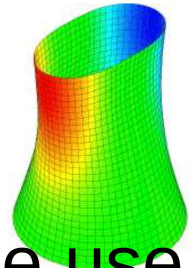

Feedback in the use of MUMPS package in EDF codes

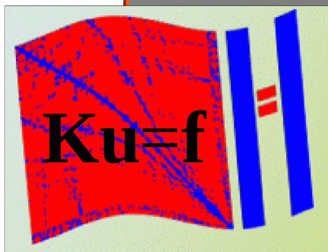

O.Boiteau (EDF Lab Paris-Saclay)
MUMPS User Days
June 1-2, 2017

1a. EDF Group: a European Electricity Utility with strong R&D involvment

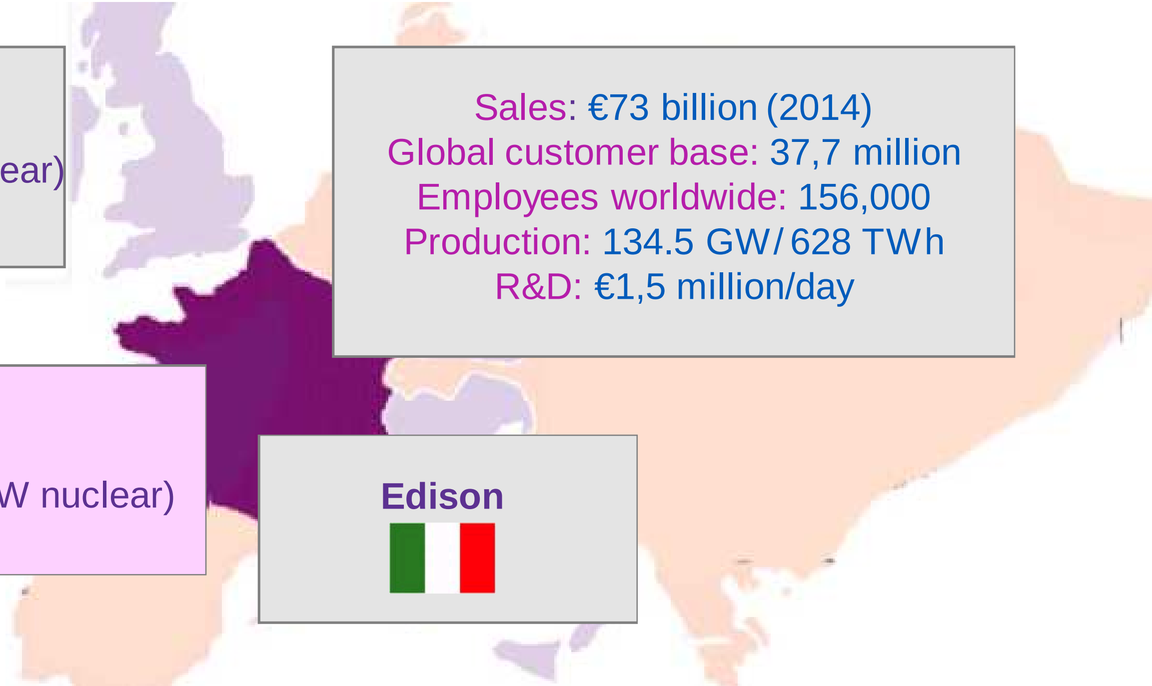

EDF Energy

Capacity: 13.8 GW (9 GW nuclear)

Sales: €73 billion (2014)
Global customer base: 37,7 million
Employees worldwide: 156,000
Production: 134.5 GW / 628 TWh
R&D: €1,5 million/day

EDF

Capacity: 101 GW (63 GW nuclear)

Edison

And many others:
Austria, Hungary, Netherlands, Poland, Slovakia, Switzerland, Brazil, USA, China, Vietnam…

1b. Operation, Maintenance & Optimization of complex systems at EDF

- **Permanent objectives**
  - ✓ Guarantee safety,
  - ✓ Improve performances/costs,
  - ✓ Maintain assets.
- **Changing operating conditions**
  - ✓ Face unexpected events, ageing issues, maintenance,
  - ✓ Improve performance through new technologies, new operating modes and system-wide optimization,
  - ✓ Adapt to evolving sets of rules (safety, environment, regulatory).

