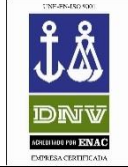



*MOHID implementation in parallel mode  
following a domain decomposition approach  
to the Rias Baixas area*



Unión Europea  
FEDER  
Invertimos en su futuro



Document		
<b>Classification:</b> Public	<b>Distribution</b> NVIST	<b>Observations</b> n/a
<b>Title</b> MOHID implementation in parallel mode following a domain decomposition approach to the Rias Baixas area		
<b>Keywords</b> MOHID, Rias Baixas, parallelization, domain decomposition, MPI		
<b>Entity that produced the document</b> Hidromod, Modelação em Engenharia, Lda Rua Rui Teles Palhinha, 4, 1º, Leão, 2740-278 Porto Salvo Tel: 214211373 – Fax: 214211272 Email: <a href="mailto:hidromod@hidromod.com">hidromod@hidromod.com</a>		Certificação de Qualidade ISO 9001:2008 
<b>Client</b> Consellería de Medio Ambiente, Territorio e Infraestruturas Rúa Roma, nº 6. 15707 Santiago de Compostela. A Coruña. Teléfono: (+34) 881 999 654		
<b>Authors</b> Paulo Chambel Leitão	<b>Title</b> Civil Eng., PhD Environment	
<b>Verification</b> Pedro Galvão	<b>Title</b> Master and degree Environment Eng.	
<b>Production date</b> 23/12/2014	<b>Page number</b> 20	<b>Reference number</b> Version 1

## Index

1	Introduction .....	5
2	Implementation of the MOHID Rias Baixas Model .....	6
2.1	MOHID Rias Baixas project .....	6
2.2	Model speedup using parallelization .....	7
2.2.1	Model surface results .....	9
2.2.2	Computational effort distribution .....	13
3	MOHID parallelization scalability .....	14
4	Conclusions .....	16
5	Future steps .....	17
6	Annex – Standard PC (GulfTown) vs CESGA (SandyBridge) .....	18

## Figure Index

Figure 1 – Time necessary to simulate 3 days function of the number of threads used.	8
Figure 2 – Speedup factor of MOHID Rias Baixas implementation in the CESGA-SVG Sandy Bridge nodes architecture.....	9
Figure 3 –Surface velocity for Rias Baixas generated modelled by MOHID with 16 sub-domains (19-12-2014 : 0h0m). .....	10
Figure 4 – Sea surface temperature for Rias Baixas generated modelled by MOHID with 16 sub-domains (19-12-2014 : 0h0m).....	11

Figure 5 - Sea surface salinity for Rias Baixas generated modelled by MOHID with 16 sub-domains (19-12-2014 : 0h0m). .....	12
Figure 6 – Division done by the domain decomposition of Rias Baixas model in 32 sub-domains : left panel - using the default procedure done automatically; right panel – division define by the user via an input file. ....	13
Figure 7 - Speedup factor of MOHID schematic 80 million computational cells implementation in the CESGA-SVG Haswell nodes architecture. ....	15
Figure 8 – Characteristic of the processors used by standard pc (middle column) and CESGA-SVG Sandy Bridge nodes (source : ceeserver.cce.cornell.edu). ....	18
Figure 9 - Run period of 3 days forecast function of the number of threads used: red line - CESGA-SVG Sandy Bridge nodes; green line - standard pc with GulfTown microarchitecture. ....	19
Figure 10 – Speedup factor of MOHID Rias Baixas implementation: red line - CESGA-SVG Sandy Bridge nodes; green line - standard pc with GulfTown microarchitecture. ....	20

## 1 Introduction

In this document the implementation of MOHID in parallel mode to the Rias Baixas area by Hidromod in the framework of the RaiaCO project is described. The work was sub-contracted by MeteoGalicia, one of RaiaCO partners.

This MOHID Rias Baixas implementation main objective was to simulate 3 days of 3D circulation for the entire Rias Baixas area with a spatial resolution of ~300 m in ~2h30 m.

