www.bsc.es

Toulouse, 07 April 2016

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Reproducibility of Earth System Models
A computational Point of View

M. C. Acosta, X. Yepes-Arbós, F. Massonnet, M. Ménégoz, K. Serradell, O. Mulla-Valls and F. Doblas-Reyes

Introduction

Introduction

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

- Climate science has a strong computational component, and the climate codes used in this discipline are typically complex and large in size.
- These models can support a variety of spatial resolutions and timescales, simulations can be run on supercomputers as well as on individual scientist's personal computers.
- Scientific codes are often in a near-constant state of development as new science capabilities are added and requirements change.

Introduction

- Climate science has a strong computational component, and the climate codes used in this discipline are typically complex and large in size.
- These models can support a variety of spatial resolutions and timescales, simulations can be run on supercomputers as well as on individual scientist's personal computers.
- Scientific codes are often in a near-constant state of development as new science capabilities are added and requirements change.

Due to the complexity of climate software, the evolution of the code requires a strict control of accuracy, reproducibility and software quality.

Introduction

- EC-EARTH is a project, a consortium and a model system.
  - The EC-EARTH consortium consists of several academic institutions and meteorological services from different countries in Europe.
  - The EC-EARTH model is a global, coupled climate model that consists of two main components: **IFS** for the atmospheric model and **NEMO** for the ocean model. They are coupled using **OASIS3-MCT.** It has other sub-components: **LIM** for the sea ice, **XIOS** for NEMO's input/output, and **Run-off mapper** for ice coupling.
  - For high resolution modeling, which needs to run on modern supercomputers with a distributed memory system, EC-Earth uses the **MPI paradigm**, using a specified number of tasks for both NEMO and IFS models, and one process for XIOS and another for Run-off mapper.

Model development has the following objectives

Accuracy
(be close to a reference)

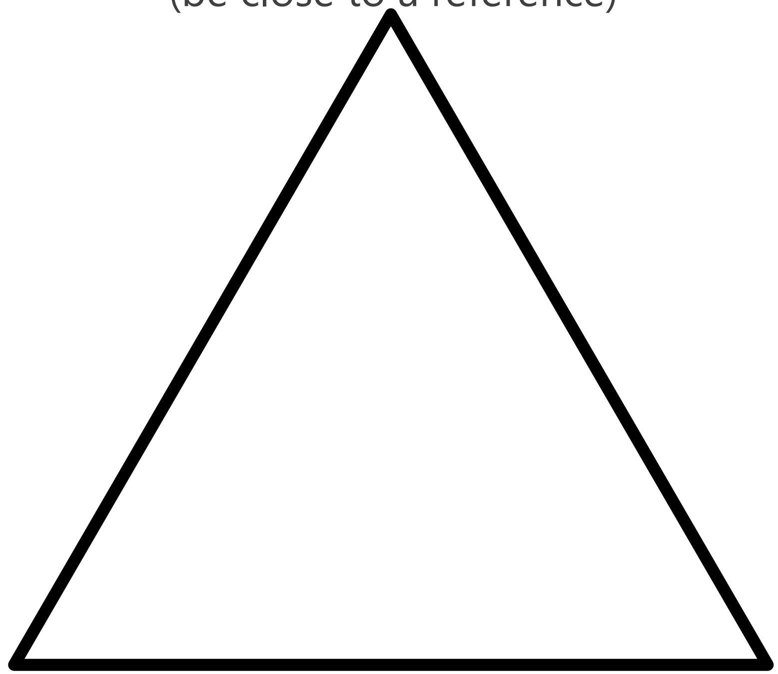

Reproducibility
(be similar across configurations)

Performance
(use resources efficiently)

Introduction

- **Parallel computing errors**
  - Finite precision addition is not associative.
    - Variables have finite resolution   1.777777777 → 1.77778
    - Rounding can change intermediate results

      A+B+C =/ A+C+B
  - Order of operations and solver iterations change with number of processors.
- **Floating Point (FP) errors are caused by:**
  - Algorithm
    - Different systems and/or input data can have unexpected results.
  - Non-deterministic task/process scheduler
    - Asynchronous task/process scheduling can change the order of some operations between reruns.
  - Alignment (heap & stack)
    - If alignment is not guaranteed, the results could be computed differently between reruns.
  - Compiler optimization options
    - Simplification of operations to reduce the computational cost (e.g. vectorization).

Introduction

- Possible solutions for not associative additions (e.g. MPI reduction)

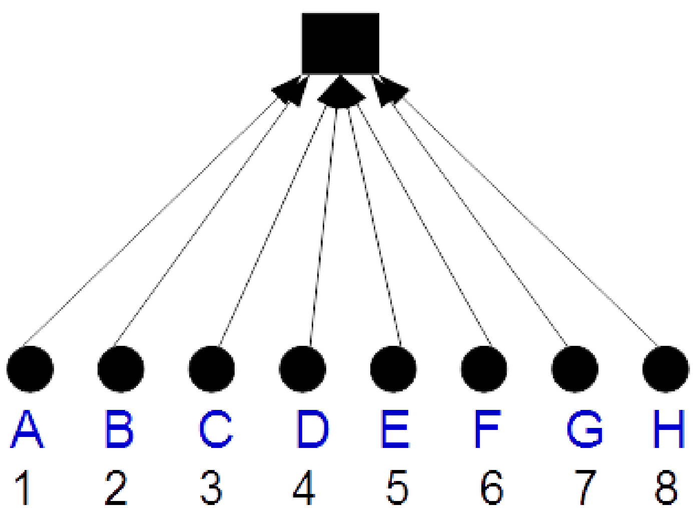

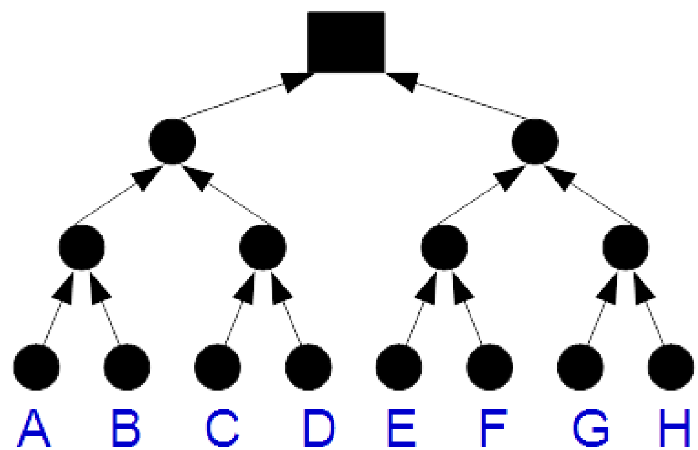

These solutions could reduce performance of the parallel application

- Increase the precision of the variables
  - Some works show how the reproducibility/accuracy of parallel numerical models improves using long double (80 bits) or two doubles (128 bits) instead of 32 or 64 bits.
    - 80 bits is not portable in all machines and two doubles increase the computational cost.

