

OASIS3-MCT, a coupler for climate modelling

S. Valcke, L. Coquart (CNRS/CERFACS), T. Craig

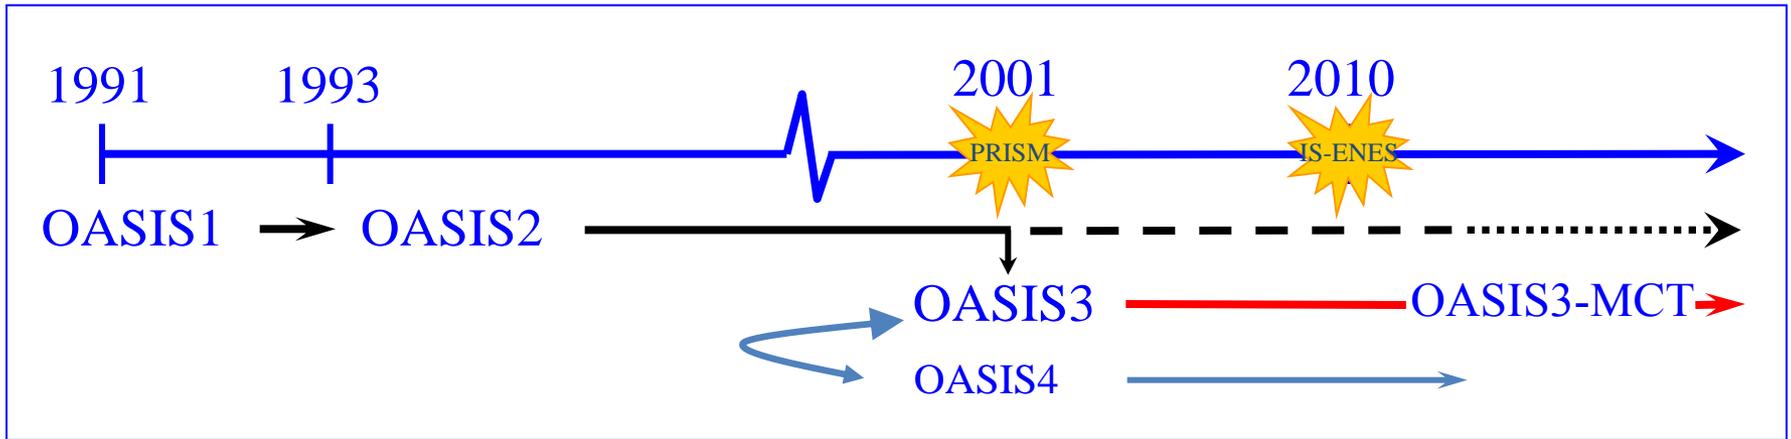
- A bit of history: what worked, what didn't
- Application Programming Interface
- Communication
- Interpolations et transformations
- Users
- Performance
- Conclusions and perspectives



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



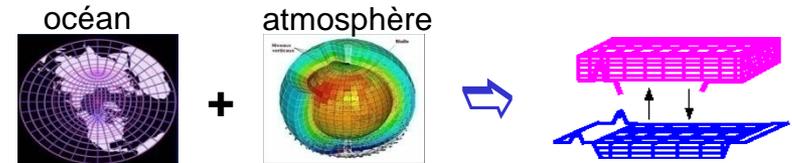
OASIS historical overview



- **OASIS1 -> OASIS2 -> OASIS3:**

2D ocean-atmosphere coupling
low resolution, low frequency

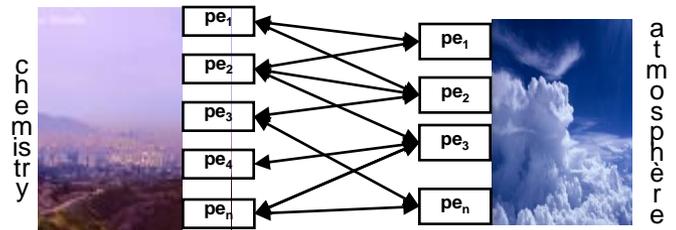
→ flexibility, modularity, 2D interpolations



- **OASIS4 / OASIS3-MCT:**

2D/3D coupling of high resolution parallel components on massively parallel platforms

→ parallelism, efficiency, performance





OASIS historical overview: OASIS(3) users

About 35 groups world-wide (climate modelling or operational monthly/seasonal forecasting):

