



Phase inversion problem: performances on EOS

Annaïg PEDRONO

IMFT

Service Codes et Simulations Numériques




IMFT and CALMIP

- IMFT and CALMIP : a partnership to improve code performances since 2004
- 2009-2014 : IMFT consumed 24 millions hours on Hyperion
- Involvement in the renewal of CALMIP supercomputers, 1 CFD benchmark with 2 codes (JADIM & Neptune) : 2009 (hyperion) and 2013 (eos)
- Benchmark permits to take into account specificity of CFD in the supercomputer choice
- JADIM : from a serial code (2005) on « Soleil » to 4096 cores on « EOS » (2014)

Phase Inversion problem

- 3 codes (PARIS, Thetis, JADIM)
- a mesh of 134 millions of cells (512x512x512)
- performance tests from 64 to 2048 cores
- ~36 hours for each run using ~2000 cores
- about 1 Tb of generated data per simulation

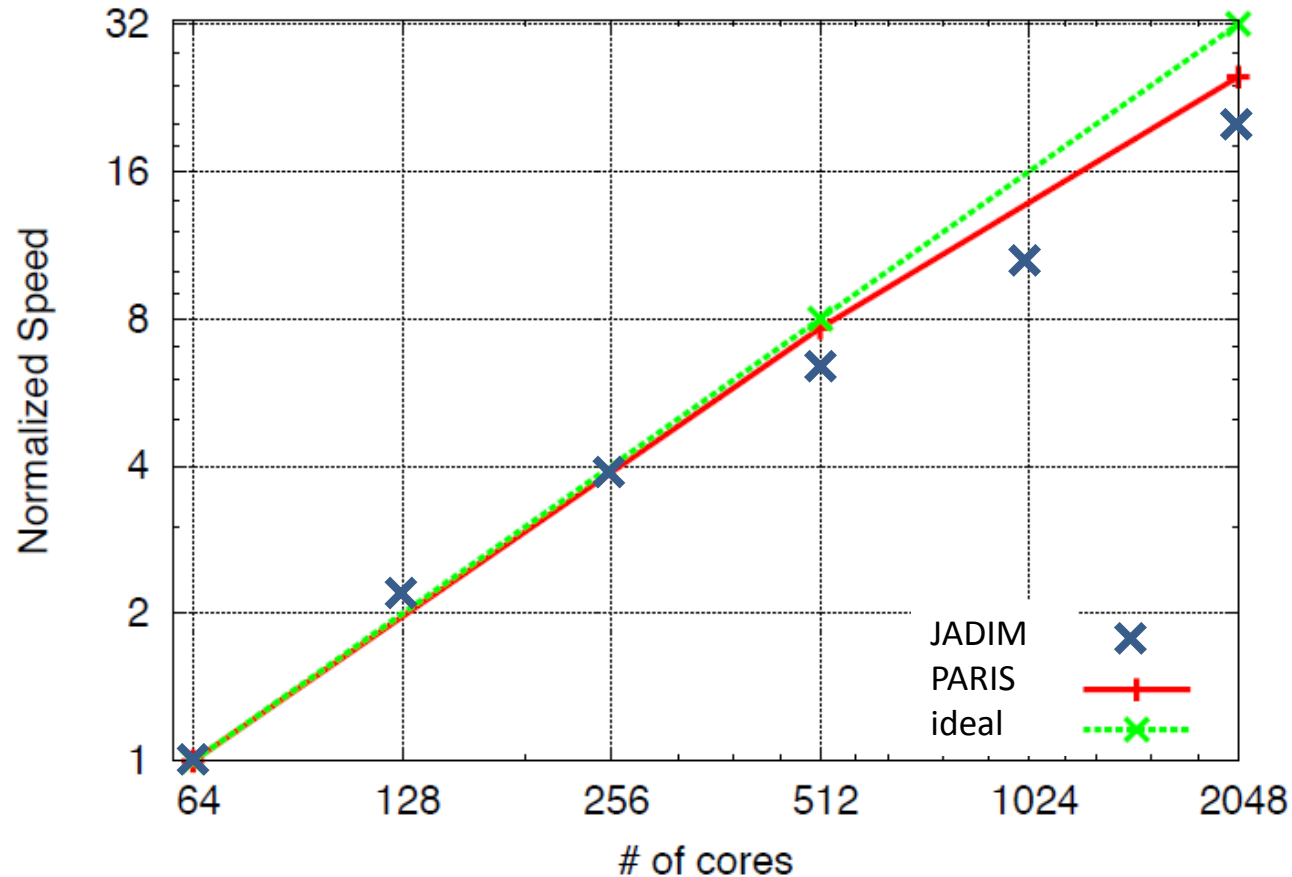


Code performances

How to compare the performances of different codes?


