OASIS3-MCT development plan
Sophie Valcke, Anthony Craig, Laure Coquart
November 2017

Introduction

This document proposes a draft for a plan of developments for OASIS3-MCT for the next 3 to 4 years and also introduces some elements to start a discussion on a longer-term vision for OASIS. It is based on the experience and analysis of the current OASIS developers and on a user survey sent to the OASIS User Group. It also builds on the IS-ENES2 reflection regarding the coupler governance and funding presented in the deliverable D4.3 "Coupler Governance model"[1], which the reader is strongly invited to read.

In the survey[2], users were asked to describe and prioritize the developments (improvements or new functionalities) they would like to see in OASIS3-MCT. 14 replies were received, 9 by official User Group members (from LOCEAN, Météo-France, LSCE, Cerfacs in France, SMHI in Sweden, Met Office in the UK, MPI-M in Germany, BSC in Spain, A. Sinica in Taiwan) and 5 by other users (from IPSL in France, DMI in Denmark, SMHI in Sweden, Fujitsu Lab EU in the UK, and HZG in Germany).

There were about 50 unique tasks identified and they were classified into 7 categories s1-Interpolation & transformations, s2-Communication, s3-Configuration, s4-Other functionalities, s5-Support & training, s6-Bugs. Development efforts already planned by OASIS developers were merged into these categories. A few requests from the survey were rated as not relevant when the feature was already implemented (s7-Not relevant). The developments were further classified by priority (1-high, 2-medium, 3-low, 4-very low) and by complexity, directly related to the manpower needed to achieve it, on a 1 to 4 scale, C1 being immediate and C4 being very complex. The results are available in the file 201711_OASIS3-MCT_developments_UserSurvey.xlsx .

Here below, we first present the general context into which OASIS developments should be considered with an analysis of the funding situation and of the longer-term possible evolution. Then we detail the possible developments for OASIS3-MCT classified by priority.

---

[1] https://portal.enes.org/oasis/ISENES2/documents/deliverables/is-enes2_d4-3_coupler-governance-model-document

[2] The results of the user survey are available at https://portal.enes.org/oasis/users/oasis-governance .

1. Context, longer term and funding

OASIS development started in 1991 in Cerfacs. The current version of the coupler, OASIS3-MCT (Craig et al. 2017), offers fully parallel regridding and distributed exchanges of the coupling fields, thanks to MCT (Model Coupling Toolkit, www.mcs.anl.gov/mct) developed at Argonne National Laboratory in the USA. Today, about 45 climate-modelling groups around the world use the coupler. Since its first release, more than 700 OASIS3-MCT downloads were registered from groups in Europe but also in Canada, USA, Colombia, India, Japan, China, Saudi Arabia, etc. The latest official release of the coupler, OASIS3-MCT_3.0, was distributed in August 2015 (Valcke et al 2015) and the next release is planned for February next year (2018). Given the latest performance improvements and the evolutions proposed in this document, we consider that OASIS3-MCT will provide an efficient and easy-to-use coupling solution for many climate modelling groups for at least the next 5 years, even if no revolutionary development is planned.

As written above, OASIS3-MCT offers fully parallel regridding and distributed exchanges of the coupling fields but does not perform the parallel calculation of the regridding weights and addresses per se. Today, different tools (like the Earth System Modelling Framework –ESMF, Theurich et al. 2016) can be used off-line to pre-calculate these weights and addresses in parallel. As is described in section 2., high priority should now be given to interfacing one of these libraries in the initialisation phase of the coupler, so that an efficient calculation of the interpolation weight-and-address file can be performed on-line, to simplify the set-up of a coupled configuration. We have, however, a priori no plan to transform OASIS3-MCT into a dynamic coupler, i.e. allowing the recalculation of the weights at the coupling frequency during the simulation, which is a feature needed to support models with grids evolving in time (adaptive mesh), such as ice-sheet models. This development would represent a step change for OASIS and we estimate that we do not have the resources to address this evolution. Furthermore, other libraries, such as the OpenPALM (Duchaine et al. 2013) coupler and the XIOS I/O server, offer this feature. OpenPALM is a dynamic coupler used in diverse multi-physics applications in computational fluid dynamics but also in several geophysical applications (hydrology, hydraulics, atmospheric chemistry, etc.). XIOS, the I/O server developed at IPSL, is used in many European climate components and already includes many coupling functionality such as communication of data (between a first and a second level of servers) and regridding of the data (to store the output data on a grid other than the model grid). Indeed it has been recognized for many years that coupling and I/O functionalities are very similar. Therefore, instead of "reinventing the wheel", we propose that converging with one of these libraries should be considered, especially in the longer-term (5-10 years) context. Working toward the convergence of XIOS and OASIS3-MCT was proposed by IPSL in the survey (s4 l19).

