Interpolation weights calculation distributed with OpenMP®

Eric Maisonnave

WN-CMGC-15-5

Table of contents

The OASIS[1] coupling library includes an automatic calculation of weight-and-address data associated to interpolations between two grids in spherical coordinates. This option is widely used by OASIS user community to quickly set up coupled model configurations[2]. Unfortunately, the increasing resolution of model grids and the sequential behavior of algorithms leads to important or even unreasonable (longer than a day) response times.

This document describes and evaluates performances of the OpenMP distribution of these calculations on standard and large shared memory node computers.

The author appreciated the help of Romain Bourdalle-Badie (MERCATOR) in setting the test case suite for performance measurements.

1. Rationale

The SCRIP[3] "Distance-weighted averaging" interpolation algorithm aims to calculate addresses and weights for interpolating variables between grids in spherical coordinates. It associates the inverse-distance-weighted average of an user-specified number of nearest neighbor values from a source grid to each non masked grid point of a target grid. The operation that consists in finding the n-nearest neighbors and calculating the associated weights is originally done sequentially for each unmasked target grid point. This outer loop is easily distributed between memory-shared processes using the standard OpenMP[4] library:

```
!$OMP DO SCHEDULE(RUNTIME)
do target_grid_point = 1, target_grid_size
 ! find nearest grid points on source grid
 ! compute weights based on inverse distance
 ! normalize weights and store the link
end do
!$OMP END DO
```

Included in the OASIS coupling library, a few intermediate and allocatable arrays has to be declared as private (and the initial values, previously set in a non parallel section, has to be copied before the parallel section into each private array):

```
!$OMP PARALLEL DEFAULT(FIRSTPRIVATE)
     & SHARED(coslat,coslon,sinlat,sinlon)
```

[1] Valcke, S., T. Craig, and L. Coquart (2013): *OASIS3-MCT User Guide OASIS3-MCT 2.0.*, Technical Report, TR/CMGC/13/17, CERFACS/CNRS SUC URA 1875

[2] Maisonnave, E., Valcke, S. and Foujols, M.-A., 2013: OASIS Dedicated User Support 2009-2012, Synthesis ,Technical Report, **TR/CMGC/13/19**, SUC au CERFACS, URA CERFACS/CNRS No1875, France

[3] Jones PW (1998) A user's guide for SCRIP: a spherical coordinate remapping and interpolation package, Los Alamos National Laboratory. http://oceans11.lanl.gov/trac/SCRIP

[4] *Using OpenMP: Portable Shared Memory Parallel Programming*, by Barbara Chapman, Gabriele Jost and Ruud van der Pas, MIT Press, Cambridge, MA, 2008

Those two instructions allows to distribute the interpolation weight calculations to as much processes as held by a computing node.

2. Test case and performances

The “Distance-weighted averaging” interpolation is performed between 3 set of increasing resolution grids (Table 1). The environment variable OMP_STACKSIZE must be set to 16m (Mbytes) to avoid memory limitations.

Source grid	Target grid
ORCA2, 2°, 182x149,	Regular3, 4°x2.5°, 96x72
ORCA025, 0.25°, 1442x1021,	Regular0.5, 0.5°, 721x361
ORCA12, 1/12°, 4322x3059	ARPEGE-t359, Gaussian 0.25°, 181724

Table 1: Name, resolution and grid points number of source/target grids test case

An Intel Xeon Sandy Bridge processor is used first to measure the parallel efficiency of our decomposition. As shown on figure 1 (red lines), calculations can be speed up by a factor 10 (from a theoretical maximum of 24) for ORCA2 grid and by a factor 5.8 for ORCA025 grid. At full decomposition (Figure 2), they last less than 1 second in the first case, and about 11 minutes in the second one. To fully benefit of the OpenMP parallelism on one node, the same measurements are lead on an Intel Xeon Phi Knights Corner processor, which includes 60 physical cores (and 240 virtual cores with multi-threading). On the same figure 1 (blue line), we see that speed up is smaller (4.38 on ORCA2 grid, from a theoretical maximum of 240).

