The IS-ENES coupling technology benchmark

S. Valcke (CERFACS)
R. Ford (STFC), M. Hanke (DKRZ), M. Hobson (MetOffice),
G. Jonville (Cerfacs), A. Porter (STFC), G. Riley (U. Manchester)

Introduction

- IS-ENES: The InfraStructure for the European Network of Earth System modelling
- Coupling the components of an Earth Climate Model

Objectives

Implementation

First results

Conclusions and perspectives

The InfraStructure for the European Network of Earth System Modelling (IS-ENES)

http://is.enes.org/

FP7 project « Integrating Activities »

1rst phase: March 2009- Feb 2013 (7.6 M€), 18 partners
2nd phase: Apr 2013- March 2017 (8 M€), 23 partners

Better understand and predict climate variability & changes by fostering

- The integration of the European Earth System Model community
- The development of Earth System Models and their environment
- High-end simulations
- The application of Earth System Model simulations for climate change impacts

Coupling components of an Earth Climate Model

An Earth Climate Model: system coupling (pre-existing) numerical representations of the Earth Climate components

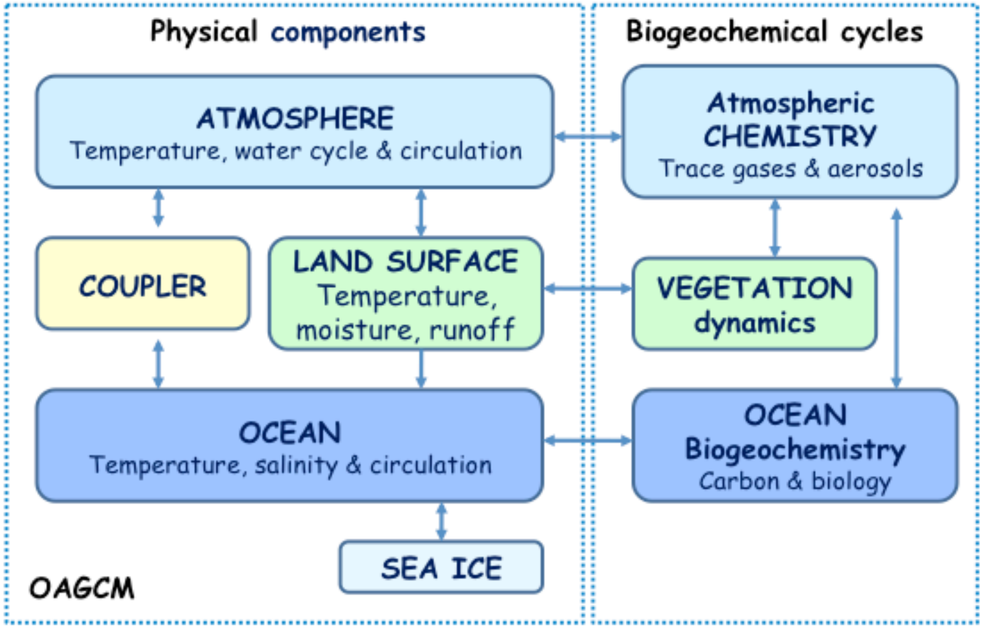

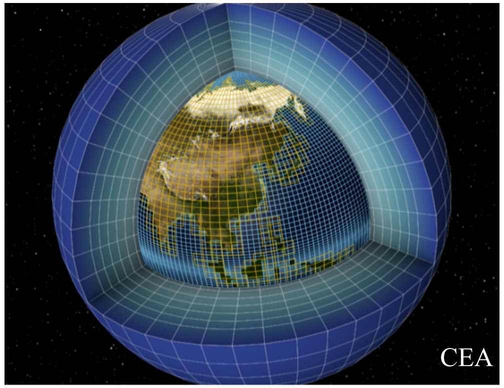

Strong needs of HPC:
Component resolution, simulation duration, system complexity, ensembles
Capability & Capacity

Subroutine coupling

- Transform one code into a subroutine
- Make one code call the other code as a subroutine

```
program prog1
…
call prog2 (fieldin, fieldout,…)
…
end prog1
```

~~program~~ subroutine prog2
…
end prog2

☺ efficient
☺ sequential components

☹ hard coded coupling
☹ no flexibility in the component layout

Coupling framework integrated approach

- Split code into elemental units
  at least init/run/finalize
- Write or use coupling units

- Adapt data structure and calling interface
- Use the framework (to build) a hierarchical merged code