Additionally scalability tests of MOHID run in parallel mode for a schematic case were done. The goal was to test the behaviour of MOHID run in baroclinic mode in a grid with 80 million computation cells (20.000x400x10) run in parallel using 384 cores (or threads).

## 2 Implementation of the MOHID Rias Baixas Model

In this chapter the characteristics of the MOHID project implemented for Rias Baixas is presented. Additionally the performance of MOHID run in parallel mode is also analysed in detail.

### 2.1 MOHID Rias Baixas project

The MOHID project implemented for Rias Baixas has the follow characteristics:

- 305 (lines) x232 (columns) horizontal cells is with a spatial resolution of ~300 m. A total of 40 layers (z-level). No nesting;
- Run in parallel mode using MPI and a domain decomposition approach;
- The open boundary and initial conditions are imposed using the ROMS (initial and open boundary condition) and WRF (surface boundary condition) operational solutions computed by MeteoGalicia;
- Boundary and initial conditions are interpolated in run time. The ROMS and WRF are converted from NetCDF to HDF5 with no interpolation;
- Input and output files are all in ASCII or HDF5 formats. No binary formats are used;
- The project is prepared to be run directly in windows and linux. The user can copy the project to any computer (in linux or windows) and execute the respective executable (MohidWater\_r4\_ windows.exe or MohidWater\_r4\_ linux.exe);

- MOHID output files are write separately for each sub-domain due to the domain decomposition procedure. In the end of the run the output files are merged using a tool developed by Ricardo Miranda (IST). The tool was developed using an Actor Model paradigm<sup>1</sup>. This tool uses two executables: DDCParcer and DDCWorker.

The MOHID Rias Baixas project is available in the follow svn site:

[https://hidromod.com:8443/svn/mohid\\_hidromod/Samples/meteogalicia\\_parallel](https://hidromod.com:8443/svn/mohid_hidromod/Samples/meteogalicia_parallel)

The project is ready to be run (windows or linux) from 16<sup>th</sup> to 19<sup>th</sup> of December 2014 in parallel mode using 3 threads.

## 2.2 Model speedup using parallelization

The Rias Baixas MOHID implementation was run in CESGA<sup>2</sup> - SVG<sup>3</sup> Sandy Bridge<sup>4</sup> nodes using several threads. It was possible using 16 and 32 to simulate 3 days of circulation in 2h33 m and 1h21m respectively. In conclusion, it is possible to run 3 days forecasts with Rias Baixas MOHID implementation in CESGA-SVG Sandy Bridge nodes in less than 2h30m (target request by MeteoGalicia see Figure 1).

---

<sup>1</sup> [http://www.mohid.com/wiki/index.php?title=Actor\\_Model](http://www.mohid.com/wiki/index.php?title=Actor_Model)

<sup>2</sup> The Supercomputing Centre of Galicia (CESGA) is the centre of computing, high performance communications systems, and advanced services of the Galician Scientific Community, the University academic system, and the National Scientific Research Council (CSIC).

<sup>3</sup> SVG - Supercomputador Virtual Gallego

<sup>4</sup> Sandy Bridge is the codename for a microarchitecture developed by Intel beginning in 2005 for central processing units in computers to replace the Nehalem microarchitecture.

For this specific MOHID implementation and hardware architecture it was possible with 32 cores to achieve a speedup factor of 15 (Figure 2).

Originally MOHID was developed to be run in standard personal computers. In *Annex – Standard PC (GulfTown) vs CESGA (SandyBridge)* is done a comparison between running the MOHID Rias Baixas in CESGA-SVG Sandy Bridge nodes and in a standard pc (windows operative system) with GulfTown processors.




Figure 1 – Time necessary to simulate 3 days function of the number of threads used.



