www.bsc.es

Barcelona , 17 April 2018

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Performance activities in Earth Science Department

Mario C. Acosta & Miguel Castrillo

Outline

- Introduction
- Next generation of HPC Challenges
- Earth System Models: EC-Earth
- Methodology
- Some successful examples
- Discussion

Introduction

Introduction

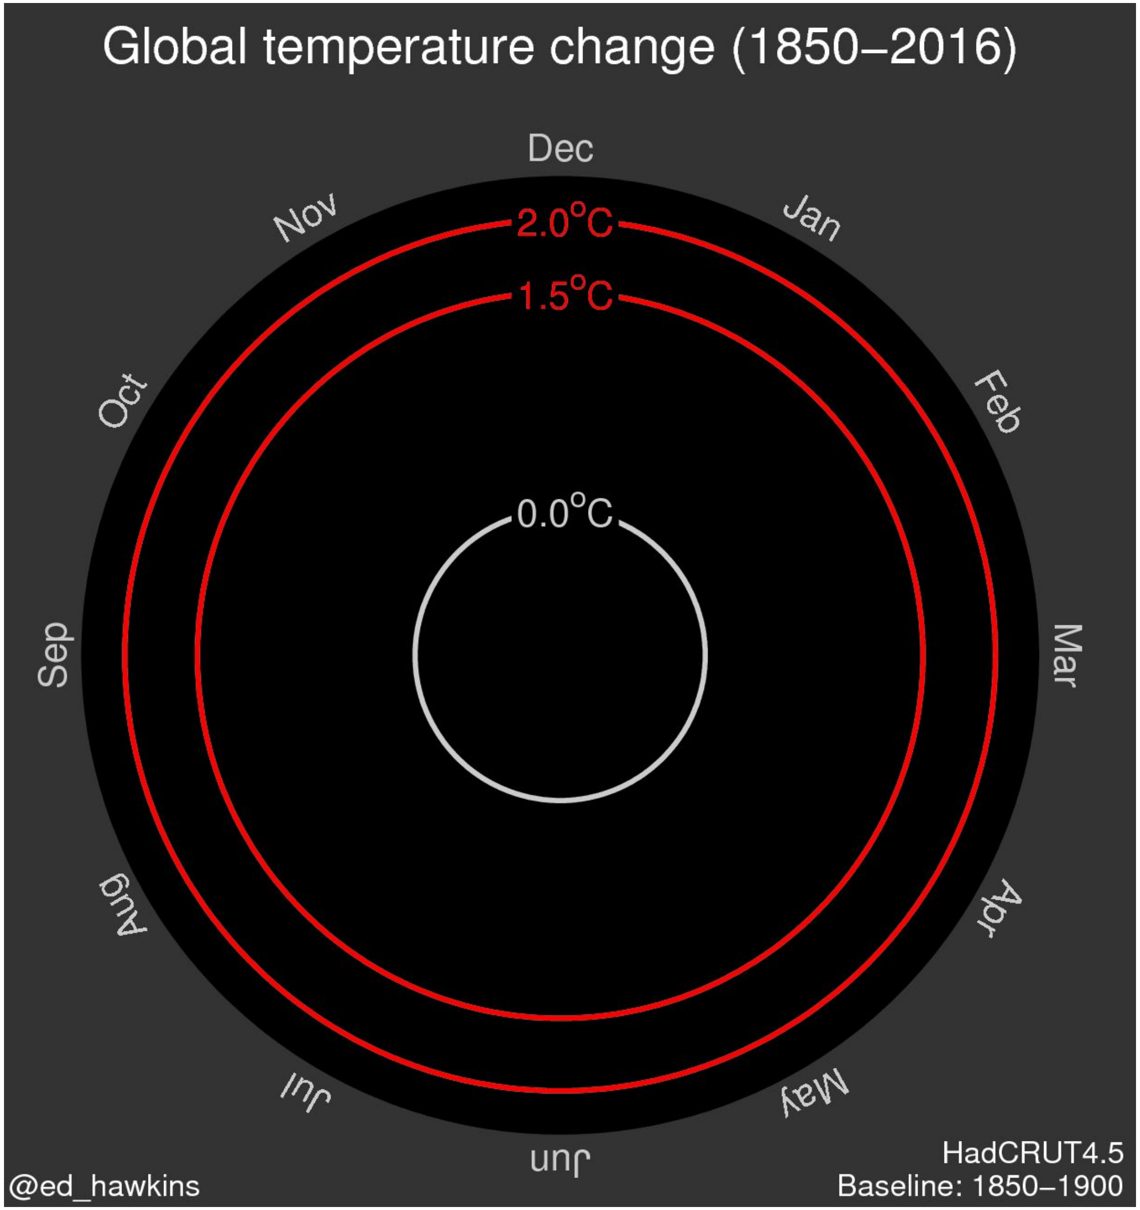

Introduction

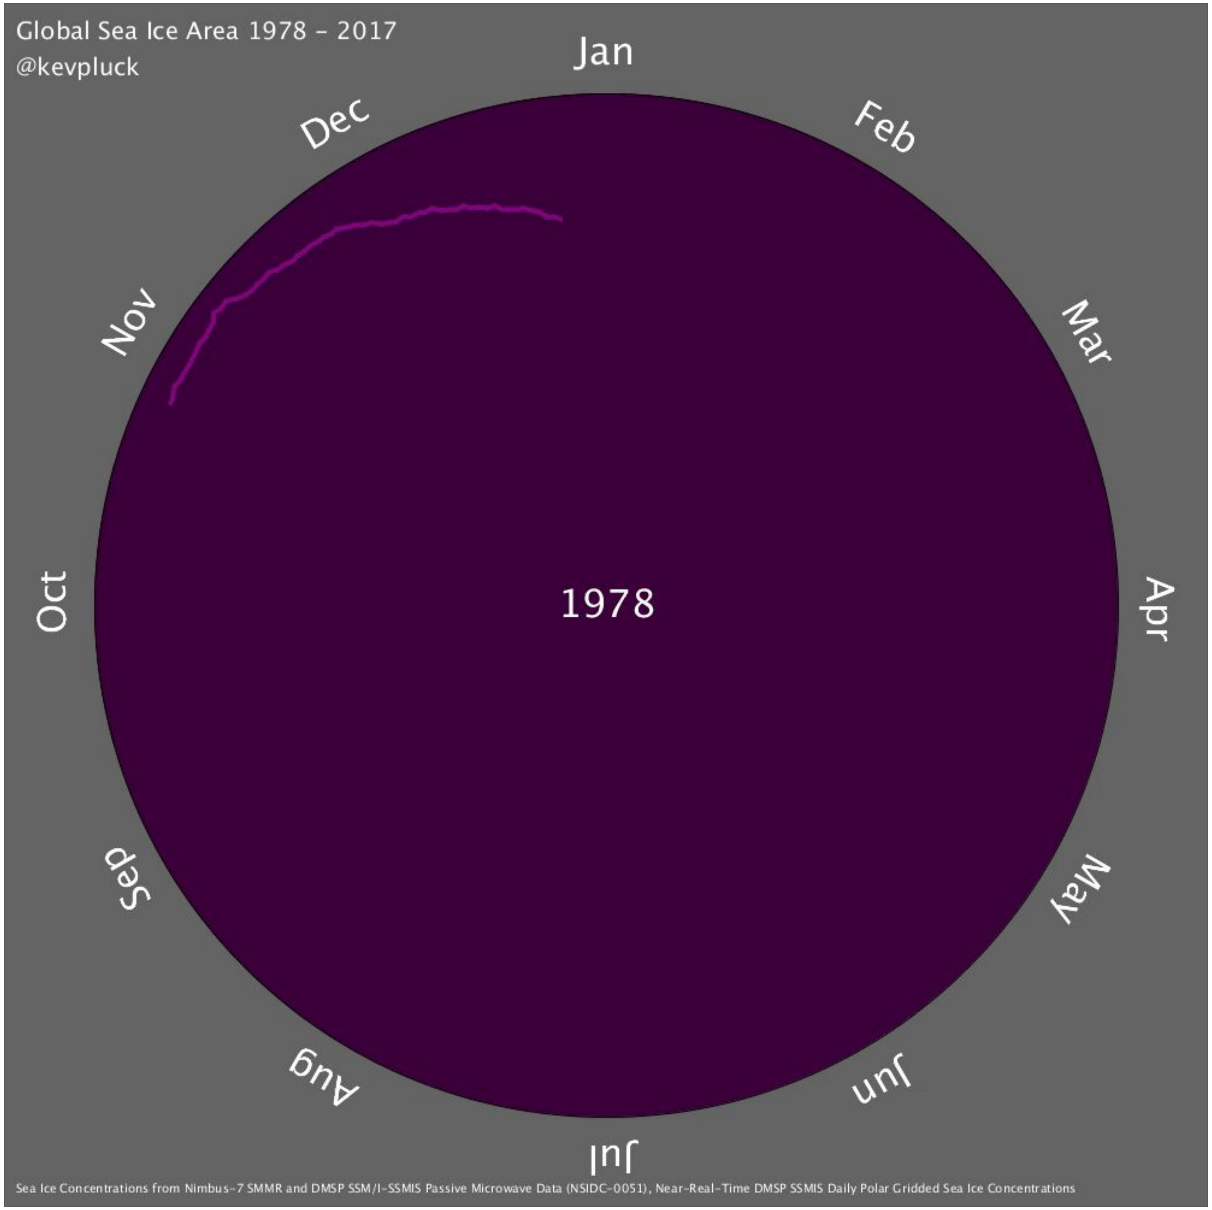

- Climate is changing.
- The potential impact of that change is huge.
- The effects would be **catastrophic** for humankind.

IOP Publishing Environ. Res. Lett. 11 (2016) 048002 doi:10.1088/1748-9326/11/4/048002

Environmental Research Letters

OPEN ACCESS

RECEIVED 28 April 2015
REVISED 27 November 2015
ACCEPTED FOR PUBLICATION 1 March 2016
PUBLISHED 13 April 2016

REPLY

Consensus on consensus: a synthesis of consensus estimates on human-caused global warming

John Cook[1,2,3,16], Naomi Oreskes[4], Peter T Doran[5], William R L Anderegg[6,7], Bart Verheggen[8], Ed W Maibach[9], J Stuart Carlton[10], Stephan Lewandowsky[11,2], Andrew G Skuce[12,3], Sarah A Green[13], Dana Nuccitelli[3], Peter Jacobs[9], Mark Richardson[14], Bärbel Winkler[3], Rob Painting[3] and Ken Rice[15]

Joint science academies' statement: Global response to climate change

Climate change is real

There will always be uncertainty in understanding a system as complex as the world's climate. However there is now strong evidence that significant global warming is occurring[1]. The evidence comes from direct measurements of rising surface air temperatures and subsurface ocean temperatures and from phenomena such as increases in average global sea levels, retreating glaciers, and changes to many physical and biological systems. It is likely that most of the warming in recent decades can be attributed to human activities (IPCC 2001)[2]. This warming has already led to changes in the Earth's climate.