- France: CERFACS, METEO-FRANCE, IPSL (LOCEAN, LMD, LSCE), OMP, LGGE, IFREMER
- Europe: ECMWF + Ec-Earth community
- Germany: MPI-M, IFM-GEOMAR, HZG, U. Frankfurt
- UK: MetOffice, NCAS/U. Reading, ICL
- Denmark: DMI
- Norway: U. Bergen
- Sweden: SMHI, U. Lund
- Ireland: ICHEC, NUI Galway
- The Netherland: KNMI
- Switzerland: ETH Zurich
- Italy: INGV, ENEA, CASPUR
- Czech_Republic :CHMI
- Spain: U. Castilla
- Tunisia: Inst. Nat. Met
- Japan: JMA, JAMSTEC
- China: IAP-CAS, Met. Nat. Centre, SCSIO
- Korea: KMA
- Australia: BoM, CSIRO
- New Zealand: NIWA
- Canada: RPN-Environment Canada, UQAM
- USA: Oregon State U., Hawaii U., JPL, MIT
- Peru: IGP

➤ OASIS3 is used in 5 of the 7 European ESMs that participate in IPCC AR5



OASIS historical overview

OASIS4

- Developed by CERFACS, CNRS, DKRZ and MPI-M (PRISM & IS-ENES EU projects)
- Fully parallel generation of regridding weights & addresses
- Specific OASIS4 versions used in:
 - 3D coupling between atmosphere and atmospheric chemistry at ECMWF, KNMI and Météo-France (EU GEMS & MACC projects)
 - SMHI (Sweden) for regional ocean-atmosphere in the Arctic
 - BoM (Australia) for global ocean-atmosphere coupling
- Development was stopped in July 2011:
 - Lack of resources for development of real-time parallel calculation of regridding weights and addresses for all type of grids
 - OASIS4 code too complex to evolve easily and original design did not include unstructured grids
- Take a step back:
 - Suppose that (static) weights & addresses are pre-calculated before the run
 - Use MCT in OASIS3 for parallel matrix multiplication (regridding) and parallel data transfer & redistribution

- First official version released in August 2012, second in May 2013
- Uses the Model Coupling Toolkit (MCT) from Argonne National Lab
- Open source, LGPL license, F90 and C
- External libraries: public domain (MPI, NetCDF) or open source (SCRIP, MCT)
- Current developers are:
 - 2 permanent FTEs (CERFACS, CNRS)
 - 1 consultant (T. Craig, also NCAR CESM cpl7 developer)



IS-ENES (InfraStructure for ENES) EU FP7 project
2009-2012 - 18 partners - 7,6 MEuros ; coord: IPSL

- 93 pm for OASIS development and support



IS-ENES2, EU FP7 project 2013-2016 - 25 partners -
8MEuros; coord: IPSL

- 27 pm for OASIS3-MCT development and support



OASIS3-MCT Application Programming Interface

- ◆ Component models remain separate executables
- ◆ OASIS3-MCT acts as a communication library linked to the models
- ◆ OASIS3-MCT API is strictly backward compatible with OASIS3.3 API

```
•Initialisation:           call oasis_init_comp(...)  
•Grid definition:         call oasis_write_grid (...)  
•Local partition definition: call oasis_def_partition (...)  
•Coupling field declaration: call oasis_def_var (...)
```

•**Coupling field exchange** (in model time step loop) :

```
    call oasis_put (... , date, var_array. ...)  
    call oasis_get (... , date, var_array, ...)
```

- **all coupling parameters defined externally by the user** in a text file
 - source or target (end-point communication)
 - sending or receiving at appropriate time only
 - automatic averaging/accumulation if requested
 - automatic writing of coupling restart file at end of run

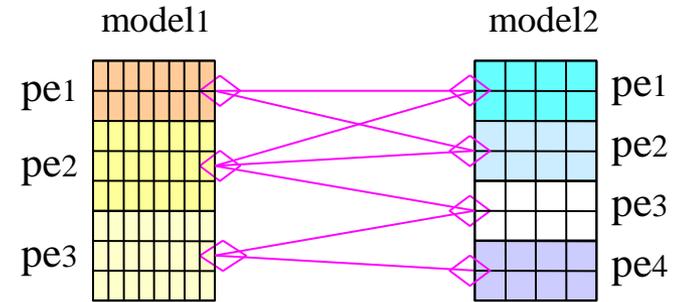
```
•Termination:           call oasis_terminate (...)
```