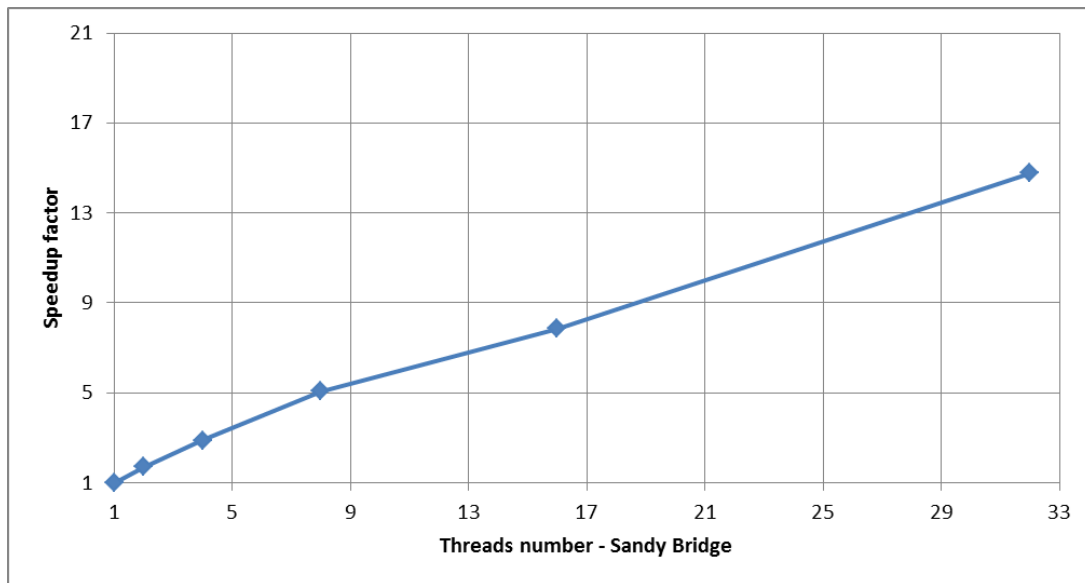


Figure 2 – Speedup factor of MOHID Rias Baixas implementation in the CESGA-SVG Sandy Bridge nodes architecture.

### 2.2.1 Model surface results

In the follow figures the results of MOHID surface results using 16 threads (or sub-domains) for the last instant of the 3 days forecast simulation mention above are presented:

- Surface velocity (Figure 3);
- Sea surface temperature (Figure 4);
- Sea surface salinity (Figure 5).

These results show in a clear way the strong interaction between the several Rias.