The existence of greenhouse gases in the atmosphere is vital to life on Earth – in their absence average temperatures would be about 30 centigrade degrees than they are today. But human activities are now cau atmospheric concentrations of greenhouse gases – including carbon dioxide, methane, tropospheric ozon and nitrous oxide – to rise well above pre-industrial le Carbon dioxide levels have increased from 280 ppm i 1750 to over 375 ppm today – higher than any previo levels that can be reliably measured (i.e. in the last 42 years). Increasing greenhouse gases are causing temperatures to rise; the Earth's surface warmed by approximately 0.6 centigrade degrees over the twenti century. The Intergovernmental Panel on Climate Cha (IPCC) projected that the average global surface temperatures will continue to increase to between 1.4 centigrade degrees and 5.8 centigrade degrees above levels, by 2100.

Reduce the causes of climate change

The scientific understanding of climate change is now sufficiently clear to justify nations taking prompt actio is vital that all nations identify cost-effective steps tha can take now, to contribute to substantial and long-te reduction in net global greenhouse gas emissions.

Action taken now to reduce significantly the build-up greenhouse gases in the atmosphere will lessen the magnitude and rate of climate change. As the United Nations Framework Convention on Climate Change (UNFCCC) recognises, a lack of full scientific certainty about some aspects of climate change is not a reason delaying an immediate response that will, at a reason cost, prevent dangerous anthropogenic interference w the climate system.

potentially cost-effective technological options that could contribute to stabilising greenhouse gas concentrations. These are at various stages of research and development. However barriers to their broad deployment still need to be overcome.

Carbon dioxide can remain in the atmosphere for many decades. Even with possible lowered emission rates we will be experiencing the impacts of climate change throughout the 21[st] century and beyond. Failure to implement significant reductions in net greenhouse gas emissions now, will make the job much harder in the future.

Prepare for the consequences of

Next generation of HPC challenges

- Projections
- Impact analysis
- Adaptation to climate change.

Currently, **only computational models have the potential** to provide geographically and physically consistent estimates.