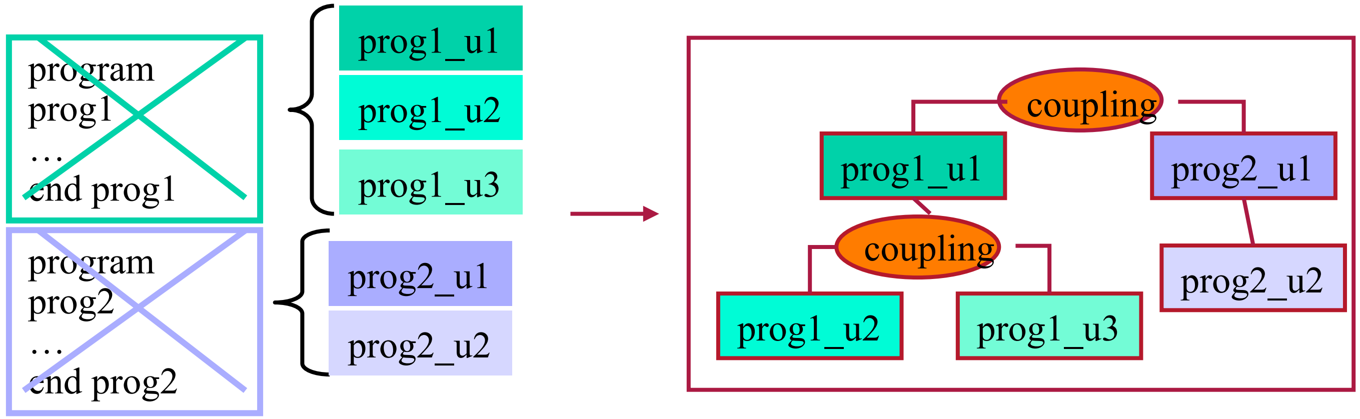

☺ efficient
☺ sequential and concurrent components
☺ use of generic utilities (parallelisation, regridding, time management, etc.)

☹ existing codes
😐 (easy)

ESMF

FMS (GFDL)

CESM (NCAR)

→ probably best solution in a controlled development environment

Coupler or coupling library approach

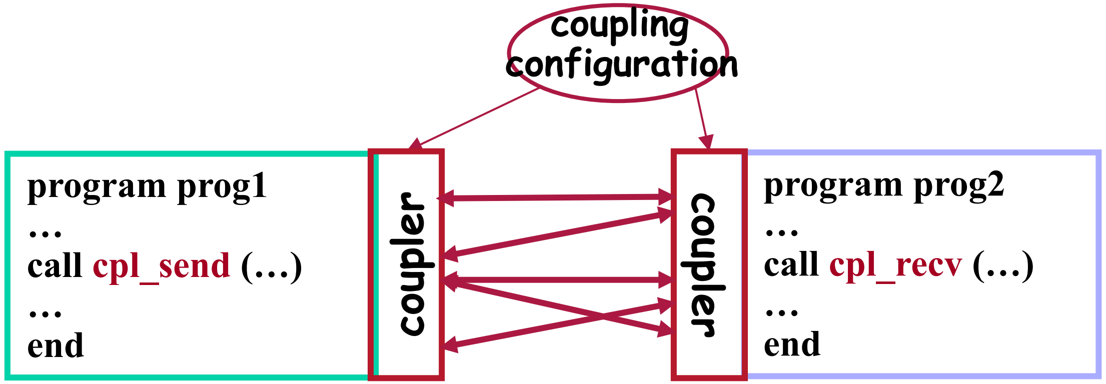

- ☺ existing codes
- ☺ use of generic transformations/ regridding
- ☺ concurrent coupling (parallelism)

- ☹ Sequential components: waste of resources?
- ☹ multi-executable: more difficult to debug; harder to manage for the OS
- 😐 efficient

→ probably best solution to couple independently developed codes

The coupling technology benchmark objectives

Define a suite of coupled benchmarks based on simplified components, which capture the essence of the coupling in Earth System Models without the science complexity

1. Capture functional and performance coupling characteristics of Earth System Models
2. Code a set of simplified components reproducing these coupling characteristics
3. Implement the coupling with different coupling technologies
4. Run the benchmark suite on specific platforms
5. Analyse results and present them to the community

Coupling characteristics of Earth System Models

Working groups at the 2nd Workshop on Coupling Technologies for Earth System Models (CW2013), Boulder, 2013:

=> exhaustive list of coupling system characteristics

US project Earth System Bridge + 2014 IS-ENES2 Exeter workshop:

=> mindmaps (https://earthsystemcog.org/projects/es-fdl/mindmaps)

