OASIS Developer Meeting
May 27th – 28th , 2009, CERFACS, Toulouse
S.Valcke
CERFACS Working Notes WN-CMGC-09-153

Participants:

NLE-IT: Hubert Ritzdorf, René Redler
AWI: Kerstin Fieg
CERFACS: Laure Coquart, Jean Latour, Sophie Valcke, Jean-Marie Epitalon

I. Administrative and management issues

1. CERFACS-NLE-IT collaboration
   - NLE-IT will be closed Dec 31st.
   - Some support for OASIS4 can be provided until then (30% René, freetime from Hubert).
   - In IS-ENES, collaboration with DKRZ will be established (DKRZ has 36 pm for OASIS4 development in IS-ENES)

2. Source management
   - The development branch has been recently merged in the trunk.
   - The developments should be done on development branch(es) but the merging on the trunk should be done frequently. Major bug fixes should be done both on the development branch(es) and on the trunk.
   - The users should get the sources from the trunk only.
   - Kerstin should download the development branch, update her sources regularly and check in her developments back on the development branch.

II. Current developments

1. Global parallel search:
   - The information about the local partition localisation in the global index space provided by the prism_def_partition is currently mandatory for the point based and cell based global parallel search. It is used to identify the missing source points for the interpolation star near the source border and to identify source cells across process borders. Otherwise a search based on geographical data would be required but this would be costly and is not implemented.
   - The halo info which is provided by the user (actual_shape larger than valid shape) is currently ignored (see prism_get_halo_points and the comment made by Hubert therein). If we would want to use the halo definition provided by the difference between the actual_shape and the valid_shape, we would need an additional indication from the application programmer that this is the halo we need. The application code might have other reasons (for example performance reasons) to use these different shapes.
   - Note for conservative remapping: as soon as at least one corner of a target cell falls into the domain of a source process (defined by the corners of its cells), the cell-based search will be done for that target cell.
   - The search is initially done for all source cells without considering the mask in case the mask changes during the run.

2. Treatment of clone fields
   - Currently, clone fields of a source process (i.e. coupling fields sharing the same source and target components, source and target grids, source and target masks, same

interpolation) are associated to the same EPIO and are treated by the same Transformer process (so that same weight calculation is not duplicated on different Transformer processes). This is OK.
- The case described by Sophie which would result in a Transformer load imbalance (a coupled configuration with 10 clone fields, a source and a target launched on one pe each, and a Transformer launched on 10 pes) is not likely to happen.

3. Timestamp on coupling fields within OASIS4 PSMILe and T
    - Currently, there is no timestamp associated to the coupling field instances in the Transformer. The Transformer applies a "first in, first out" (FIFO) algorithm.
    - This does not lead to any problem if the rule that the date_bounds of the prism_put/prism_get calls cover exactly the whole run duration without any gap and any overlap is strictly respected.
    - Routines prism_put/prism_get should be modified to ensure that the rule is respected (note after the meeting: this was done and checked in with r1958).
    - However, this rule could be relaxed if the date_bounds were transferred with the data to the Transformer and if the Transformer would check the consistency of the prism_put and prism_get date_bounds.

4. Collaboration with the BMBF funded ScalES project (unstructured grids and mixed parallelisation)
    - In a first step, the coupling between ECHAM5 and FEOM (AWI finite element ocean model) will be done considering a structured grid in FEOM. The interpolation from this structured grid to FEOM unstructured grid will be done internally to FEOM with its interpolation routines.
    - The work to fully integrate a search on unstructured grid in OASIS4 will be started in a second step. It should be done similarly to what is done for Gaussian grid, i.e. with a mapping from the unstructured grid to an auxiliary regular grid on the source side and with the search done on the auxiliary source grid. The prism_set_corner routine will have to be adapted to unstructured information. A prism_set_connectivity routine will have to be developed to provide the information about the neighbour nodes for each node; it has to be investigated how the connectivity in partitioned domain is expressed in current models using unstructured grid.

5. Possibility to use user-defined weights-and-addresses:
    - On-going work at CERFACS by Jean Latour
    - Sophie has to sent Jean's report to all developers

6. GUI for SMIOC XML constitution (demo by J.-M. Epitalon)
    - Jean-Marie demonstrated the GUI he is currently developing for automatic constitution of SMIOCs given pre-defined PMIODs. His demonstration was very convincing and this GUI will answer a long-time demand. Jean-Marie is currently finalizing the GUI and developing an easy installation procedure.
    - Jean-Marie is also working on a separate GUI that will help the users build their PMIOD (interface to the XML structure).
    - The XML structure needs revision and simplification (to be done within METAFOR).