Next generation of HPC challenges

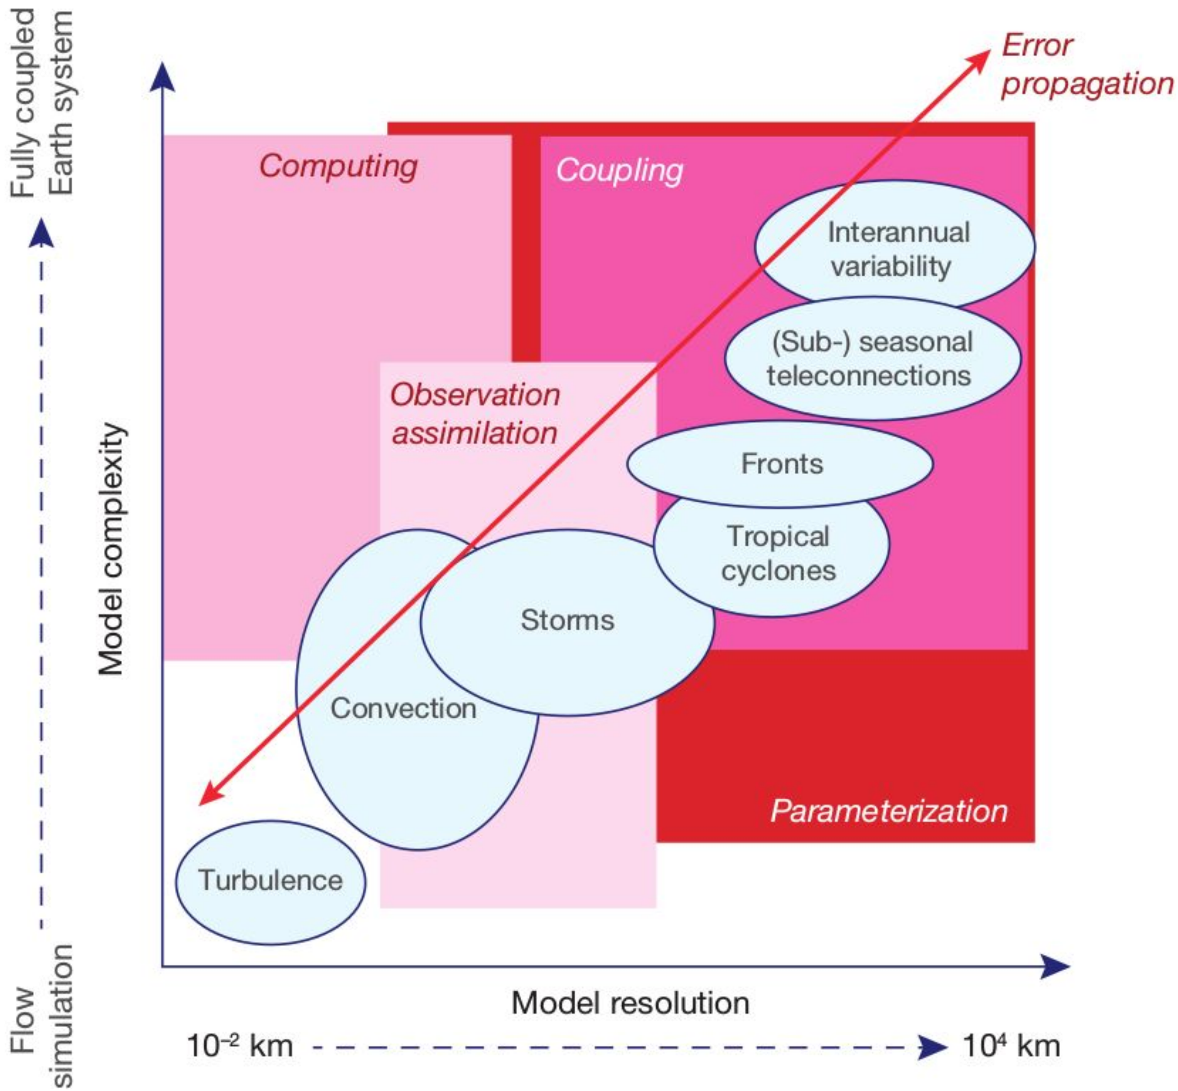

Earth System Models: EC-Earth

EC EARTH

Atmosphere:
IFS
ECMWF

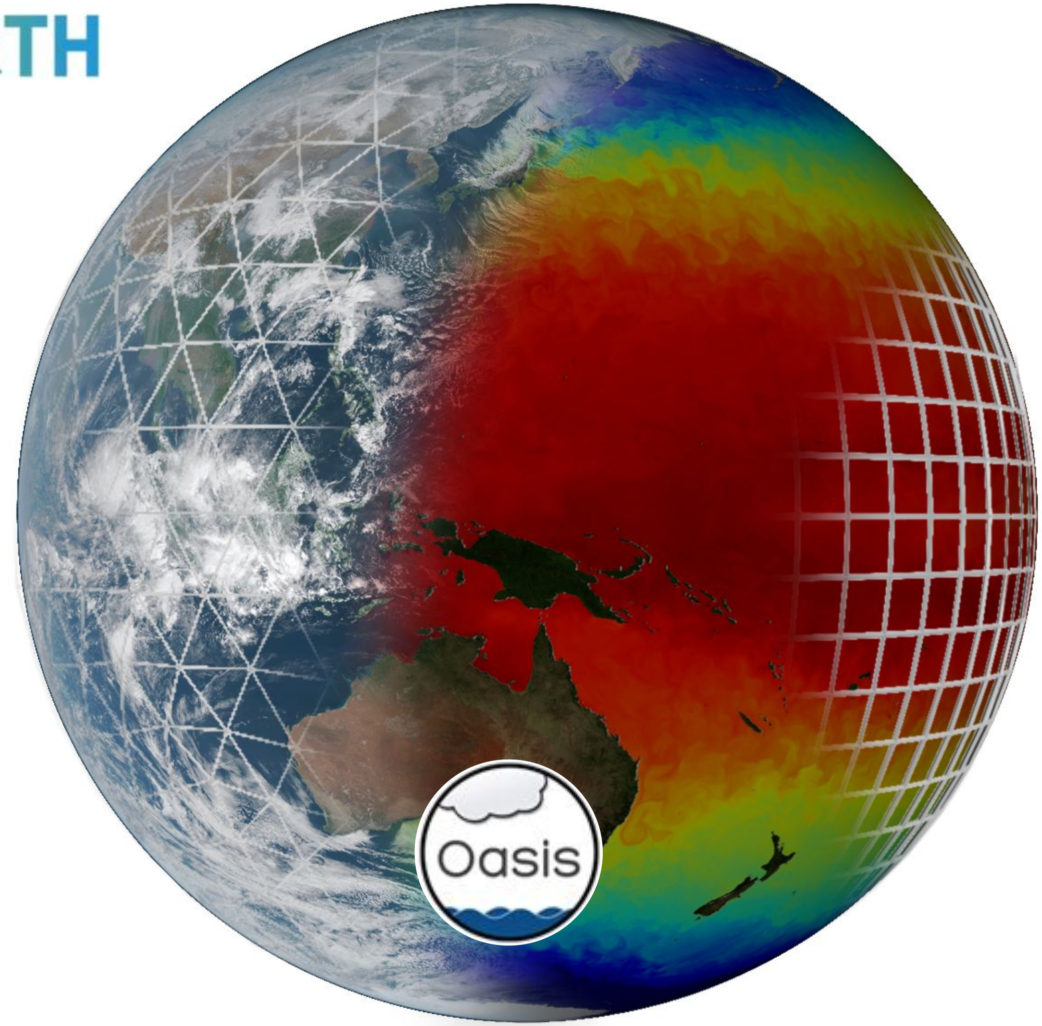

Ocean:
NEMO

Coupler:

MareNostrum IV

- MareNostrum IV in operation since July 2017
- One of the first HPCs featuring new Intel Scalable Processors

	MareNostrum III	MareNostrum IV
Processor	Intel Xeon E5-2670 2.6 GHz	Intel Xeon Platinum 8160 2.1 GHz
#Cores per socket	8	24
#Sockets	2	2
Memory	32Gb DDR3-1600 2 GB/core	96Gb DDR4-2667 2 GB/core
Interconnection	Infiniband FDR10 10Gb	Intel Omni-Path 100Gb

MareNostrum III - 11,15 petaFLOPS

MareNostrum IV - 11,15 petaFLOPS

www.bsc.es

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Methodology

Methodology

- To be able to use the computing power of modern supercomputers, applications must exploit parallelism.
- Parallelism produce overhead (extra computation and communications)
  - We need to evaluate and optimize the parallelism efficiency of our computational models.
    - When the hardware change
    - When the number of resources change
    - When the model complexity increase
    - When the resolution increase
    - …

Methodology

High Performance Computing is an essential part of Weather and Climate models nowadays

- Path to exascale is coming → But is not automatic, the free lunch is over
- Processors can not be faster → But supercomputer parallel processing units can increase
- Compilers work great at low level optimizations → But human decisions in the development will be more and more critical to achieve optimizations
- Overhead does not look a problem in my model → But if the needs increase (i.e. higher resolutions), a bad implementation will be a problem in some point.

Although there are physical limitations, performance work help to achieve the modelization dreams of the community

Methodology

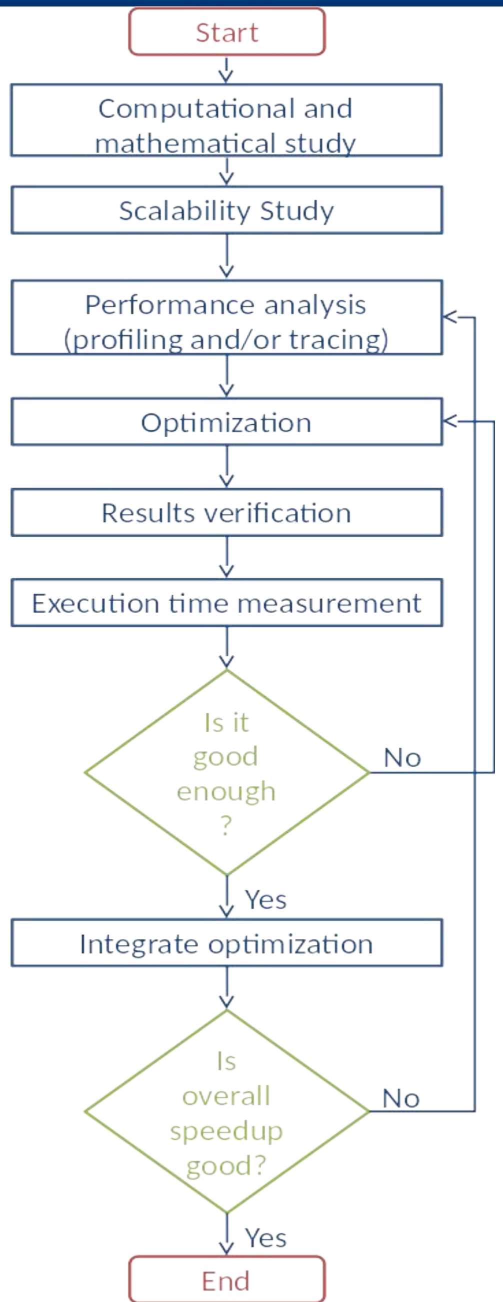

BSC Tools (Profiling Study)

- Since 1991
- Based on traces
- Open Source: http://www.bsc.es/paraver
- **Extrae**: Package that generates Paraver trace-files for a post-mortem analysis
- **Paraver**: Trace visualization and analysis browser
  - Includes trace manipulation: Filter, cut traces
- **Dimemas**: Message passing simulator

