EUROPEAN COMMISSION

Community research

is-enes
INFRASTRUCTURE FOR THE EUROPEAN NETWORK FOR EARTH SYSTEM MODELLING

IS-ENES2 MILESTONE (M -N°: 4.3)

Draft governance model document for the community coupler

CERFACS Working Notes, WN/CMGC/14/105

File name: IS-ENES2_M4.3_Coupler_Governance_Draft.docx or .pdf

Reporting period: 01/04/2013 – 30/09/2014

Author(s): Sophie Valcke

Release date: **15/11/2014**

Abstract

After a historical overview of the OASIS coupler development and some elements on its possible future evolution, this document presents the current status of the community governance process for the coupler and how to develop it. A "straw man" governance structure including a stakeholder group, a technical advisory group and a user group is then proposed as a basis for further engagement. Some questions we propose to ask shortly to the OASIS user institutions to better understand their position on how they envisage, or not, a possible governance process around a community coupling software are then formulated.

Project co-funded by the European Commission's Seventh Framework Programme (FP7; 2007-2013) under the grant agreement n°312979		
Dissemination Level		
PU	Public	X
PP	Restricted to other programme participants including the Commission Services	
RE	Restricted to a group specified by the partners of the **IS-ENES2** project	
CO	Confidential, only for partners of the IS-**ENES2** project	

Table of contents

Executive Summary

Establishing a community governance strategy for the coupler is needed on the long term to address the foreseen jump in complexity of the coupling problem on exascale platforms. After a historical overview of the OASIS development and some elements on its possible future evolution, this document presents the current status of the governance process and how to develop it. Based on first discussions organized in June 2014 during the IS-ENES2 General Assembly and subsequent interactions, we propose a "straw man" governance structure as a basis for further engagement. This governance could include a stakeholder group, a technical advisory group and a user group. However, before pushing forward such a structure, we think that the OASIS user institutions should be more formally consulted on how they envisage, or not, a possible governance process around a community coupling software and we formulate some questions we propose to ask them in the coming months to better understand their position on this issue.

1. Objectives

At the time the IS-ENES2 project was written, the short-, mid- and long-term evolution of the OASIS coupler was unclear and it was considered urgent to actively involve its user community in the development and governance of the software. Even if this seems today less urgent thanks to the recent success of OASIS3-MCT, establishing a community governance strategy for the coupler is needed on the longer term to address the foreseen jump in complexity of the coupling problem on exascale platforms. First discussions on this issue organized in June 2014 during the IS-ENES2 General Assembly (see the appendix) and other contacts with OASIS users led to the conclusion that the main OASIS user institutions should be more formally consulted to further input on the need and form of a governance process for OASIS. The document presents the current status of the governance and the process to develop it.

2. Context

The OASIS coupler, developed since 1991 at CERFACS, is used today by more than 40 climate-modelling groups around the world. OASIS3 was, for example, used in 5 of the 7 European Earth System Models (ESMs) that participated to international Coupled Model Inter comparison Project (CMIP5) that provided the scientific basis for the IPCC Fifth Assessment Report (AR5).

OASIS was first developed for 2 different French climate models at a time where ocean-atmosphere coupling was in its infancy. In fact, OASIS development started with a top-down decision taken by two French climate modelling groups, at Météo-France and IPSL/CNRS, that commissioned CERFACS to develop a modular coupling software. One can therefore consider that upper level governance existed at that time at the French national level even if CERFACS was fully responsible for making the relevant technical choices. OASIS became progressively used by more and more groups and proved to be the right tool developed at the right time to offer a concrete answer to specific climate modelling needs. The wide use of the OASIS coupler in climate models was not imposed by any upper authority in a top-down approach but naturally emerged as a bottom-up process. Other important aspects of OASIS success are the great care taken by CERFACS to provide active user support and to constantly integrate community developments in its official version. This is how the user community partially guided the coupler development even if no formal governance existed outside CERFACS. OASIS is therefore a great example of successful community software as the total time invested by the different groups in its development is certainly much less than the time it would have taken for each group to develop its own coupler. Furthermore, using a common coupling tool also facilitates sharing of component models and favours scientific interactions.