Introduction

Model development has the following objectives

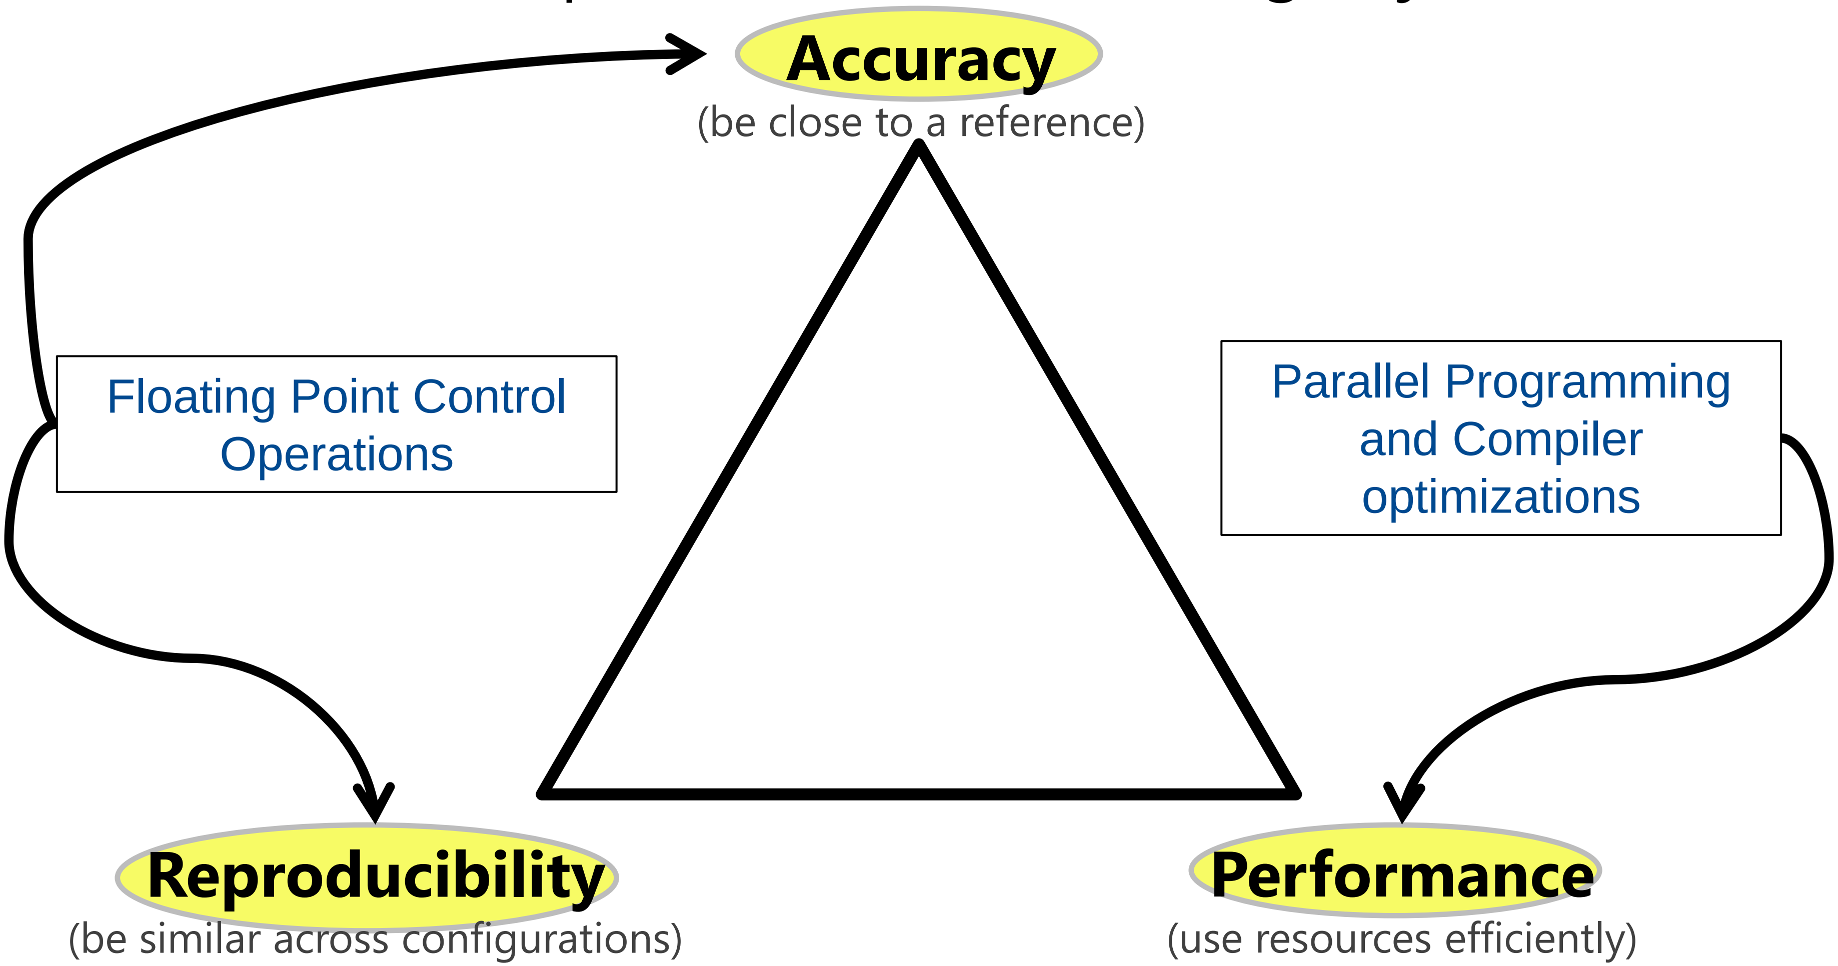

Introduction

Relation between Performance and Accuracy & Reproducibility

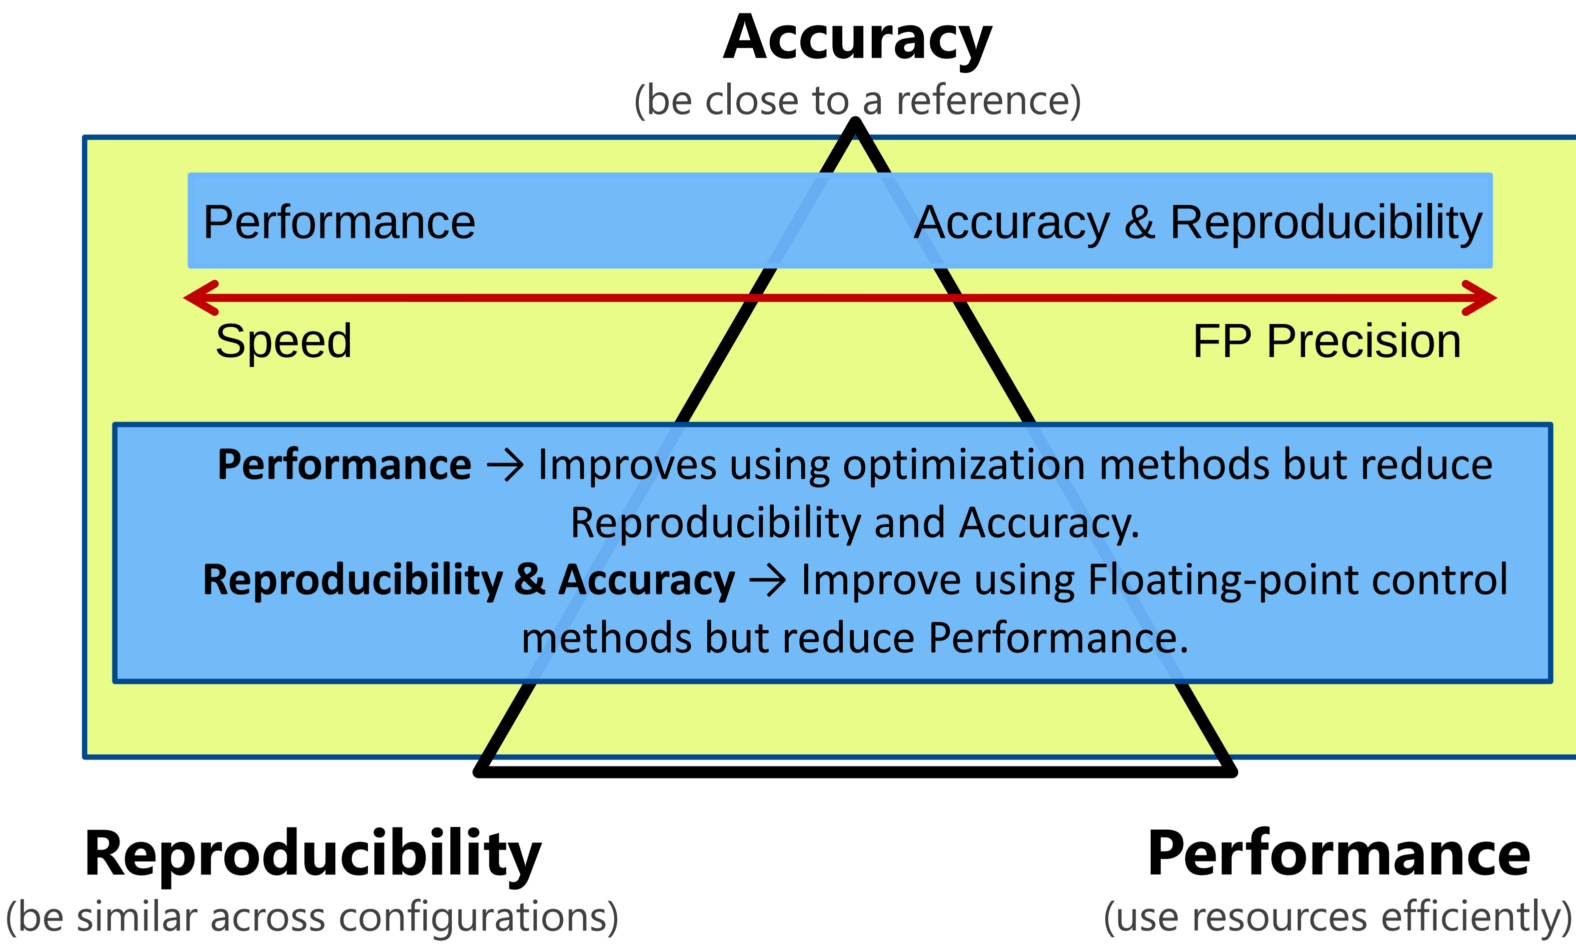

Compiler options let you control the tradeoffs among accuracy, reproducibility and performance.

Methodology

Methodology

- Different compilation flags can be used to control the tradeoffs between accuracy, reproducibility and performance.

  - To control Floating-Point (FP) operations
    - fp-model precise, fimf-arch-consistency, fpe0, fma, ftz ...

  - To control optimization options
    - O1, O2, O3, xHost, ipo ...

Methodology

- To control FP operations

Flag	Description
-fp-model precise	Allow value-safe optimizations only
-fp-model source/double/extended	Intermediate precision for FP calculations (For Fortran only source)
-fimf-arch-consistency=true	Math library functions produce consistent results
-no-fma	FP contractions are disallowed
-fp-model except	Determine whether floating point exceptions semantics are used
-fp-model strict	Enables precise and except, disables contractions, and enables the property that allows modification of the floating-point environment (include -no-fma, fpe0)
-ftz	Flush denormal results to zero
-fpe0	Unmask floating point exceptions and disable generation of denormalized numbers (include ftz)

Methodology

- To control optimization options

Flags	Description
-ipo	Permits inlining and other interprocedural optimizations among multiple source files.
-prof-gen/-prof-use	Gen enable profile generation and use the profile generation during optimization
-no-prec-div	Reduce precision of floating point divides
-no-prec-sqrt	Reduce precision of square root computations
-O0	No optimizations
-O1	Enables optimizations if they do not increase code size
-O2	Enables optimizations for maximized speed, such as code inlining, loop unrolling, variable renaming...
-O3	Performs O2 optimizations and enables aggressive loop transformations such as Fusion, collapsing “if” statements...
-xHost	Generates instruction sets up to the highest that is supported by the compilation host.
-r8	Specifies the variable names to be double precision 8-byte real numbers which has 15 digits of accuracy and a magnitude range of 10 from -308 to +308

Methodology

- These flags enable or disable (FP Flags):
  - Value safety (fp-model precise, ftz)
    - Make safe some operations such as Reassociation (**(a+b)+c or a+(b+c)), Zero folding (X+0), Multiply by reciprocal (A/B = A*(1/B))…**
  - Floating-point expression evaluation (fp-model source/double, fimf-arch-consistency=true)