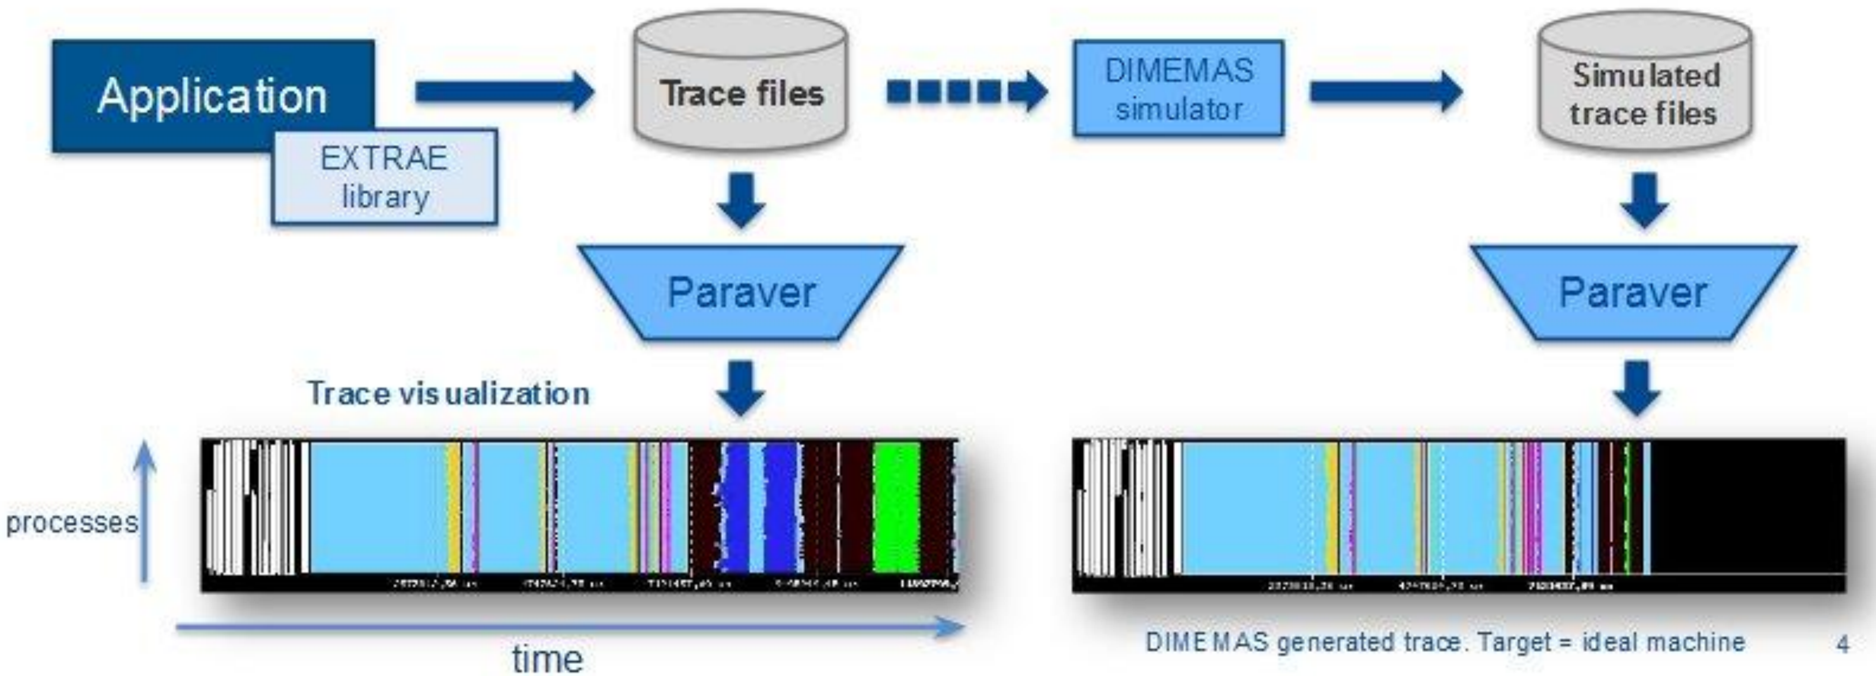

DIMEMAS generated trace. Target = ideal machine

Computational Study (IFS example)

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

- Localize each scientific phase in your model and evaluate it independently

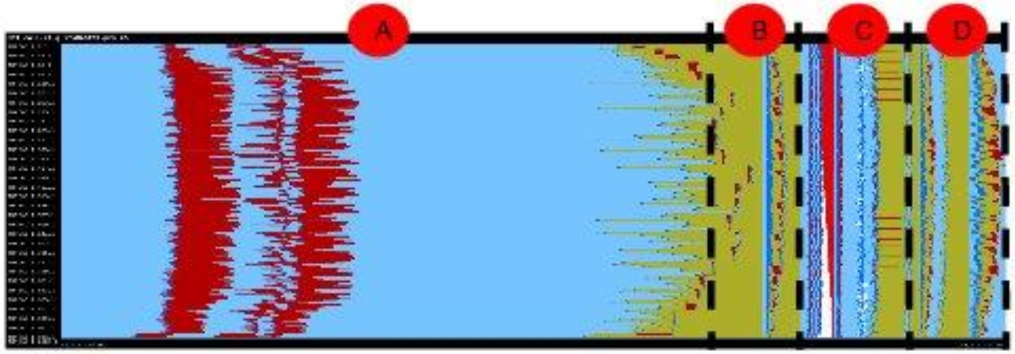

- A Grid Point Calculations
- B Transformations and Transpositions (Fourier + Legendre)
- C Spectral Calculations
- D Fourier + Legendre Inverse

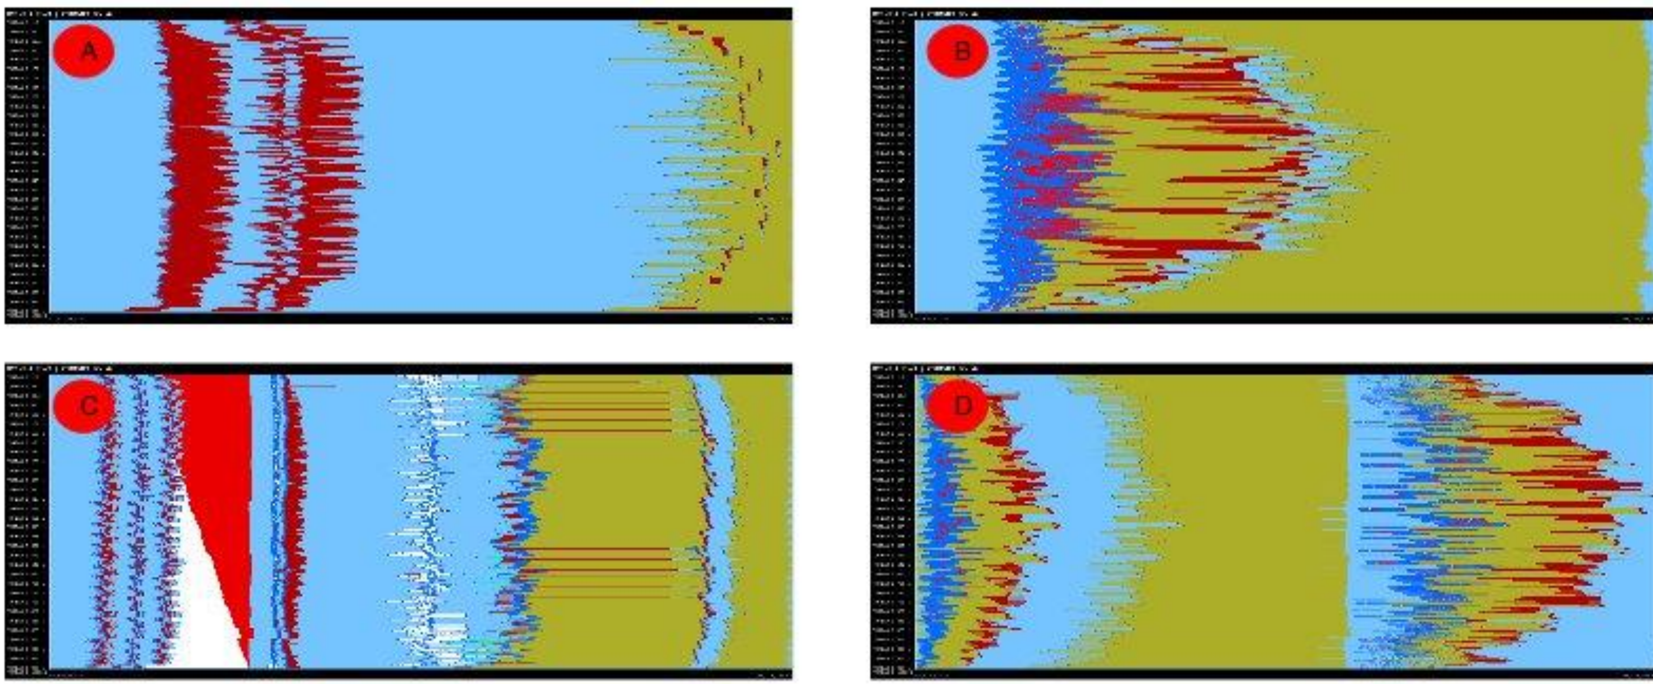

EC-Earth scalability on MN4

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

Scalability for EC-Earth trunk with default output configuration.