- Total CPU time needed to reach 25 seconds (physical time) for one configuration of phase inversion problem
- Speed-up comparison from 64 to 2048 cores
- Velocity $Z = \text{number of computed nodes} / \text{elapsed time for one iteration}$

Speed-up



Drawback : compare the scalability of each code without information on the serial performance

Velocity Z / number of cores

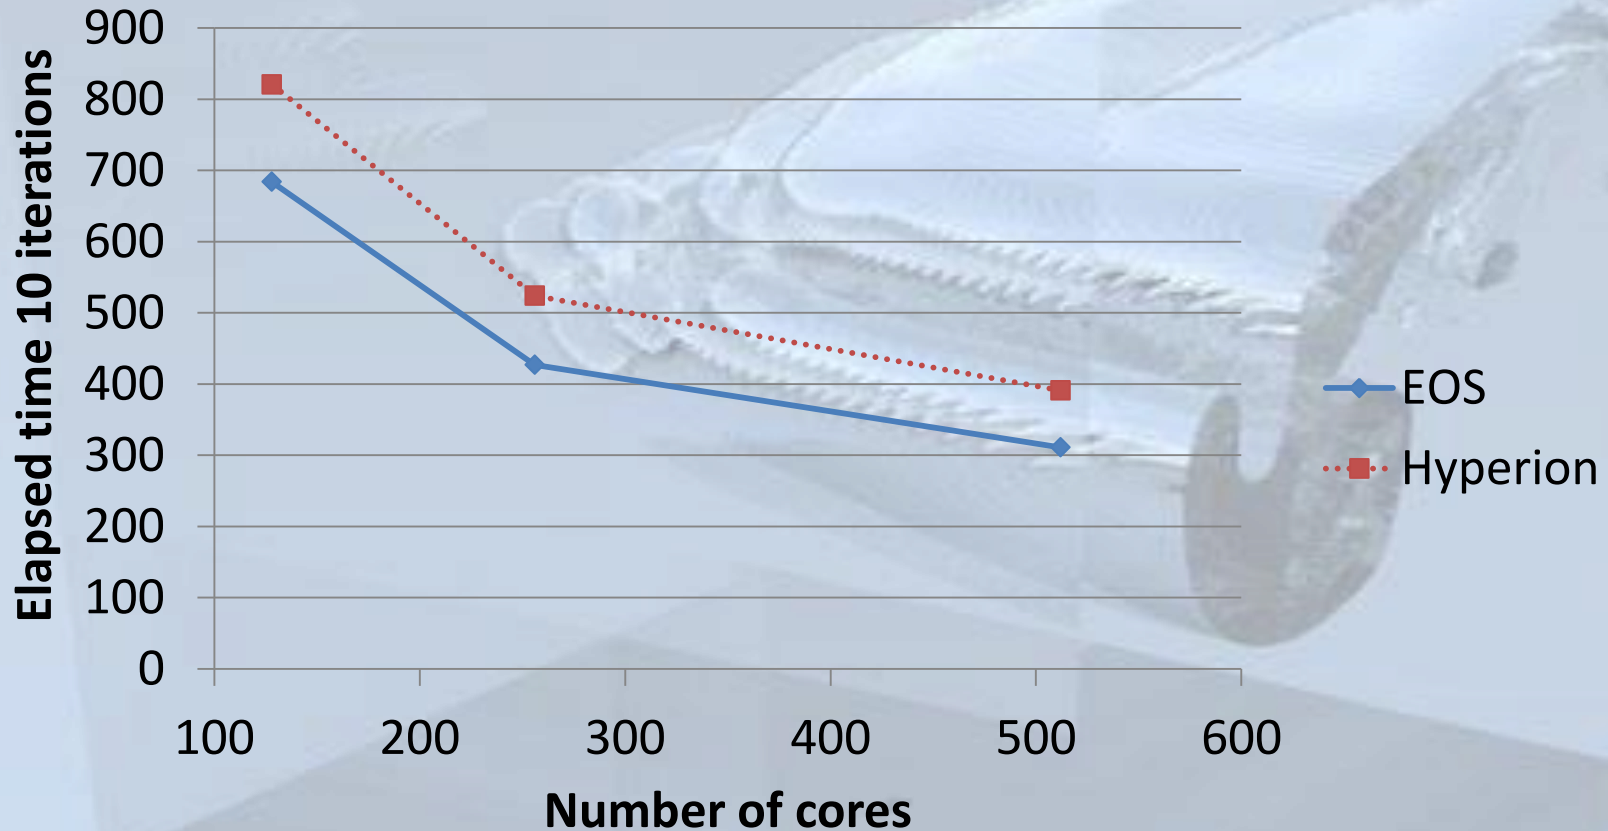


Drawback : different physical time steps, different numerical methods, different precisions, ...

JADIM : performances on EOS

- Comparison EOS / Hyperion
- Profiling
- Effect of stopping criteria value on Poisson's solver
- JADIM improvement (tuning, memory, visualisation)