    - Precision used for rounding off the intermediate results (e.g. **a=b*c+d**).
  - Precise floating-point exceptions (fp-model except, fp-model strict,fpe0)
    - FP exceptions (overflow, underflow, divide by zero…) are synchronized with the operation causing it and optionally unmasked.
  - Floating-point contractions (fp-model strict, no-fma) → **a=b*c+d**
  - Floating-point unit environment access (fp-model strict, ftz)
    - Control some options such as the rounding mode.

Methodology

- These flags enable or disable (Optimization Flags):
  - General optimization options
    - Optimizations which do not increment the size code (O1)
    - Optimizations which could change the flow and the code such as vectorization (O2)
    - Aggressive optimizations in loops such as loop unrolling (O3)
  - Instruction sets (AVX, SSE4.2, SSE3)
    - Use the same precision in all instruction sets (xHost, r8)
  - Approximation of operations (no-prec-div, no-prec-sqrt)
    - The machine solves the operation using an approximation.
  - Other optimizations evaluated
    - Inline and interprocedural optimizations among multiple source files (ipo)
    - Use of profile information during the optimization (prof-gen, prof-use)

Methodology

Performance

- **(no-prec)**-O3 -xHost -r8 -ipo -prof-use -no-prec-div -no-prec-sqrt
- **(prof-use)**-O3 -xHost -r8 -ipo -prof-gen → -O3 -xHost -r8 -ipo -prof-use
- **(ipo)**-O3 -xHost -r8 -ipo
- **(O3)**-O3 -xHost -r8
- **(O2)**-O2 -xHost -r8
- **(no-fma_fz)**-O2 -no-fma -ftz -r8
- **(fp-except)**-O2 -fp-model except -no-fma -ftz -fpe0 -r8
- **(fp-precise)**-O2 -fp-model precise -fimf-arch-consistency=true -no-fma -fpe0 -r8
- **(fp-strict)**-O2 -fp-model strict -fimf-arch-consistency=true -no-fma -fpe0 -r8
- **(O1)**-O1 -fp-model strict -fimf-arch-consistency=true -no-fma -fpe0 -r8

Accuracy

Methodology

- Marenostrum III (BSC)
  - 2x E5–2670 SandyBridge-EP 2.6GHz cache 20MB 8-core
  - 8x 4G DDR3–1600 DIMMs (2GB/core) Total: 32GB/node.
  - Infiniband Mellanox FDR10: High bandwidth network used by parallel applications communications (MPI).
  - Intel Fortran Compiler and Intel MPI library.
- Experiment
  - EC-Earth 3.2beta
    - IFS 36r4, NEMO 3.6, LIM3, XIOS, Runoff-mapper, OASIS3-MCT.
  - Standard configuration (T255L91, ORCA1L75).
  - 1 year simulated (1990) using a time-step of 2700 seconds.
  - Outputs with six-hourly, daily and monthly frequency.
  - Use MPI processes (288 for NEMO, 320 for IFS, one for XIOS and one for Run-off mapper), using 39 nodes of MN3 for the simulation.
  - Five-members ensemble with small perturbations in the initial condition.