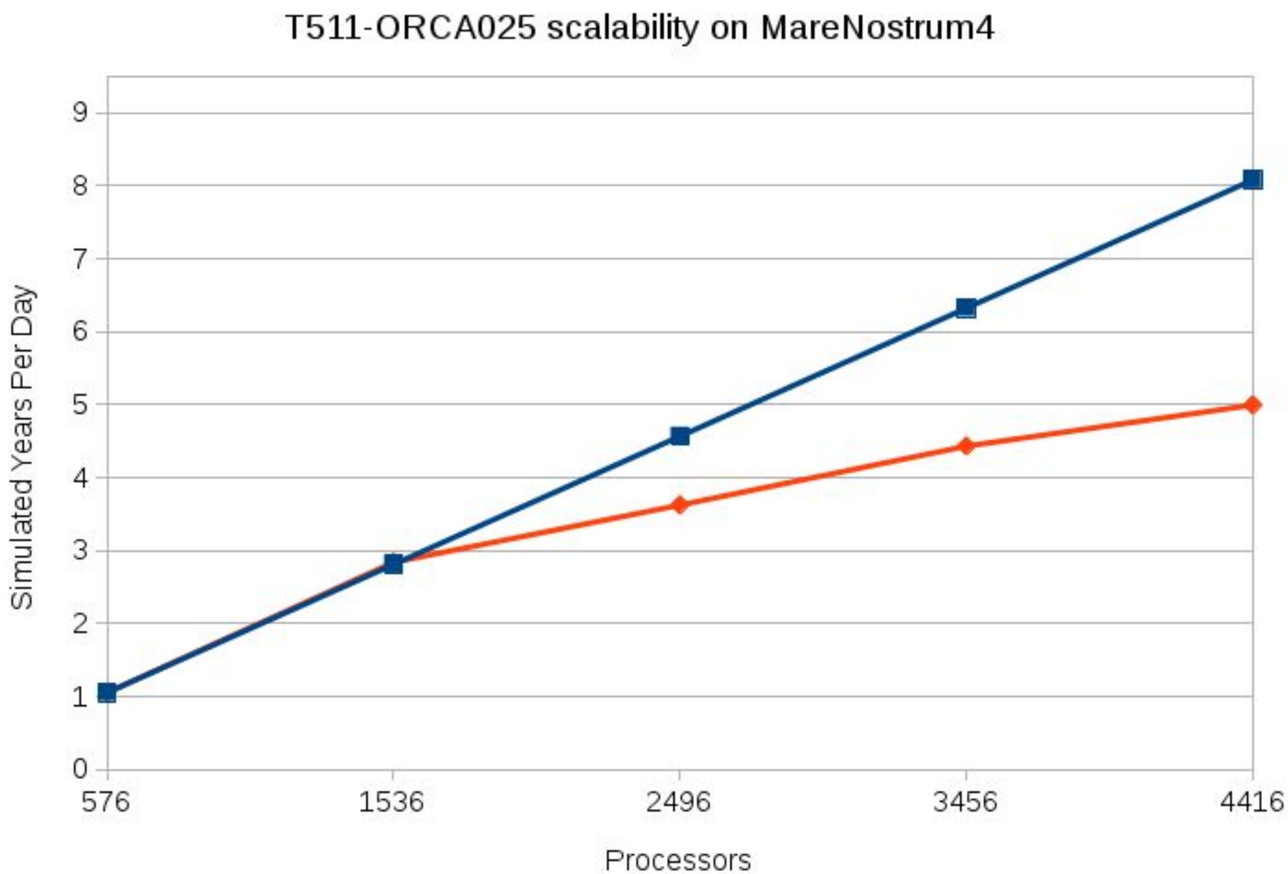

Some successful examples

ELPiN (Exclude Land Processes in NEMO)

- ELPiN allows to find proper namelist parameters to **exclude land-only** processes in NEMO simulations
- Avoids NEMO to waste resources, **speeds up** simulations.
- Included in EC-Earth production branches.

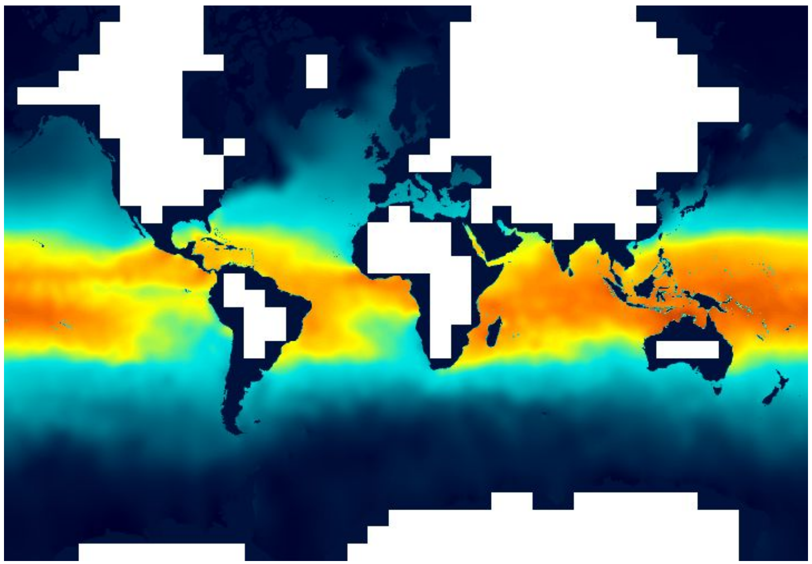

ORCA025 domain decomposed in 1287 sub-domains. 312 are land-only and therefore removed (24% of the total grid).
O. Tintó (BSC)