Comparison Hyperion / EOS

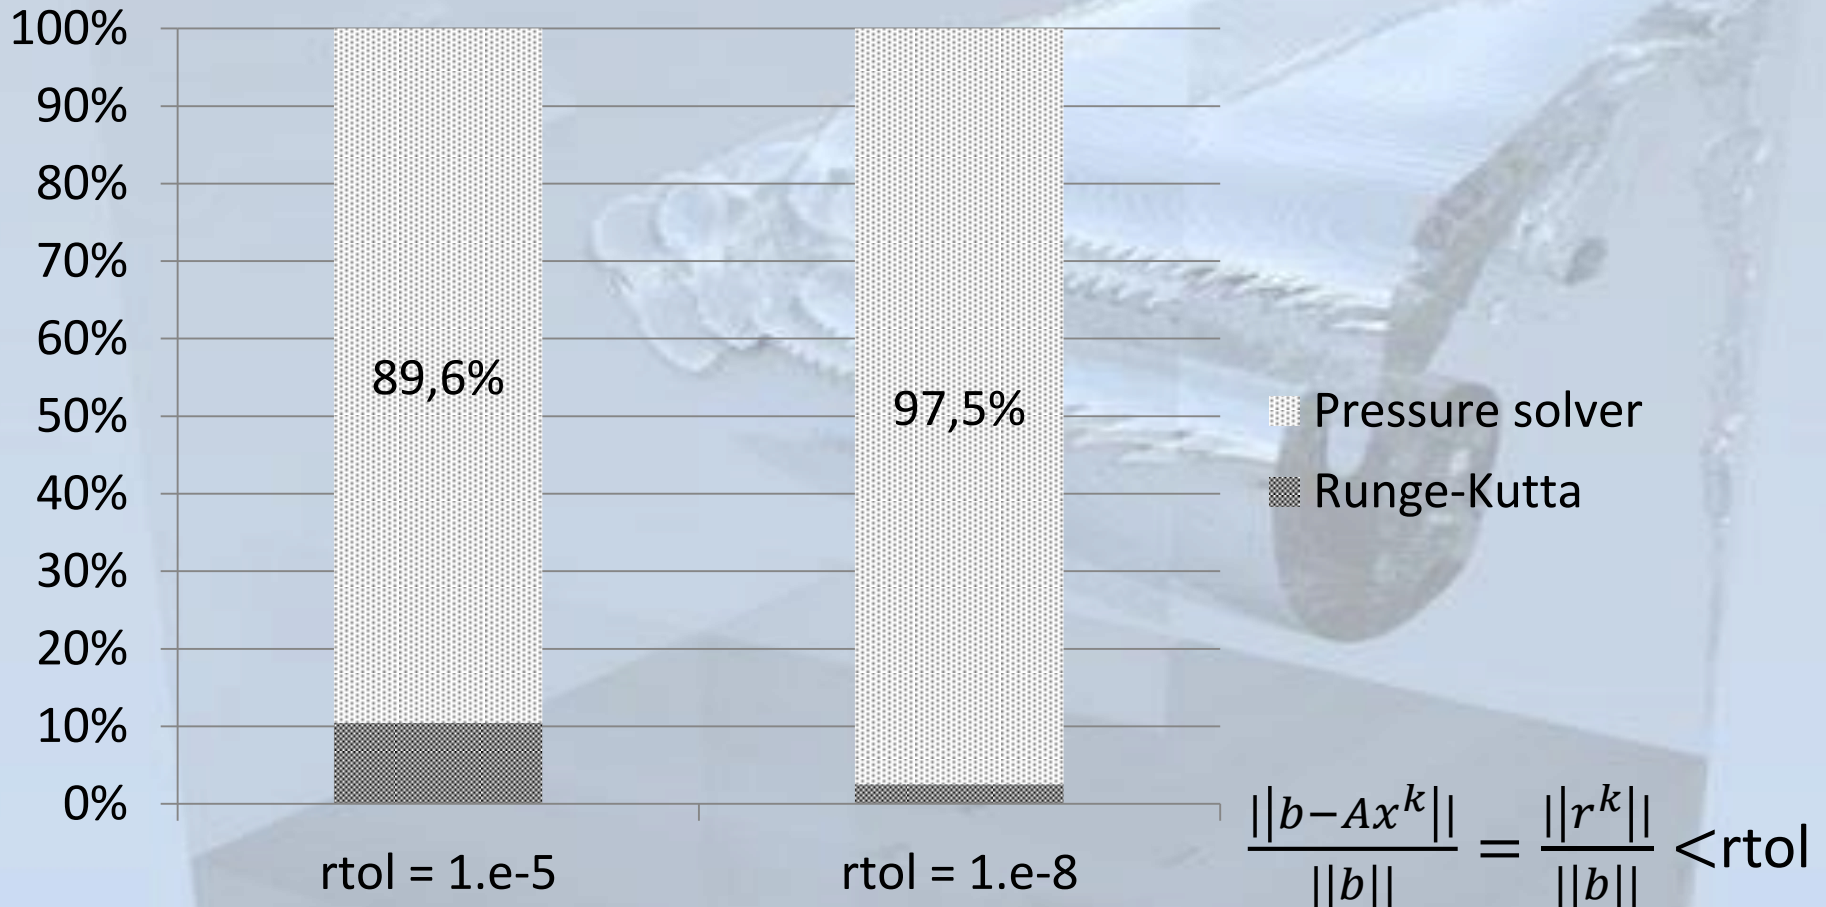


- 20% gain between hyperion and EOS
- same scalability until 512 cores

EOS : Processors Ivybridge, 