Methodology

- Performance
  - Execution time
    - Average time step.
    - Time to simulate the complete year (Total Time).
- Precision and reproducibility
  - Reichler-Kim normalized index for 13 variables
    - Calculate for each variable a normalized error variance e^2 by squaring the grid-point differences between simulated and observed climate

$$e_{vm}^2 = \sum_n \left(w_n \left(\bar{s}_{vmn} - \bar{o}_{vn} \right)^2 / \sigma_{vn}^2 \right)$$

    - Ensure that different climate variables receive similar weights when combining their errors

$$I_{vm}^2 = e_{vm}^2 \Big/ \overline{e_{vm}^2}^{m=20C3M}$$

    - Mean over all climate variables

$$I_m^2 = \overline{I_{vm}^2}^{v}$$

Methodology

Climate Variables and corresponding validation data

Variable	Domain	Validation data	Period
Sea level pressure	ocean	ICOADS (Woodruff et al. 1987)	1979–99
Air temperature	zonal mean	ERA-40 (Simmons and Gibson 2000)	1979–99
Zonal wind stress	ocean	ICOADS (Woodruff et al. 1987)	1979–99
Meridional wind stress	ocean	ICOADS (Woodruff et al. 1987)	1979–99
2-m air temperature	global	CRU (Jones et al. 1999)	1979–99
Zonal wind	zonal mean	ERA-40 (Simmons and Gibson 2000)	1979–99
Meridional wind	zonal mean	ERA-40 (Simmons and Gibson 2000)	1979–99
Net surface heat flux	ocean	ISCCP (Zhang et al. 2004), OAFLUX (Yu et al. 2004)	1984 (1981) –99
Precipitation	global	CMAP (Xie and Arkin 1998)	1979–99
Specific humidity	zonal mean	ERA-40 (Simmons and Gibson 2000)	1979–99
Snow fraction	land	NSIDC (Armstrong et al. 2005)	1979–99
Sea surface temperature	ocean	GISST (Parker et al. 1995)	1979–99
Sea ice fraction	ocean	GISST (Parker et al. 1995)	1979–99
Sea surface salinity	ocean	NODC (Levitus et al. 1998)	variable

Reichler, T., and J. Kim (2008): **How Well do Coupled Models Simulate Today's Climate?** *Bull. Amer. Meteor. Soc.*, **89**, 303-311.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Results