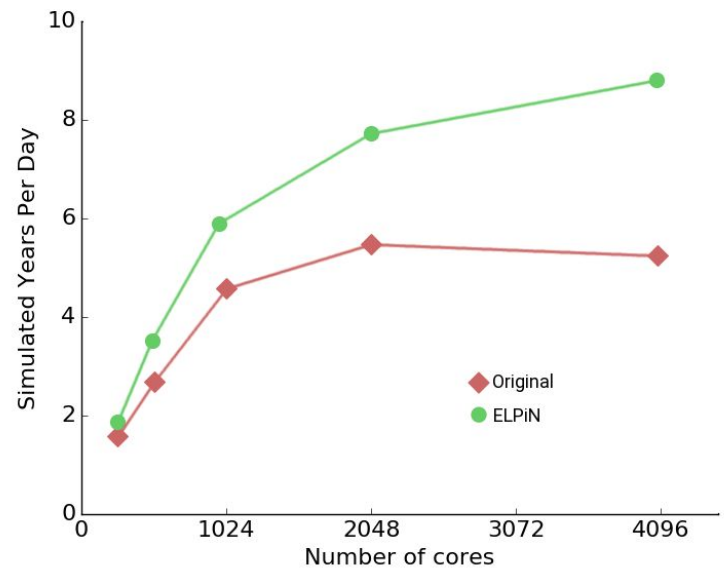

Impact of ELPiN on the NEMO model for an ORCA025-LIM3 simulation.
O. Tintó (BSC)

NEMO performance analysis

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Original code

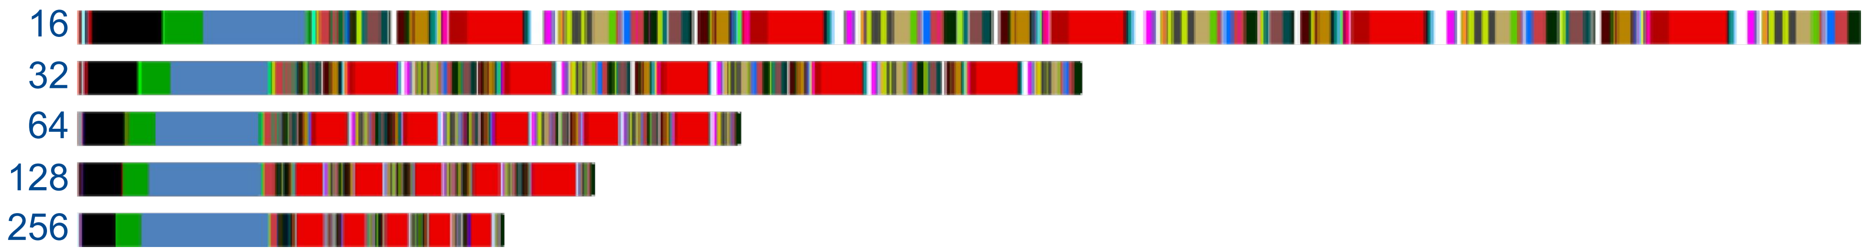

Optimized code

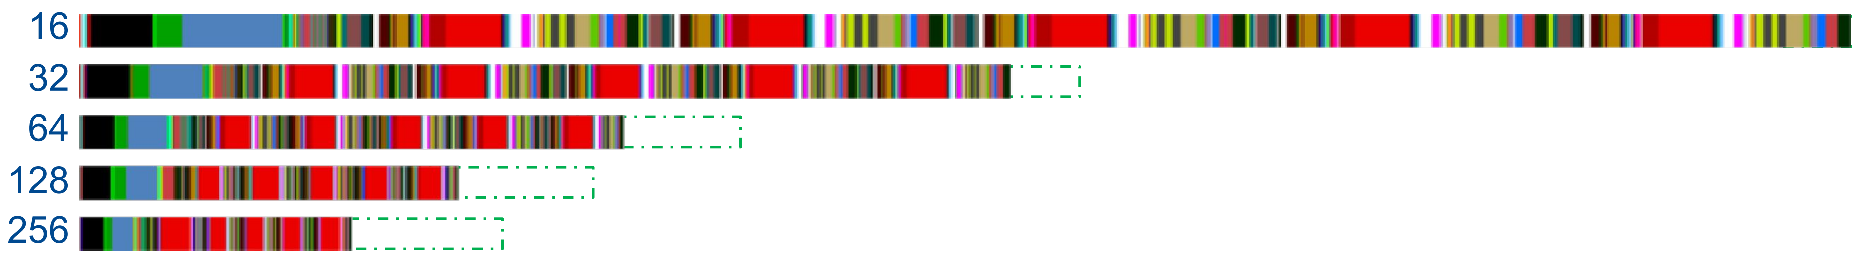

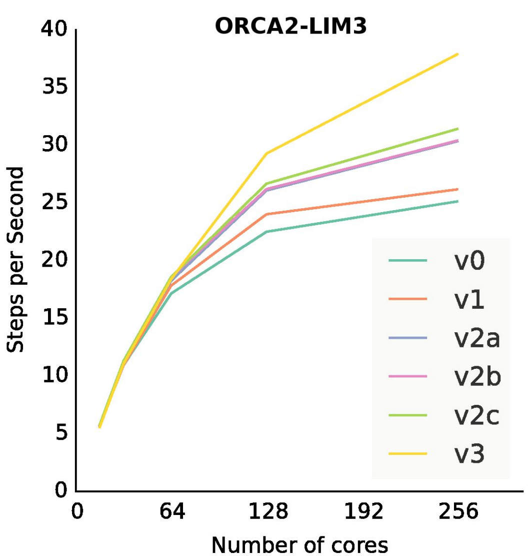

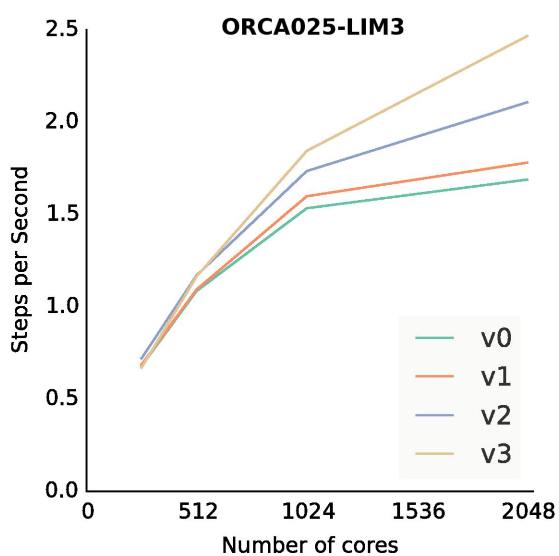

V0 → Original
V1 → Message packing
V2 → Conv. Check reduction
V3 → Reordering

Check method for IFS

- Synchronous point to point communication could be a bottleneck even for only one message from one master to hundreds of slaves
    - Sigcheck method