▪In-house technical backing

- **✓ Expertise: strong Engineering and R&D divisions,**
- **✓ Physical testing and simulation are key tools from the outset.**

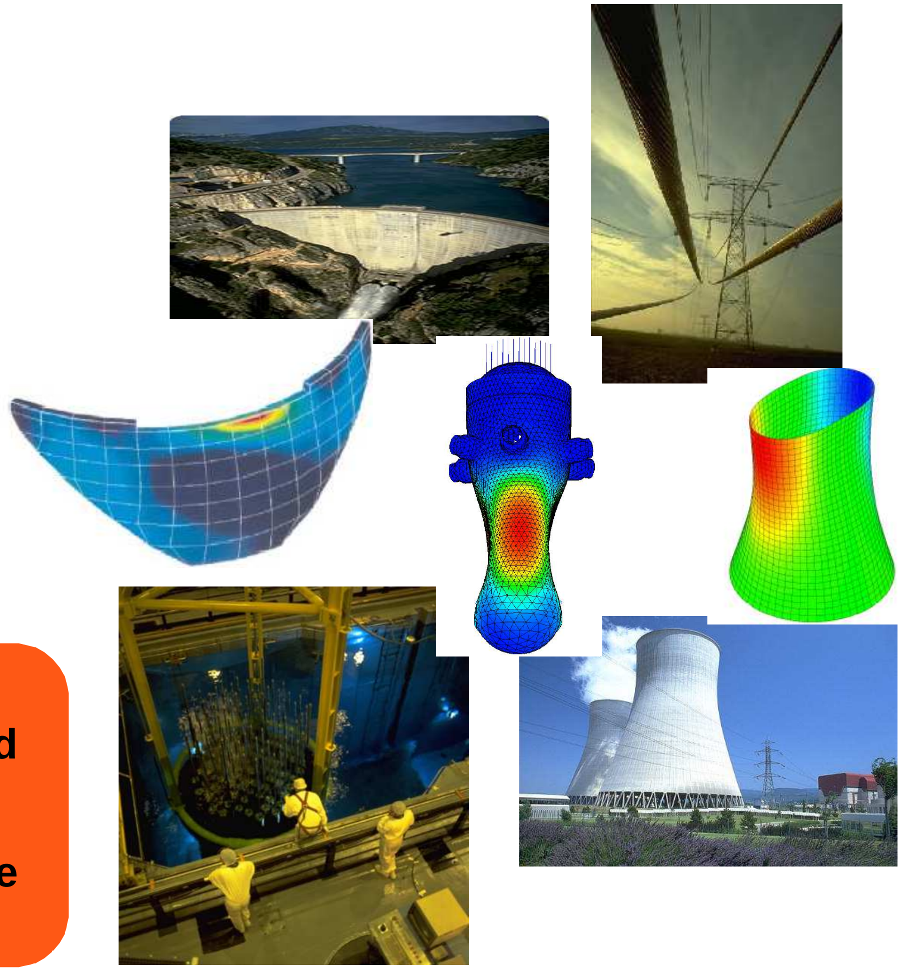

1c. EDF R&D worldwide presence

With 3 sites in France and 7 abroad, EDF conducts research both at national and international levels

2 100 people

370 doctors

180 PhD. students

200 research fellows from Universities and other higher education establishments

R&D Partners

- International academic support
- Major international partnerships
- Common laboratories & institutes
- National research bodies