2.8GHz, Intel MPI, opt = -O3 -xAVX

Hyperion : Proc. Nehalem, 2.8 GHz, SGI MPT, opt = -O3 -xsse4.2

CPU time distribution

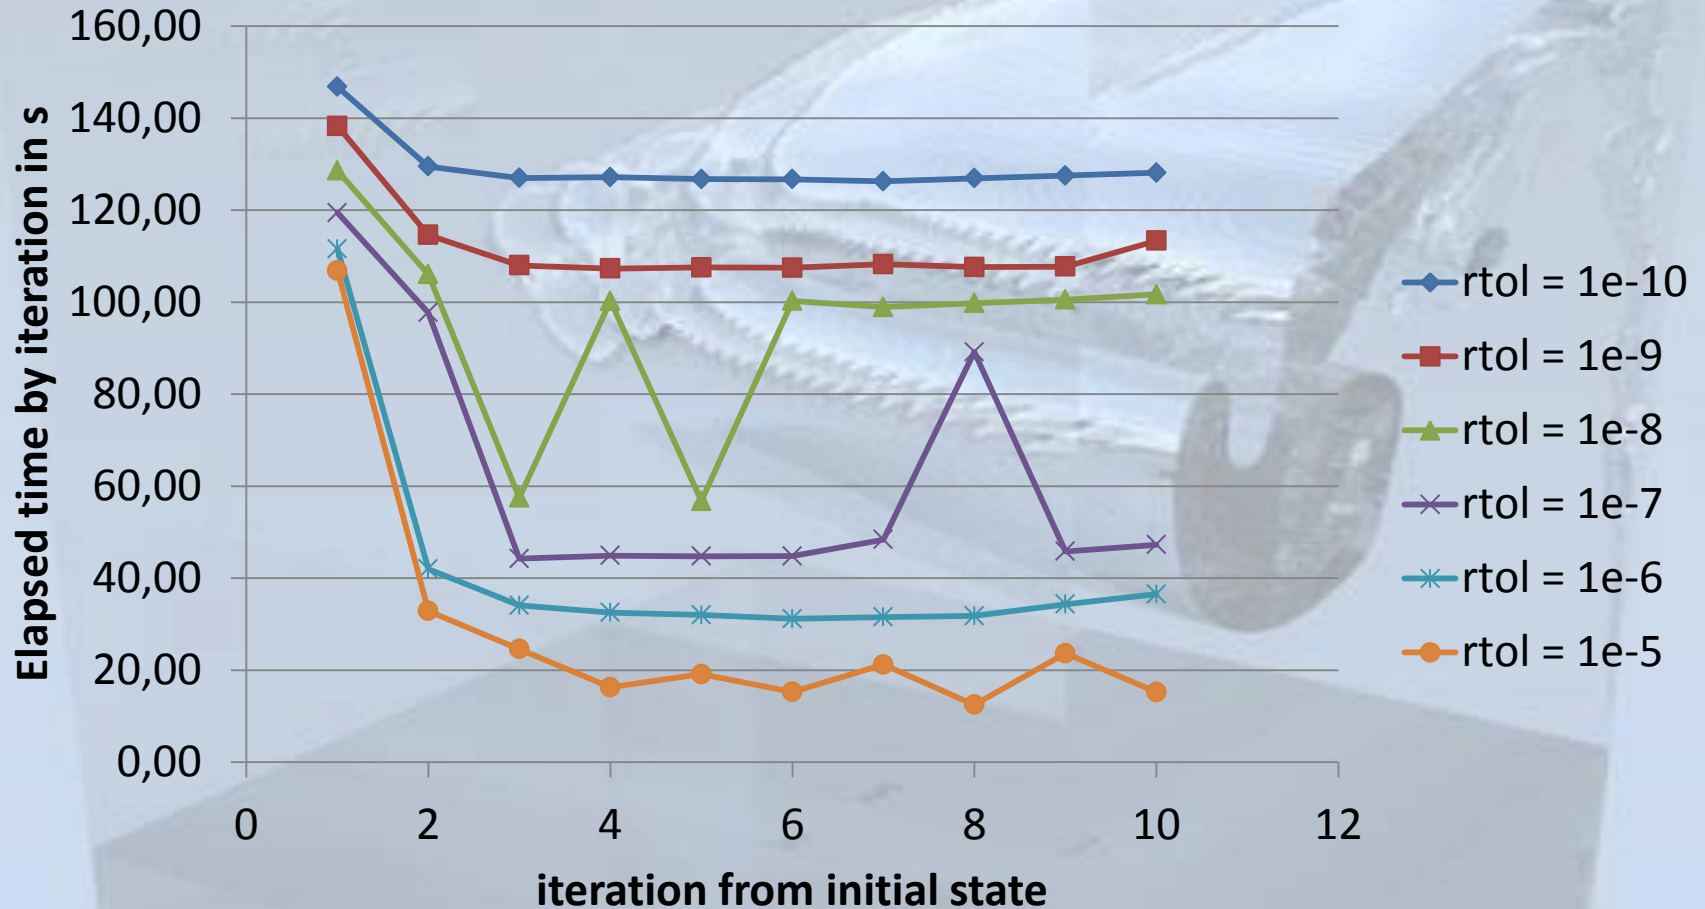