OASIS3-MCT communication

Fully parallel communication between parallel models based on MCT using MPI:

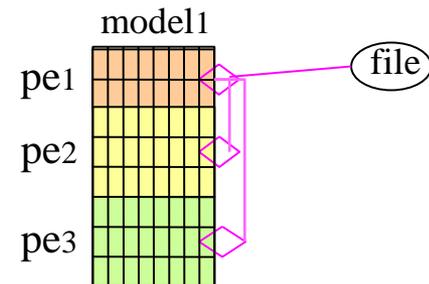


- computation of communication patterns
- matrix multiplication for regridding (on the source or target processes)
- coupling data transfer/redistribution



If specified by the user, the regridding weights and addresses are calculated onto one model process using the LANL SCRIP library

- I/O functionality (switch between coupled and forced mode):





OASIS3-MCT: interpolations & transformations

- on 2D or 3D scalar fields
- on different types of grids: lat-lon, rotated (logically rectangular), gaussian reduced, unstructured

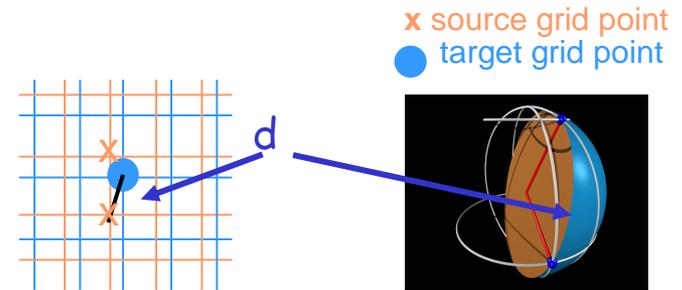
❖ Transformations

- Statistics
- Addition/multiplication by scalar
- Global conservation

❖ Interpolations/regridding

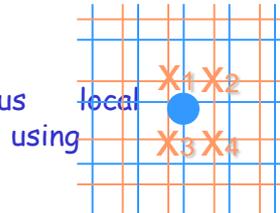
SCRIP library (Jones, 1999)

- n-nearest-(gaussian-weighted)-neighbours: $\text{weight}(x) \propto 1/d$
d: great circle distance on the sphere:



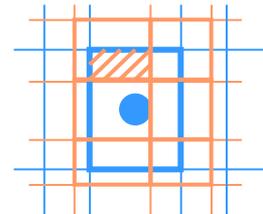
- bilinear interpolation

- general bilinear iteration in a continuous coordinate system
 $f(x)$ at x_1, x_2, x_3, x_4



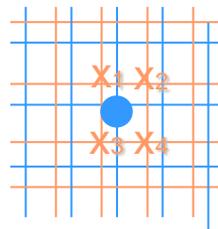
- conservative remapping

- weight of a source cell % to intersected area

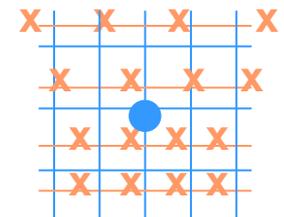


- bicubic interpolation:

- general bicubic iterations in a continuous local coordinate system:
 $f(x), \delta f(x)/\delta i, \delta f(x)/\delta j, \delta^2 f/\delta i \delta j$ in
 x_1, x_2, x_3, x_4
for logically-rectangular grids (i,j)



- standard bicubic algorithm:
16 neighbour points
for Gaussian Reduced grids



** gradients must be given as extra arguments to the oasis_put



OASIS3-MCT users

- ◆ CERFACS (France): running
 - NEMO ocean (ORCA025, 1021x1442) - ARPEGE atmosphere (Gaussian Red T359 grid, 817240 points).
 - Seasonal prediction experiments in SPRUCE PRACE project (27 Mhours on tier-0 Bullx Curie at TGCC).
 - Decadal experiments in HiResClim I & II PRACE project (38 & 50 Mhours on IBM Mare Nostrum at BSC)

- ◆ IPSL (France): running
 - WRF atm - NEMO ocean model, both with two-way nested zooms, resolution from 27 km to 9 km.
 - Used in PULSATION project funded by the French ANR, 22 Mhours on PRACE tiers-0 Bullx Curie.

- ◆ MPI-M (Germany): running
 - All MPI-ESM versions, in particular MPI-ESM-XR: atmosphere ECHAM6 T255L95 (768x384 grid points, ~50km, 95 vertical levels) - ocean MPIOM TP6ML40 (3602x2394 grid points, ~10km, 40 vertical levels); 17 coupling fields, 1h coupling frequency.