We stress here that for the 2018-2021 period the funding of these developments will come from CNRS (3/4 FTE), Cerfacs (1/2 FTE), the ESiWACE CoE (6 pms total), and possibly IS-ENES3 EU project if accepted (the precise amount is still to be defined). The commitment of CNRS and Cerfacs can be considered stable: in practice, it ensures active user support to the community and the maintenance of the actual code with the possibility of minor upgrades. Additional developments will require additional funding and the only source of such additional funding seems to be EU research infrastructure

projects (like IS-ENES3) or French national projects (like a CONVERGENCE follow-up). Other funding models were analysed and proposed to the community but none was considered practical (see the "Coupler Governance model " document"[1]). Such efforts have been successful for the past 20 years and, even if there is obviously no guarantee for the future, Cerfacs commits to keep on seeking such additional EU or national funding.

2. Developments with high priority (P1)

The developments that we propose to consider with high priority are described here. For each development, we provide an estimate of the complexity, C1 - being immediate up to C4 -being very complex, and its location in the 201711_OASIS3-MCT_developments_UserSurvey.xlsx file; for example "(C4) [s1 l12]" means that it is a very complex development and that more details can be found on line 12 of the first sheet in the excel file. When a user proposed the development in the survey, we also mention the user's institute and, when a development is already described in a ticket on OASIS Redmine site (https://inle.cerfacs.fr/projects/oasis3-mct), we provide the ticket number.

s1 - Interpolations & transformations

- Interfacing with a parallel library for the calculation of the interpolation weight-and-address file (C4) [s1 l12 ticket #1011, s1 l13 ticket #1069]

  We consider that this should be of high priority even if no user directly asked for it. Currently, the coupler is interfaced with the SCRIP library, which is sequential and suffers from many quality problems, especially for the conservative remapping. Parallel and higher-quality interpolation routines from ESMF or XIOS could be considered.

- OpenMP parallelisation of the SCRIP routines (C2) [s1 l2, Cerfacs]

  In the short term however, users could benefit from the OpenMP parallelisation of the SCRIP routines. This has been already implemented and positively evaluated for the "N-nearest-neighbour" interpolation. It could be easily implemented for the "N-nearest-neighbour weighted by a Gaussian function", bilinear and bicubic interpolations. The implementation for the conservative remapping would be more complex and should be considered with care as investing the same effort in the interfacing with another already parallel library (see above) could be more appropriate.

- Other issues requiring only minimal efforts (C1):
  - Adding the number of neighbours used in the remapping file name so to avoid confusion when the same interpolation is used between the same grids but with different number of neighbours [s1 l9, SMHI]
  - Easy detection of the target grid point that do not receive any value [s1 l11, ticket #731]
  - Better display of the diagnostics in the log file [s1 l8, IPSL, ticket #1761]

s2 - Communication

- Coupling initialisation and exchange performance (C2) [s2 l5, MPI-M]

It is of course crucial to ensure the performance of the coupling initialisation and exchanges. Recent developments (i.e. mapping the target decomposition on the source processes, or vice-versa, based on the set of remapping weights and addresses, see s2 l7, ticket #1547) have significantly improved the communication scheme. The coupling initialisation has also been significantly optimised thanks to some debugging (ticket #2039) and to the update of MCT (s2 l6, ticket #1722). We now need to better quantify and stabilise these latest developments.

- Sending/receiving simple scalars (C1)

  Two users asked for this functionality that should be implemented (s2 l3, SMHI; s2 l4, MetO).

- Control of the time evolution on a per-field basis (C1) (s2 l1, IPSL]

  OASIS3-MCT implements a check to ensure that the model is not going backward in time. This has to be relaxed or implemented on a per-field basis. Indeed, when two components are running sequentially within one executable, one component may perform different oasis_puts over a certain period before the other one starts its oasis_puts for the same period.

s4 - Other functionalities

- Further development of LUCIA, the load balance analysis tool (C2)

  LUCIA is a tool delivered with OASIS3-MCT to automatically take time measurements during the coupled run and evaluate the load balance between the coupled components. Two users asked for further improvements of LUCIA (SMHI s4 l2, BSC, s4 l3). Currently the LUCIA analysis is strictly valid for standard coupling. Taking into account more specific cases like the fact that one coupling time step may be longer than another due to some higher workload should be included. Fully integrating LUCIA analysis tool in the OASIS3-MCT finalisation step (instead of having to run if afterwards) and unifying LUCIA and the standard OASIS timer result presentation was already planned (s4 l4, ticket #1003).

- Allowing to launch a coupled model without coupling restart files (C1)

  This point was raised by Météo-France (s4 l9). We think it is useful to offer this possibility for users in development mode. In that case, a field with uniform value can be provided as coupling field. This is partly implemented already but used by OASIS developers only.