EARTH SYSTEM BRIDGE

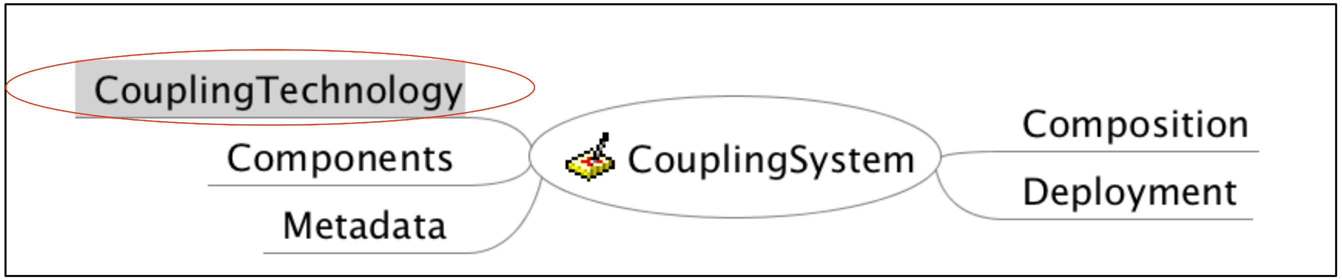

Characteristics of coupled Earth System Models

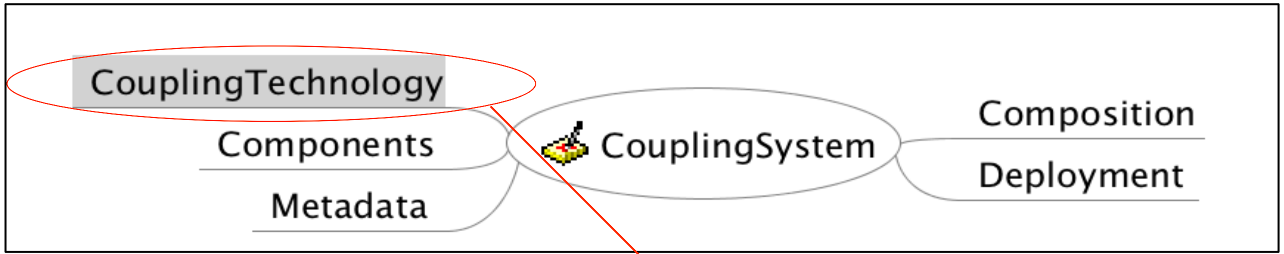

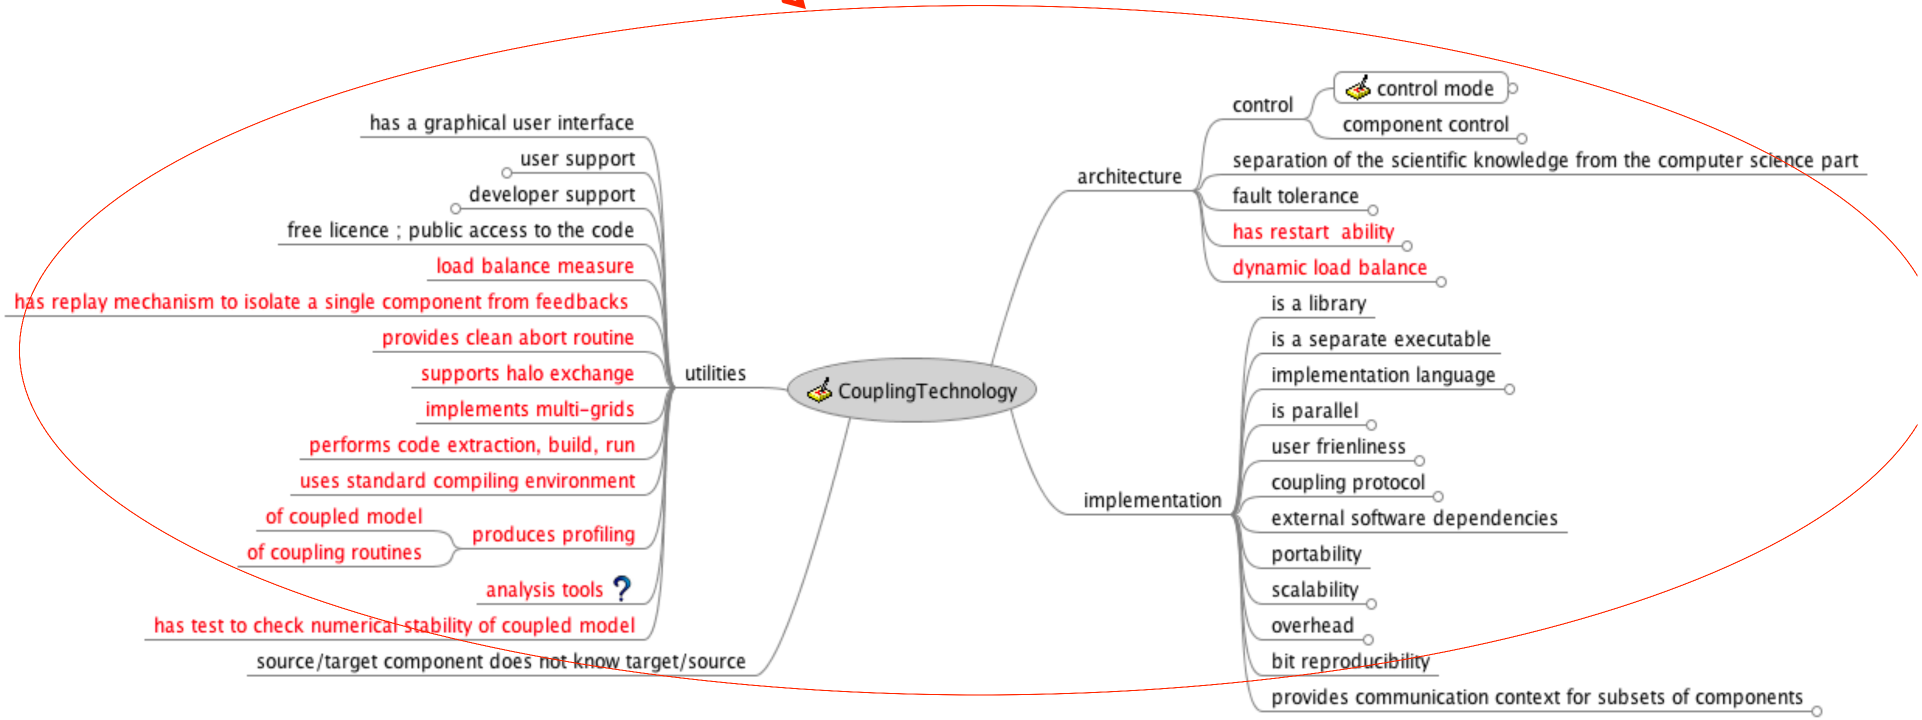

- architecture: basic design principles and other general characteristics
- implementation: how the technology is implemented (library, parallelism, language, etc.
- utilities: all the utilities offered by the technology

- ✓ Priority coupling characteristics to benchmark
  - Type of the component grid
  - Number of cores per component
  - Numbers of fields exchanged
  - Frequency of exchange
  - Size of the coupling fields
  - (Ease of use: code intrusion, development time, techniques for overcoming specific issues)

Stand alone components

✓ 4 stand alone components on 4 different grids

- MPI parallel Fortran subroutine(s) not modelling any physics or dynamics but implementing real coupling characteristics
- Potential coupling fields as IN/OUT arguments, arrays in modules, local data
- Use specific numerical grids
  - o *latlon:* latitude-longitude, arbitrary resolution
  - o *stretched:* stretched, rotated, logically rectangular, e.g. NEMO ocean ORCA grid
  - o *icosa:* quasi-uniform icosahedral, e.g. DYNAMICO (LMD, Fr) or ICON (DWD, DKRZ, De)