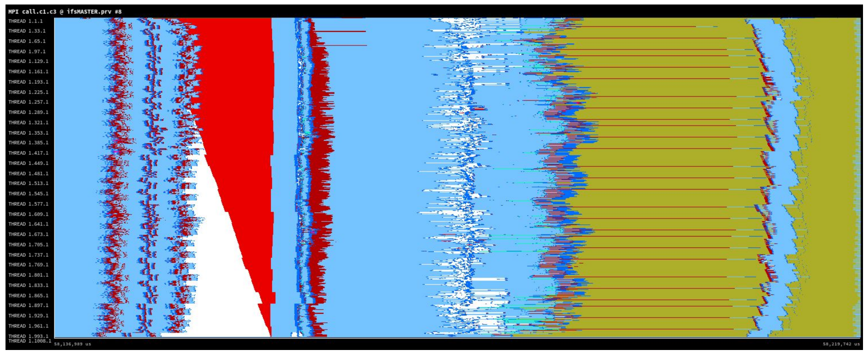

MPI call profile @ ifsMASTER.prv #8

89 (sigcheck.F90, ifsMASTER)	95 (sigcheck.F90, ifsMASTER)	111 (sigcheck.F90, ifsMASTER)
0.81 %	442.74 %	806.93 %
0.81 %	0.44 %	0.80 %
0.81 %	1.73 %	1.84 %
0.81 %	0.00 %	0.00 %
0 %	0.45 %	0.41 %
1	0.25	0.44

- Using one asynchronous collective communication this time is reduced almost to 0

Bottleneck of the coupler OASIS

- First studies showed that IFS-NEMO coupling was not a big issue

IFS: 128 cores

NEMO: 128 cores

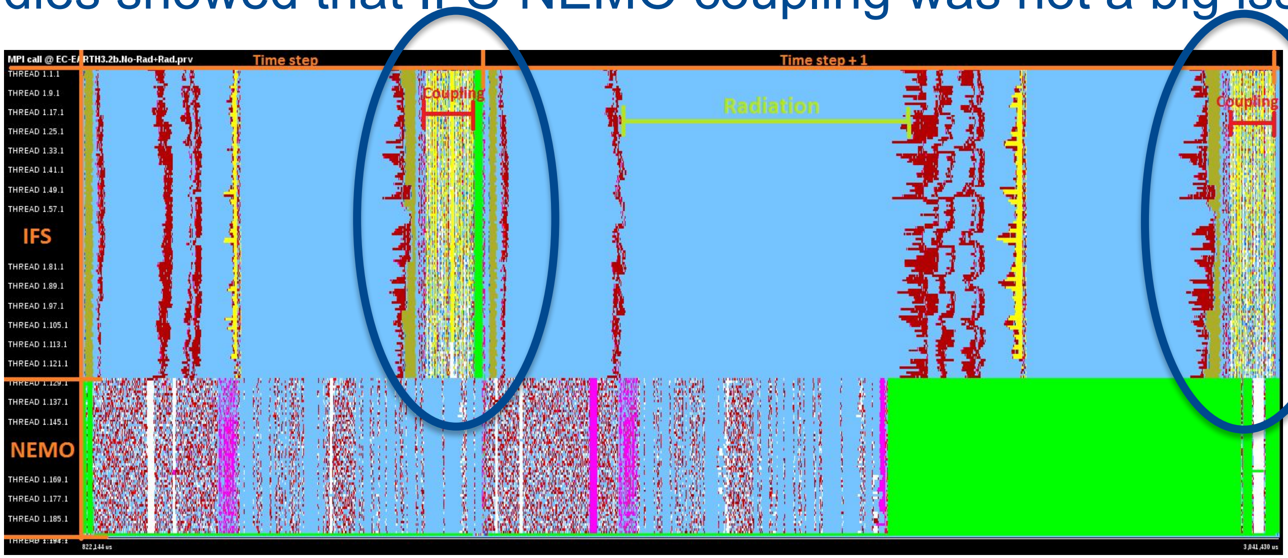

TIME

Coupling time is increased

- But it seems that it is when increasing number of cores

IFS: 512 cores

NEMO: 128 cores

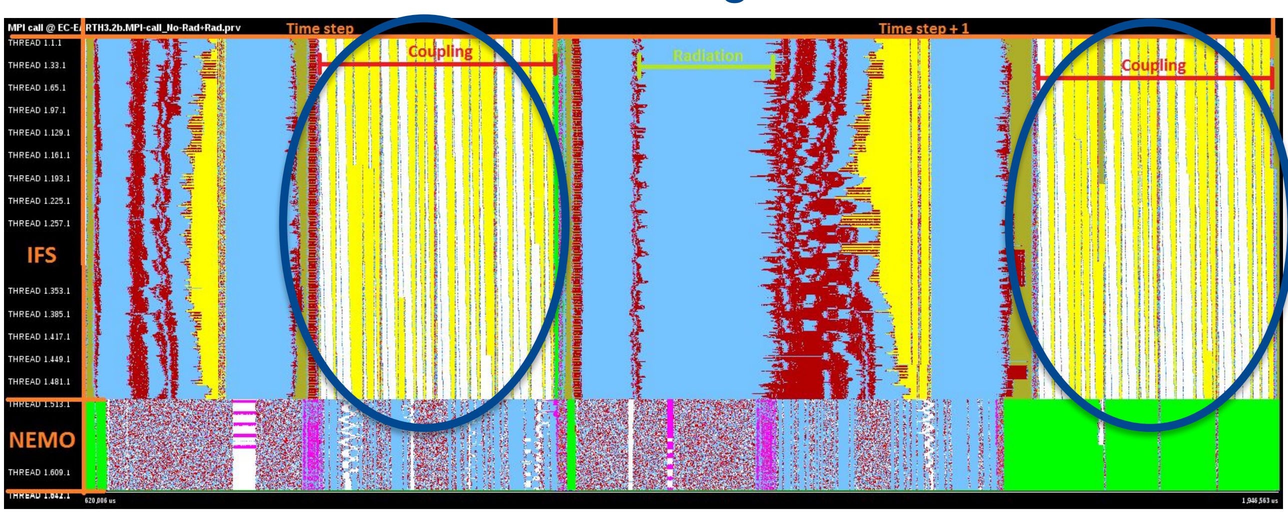

Discussion

- There are some HPC challenges in the near future for Earth System Modelling
  - **is it similar for other fields?**
  - **Is it important for you the optimization of the models or a machine with more computational power is enough?**

- There are several approaches to improve your application…

Optimization or revolution?

  - Optimizations → Are you interested in our work? Could be helpful for you?
  - Revolution → New approaches to solve the same problem (DSLs, Accelerators for heterogenous computation, new mathematical solutions...)
- Which one is better? A combination of both?

www.bsc.es

Thank you!

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

For further information please **contact to the Performance Team**

Mario Acosta: mario.acosta@bsc.es
Miguel Castrillo: miguel.castrillo@bsc.es
Oriol Tintó: oriol.tinto@bsc.es
Xavier Yepes: xavier.yepes@bsc.es