pdf>
- Eric Chassignet (HYCOM)
<http://www.cesm.ucar.edu/events/ws.2016/presentations/cross/chassignet.pdf>
- Jean-Michel Campin (MITgcm)
<http://www.cesm.ucar.edu/events/ws.2016/presentations/cross/mit.pdf>
- Alistair Adcroft (MOM6)
<http://www.cesm.ucar.edu/events/ws.2016/presentations/cross/adcroft.pdf>
- Mark Peterson (MPAS-O)
<http://www.cesm.ucar.edu/events/ws.2016/presentations/cross/peterson.pdf>
- Antonio Navarra (NEMO)
<http://www.cesm.ucar.edu/events/ws.2016/presentations/cross/navarra.pdf>
- Alexander Shchepetkin (ROMS)
<http://www.cesm.ucar.edu/events/ws.2016/presentations/cross/shchepetkin.pdf>