- Analyse and optimise the memory use in OASIS3-MCT (C2)

  This important task is already planned within the ESiWACE framework (s4 l14, ticket #1104)

- Develop, maintain, and systematically apply the buildbot test suite (C3)

  The builbot suite is an automatic testing system composed of more than 20 toy models that are run automatically over four different platforms each time a check-in is done in the SVN source manager. The results are automatically analysed and compared to a reference state to detect any bug that a new development would introduce. We consider with high priority the maintenance and further enhancement of this test suite.

s5 - Support and training

It is also a high priority for us to keep on offering high quality user support and training. In particular:

- Regular user support through the forum, by mail or by phone will continue (C2) (s5 l3).
- Face-to-face training at Cerfacs will continue to be offered based on demand (C2) (s5 l4).
- In IS-ENES3, we propose to revive the Dedicated User Support, i.e. organising the visit of an OASIS developer to a chosen institute to help setting up a coupled system, as was offered in IS-ENES1 (2-3 institutes x 1 month per year) (C2) (s5 l5).
- Toy coupled models are delivered with the coupler sources so that users can get familiar with compiling and running with OASIS3_MCT and/or test offline the quality of their interpolations in simple environments; these toy coupled environments should be revamped and updated (ticket #1801) (C2) (s5 l2).

s6 - Current bugs

Users reported few bugs either through the survey or by mail. These should be addressed with high-priority.

- Deadlock with orange partitioning with non contiguous sub partitions (C2) (Fujitsu, s6 l2)
- Arbitrary match between redundant decompositions (C2) (s6 l3, ticket #1937)
- Problem in m_MCTWorld.F90 in the routine initm (C2) (s6 l4, ticket #1321)
- Multiplication and addition by a scalar (C1) (s6 l5, ticket #756)

3. Developments with medium priority (P2)

s1 - Interpolations & transformations

- Diagnostics

  One user asked for more diverse diagnostics (e.g. for intensive variables or for fields with fractional masks) (C2) (IPSL, s1 l7, ticket #1069)

- Pre- and post-processing transformations

  Two users asked for extension of the pre- and post processing transformations (BLASOLD/BLASNEW), e.g. combination with other fields coming from disk files or user-defined transformations (C2) [A. Sinica & BSC, s1 l5, s1 l6]

- Normalisation by the “true” area for the conservative remapping

  To be truly conservative the remapping weights should be normalised by the “true” area of the grid cells i.e. the area of the cells as considered by the models themselves and not as calculated by the remapping library (C2) (s1 l10, ticket #1010)

s3 - Configuration

- Layer for “automatic” configuration of the coupling exchanges in the models based on the namcouple

  Two users asked for that functionality or something similar). The idea is to define, in the namcouple, a unique identifier for each coupling field. The identifier would have to be associated to a coupling field array in the model but, besides this, all coupling action would be automatic. For example, there would no longer be the need to hard code the symbolic name of a coupling field in the models as the layer would automatically call the coupling field declaration routine (oasis_def_var) with the symbolic name read from the namcouple for its coupling field identifiers (C3) (MetO, s3 l3; SMHI, s3 l4).

- Tool to perform offline a complete and explicit control of the namcouple

  This tool would detect for example when configuring line is not complete or when an argument is missing. We consider that this tool could be useful but OASIS3-MCT itself performs a quite extensive check of the namcouple with precise error messages in the initialisation phase. Instead of developing a separate tool, the OASIS namcouple error handling functionality could also be extended. (C2) (Météo-France, s3 l5)

s4 - Other functionalities

- Modernise OASIS code and transform F77 routines (C3) (SMHI, s4 l8)
- Debug files (EXPOUT coupling fields):

  When written out for debug, each bundle member appears in a different file. MetO would find it more practical to have all bundle members into one file (C2) (s4 l6). IPSL would like to have the field grid written out in debug file for easier plotting (C2) (s4 l7).

- Systematic tests of NetCDF returned error code:

  This was asked some time ago by IPSL and would facilitate the debugging (C2) (s4 l11, ticket #550).

- Support OpenMP models

  This was not directly asked for but we know users who couple OpenMP models with OASIS3-MCT. In a first step, only models with "full" OpenMP parallelisation (i.e. the threads are active from start to end) will be considered. Each thread will implement the OASIS3-MCT API but below the put/get API routines, the coupling field will be gathered in a shared buffer and effectively managed/sent/received only by the master thread. The development of a prototype is planned in ESiWACE (C3) (s4 l12 ticket #1223)

- Develop standard components for standard workflow tasks (C2)

  The idea is to develop (or collect existing ones in the community) and distribute "standard" components that would be coupled by OASIS3-MCT to the coupled model for basic tasks in the workflow such as data readers and writers (C2) (SMHI, s4 l15)