- ◆ MetOffice (UK) : running
 - Global ocean-atmosphere coupling between UM global atmosphere (N512, 1024x769) and NEMO ocean (ORCA025, 1021x1442).

- ◆ BTU-Cottbus (Germany): running
 - 3D coupling between COSMO-CLM regional atmosphere (221x111x47, ~2 deg) and ECHAM global atmosphere (T63, 192x96x47), + 2D coupling to MPI-OM ocean (254x220)
 - 6% coupling overhead for exchange of 6 x 3D fields every ECHAM time step

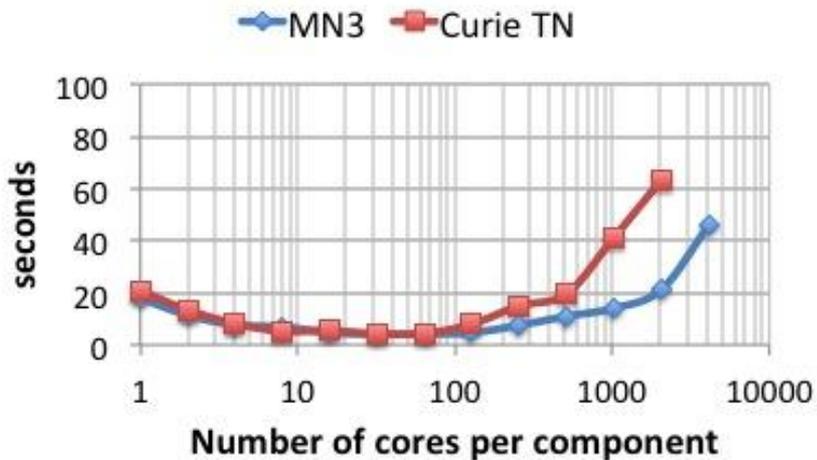
- ◆ BoM (Australia): under developments
 - limited area coupled model MOM4p1 ocean and UK Met Office UM6.4 atmosphere.

◆ ...

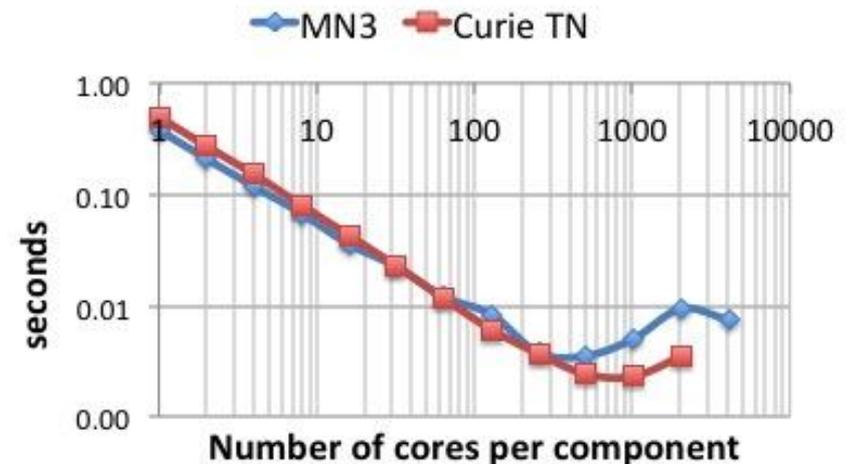
OASIS3-MCT performance

- Toy coupled model: ping-pong exchanges between one component on NEMO ORCA025 grid (1021x1442) and the other on a Gaussian Reduced T799 grid (843 000)
- **Bullx Curie** thin nodes; Intel® procs Sandy Bridge EP; IFort 12.1.7.256, Bullx MPI 1.1.16.5
- **IBM MareNostrum3**: 48,896 Intel Sandy Bridge processors, Intel MPI 4.1.0.024

Total initialisation time



Time for one ping-pong



➤ Coupling overhead for:

- one-year long simulation with one 1 coupling exchange every hour in each direction between codes with $O(1\text{ M})$ grid points running on 4000 cores/component
 - ◆ ~60 seconds for initialisation, ~30 seconds for data exchange (previous OASIS3.3 would have taken ~2600 seconds for the same data exchanges)