Up to now, CERFACS and, since 2007, CNRS (the French Centre National de la Recherche Scientifique), devoting each one person full time to OASIS, have been able to ensure development, maintenance, integration and user support of the coupler, helped by a number of collaborations on specific developments and by temporary but important funding streams (e.g. PRISM, CICLE, METAFOR, IS-ENES1&2 projects). Again, even if only CERFACS and CNRS are ultimately making all the technical decisions, the development of OASIS through

externally-funded projects has in practice provided another mean for a light community engagement in the coupler governance as the work proposed in projects needs to be agreed on with the whole project community.

From 2001 until 2004, the development of OASIS benefited from an important support by the European Commission in the framework of the PRISM project and led to the first release of OASIS3 (Valcke et al., 2012), which is still widely used in the climate modelling community. But even during PRISM time it was expected that the limited parallelism of OASIS3 would one day become a bottleneck for massively parallel computations. This is when the development of the fully parallel OASIS4 started in collaboration with NEC Laboratories Europe - IT Research Division (NLE-IT). Objectives at the time were to develop a coupling library that would efficiently perform parallel exchange and regridding of the coupling fields between the source and target codes but also parallel calculation of the source neighbour weights and addresses (also called the "neighbourhood search") needed for the regridding. A first version of OASIS4 was produced (Redler et al. 2010) and was used in few coupled systems, in particular to couple the IFS atmosphere model to different atmospheric chemistry models in the framework of the EU MACC project. However, a deep evaluation of OASIS4 performed during the IS-ENES1 project lead to the conclusion that OASIS4 presented some fundamental weaknesses in its design, in particular the lack of support of unstructured grids. At this point, CERFACS and CNRS decided to stop the development of OASIS4, as it was not practical to adapt the design to this requirement.

It was then urgent to provide to the community a parallel coupling library even if the on-line parallel calculation of the regridding weights and addresses per se was not considered mandatory (i.e. using weights and addresses pre-defined offline met the needs of the community). To fulfil this urgent need, CERFACS and CNRS decided to interface MCT (the Model Coupling Toolkit, www.mcs.anl.gov/mct developed by the Argonne National Laboratory) with OASIS3; this happened at the time the writing of the IS-ENES2 proposal was starting. Version 1 of OASIS3-MCT was released in July 2012 and a second version is now available since May 2013 (Valcke et al. 2013).

Since its first release in July 2012, more than 250 OASIS3-MCT downloads were registered from groups in Europe but also in Canada, USA, Colombia, India, Japan, China, Saudi Arabia, etc. Among these groups, many have effectively migrated to OASIS3-MCT, in general for relatively high-resolution version of their coupled model, for example all IPSL labs, ETHZ in Switzerland, SMHI in Sweden, MPI-M and the regional climate modelling community COSMO in Germany, the UK Met Office and the BoM in Australia. The latest official releases of the NEMO ocean model and the atmospheric WRF model now include the OASIS3-MCT API. Regarding user support, three-day training sessions on OASIS3-MCT are offered twice a year, the OASIS web site (http://oasis.enes.org) is continuously maintained, and additional support is provided through the forum and by mail and phone exchanges, as required.

OASIS3-MCT supports coupling of 2D logically-rectangular fields but 3D fields and 1D fields expressed on unstructured grids are also supported using a one-dimension degeneration of the structures. In OASIS3-MCT, MCT implements fully parallel regridding, as a parallel matrix vector multiplication based on pre-computed weights and addresses, and distributed exchanges of the coupling fields. Even if the underlying implementation has evolved significantly, usage of OASIS3-MCT in the codes has largely remained unchanged since the previous OASIS3 version, keeping the same Application Programming Interface (API) based

on flexibility and minimal invasiveness. Tests done with 8000 cores on the Bullx Curie machine at the TGCC near Paris and 16000 cores on the IBM MareNostrum III at BSC in Barcelona show that OASIS3-MCT presents good performance and scalability for up to O(10000) cores. We consider it is therefore very likely that OASIS3-MCT will provide an efficient and easy-to-use coupling solution for many climate modelling groups for at least the 5 years to come. Of course, the software will always evolve to fit particular user needs and even if CERFACS and CNRS manage and make decision for these developments, the users are regularly consulted through surveys to ensure that the community partially govern these choices.