1d. Examples of needs for numerical simulations

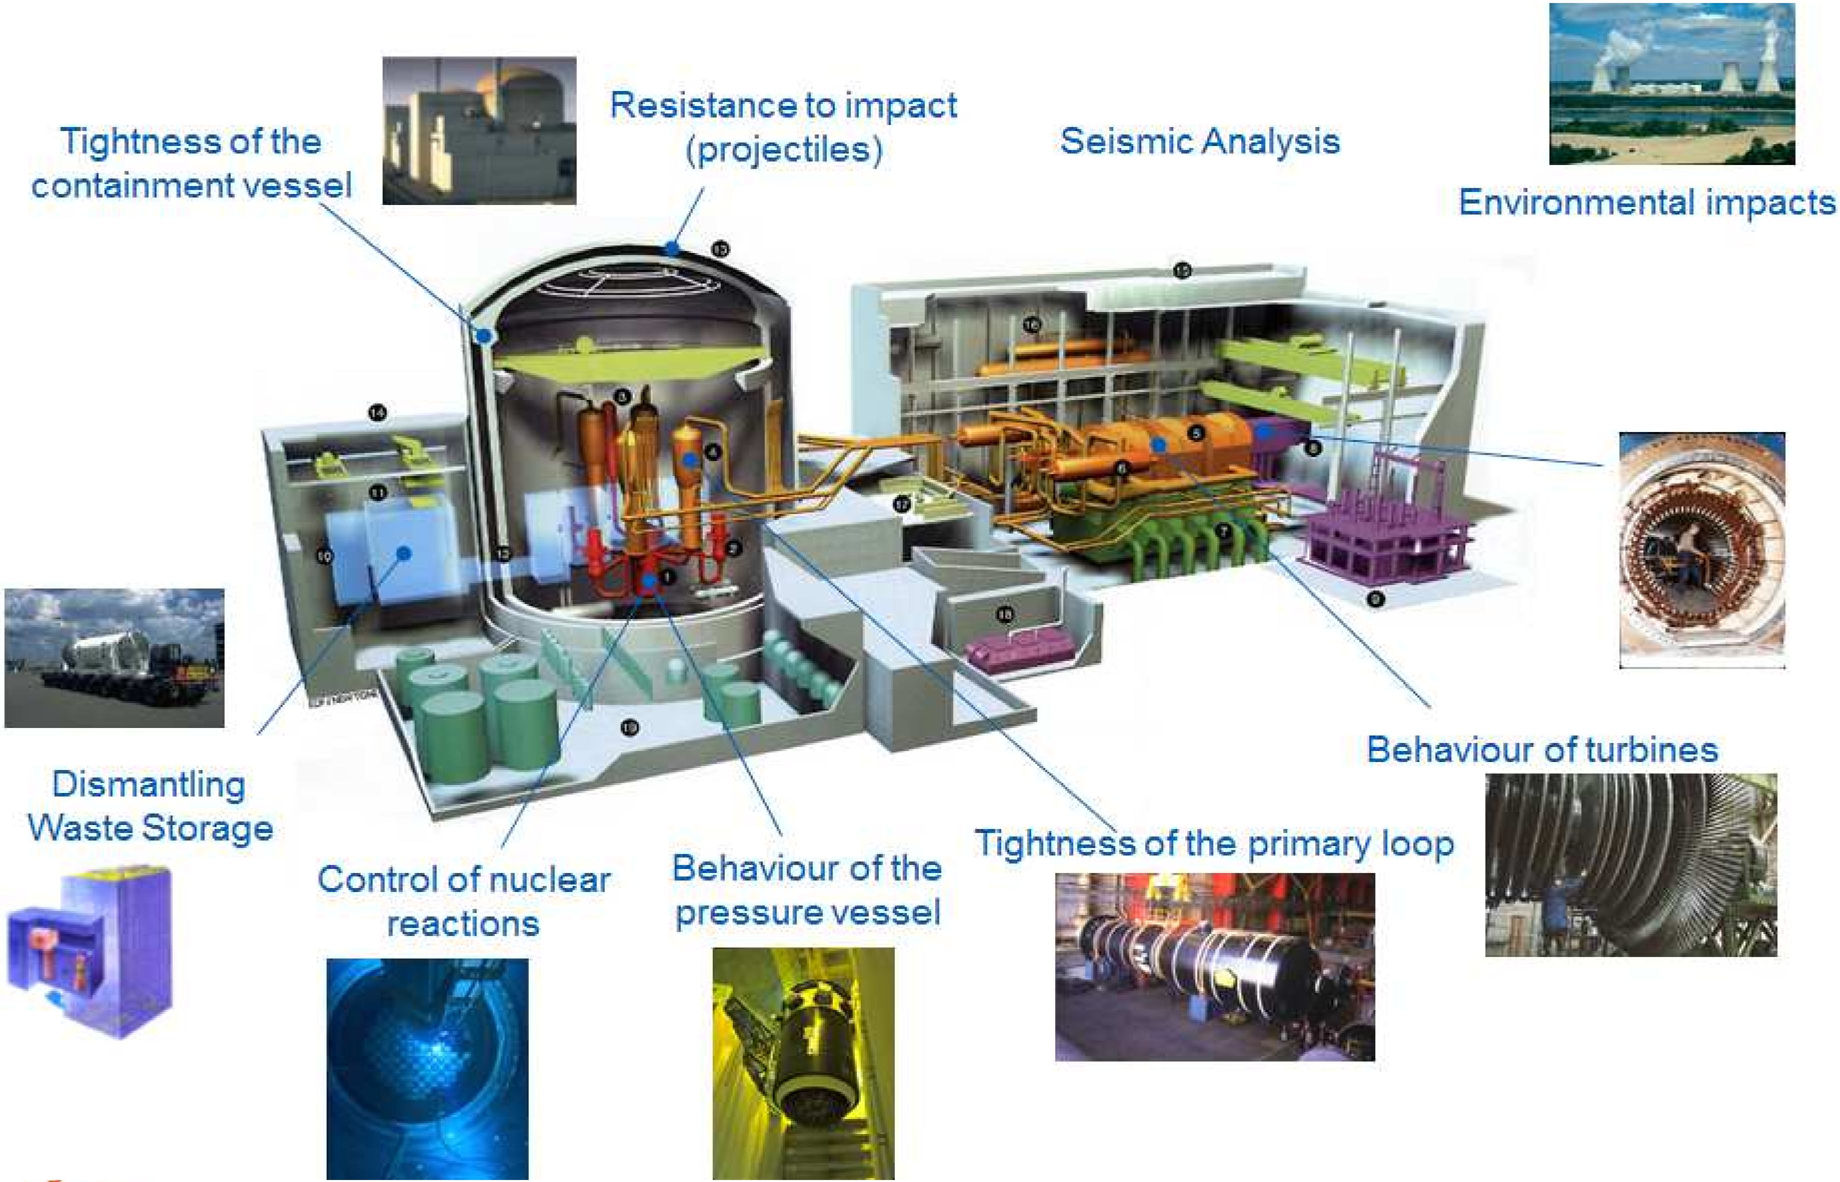

2a. Computational softwares used by engineers, experts and researchers

- All-purpose tools
  - ✓ **Studies**: user-friendly, highly versatile…
  - ✓ **Researches:** continuous integration of new models/methods, prototyping…
  - ✓ **Quality Management:** robust/reliable, tested/qualified (V&V)…
- Research codes
- Prototypes

Often in-house open-source codes,

not « black-box » closed sources ones

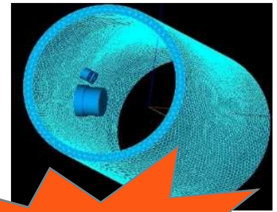

Electromagnetics,

NDT code

code_Carmel

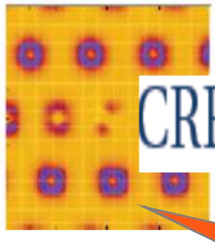

CRESCENDO

Material's structure research code

But also, CFD, neutronics for nuclear…

code_aster

Thermomechanical code

telemac

Hydraulic code