Most of CPU time is spent in Poisson's solver

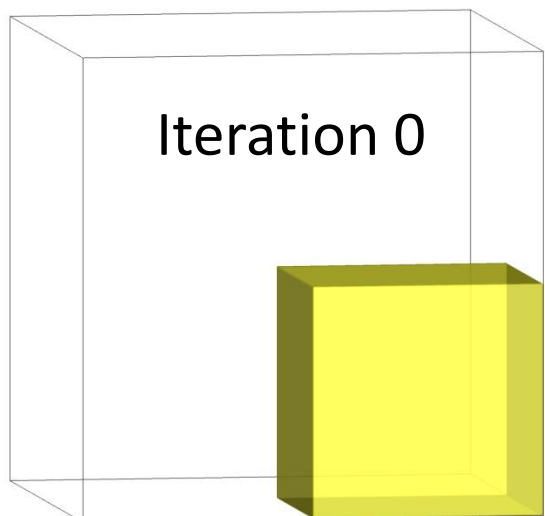
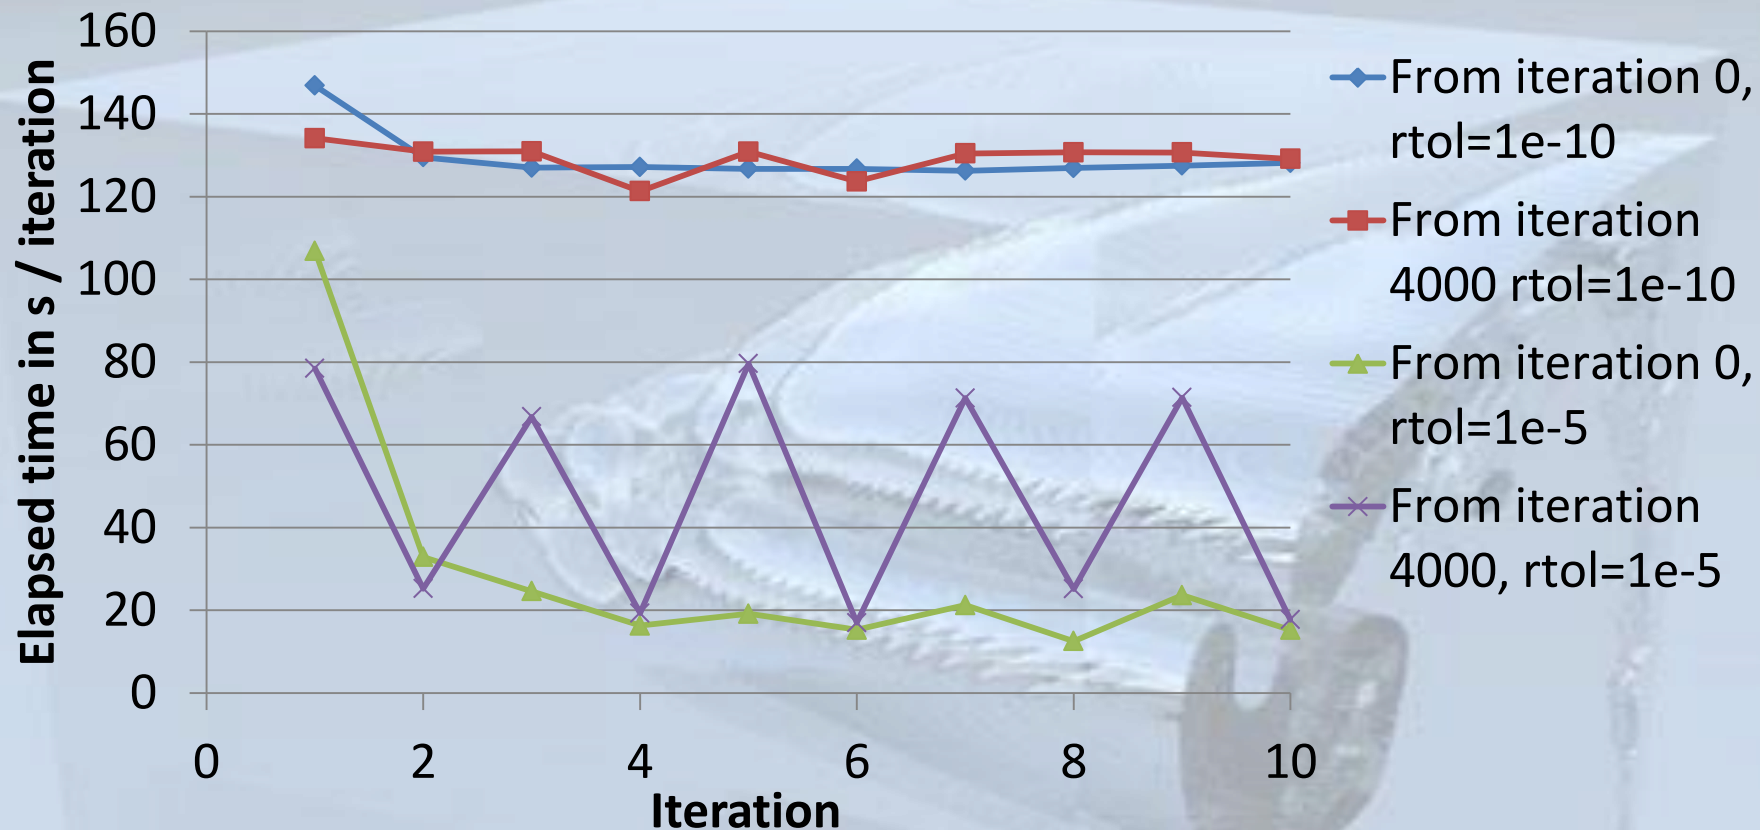
(solved with PETSc : CG preconditioned with block Jacobi)