### Surface Velocity (dx - 300 m) - 40 Layers

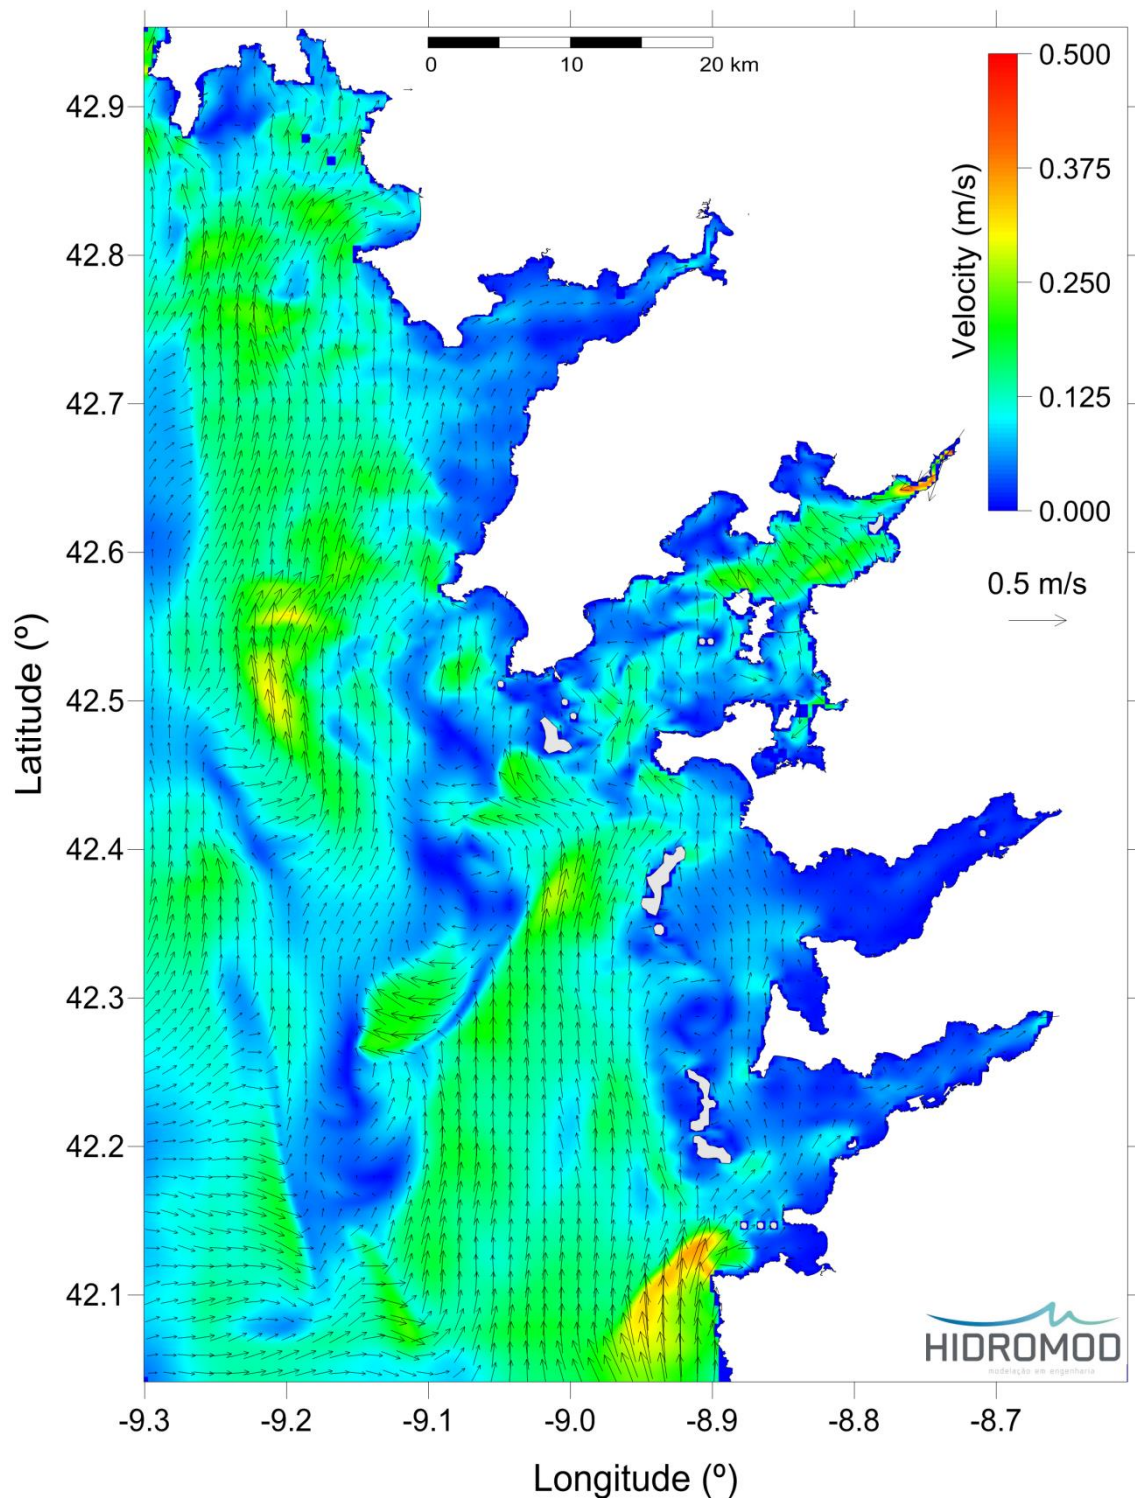


Figure 3 –Surface velocity for Rias Baixas generated modelled by MOHID with 16 sub-domains (19-12-2014 : 0h0m).