Results

- Evaluation of execution time

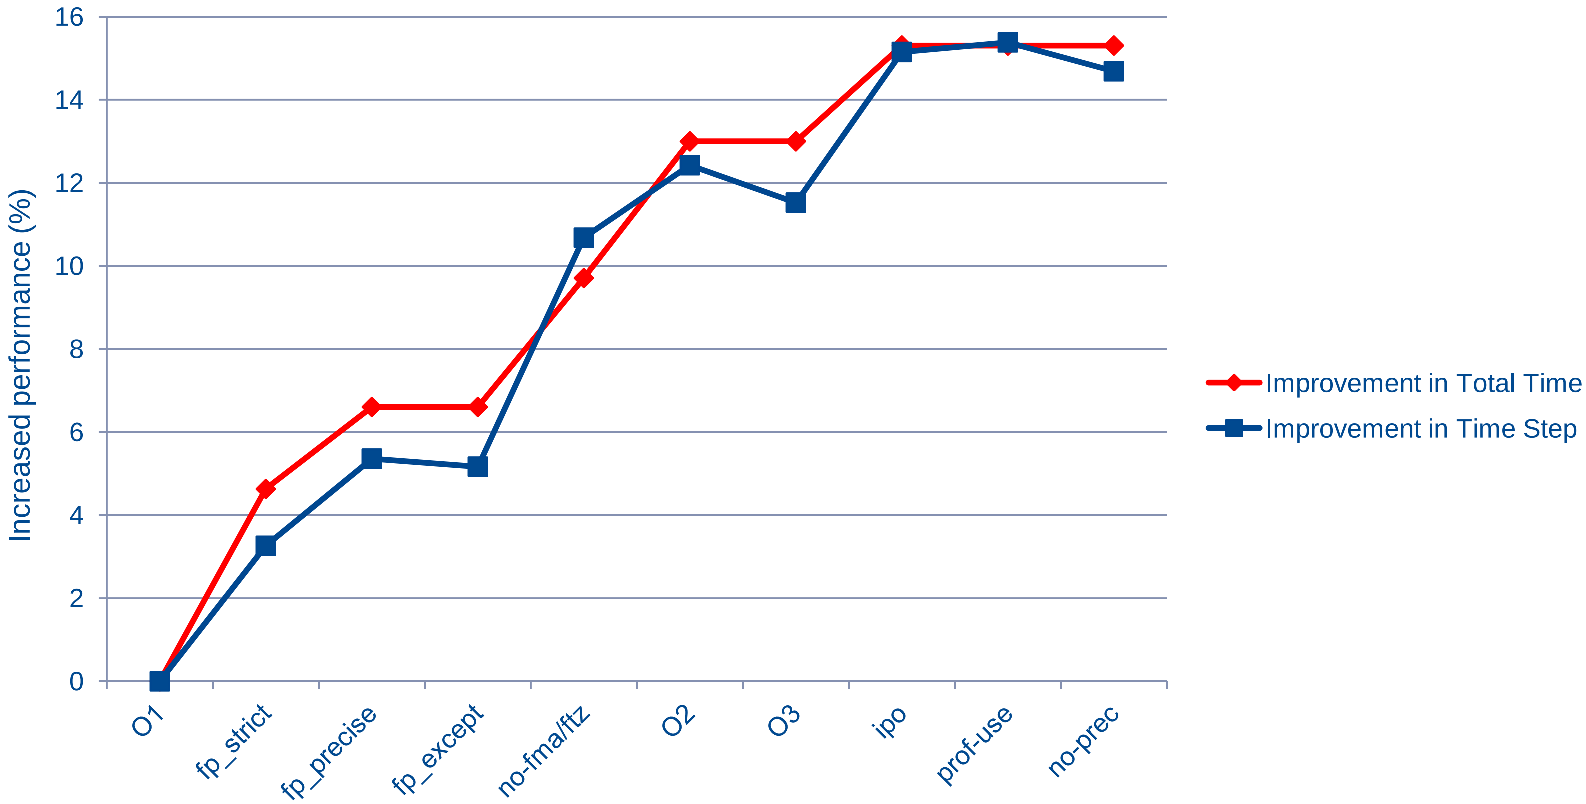

Comparing to restrictive O1 option

Results

- Evaluation of execution time

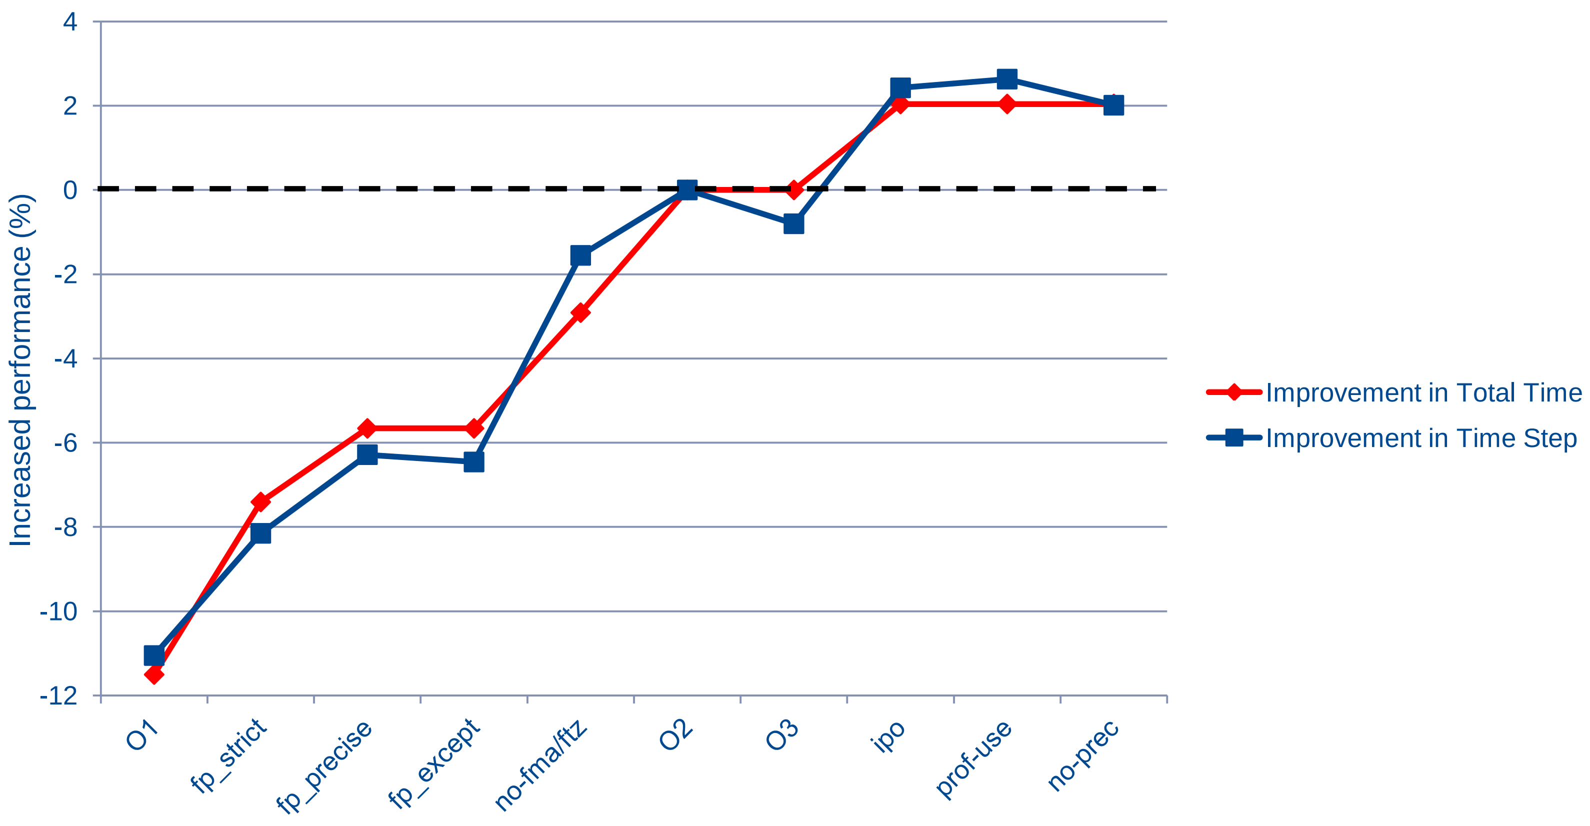

Comparing to standard O2 configuration

Results

- Evaluation of execution time

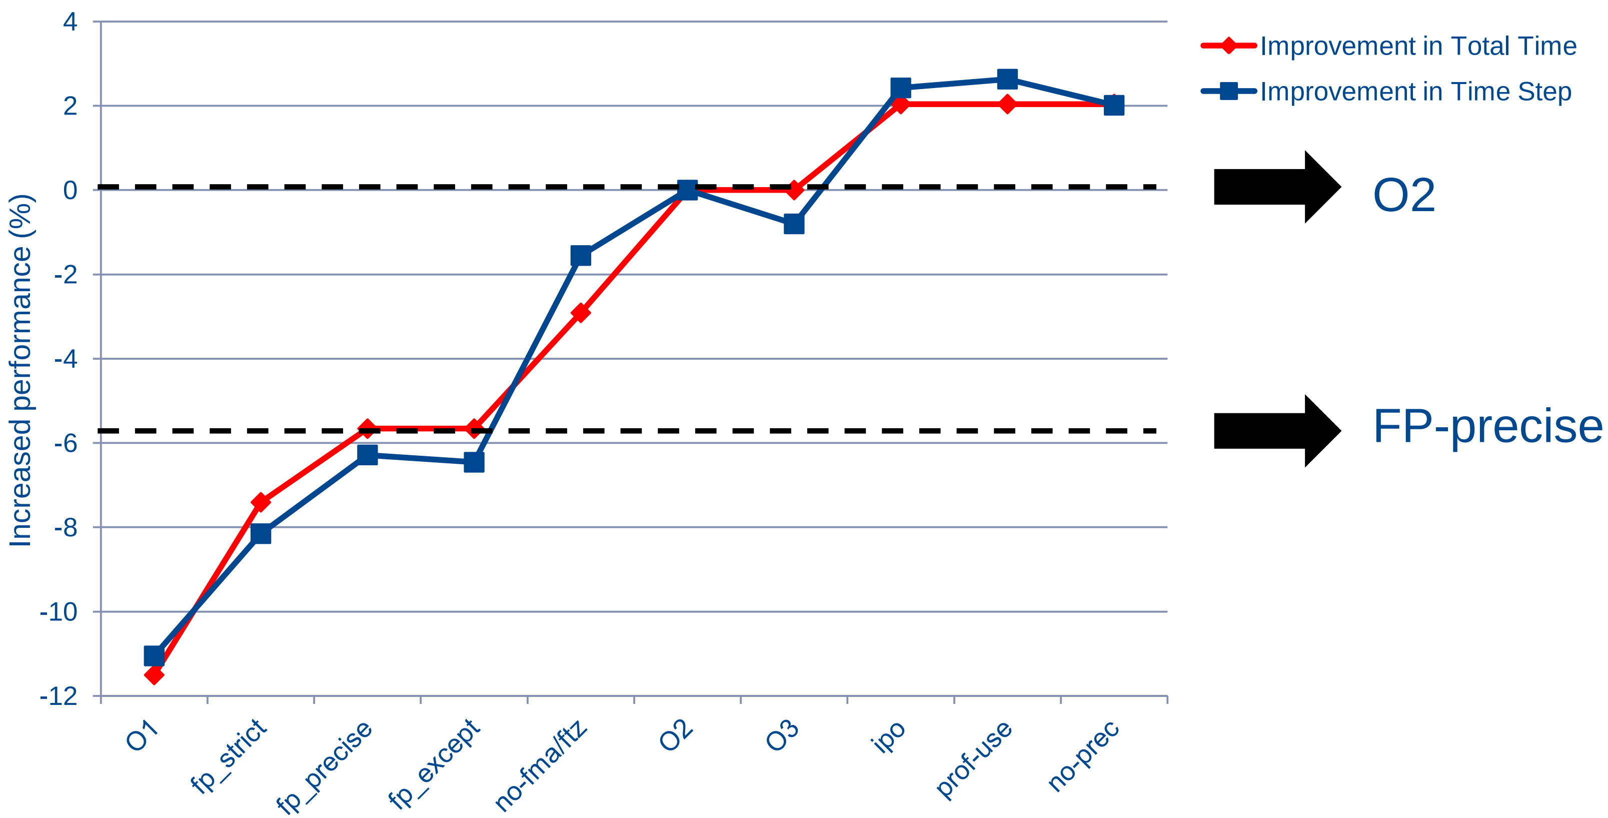