rtol = residual convergence stop criteria of CG's method


Effects of iterative solver stopping criteria




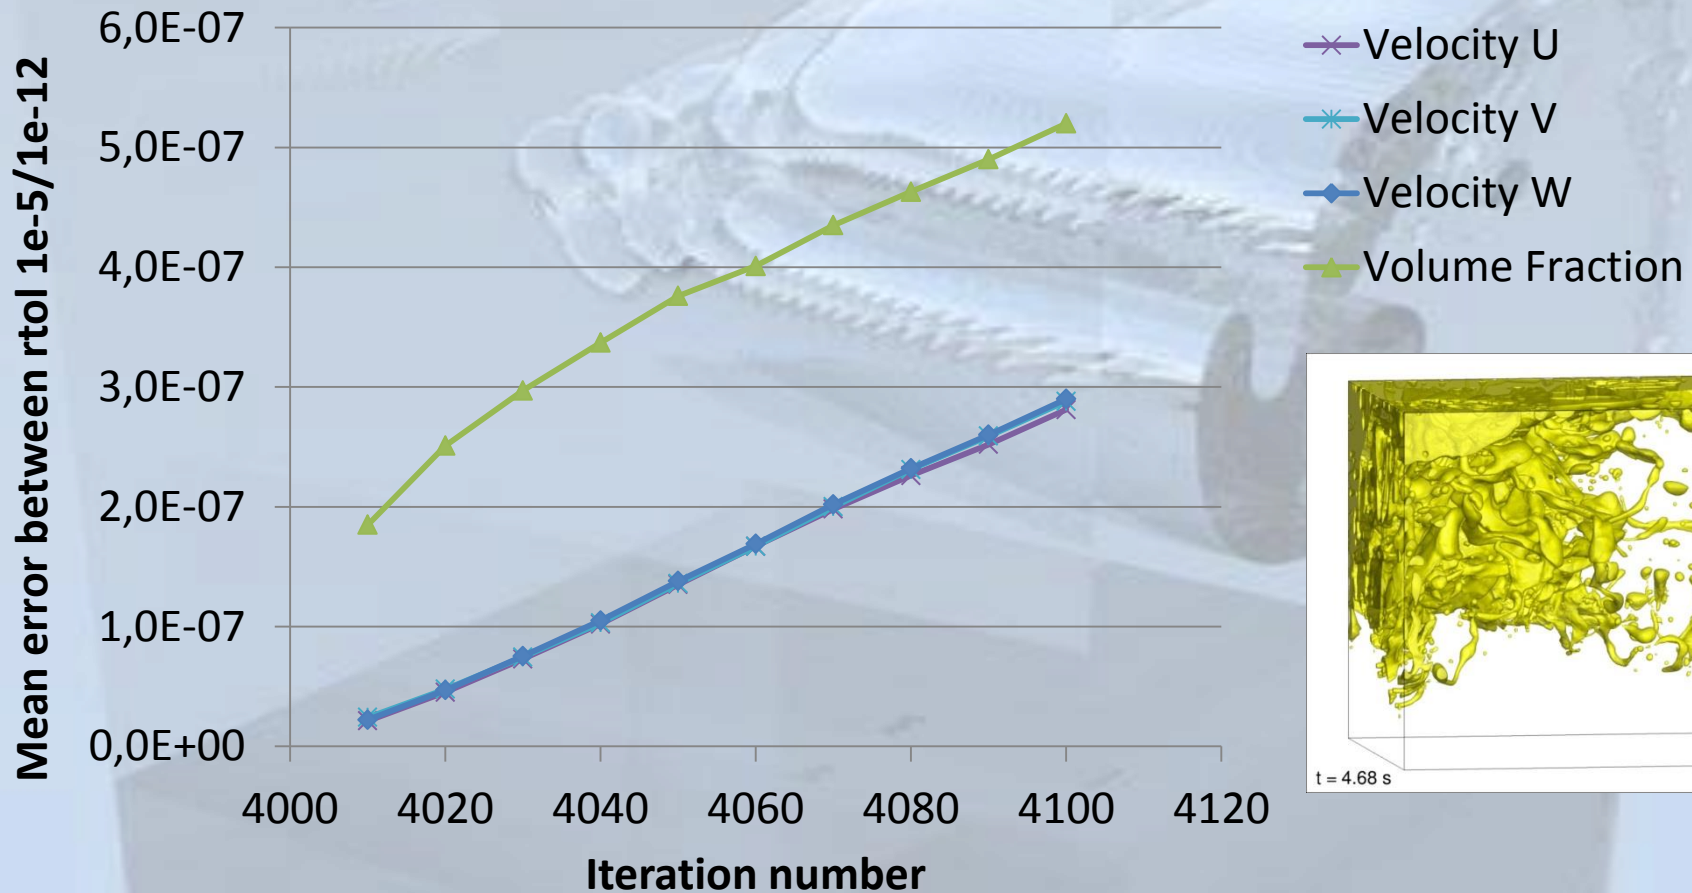
Elapsed time is not constant from one iteration to another.
It could explain variability of performance tests.



Performance tests induce different results depending on the physics.



Mean cumulative error for $\text{rtol}=1. \cdot 10^{-5}$



$\text{rtol} = 1.0 \cdot 10^{-5} \Rightarrow$ mean absolute error = $1 \cdot 10^{-4}$ on volume fraction after 20000 iterations

Jadim improvement thanks to the Workshop

- Tuning
 - Which MPI?
 - Effects of placement
- Memory
 - On 512 cores on Hyperion : Mesh 512^3 -> 4 Gb / core
 - No more possible on EOS -> work to save memory space
- Visualisation
 - Visualisation of large data with Tecplot
 - Post-treatment with Blender to get realistic pictures

BullX MPI vs Intel MPI?