  - o *cubesphere:* quasi-uniform cubed sphere, e.g. NOAA/GFDL dynamical core

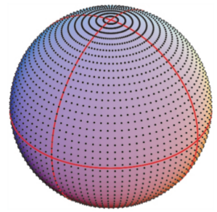
latlon

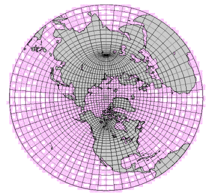
stretched

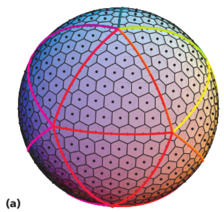

icosa

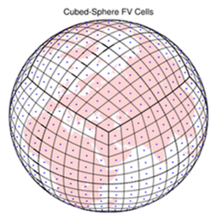

cubesphere

Coupled test case implementation

Set-up of coupled test-cases implementing ping-pong exchanges between two components running on regular lat-lon grids with:

- **OASIS3-MCT** (https://portal.enes.org/oasis)

  Legacy coupler developed at Cerfacs (France) and used by many climate modelling groups in Europe

- **ESMF**: (https://www.earthsystemcog.org/projects/esmf/)

  High-performance, flexible software for building climate and weather applications; US multi-agency governance (NSF, NASA, DoD, NOAA) with many partners

- **OpenPALM** (http://www.cerfacs.fr/globc/PALM_WEB/)

  Dynamic coupler developed by Cerfacs and ONERA (France) originally for data assimilation suites

- **MCT** (https://www.earthsystemcog.org/projects/mct/)

  Set of open-source software tools for creating coupled models

- **YAC** (https://doc.redmine.dkrz.de/YAC/html/)

  Light weight coupling software infrastructure developed at DKRZ (Germany)

Coupled test case implementation

Measure of the time for a ping-pong exchange (back-and-forth) between two components