Comparing to standard O2 configuration

Results

- Evaluation of precision and reproducibility

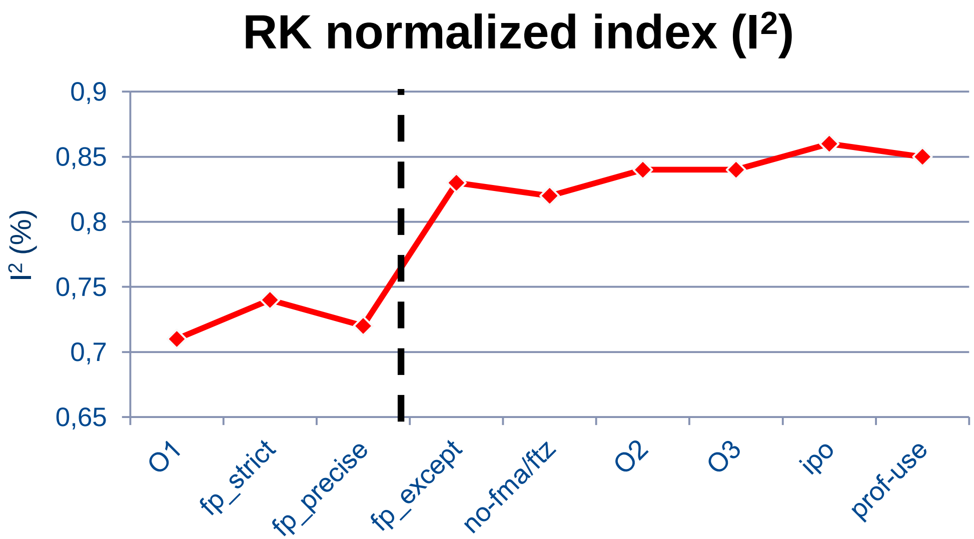

Results

- Evaluation of precision and reproducibility

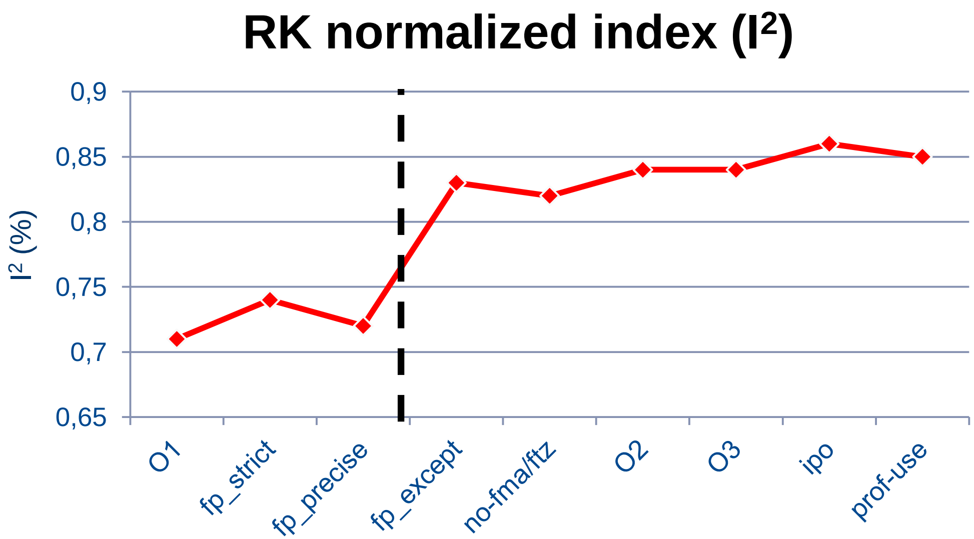

		Reproducibility Test
FP_precise →	$I^2 = 0.72$ →	$I^2 = 0.72$
O2 →	$I^2 = 0.84$ →	$I^2 = 0.85$