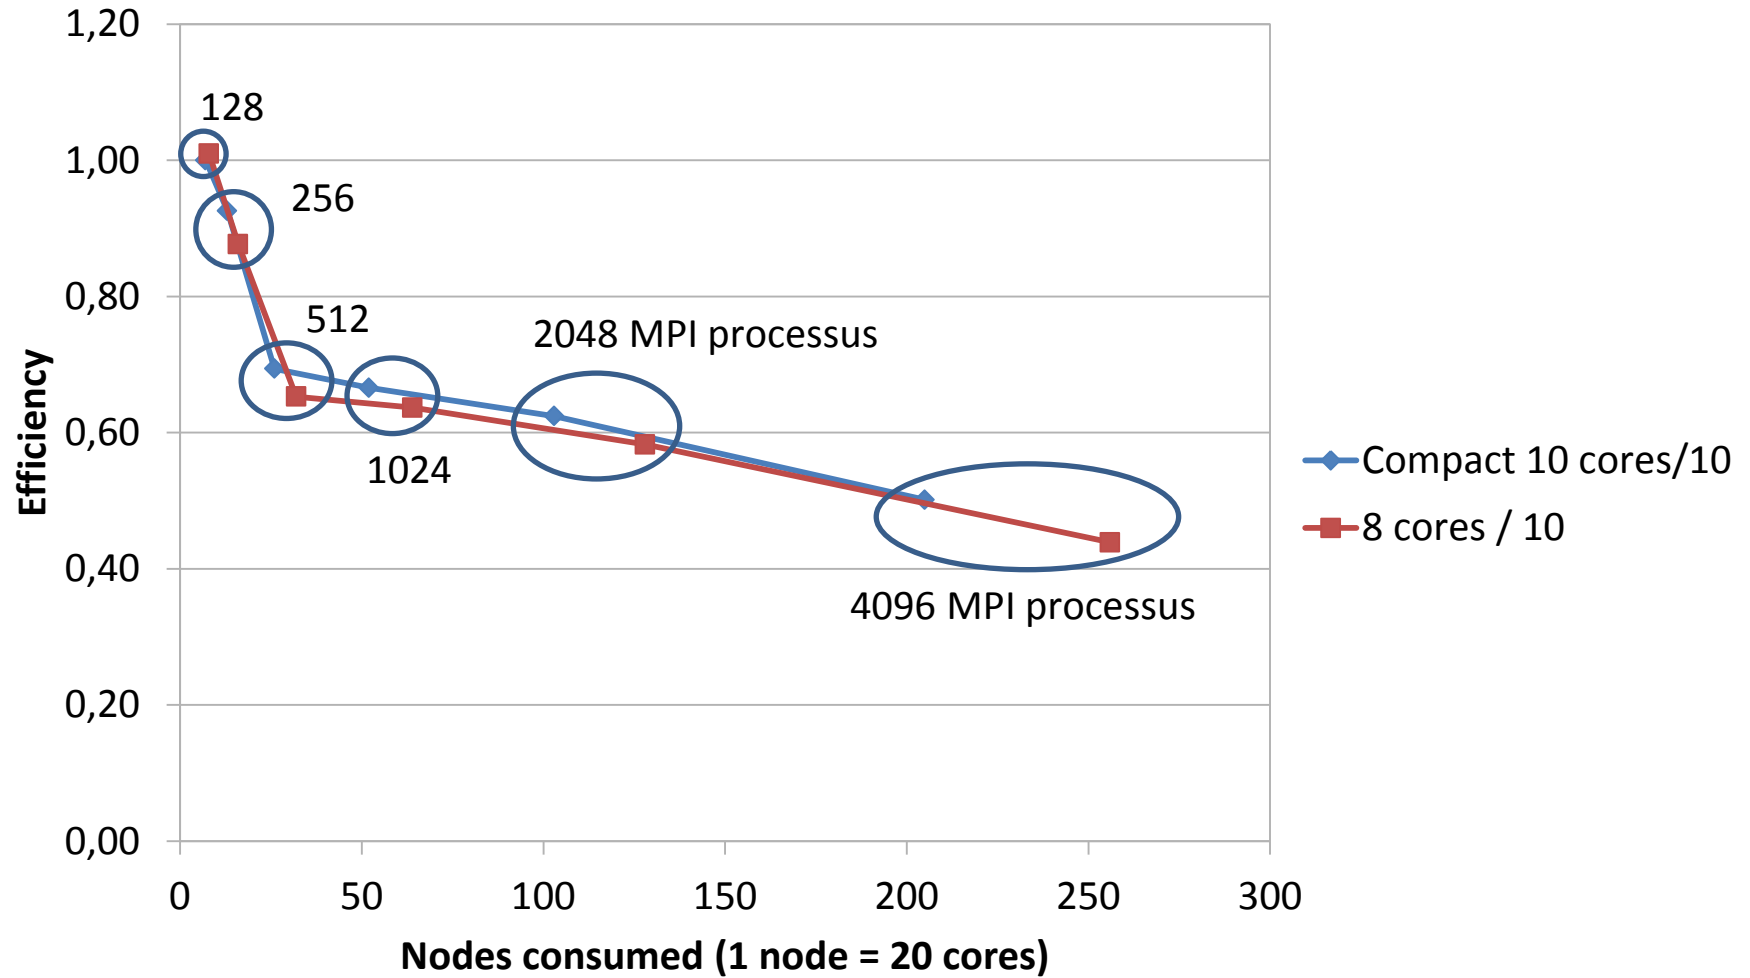
PETSc uses a lot of MPIAllReduce => Intel MPI is better than BullX MPI for this operation.

Tendency increases with the number of cores.