3. Future evolution

As detailed above, OASIS3-MCT does not offer parallel calculation of the regridding weights and addresses. This functionality will most probably not be mandatory for most of our climate models in the short and mid term (~5 years) as groups running at very high resolution can use different tools (like the Earth System Modelling Framework -ESMF) to pre-calculate the weights and addresses in parallel. However, it will be needed once the component models are running on adaptive grids (where grid point locations possibly change at each time step). We must therefore consider today the development of this functionality or the integration of an existing library offering this functionality. More generally, on the longer term, we have to ask ourselves how to provide a coupling library efficient on extreme scale computing platforms.

To answer this question, we have first to evaluate different libraries offering possibly part of the solution. For example, ESMF is already used offline by some groups to calculate in parallel the regridding weights and addresses for any set of source and target grids. The other coupler developed at CERFACS in collaboration with ONERA, OpenPALM, also needs to be considered. The unique feature of OpenPALM is that it supports dynamic coupling, which means that codes to be coupled can be launched at any time during the simulation, in a loop or under a specific condition. With the more recent interfacing of CWIPI, OpenPALM can now perform online parallel calculation of regridding weights and addresses. OpenPALM has proven to be a flexible and powerful dynamic coupler and it is now used by about 50 different groups in France and Europe for different multi-physic couplings in different application domains, such as aeronautics and space, computational fluid dynamics, combustion but also atmospheric chemistry or hydrology. Finally, on a the longer term a possible convergence with XIOS, the I/O server developed at IPSL and used in more and more climate component codes, has to be evaluated. In fact, since PRISM times, it has been recognized that the main function of an I/O server and the one of a coupler, i.e. communication of data, are extremely alike; as interpolation is now also envisaged in I/O servers to store data on other grids than their native grids, the two types of software are getting even closer allowing the projects to share underlying libraries or even to merge with one another.

To ensure a proper evolution of the coupler, we estimate that the effort needed is of the order of 4 FTEs, i.e. 1 for project management, 2 for developments and quality assurance and 1 for user support. Even if additional EU or other projects could cover part of this effort, a wider community commitment, beyond CERFACS and CNRS current one, should therefore be considered.

It is agreed within the ENES community that any such investment should be underpinned by a governance process.

4. Community governance

Before finalising any governance model, we need to evaluate how the climate community would provide its engagement in the development and governance of a next-generation community coupling software. In 2011, it was already proposed to the OASIS community to establish an Executive Board, setting strategic directions and providing technical guidance, but in the absence of any clear funding framework for the required resource, it was not possible to take this further. There is now an opportunity to seek funding for a community-centric investment into the development and support of the OASIS coupler through a Centre of Excellence and, based on discussions at the ENES general assembly and subsequent interactions, a governance structure for OASIS could be as follows:

- A stakeholder group that would oversee the investment of resources into the community coupler. This would be advised and informed by:
- A technical advisory group that would advise on technical issues, including quality assurance and,
- A user group that would collect the input of end users to inform both technical and strategic activities.

If considered too complex, this governance structure could be simplified with the technical group also fulfilling the role of the stakeholder group or with the technical group being also responsible for collecting information locally, removing the need for a separate user group.

To further define the terms and structure of a possible community governance strategy for a common coupling software, we propose here to ask the different groups the following questions.

- **The strategic requirement**. Do you think that sharing common coupling software, like the OASIS coupler used today by a majority of the climate modelling groups in Europe, is a good thing? Do you think the different groups should keep on sharing their coupling software in the mid and long term?
- **The appetite for engagement.** Do you think community governance should be established for the OASIS coupler on the short term and for its evolution or the use of alternative software in the mid and long term?
- **The structure.** If so, what type of governance should be established? Should it include an OASIS user group, providing feedbacks and advices from the users to the developers? Should it include higher-level more strategic groups providing directions for the technical development, quality control, resources and funding of the software, as described above?
- **Funding model.** If so, what should be the funding model of community coupling software? Should one group develop the coupling software with the different partners contributing money to this group? Should the different partners contribute efforts to the software development, this resulting in a geographically distributed software development? Should sustainable funding be sought from some form of research infrastructure?
- **Need for more software sharing before possible governance.** Alternatively, do you think that sharing more software, naturally leading to distributing the software development effort over more groups, is a mandatory condition before trying to establish some community software governance for the coupler?