Evaluation of differences according to the Kolmogorov-Smirnov test

Reproducibility test of O2: 1% of grid points show a significant difference

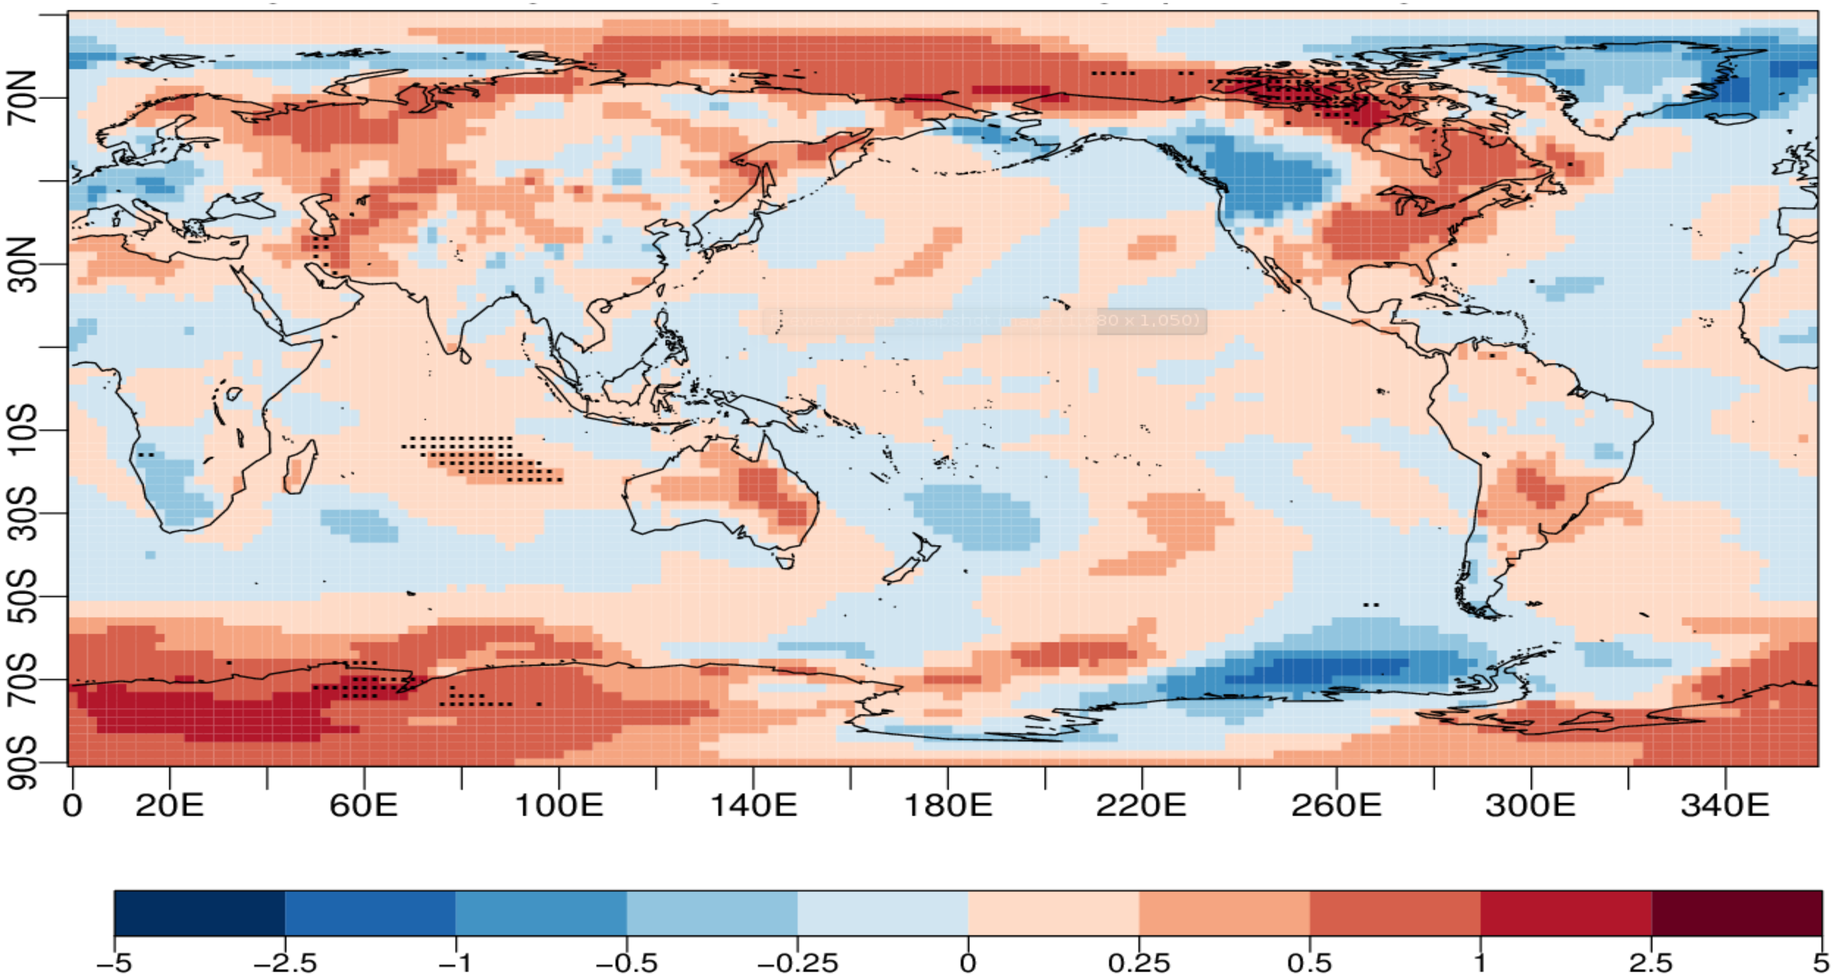

Temperature difference of the ensemble means (five-members) between FP_precise and O2. Black doted regions indicate where the differences are significant.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

Evaluation of differences according to the Kolmogorov-Smirnov test

Reproducibility test of FP_precise: 0% of grid points show a significant difference

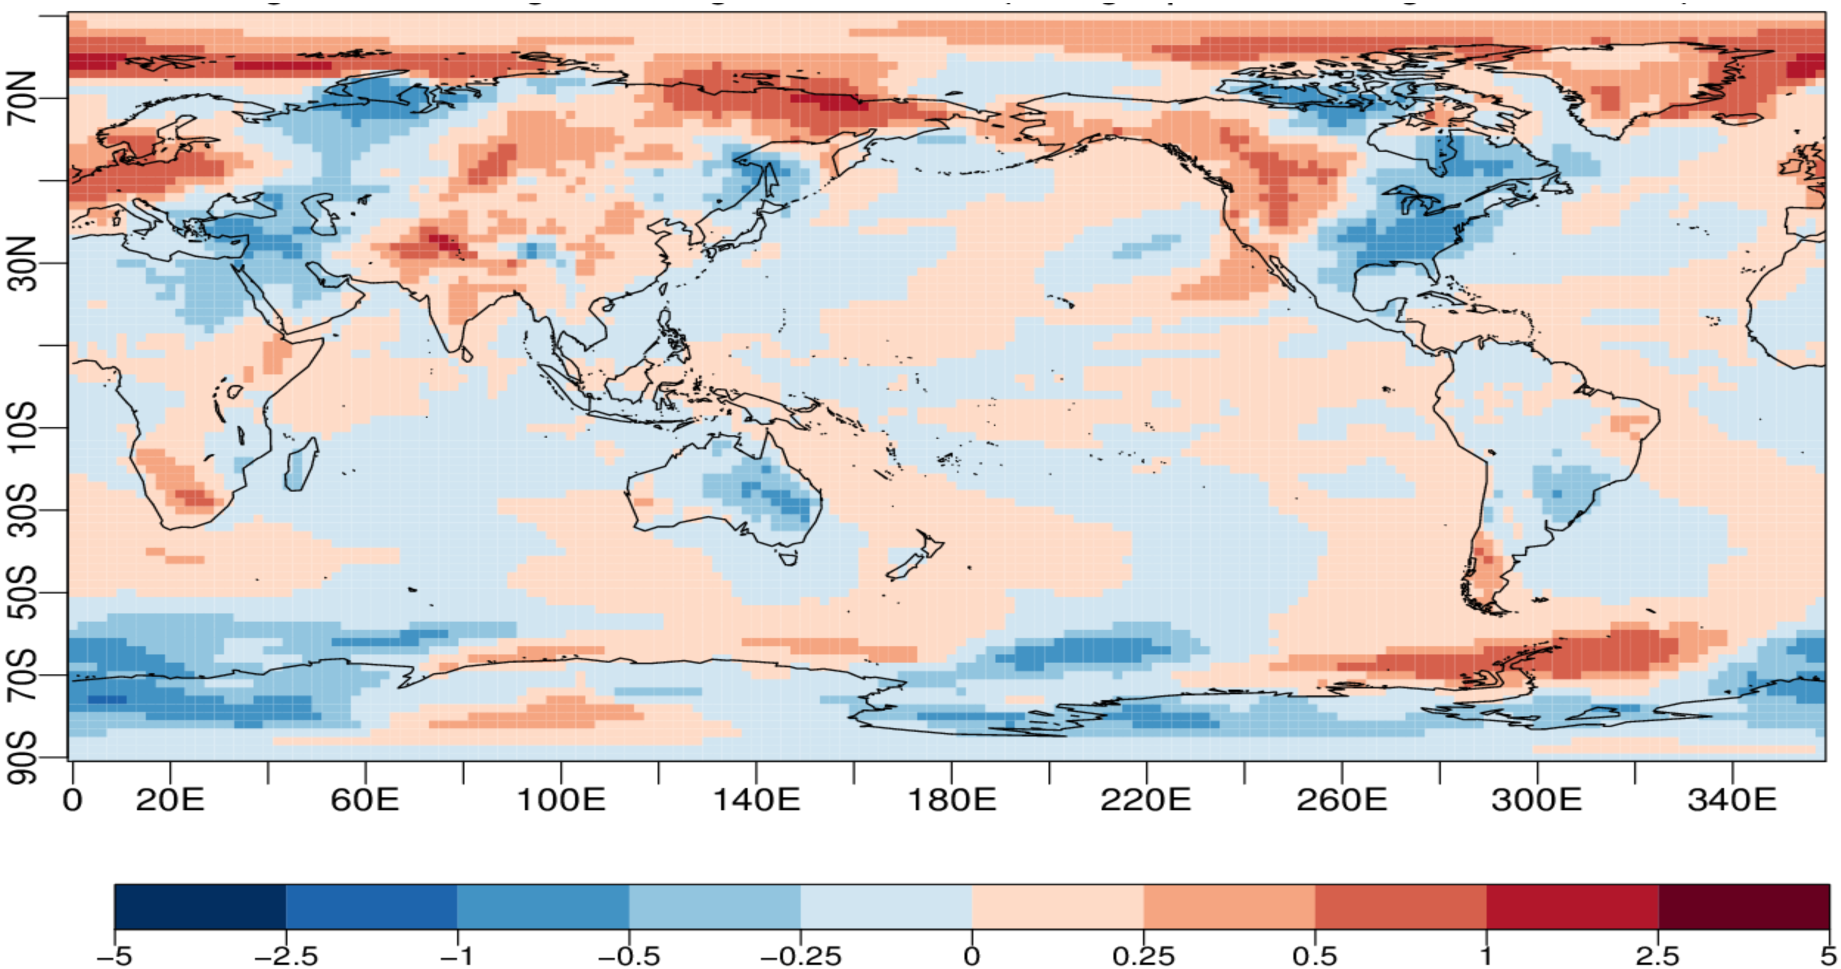

Temperature difference of the ensemble means (five-members) between FP_precise and O2. Black doted regions indicate where the differences are significant.