## Sea Surface Temperature (dx - 300 m) - 40 Layers

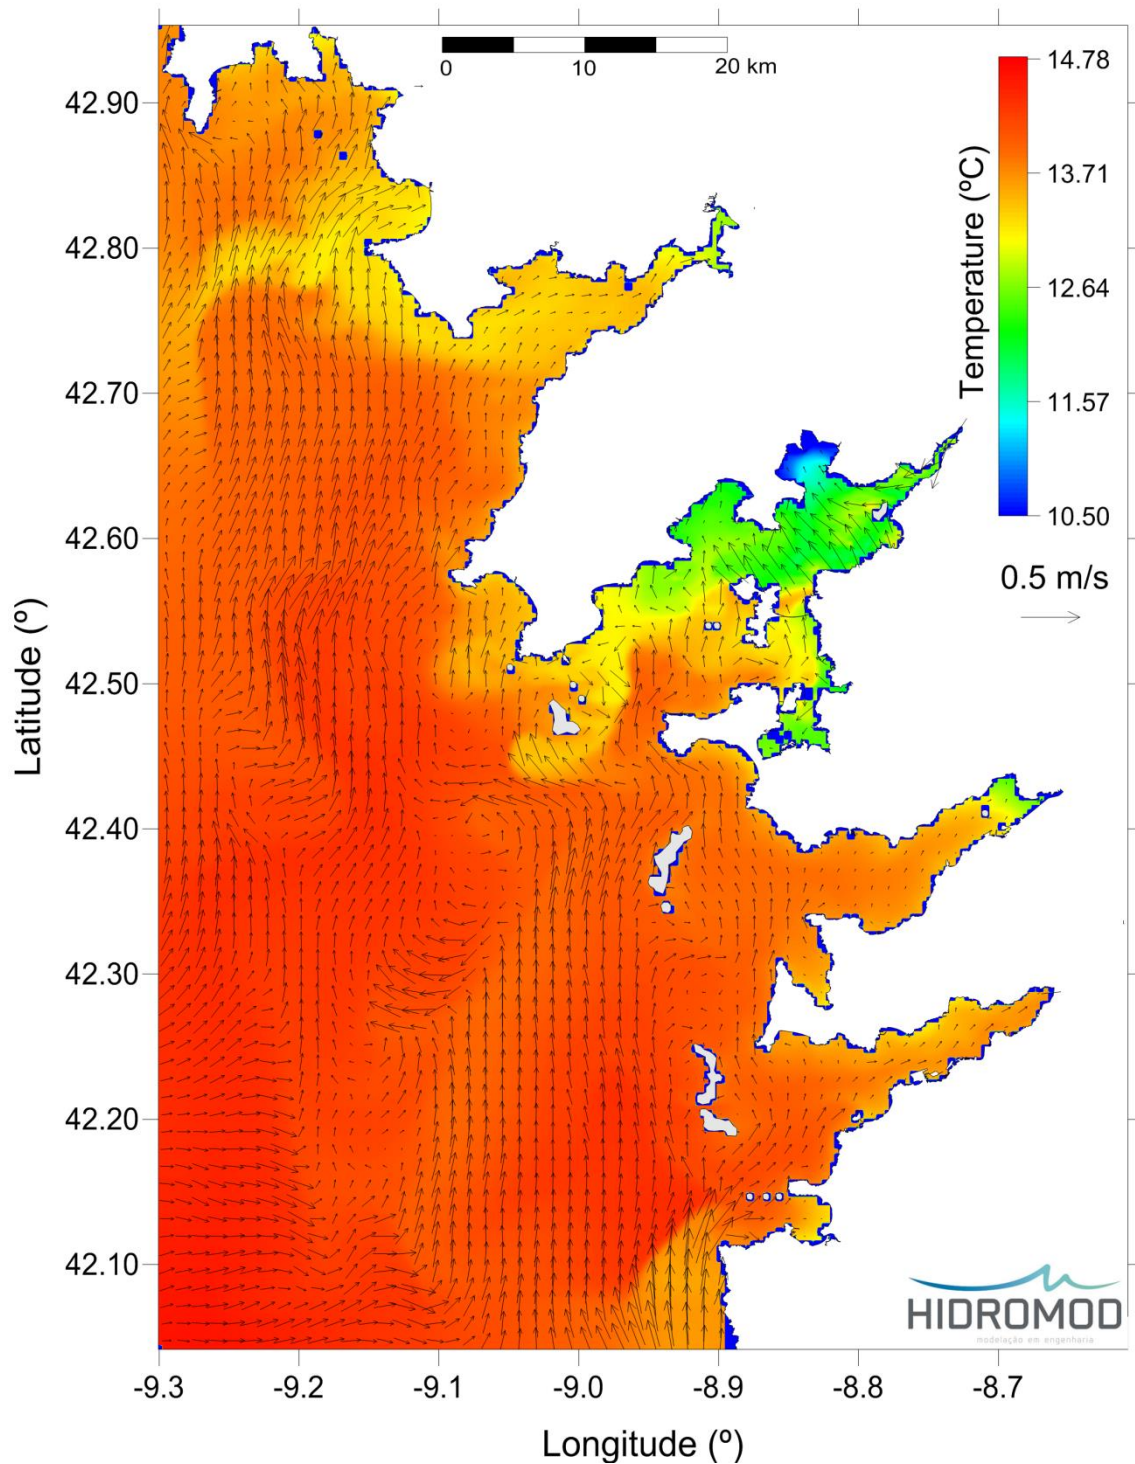


Figure 4 – Sea surface temperature for Rias Baixas generated modelled by MOHID with 16 sub-domains (19-12-2014 : 0h0m).



## Sea Surface Salinity (dx - 300 m) - 40 Layers




Figure 5 - Sea surface salinity for Rias Baixas generated modelled by MOHID with 16 sub-domains (19-12-2014 : 0h0m).

## 2.2.2 Computational effort distribution

By default the MOHID domain decomposition procedure breakdowns the model horizontal grid assuming sub-domains with number of columns equal to the original domain (see Figure 6 – left panel). However, this is a simplistic way of distributing the computation effort. For the 32 threads test case it was possible to improve the speedup factor 31% (from 11.3 to 14.8) implementing a domain division that aims to have approximately the same number of water points in each sub-domain. In this MOHID version this optimization is not done automatically need to be specified by the user. In the future an algorithm able to automatically define optimal sub-domains from a computational efficiency point of view will be developed.