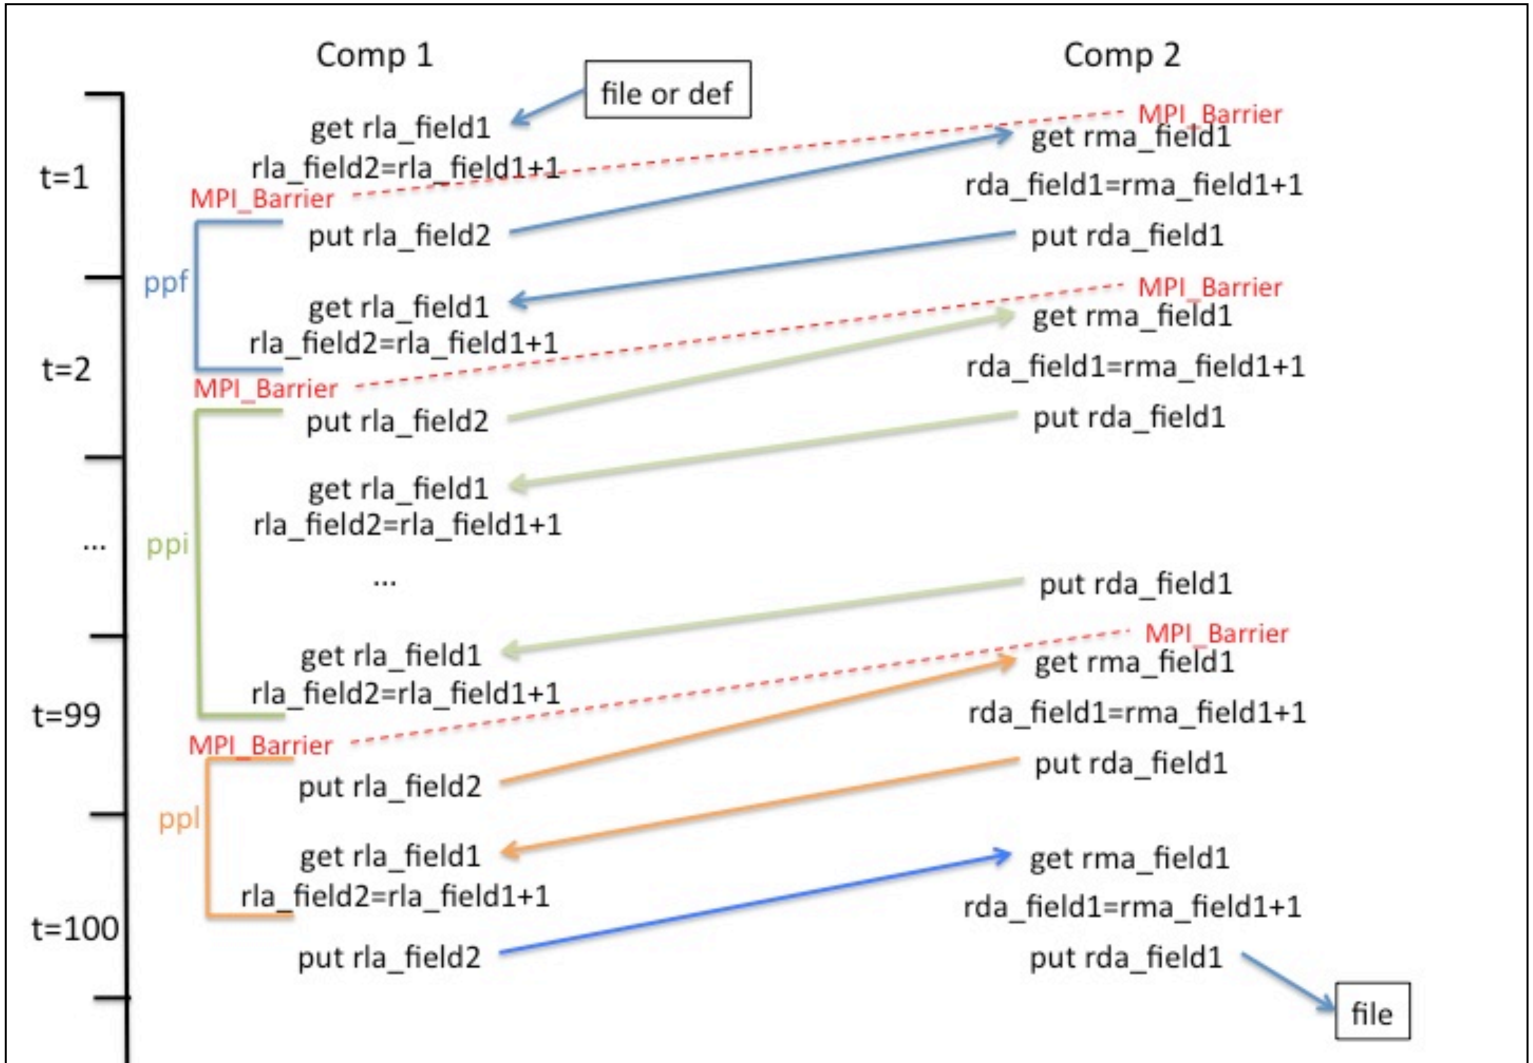

ppf: first ping-pong   ppi: 97 intermadiate ping-pongs ppl: last ping-pong

First results of specific coupled cases on specific platforms

We tested the impact of

- The coupling technology: OASIS, ESMF, Open-PALM, MCT, YAC
- The platform: Bullx Occigen (Fr), MetOffice Cray (UK), Marconi Broadwell (It)
- The grid size : HR-HR, VHR-VHR, LR-HR

  with LR:100x100, HR:1000x1000, VHR: 3000x3000
  regular lat-lon grids with same decomposition on both sides