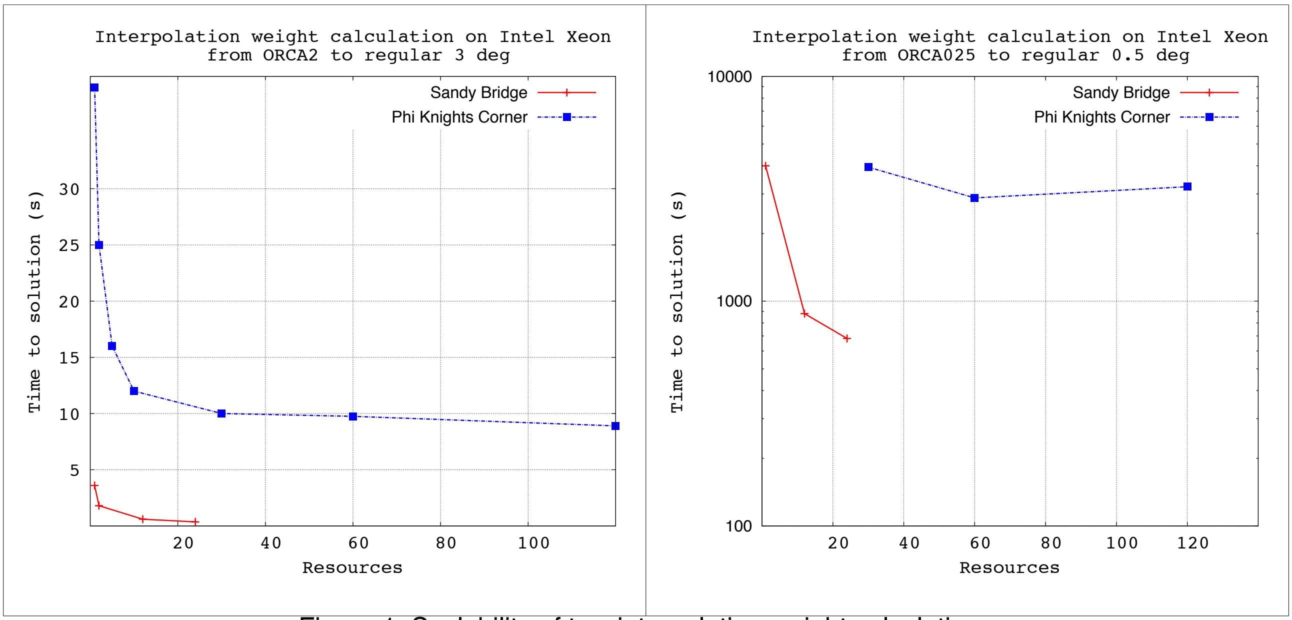

Figure 1: Scalability of two interpolation weight calculations

In addition, the minimum response time is longer (respectively 10 seconds and less than 1 hour). As shown on Figure 2, best performances on a single node are obtained on Xeon Sandy Bridge in any case. It has not be possible to launch calculations for ORCA12 grid on Xeon Phi. A more precised analysis reveals that parallel calculations are strongly unbalanced on more than 20 cores. It suggests that finer algorithm improvements are needed for any further improvement.

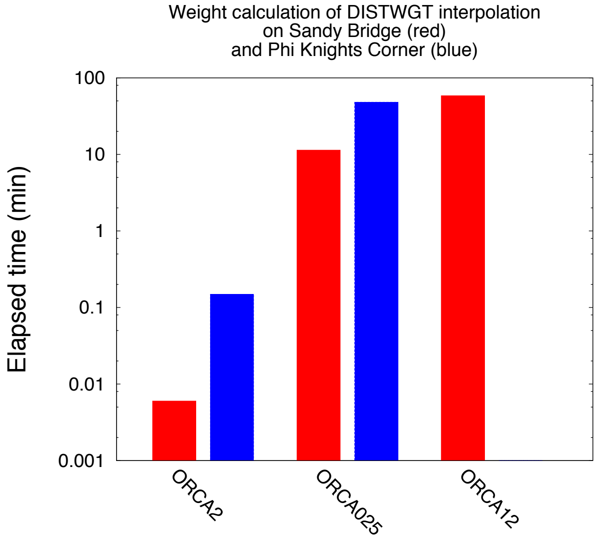

Figure 2: best performances for 3 interpolation weight calculations

3. Conclusion

The author warmly acknowledges Isabelle d'Ast, who helps him improving the “Distance-weighted averaging” calculation performances in a couple of hours. A few more efforts would lead to the same result for all OASIS interpolations available. A so-called “parallel processing of interpolation weight calculation” would be then implemented at a good price in the famous coupler. For the moment, the most demanding calculations become possible (less than one hour for an interpolation involving the ORCA12 grid) on standard machines. Even though performances on Intel Xeon Phi Knights Corner failed to match up to initial expectations, a possible increase of core number (and memory) per node on near future supercomputers should lead to better results, assuming that some time would be spent on algorithm improvements to reduce unbalanced computations.

This work was awarded by the FRIISM initiative committee. We also acknowledge funding received in the framework of the EU IS-ENES2 - Infrastructure for the European Network for Earth System modelling - Phase 2 project (GA no 312979) and in the framework of the French ANR Convergence project (GA no ANR-13-MONU-0008).