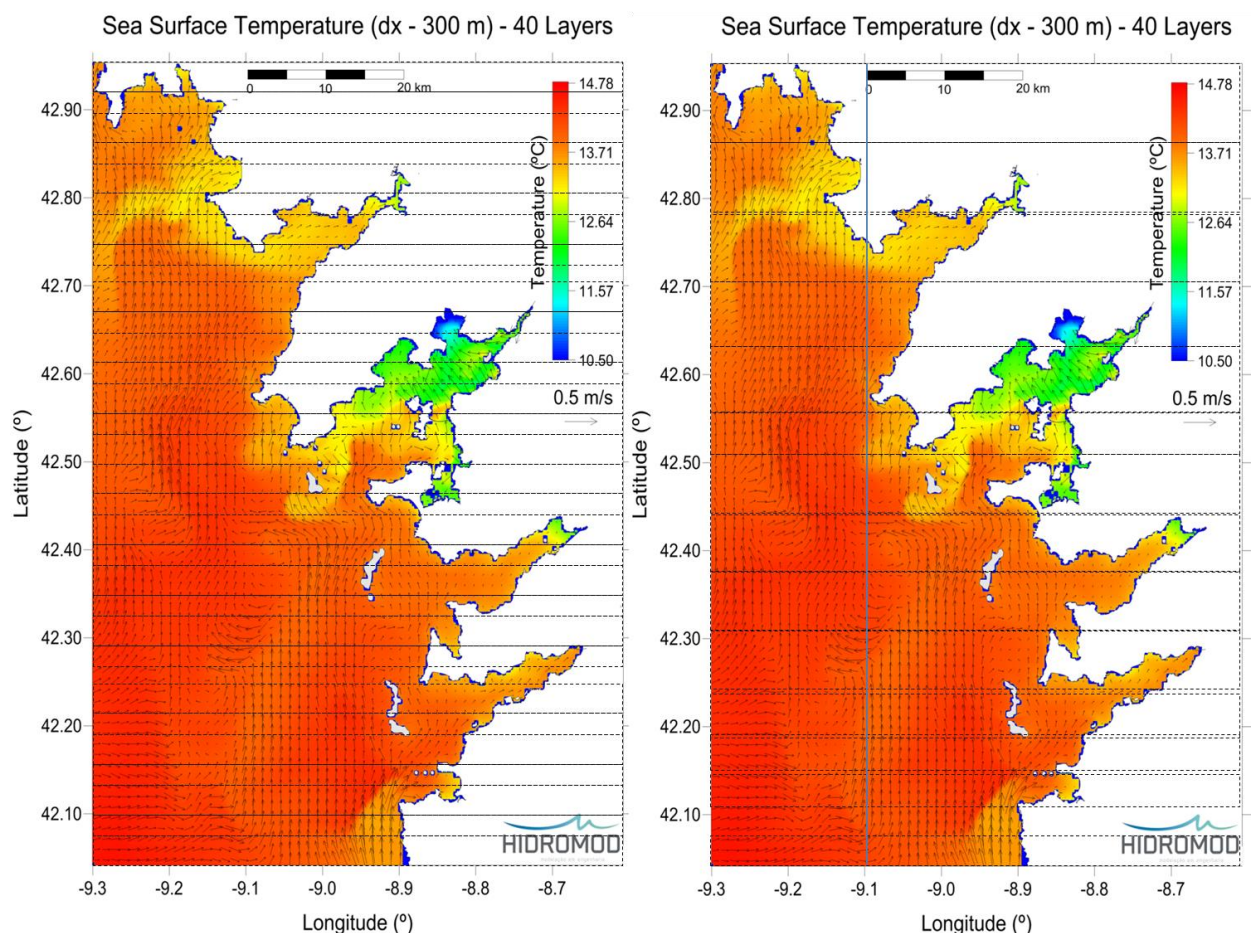


Figure 6 – Division done by the domain decomposition of Rias Baixas model in 32 sub-domains : left panel - using the default procedure done automatically; right panel –division define by the user via an input file.

### 3 MOHID parallelization scalability

CESGA-SVG was upgrade with 18 new servers with 24 cores and 64GB of memory per servers, with Infiniband network and focused for parallel workloads. This server use Haswell processors (more modern than the Sandy Bridge ones).

Together with MeteoGalicia technicians has done in these new servers a test with the goal of check MOHID performance using hundreds of threads.

A schematic MOHID project with the follow characteristics was implemented:

- Schematic bathymetry with a constant depth of 100 m with no land points;
- 20.000 (lines) x400 (columns) horizontal cells with a spatial resolution of ~100 m. Vertically were used 10 layers of equal thickness. A total number of 80 million computational cells. No nesting;
- Run in parallel mode using MPI and a domain decomposition approach;
- No forcing and constant salinity and temperature in space;
- Compute 3D baroclinic hydrodynamic, transport of salt and temperature and  $k-\epsilon$  turbulence model. The model final result is a null velocity field because there is no forcing and density gradients are null;
- Input and output files are all in ASCII or HDF5 formats. No binary formats are used;
- No transient output files are write only the final outputs,
- This implementation allocates in RAM memory ~60 Gb.

The schematic MOHID project used is available in [https://hidromod.com:8443/svn/mohid\\_hidromod/Samples/CesgaExercise](https://hidromod.com:8443/svn/mohid_hidromod/Samples/CesgaExercise).

The speedup factor was computed for 24, 32, 48, 64, 96, 128, 192, 256, 312 and 384 threads. The speedup factor was computed not considering the input and output procedures. It was not possible to run the model with less than 24 threads due to hardware limitations. For 24, 32, 48 and 64 threads the model performance was equal to the theoretical limit, in other words doubling the number of threads was possible to double the speedup factor (Figure 10). The maximum speed factor of 192 was obtained with 256 threads. Above this number of threads the speedup factor stabilizes and can be assumed constant.




Figure 7 - Speedup factor of MOHID schematic 80 million computational cells implementation in the CESGA-SVG Haswell nodes architecture.