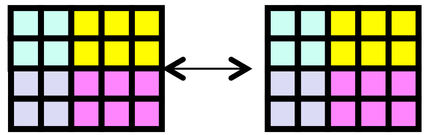

- The number of cores/component: from O(1) to O(10000)

- Each run is repeated 3 times to analyse the spread of the results

- *Allocation of 480 000 core-hours on Bullx Occigen at CINES*
- *Allocation of 19 000 core-hours on Marconi Broadwell at CINECA*

First results of a specific coupled case on specific platforms

Average time for one ping-pong exchange between components on regular latlon grids with same decomposition on both sides, wrt number of cores/component:

1000x1000 – 1000x1000 Occigen Bullx

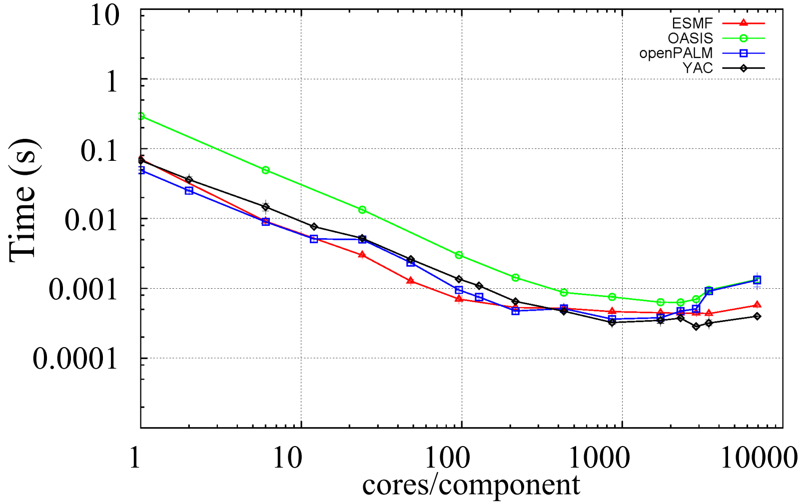

3000x3000 – 3000x3000 Occigen Bullx

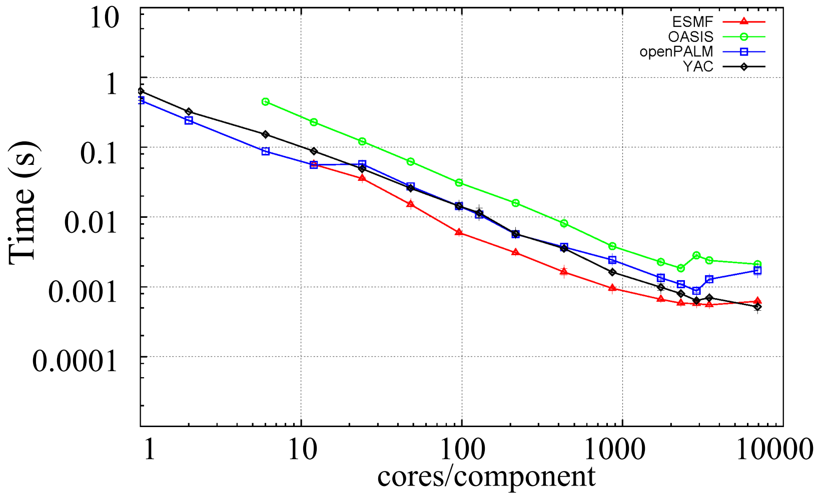

1000x1000 – 1000x1000 Cray XC40