Conclusions and perspectives

Conclusions on OASIS3-MCT

- Good performance thanks to MCT, removes OASIS3.3 bottleneck, very simple to use for traditional OASIS3.3 users (same API)
- > OASIS3-MCT most likely provides a satisfactory solution for fully parallel coupling in our climate models at the resolutions targeted operationally for the next ~3-5 years.
- ** A good example of shared software and Coupling Technology Workshops benefits **

Perspectives:

- Refactoring of *namcouple* configuration file + Graphical User Interface
- Evaluation of ESMF for off-line precomputing of interpolation weights (on going)
- Evaluation of Open-PALM (including ONERA CWIPI library) for on-line parallel calculation of interpolation weights (designed for unstructured grid)
- IS-ENES2:
 - Coupling technology benchmark launched during the Boulder Coupling Technology Workshop last February
 - Should we (Europeans) evolve to the "integrated mono-executable" approach (e.g. CESM, GFDL, ESMF): more constraints but more opportunities for performance, e.g. exchange of coupling data through the memory?



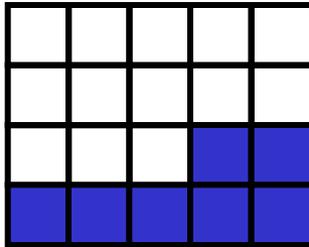
The end



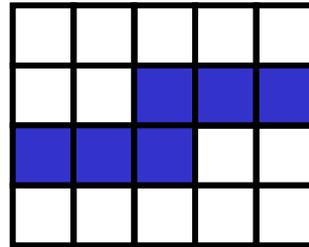
OASIS3-MCT: parallel decomposition supported

Apple

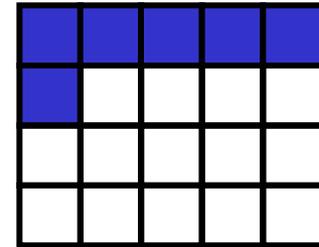
part1



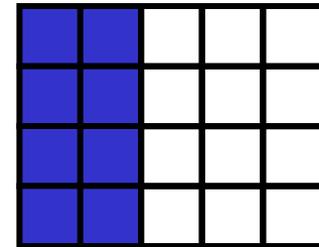
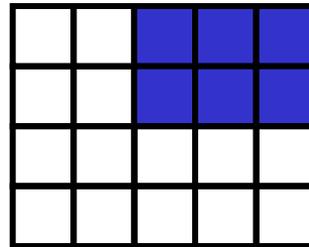
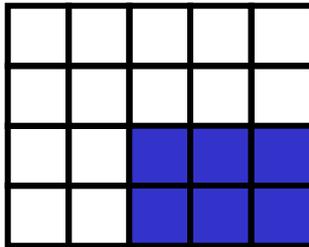
part2



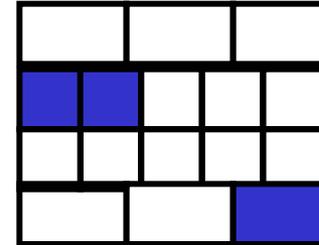
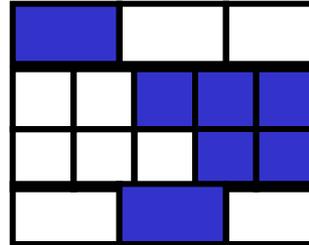
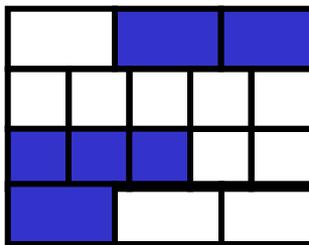
part3



Box



Orange



Apple and orange applicable to unstructured grids



Use of OASIS3-MCT

At run time, the component models remain separate executables and OASIS3-MCT acts as a communication library linked to the models.

To use OASIS3-MCT:

- Register, download the sources, compile and run the tutorial toy model on your platform
 - Identify your component models, their grids, the coupling fields to be exchanged between those models
 - Identify the interpolations/transformations needed to go from the source to the target grid
 - Use the "test_interpolation" environment (offline interpolations/transformations) to test their quality
 - Adapt your model i.e. insert calls to OASIS3-MCT communication library
 - Choose the other coupling parameters (source and target, frequency, field transformations, etc.) and write the full *namcouple* configuration file
 - Compile OASIS3-MCT and link the components models with it
 - Start the models and let OASIS3-MCT manage the coupling exchanges
-