4. Developments with low priority (P3)

s1 - Interpolations & transformations

- Vertical interpolation

  Simple level-to-level or pressure-to-pressure interpolation would be relatively simple to implement but more sophisticated pressure-to-level or hybrid-to-level would be much more complex (because the interpolation then depends on auxiliary field that changes at each coupling timestep). We consider that this is better done in the model themselves and give only a low priority to that task. (C4) (Cerfacs, s1 l4)

s3 - Configuration

- Modernisation and extension of namcouple syntax

  The priority of this task is debatable. We have considered transforming the namcouple format (xml, standard FORTRAN namelist). Finally we have always come back to the conclusion that the current ASCII format, even if not extremely friendly, fulfils its mandate and changing it, while keeping the same quality of error handling, would require a non negligible effort for low benefits (C3) (SMHI, s3 l2)

s4 - Other functionalities

- Python bindings for the OASIS3-MCT library

  This was suggested by R. Ford as a possible development for IS-ENES3. Python bindings would allow the coupling of models written in Python. It would bring OASIS to a much wider ecosystem and be a good way to attract countries less advanced in HPC (C3) (STFC, s4 l10)

- Further OpenMP implementation

  The possibility of going further in the OpenMP multithreading of OASIS3-MCT, i.e. each thread sending and receiving its part of the coupling fields, could be evaluated although we are not convinced that there is any demand for this (C4) (s4 l13, ticket #1223)

- Interoperability with ESMF

  This is an interesting proposition but we think it should be better addressed by asking for dedicated support. An OASIS developer could visit the group to better understand the technical implications and develop a first prototype. It would then be easier to evaluate the technical implication of a generalisation of this feature (C2) (HZG, s4 l16).

5. Developments with very low priority (P4)

s1 - Interpolations & transformations

- Dynamic weight calculation

  It is currently out of our scope to transform OASIS3-MCT into a dynamic coupler, i.e. supporting grids evolving in time (see section 1.). Even the user mentioning this functionality writes “there seems little point is requesting this until we can be more precise about requirements” (C4) (MetO, s1 l3)

s4 - Other functionalities

- Saving of coupling restart files in the middle of a coupling period

Currently the write_restart argument of the oasis_put routine allows the coupling fields to be written in intermediate coupling restart files only at coupling time steps. This would not be straightforward to support, especially as writing of coupling restart take care of time transformation over the coupling period such as accumulation or averaging. Further discussion with the user convinced us that there might be no real user case for this (C3) (BSC, s4 l5)

- OASIS as a command line interpolator (C3)

  The "test_interpolation" environment, that can be used offline to test offline the interpolation quality, will be revised to be more user friendly but we think it is not an objective of OASIS to be usable as a command line interpolator. Other tools exist (CDO, ESMF, NCL).

Conclusions

This document proposes a first development plan for OASIS3-MCT, addressing more or less precisely the short, medium and longer terms. Based on the user survey and on the developer experience, we have assigned a high (section 2), medium (section 3), low (section 4) or very low (section 5) priority to the possible developments.

For the short and medium term (3-4 years), with limited funding confirmed today (i.e. 0.75 FTE from CNRS, 0.5 FTE by Cerfacs, 6 pms by ESiWACE), focus will be on highest-priority and lowest-cost tasks. If IS-ENES3 is funded, which we certainly hope will be the case, we will most likely be able to address all developments with high priority and several with medium priority as well as most of the easier tasks. At that time, priorities could be revised, but it is unlikely that the more expensive developments with low or very low priority will be considered unless their priority is upgraded.

Some thought on the longer-term evolution are also proposed in the first section of the document.

We are now looking forward receiving feedback from the Advisory Board on our propositions and long-term vision.

References

Craig, A., Valcke, S. Coquart, L. 2017: Development and performance of a new version of the OASIS coupler, OASIS3-MCT_3.0, Geosci. Model Dev., 10, 3297-3308, https://doi.org/10.5194/gmd-10-3297-2017, 2017.

Valcke, S. Craig T. and Coquart, L. 2015. OASIS3-MCT User Guide, OASIS3-MCT_3.0, Technical Report TR/CMGC/15/38, Cerfacs, France.

Theurich, G., Deluca, C., Campbell, T., Liu, F., Saint, K., Verten- stein, M., Chen, J., Oehmke, R., Doyle, J., Whitcomb, T., Wall- craft, A., Iredell, M., Black, T., Da Silva, A. M., Clune, T., Fer- raro, R., Li, P., Kelley, M., Aleinov, I., Balaji, V., Zadeh, N., Jacob, R., Kirtman, B., Giraldo, F., McCarren, D., Sandgathe, S., Peckham, S., and Dunlap IV, R.: The Earth System Prediction Suite: Toward a Coordinated U.S. Modeling Capability, B. Am. Meteor. Soc., 97, 1229–1247, https://doi.org/10.1175/BAMS-D- 14-00164.1, 2016.