In parallel with this, we need to establish some funding model for the resource as the governance structure needs to support both the community needs and the funding model simultaneously.

5. Conclusions

One objective of IS-ENES2 is to establish a community governance model for the coupler software. Based on first discussions on this issue organized in June 2014 during the IS-ENES2 General Assembly and subsequent interactions, we propose here a "straw man" governance structure as a basis for further engagement. At this point, we think that the OASIS user institutions should be more formally consulted on how they envisage, or not, a possible governance process around a community coupling software. After a historical overview of the OASIS development and some elements on its possible future evolution, this document presents the current thinking on a governance structure and proposes some questions to ask to the OASIS user institutions to better understand their position on this issue in the coming months. In the next stage, we will need to engage with funding opportunities to underpin the governance of a community coupler. This document represents an essential step in the establishment of a community governance model for the coupler software.

6. References

Redler R., S. Valcke and H. Ritzdorf, 2010: OASIS4, A Coupling Software for Next Generation Earth System Modelling. Geosci. Model Dev., 3, pp. 87-104, doi: 10.5194/gmd-3-87-2010

Valcke S., 2013: The OASIS3 coupler: a European climate modelling community software, Geosci. Model Dev., 6, 373-388, doi: 10.5194/gmd-6-373-2013.

S. Valcke, T. Craig and L. Coquart, 2013: OASIS3-MCT User Guide, OASIS3-MCT_2.0, Technical Report TR/CMGC/13/17, Cerfacs, France.

7. Appendix 1 – Discussion on software governance – IS-ENES2 GA, June 2014

1. Governance definition
   - What do we mean by software governance?
     - RB: learnersdictionary.com: “the way that a city, company, etc., is controlled by the people who run it”
     - RB: Wikipedia: G …” refers to "all processes of governing, whether undertaken by a government, market or network, whether over a family, tribe, formal or informal organization or territory and whether through laws, norms, power or language.”
     - Governance could be as light as community agreement on some principles when we develop software
     - Management is the expression of governance ; agreeing on the level of quality is governance ; how you do this is management
     - For IS-ENES2, governance should be applied to infrastructure software that everybody needs; this software should be identified and resource should be provided to allow for proper governance
     - Discussion on top-down vs bottom-up software development:
       - Identify infrastructure software that everybody needs, and govern that. Wrt develop what you need and see if used by others for their institution
       - Top-down: identify software needed and govern the development of that software
       - Bottom-up: each group develops what it needs, and with some networking and communication between the groups, the good software will naturally emerge
       - A mix is also possible: different developments start in the community and at one point there is a transfer of ownership (some governance) from the developing group to the community via a committee (e.g. MPI, but this was possible because of incentive of the vendors and large community)
       - Going back and forth between two approaches is probably what is needed
       - Standard have to emerge from bottom-up ; but at one point, when the community is mature, than the community should get together and recognize the standard and recommend to use it.
       - If the resources are scarse, top-down should be preferred (more risky but less costly)
       - What should be ENES approach regarding the development of software needed by all: Let it go and see what happens or adopt a top-down approach? (Does the scale of the problem should impact the way we are considering the governance?)
     - To allow proper development of a software, sustained funding is needed; a gate keeper is also needed.
     - Sylvie: 2 discussions: More common software? How do we manage shared software.
   - What does software governance imply?
     - Quality and discipline of software development
     - Community able to input to direction
     - A formal process for decision making on what is done and how
     - A formal agreement - consortia agreements
   - Relation between management and governance?
   - Does software governance help make a software package : a standard? be shared?
   - Does governance come for free? If not, who should pay? How to fund?
     - In what ways do we pay?
     - Do we want community development?
     - Testing, other work
     - Do we want community funding?
     - In what ways should we pay?
     - Do we need financial support for governance?
     - What options do we have and how effective are they? Could they be FP6/7/h2020, ENES, EU research infrastructures, Payment in kind - you do this and we do that, other funding Options?