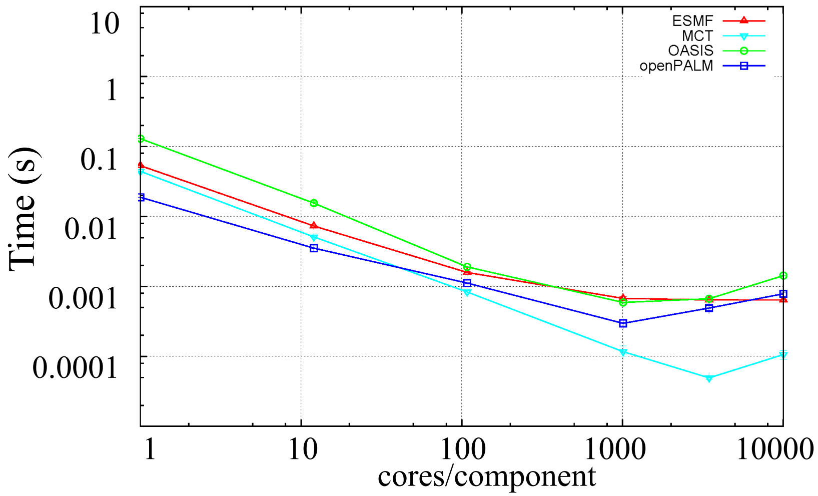

3000x3000 – 3000x3000 Cray XC40

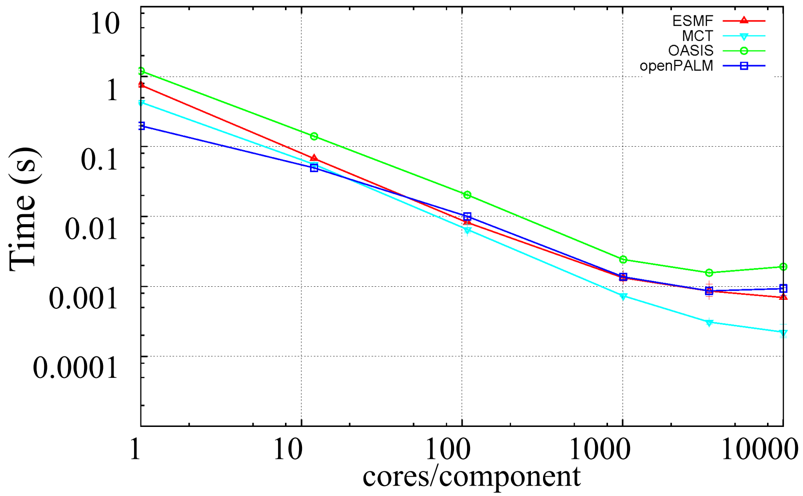

First results of specific coupled cases on specific platforms

Average time for one ping-pong exchange between components on regular lat-lon VHR grids, with same decomposition on both sides, wrt number of cores/component, for each coupling technology
for all 3 platforms

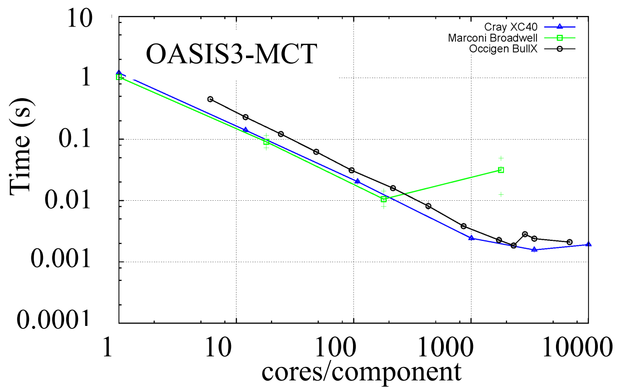

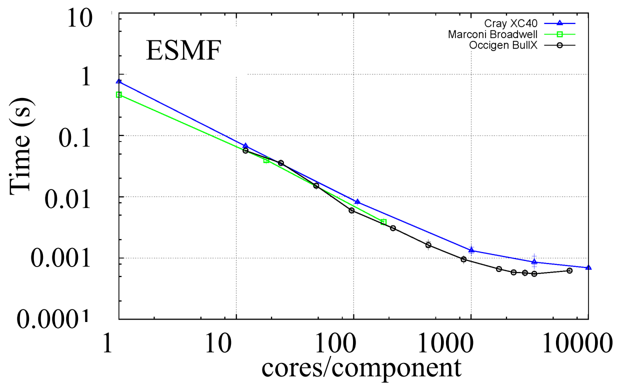

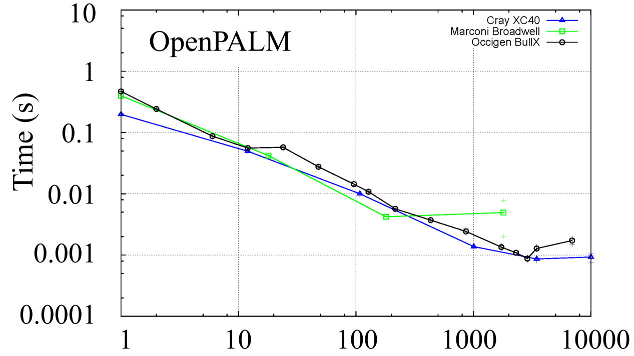

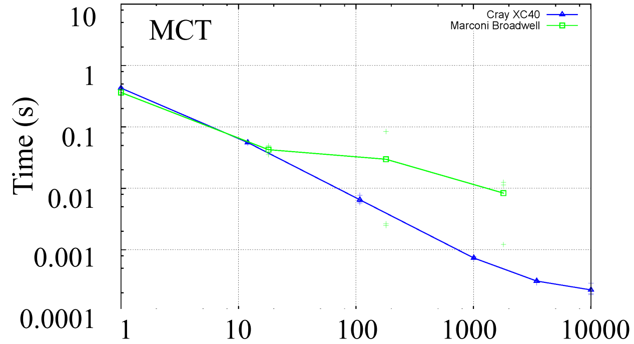

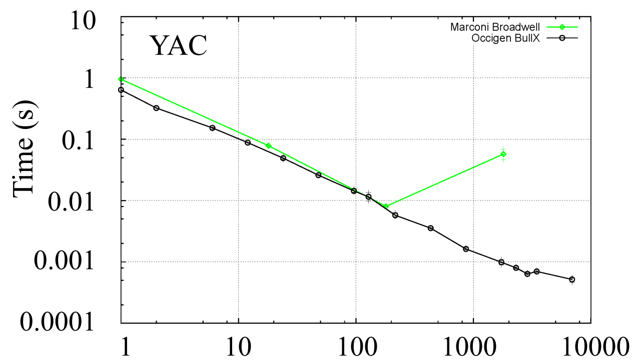

First results of specific coupled cases on specific platforms

Impact of opposite decompositions :