III. Potential or future developments

1. Important OASIS4 developments to consider
    - OASIS3 reproducibility is important; time has to be spent on this task.
    - Proactive development for the next generation of HPC
    - Analyse the efficiency of XML ingestion in the initialisation phase

- Conservative remapping improvement (also in OASIS3)
    - Use of rotation when near the pole
    - IPSL algorithm
- Support of regional domain (i.e. calculate a nearest-neighbour value for target points falling outside the source domain): to be re discussed with Hubert
- Verification of current vertical interpolation
- Development of a battery of automatic tests to be run and analysed automatically (e.g. to check if a modification has no side effect on different machines)

2. Support of sequential components into one application/executable
    - This functionality might be not so difficult to integrate.

3. Storage and reuse of weights-and-addresses
    - The weights-and-addresses (or at least the nearest-neighbour or "reference source point") could be stored during the first run and read and reused in the subsequent runs.
    - This functionality would be difficult to implement because the order of the points in the EPIOs is not ensured (because of parallel calculation); a unique index would have to be assigned to each point and a reordering of the points according to this index would have to be done.
    - We first need to compare the time needed for the calculation of the weights with-respect-to the time needed for the storage and rereading before considering implementing this functionality.

4. Potential transfer of Transformer functionality in the source PSMILe:
    - One benefit would be that the total memory required would be smaller. Currently the source data (memory X) are sent to the Tranformer (memory X) where they are interpolated (memory Y) and then sent to the target model (memory Y). The total memory is therefore X+X+Y+Y. If the transformer functionality was transferred in the source PSMILe, the total memory required would be X+Y (in the source PSMILe) +Y (in the target model).
    - Another advantage is that there would be one less executable in the coupled system (this might be easier to manage for batch scheduler).
    - One disadvantage is that the memory required would be less distributed. Currently, the source PSMILe does not need to store the interpolated field because the Transformer acts as an intermediate buffer.

5. Mixed openMP - MPI parallelisation
    - This mixed parallelisation could be interesting for OASIS4, in particular for systems for which all MPI tasks of the coupled model have to use the same number of openMP threads (e.g. loadleveler on IBM)
        - Q? Is this a peculiarity of loadleveler or is this a standard restriction?
    - OpenMP parallelisation is also adapted to multicore processors
    - The mixed parallelisation of OASIS4 using both MPI and openMP is described as task 7 (deliverable A3) in SCALEs work package 4 lead by University of Karlsruhe). 6 pm are devoted to U. Karlsruhe and 3 pm to AWI for this task.

6. Support of vector fields
    - Only vector fields provided in the spherical system (lon, lat, z) will be supported (no local rotation); transforming a vector field from the local coordinate system to the spherical system (or from the spherical system to the target local coordinate system) can be easily done in the source model (or in the target model).
    - Projection in a Cartesian coordinate system will be included.

7. Global conservation:
    - A routine prism_set_scale_factor or prism_set_areas has to be developed.
    - The area integral of the source field has to be calculated on the source grid, transferred to the target psmile which will then calculate the area integral of the interpolated field and the difference between the two integrals.

8. Dynamic grids:
    - Very difficult to implement because the search involves interaction between the source and the target domain and there is no way currently to let the target component know when the source component changes its grid.
    - The target information should be stored on the source side to be able to redo the search calculation when the source model grid changes.
    - We do not plan to support this within IS-ENES.

9. Other developments with medium-low priority
    - Extrapolation on the source grid points
    - Use of autoconf or FCM for compilation
    - Consider integrate CISL (AWI) interpolations in OASIS4
    - Support other exchanges dates than at a regular frequency.
    - Combination of more than one source fields for one target field: it should be not so difficult to include this functionality in the target PSMILe, but we will wait for an expression of need.

10. Other notes:
    - GPU could be used for some parts of the code; it would mean duplicating some parts of the code under CPP keys with one part calling libraries developed for GPU
    - ORCA grid and point or cell based search (j=147,148,149, i=1,2,181,182): see closed ticket #5)
    - Strategy for modifying PSMILe API: it is better in general to add new routines than to modify the API of existing ones.
    - Block structured grids should be supported by associating more than one grid id to the same field id (to test).