2. Examples of common software governance
    - What can we learn from effective examples of common software governance (NetCDF, NEMO, OASIS, CDO, Cylc, MCT, Radiation/COSP codes, …)
        - OASIS: started as a top-down initiative (in France) but then spreading and buy-in by the EU climate community was more bottom-up
        - CDO: bottom-up, open source, informed by other developments, freedom of developer, point of failure because too much work,
        - NEMO: two groups were initially using OPA, MetOffice wanted a large community (top-down decision); one ETP to be part of the consortium and steering committee. Steps to add something in NEMO : 1) ask other people if proposed development is useful; 2) ask the system team if the concept is OK; 3) develop; 4) prove that it works fine; 5) provide the code to the System team; 6) the Software team rewrites and includes it
        - GRADs: over because no sustained funding
        - Vis5D: " " " " "
        - PRISM: decision by top-people not informed enough by experts; over ambitious; not corresponding to the needs
    - What are the key factors for using these programmes for "other" centres? Was software governance a factor in selecting the software?
    - What effective governance structures are in place in these examples? Strengths, weaknesses
    - Are the requirements for governance common or all different?
    - Are there different views from providers and users?

3. Software sharing
    - How do we encourage more software sharing?
    - Is more shared software the right thing?
    - If so, what are the barriers?
        - Do not know enough
        - The effort of evaluating
        - Performance or functionality
        - Learning curve
        - Initial cost high
        - Benefit not clear or not big enough (already have something)
        - Risk you cannot solve problems
        - Documentation, training
        - Lack of support
        - Lack of consensus? Do not want to be the first.
        - Software not mature
        - Cannot control what happens
        - Licence issues
    - Do we want standard software?
    - If we think it is a good thing, how can we get what we want? Can we identify the leadership needed?

4. Next steps
    - Which groups in the ENES community are interested in participating in a common software governance?
    - Is ENES the right vehicle for setting up governance structures?
    - How do we make the ENES community strong enough to make a difference here? We could start by getting grass-roots support and evidence that it should be done
    - Missing : way to evolve from netwoking to community endorsement : no appetite within ENES to go further than bottom-up ; it is our role to say that it is needed. Document to the ENES board. More intrusive coupling approach needs to be discussed.

5. First thoughts on OASIS Governance
    - What are the likely aims of governance for OASIS?
        - To increase the chances of take-up via wider involvement;
        - To ensure OASIS developments keep it relevant to future needs;
        - To set priorities;

    - To define quality standards;
    - To provide technical advice and review.
- Would it be appropriate to set up an OASIS user and governance groups? If so, what would be the mandate of the user group?
    - Advisory or setting the direction?
    - Reporting for information to CERFACS.
    - Reporting to a strategic group below.
- Do we want a higher level, more strategic group to consider such things as:
    - Resources (see also 3 below)
    - Technical direction (review of technical plans on all timescales)
    - QC issues
- If we do not want these things, what are the reasons and could those reasons be overcome to make this attractive?
- How can we have governance of OASIS if it does not come with funding – why would the existing funders want to have others tell them what to do to in this way? Are we looking for long term funding for OASIS development and support? At one time this seemed very important but not so much of an issue now, maybe? If we had this, how would things be different?
    - What funding models would work?
    - Partners contributing money?
    - Partners contributing effort?
    - Sustainable funding from some form of research infrastructure fund?
    - Can the IS-ENES board help with this. Could it be the task of the new permanent post?
- What are the strategic drivers for long term investment in OASIS?
    - Could a user group and a strategic board help answer these questions?
    - Need for much higher scalability?
    - Need for optimisation?
    - Need for better end user support, documentation etc.
    - Functionality?
    - Ease of use features?
    - Developments to aid use alongside (complementary) with other frameworks such as ESMF?
- Can we invest in more code sharing between OASIS and other software, as has happened with OASIS MCT?
    - Anything around XIOS?
    - Anything around ESMF?