## 4 Conclusions

A MOHID model for the entire area of Rias Baixas with a resolution  $\sim 300$  m with 40 layers (z-level) was implanted. The model is ready to be run in operational mode forced by the ROMS and WRF models run operationally by MeteoGalicia. The model can be run in parallel mode using MPI and a domain decomposition approach. Computational efficiency tests were done for the CESGA-SVG Sandy Bridge architecture. It was possible to run 3 days forecast in 2h33 m and 1h21m using 16 and 32 threads respectively.

MOHID Rias Baixas implementation computational efficiency was also tested in a standard pc with a windows operative system. A similar performance was achieved until 4 threads. However, for 8 threads the CESGA-SVG Sandy Bridge showed a speedup factor 50 % higher than the standard pc.

Together with MeteoGalicia technicians, using new servers (Haswell processors), an test with the goal of check MOHID performance using hundreds of threads has done. A schematic MOHID 3D baroclinic project with 80 million cells was tested. A maximum speedup factor of 192 was achieved using 256 threads.



## 5 Future steps

The proposed future steps are:

- Develop an algorithm able to automatically define optimal sub-domains from the computational efficiency point of view;
- Run MOHID Rias Baixas with biogeochemical properties connected (e.g. nutrients, chlorophyll, zooplankton, oxygen);
- Extended the MOHID Rias Baixas to the entire Galician coast increasing the number of horizontal grid cells one order of magnitude;
- Develop an AquaSafe service where MOHID users could submit to CESGA parallel runs.

## 6 Annex – Standard PC (GulfTown) vs CESGA (SandyBridge)

This annex compares the computational efficiency of running the MOHID Rias Baixas in CESGA-SVG Sandy Bridge nodes and in a standard pc with GulfTown processors (windows operative system and with a maximum of 12 threads see Figure 8). The speedup factor was compute for 1, 4 and 8 threads in a standard pc. For 1 and 4 threads booth architectures had the same performance (Figure 9 and Figure 10). However for 8 threads the CESGa-SVG Sandy Bridge nodes (speedup factor of 5.1) had a performance 50% better than the standard pc (speedup factor of 3.4).

Product Name	Intel® Core™ i7-980X Processor Extreme Edition (12M Cache, 3.33 GHz, 6.40 GT/s Intel® QPI)	Intel® Xeon® Processor E5-2670 (20M Cache, 2.60 GHz, 8.00 GT/s Intel® QPI)
Code Name	Gulftown	Sandy Bridge EP
<b>Essentials</b>		
Status	End of Interactive Support	Launched
Launch Date	Q1'10	Q1'12
Processor Number	i7-980X	E5-2670
Cache	12 MB Intel® Smart Cache	20 MB Intel® Smart Cache
Bus Type	QPI	QPI
System Bus	6.4 GT/s	8 GT/s
# of QPI Links	1	2
Instruction Set	64-bit	64-bit
Instruction Set Extensions	SSE4.2	AVX
Embedded Options Available	No	No
Lithography	32 nm	32 nm
Scalability		2S Only
VID Voltage Range	0.800V-1.375V	0.60V-1.35V
Recommended Customer Price	BOX : \$1059.00	BOX : \$1556.00 TRAY : \$1552.00
Datasheet	<a href="#">Link</a>	<a href="#">Link</a>
<b>Performance</b>		
# of Cores	6	8
# of Threads	12	16
Processor Base Frequency	3.33 GHz	2.6 GHz
Max Turbo Frequency	3.6 GHz	3.3 GHz
TDP	130 W	115 W

Figure 8 – Characteristic of the processors used by standard pc (middle column) and CESGA-SVG Sandy Bridge nodes (source : ceeserver.cee.cornell.edu).




Figure 9 - Run period of 3 days forecast function of the number of threads used: red line - CESGA-SVG Sandy Bridge nodes; green line - standard pc with GulfTown microarchitecture.




Figure 10 – Speedup factor of MOHID Rias Baixas implementation: red line - CESGA-SVG Sandy Bridge nodes; green line - standard pc with GulfTown microarchitecture.