Evaluation of differences according to the Kolmogorov-Smirnov test

Differences between FP_precise and O2: 3% of grid points show a significant difference

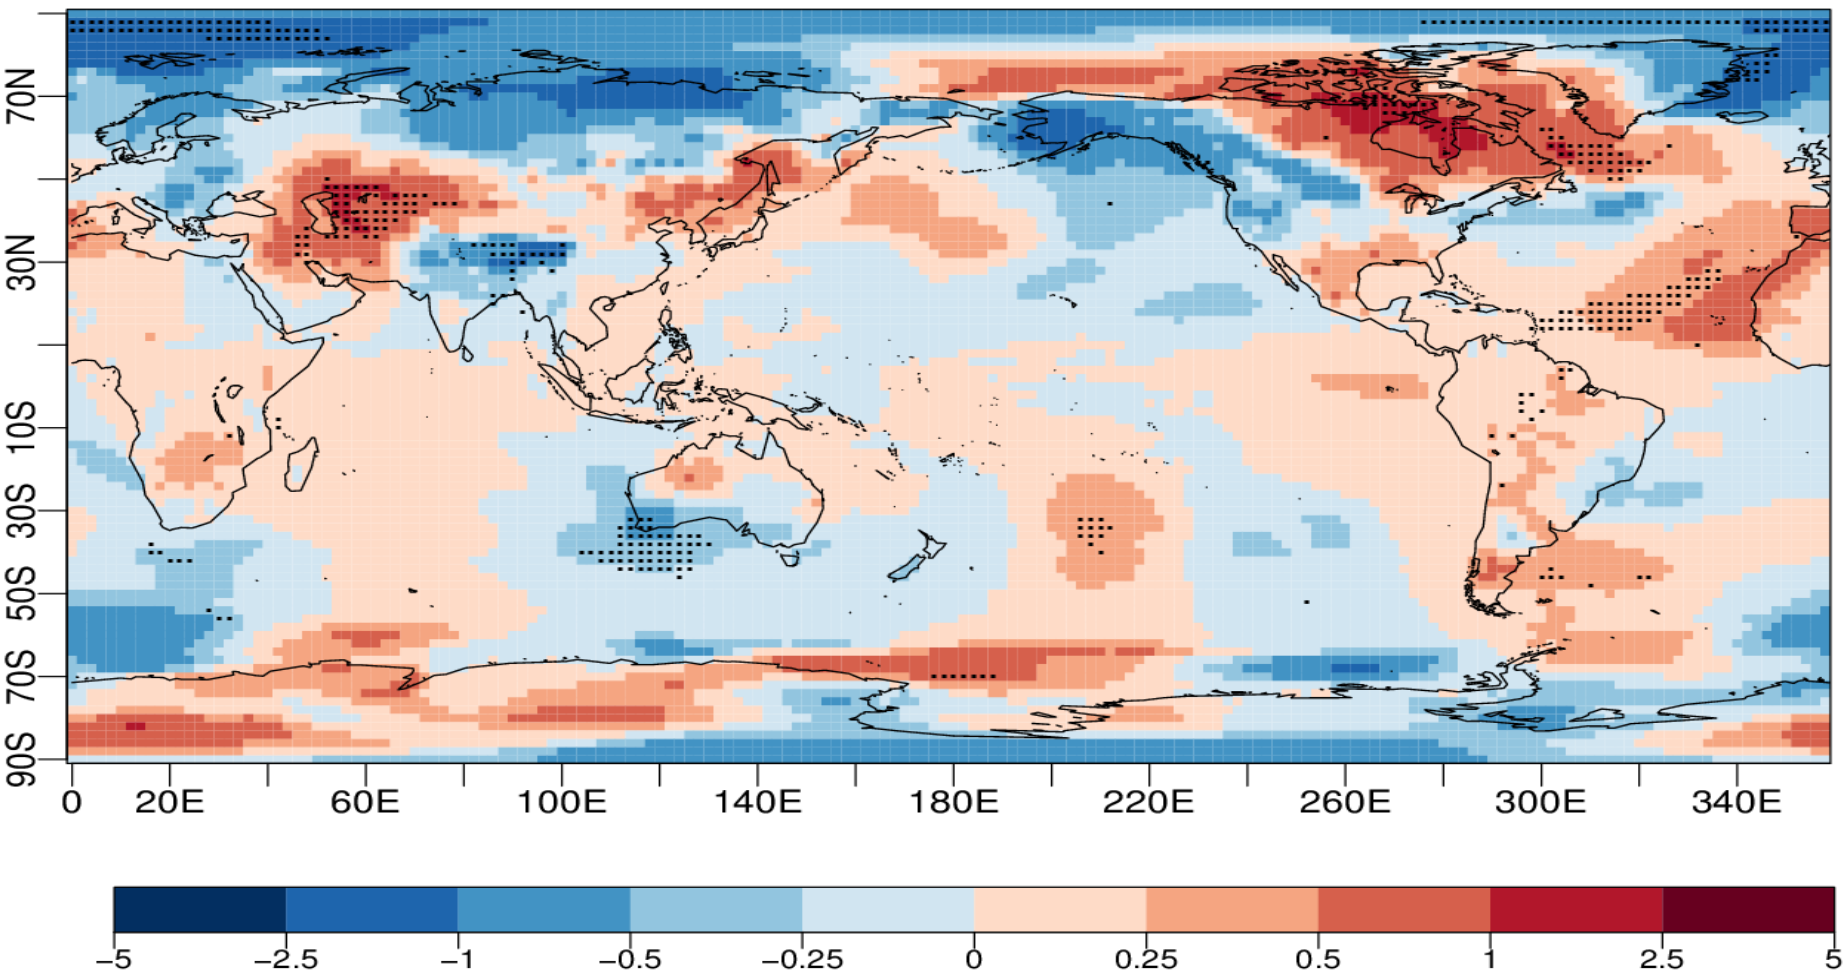

Temperature difference of the ensemble means (five-members) between FP_precise and O2. Black doted regions indicate where the differences are significant.

www.bsc.es

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Conclusions

Conclusions

Barcelona Supercomputing Center
Centro Nacional de Supercomputación
EXCELENCIA SEVERO OCHOA

– Standard flag configurations for performance and FP precision obtain the best results.
  - **-fp-model precise -fpe -no-fma -O2 -xHost -r8**
    – Good performance, better precision (3%), better reproducibility (differences close to 0), catch fpe exceptions).
  - **-O2 -xHost -r8**
    – Better performance (6%), good precision, good reproducibility (differences less than 1%).

– Aggressive optimizations (O3, ipo, prof-use) do not improve the performance.
  - Other issues could avoid additional optimizations (loop dependences, non vectorization, MPI overhead …).

– Strict FP control does not improve the precision and reduce the performance up to 6%-12%.

– Using approximations for FP operations (no-prec-div/sqrt) does not improve the performance and reduces the precision and reproducibility dramatically.

Future work

- Evaluate long simulations in time
- Evaluate diverse hardware configurations
  - Similar platforms
  - Different platforms
- Evaluate different software configurations
  - Different version of libraries
  - Same version of libraries compiled with different compilers
- Other parallel issues
  - Different affinity configurations
  - Different domain decompositions

www.bsc.es

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

PRIMAVERA

PRocess-based climate sIMulation: AdVances in high resolution modelling and European climate Risk Assessment

The research leading to these results has received funding from the EU H2020 Framework Programme] under grant agreement PRIMAVERA (H2020 GA 641727)

www.bsc.es

Toulouse, 07 April 2016

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Reproducibility of Earth System Models A computational Point of View

M. C. Acosta, X. Yepes-Arbós, F. Massonnet, M. Ménégoz, K. Serradell, O. Mulla-Valls and F. Doblas-Reyes