- The coupling technology: OASIS, ESMF, Open-PALM, YAC
- The platform: Bullx Occigen (Fr)
- The grid size : VHR-VHR (3000x3000 - 3000x3000)
- The number of cores/component: from O(1) to O(10000)

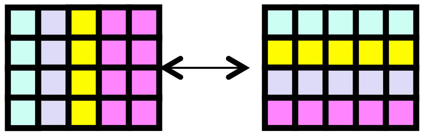

- Each run is repeated 3 times to analyse the spread of the results
- Minimal number of runs for this test case on Bullx: 4x1x1x5x3 = 60

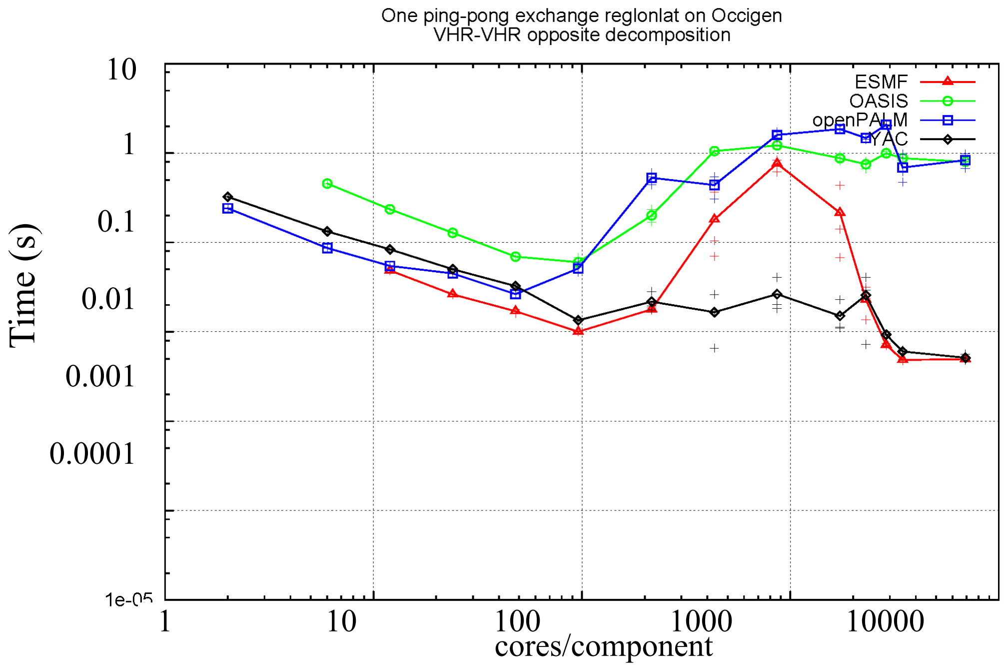

Conclusions and perspectives

Main results:

- ✓ First version of IS-ENES coupling technology benchmark available on ENES portal at https://verc.enes.org/computing/performance/benchmarks; will be made available through github soon:
  - 4 stand alone components on 4 different grids available
  - Coupled systems for regular lat-lon grids with OASIS, ESMF, OpenPALM, MCT & YAC
- ✓ This first benchmark version allows **for that very specific coupled test-case** :
  - a comparison between different coupling technologies
  - a comparison between different platforms
  - evaluate the impact of the number of cores/component
  - evaluate the impact of the decomposition
  - evaluate the impact of the size of the coupling grids / coupling fields
  - ~~Ease of use: code intrusion, development time, …~~
- ✓ Results available for Bullx Occigen, MetOffice Cray, Marconi Broadwell

Warning: more work is needed to analyse the results and understand the differences

Conclusions and perspectives

- Very interesting international collaboration for capturing the characteristics of coupling systems
- Implementing this first (simple) coupling technology benchmark version:
  - far more complex than expected
  - difficult to define standard and unbiased specifications
  - difficult to implement the same coupling with different coupling technologies and make sure we are measuring the same thing!
- Benchmark extensions (other grids, other decompositions, etc.) are needed but difficult to plan without specific funding; IS-ENES3?
- Interested? Everyone is welcome to run the benchmarks on additional platforms and we commit to publish additional results. Join the Coupling Technology Benchmark Group!

Many thanks to all the group for the energy put in this very challenging collective work

The end