	Intel MPI	BullX MPI
512 cores 26 nodes (20 cores/node)	413 s	450 s
1024 cores 52 nodes (20 cores/node)	227 s	268 s

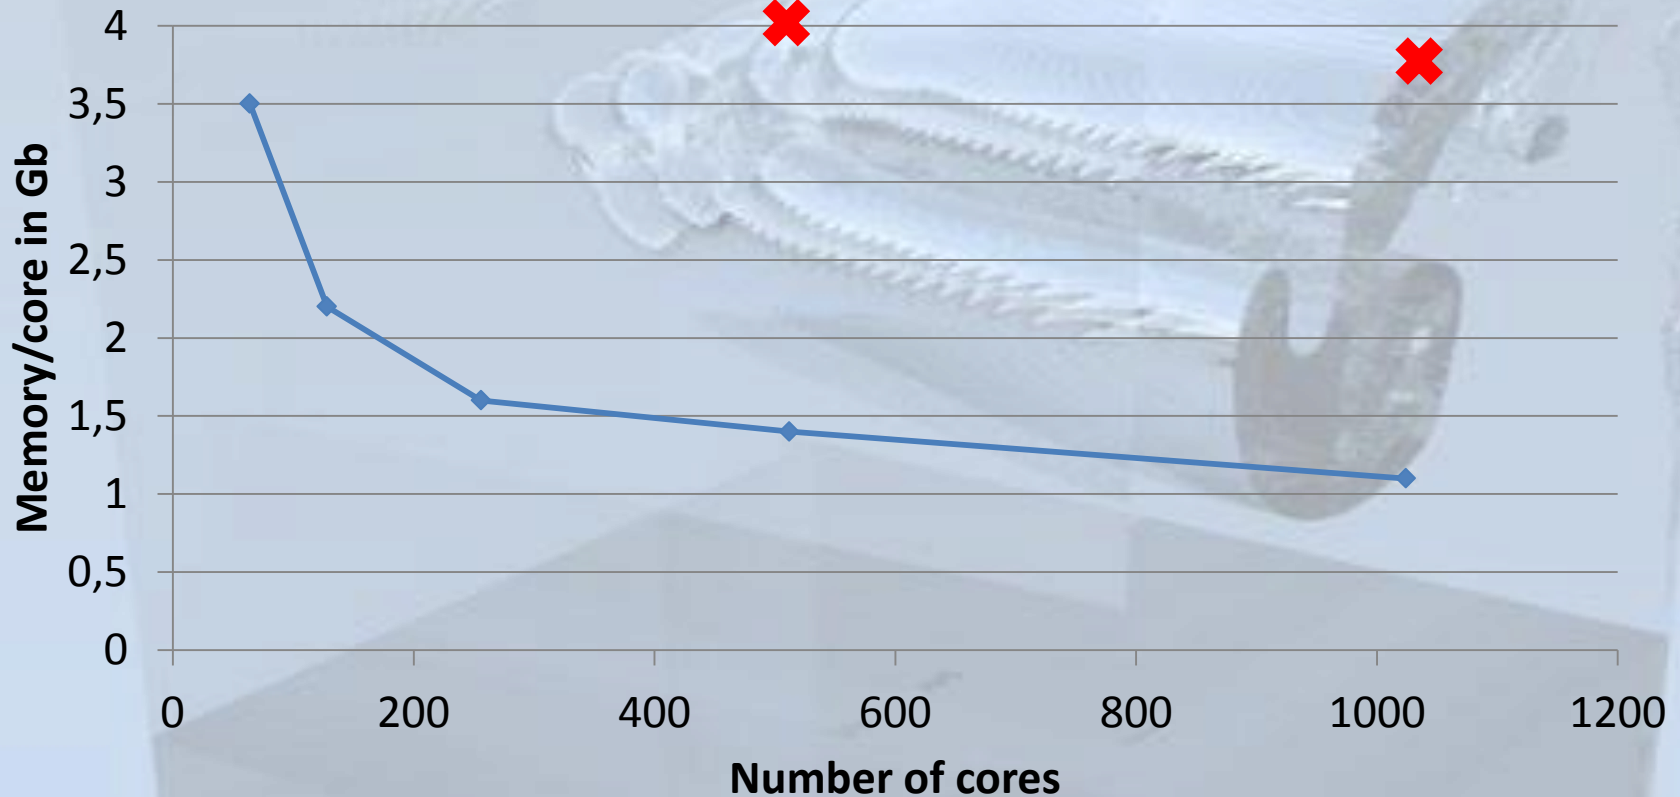
Elapsed time for 20 iterations of inversion phase problem.

Effects of placement



Memory /core

Residual memory/core




❌ 512³ before improvement (hyperion) ◆ 512³ after improvement (EOS)

But we can do better : Thetis uses only ~400 Mo by core on 1800 cores!

Visualisation

- Before this workshop, 512^3 mesh visualisation was realised on a coarse mesh (1 point on 2)
- All points are now visualised thanks to a new version of Tecplot and a new binary format (F. Auguste Tecplot Webminar coming soon...)
- Export stl file with Paraview in Blender
- Test for a new solution of a remote visualisation session on EOS (TurboVNC + VirtualGL) next week




OIL/WATER PHASE INVERSION

JADIM code

Groupe Interface :
F. Auguste
J. Magnaudet

Service CoSiNus :
A. Pedrono
H. Neau





Thank you for your attention.
Thanks a lot to CALMIP team!