epx europlexus software

Fast transient dynamics

Code_Saturne

NEPTUNE

eDF

2b. Typical workflow

Mesh generation, material, loadings…

Data distribution

Discretization time, space

Linear solver

Post-processing, visualization

e.g. Non linear loops

e.g. Time loop

Bottleneck in CPU time and RAM peak

Needs for robust, user-friendly, versatile and efficient…

sparse linear and eigenvalue solvers

Less in-house developments/researchs at EDF

More and more links to external packages

Lot of works in test/benchmark, integration, optimization, validation

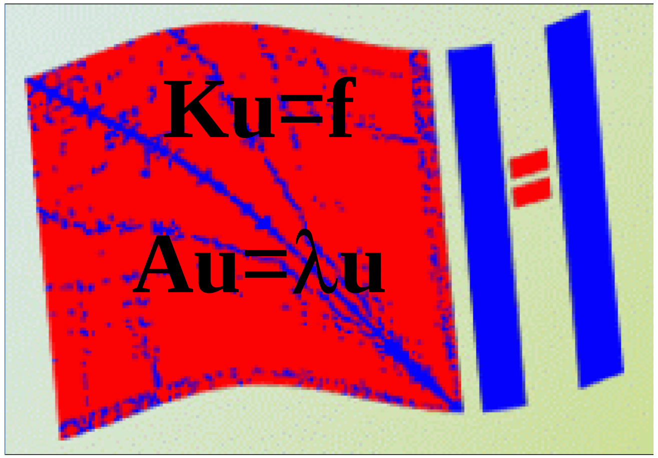

Components may be crucial for performance and feasibility of a study

2c. Question:
Is it a good idea to use external libraries ?

- **Search of optimal performances**
  - ✓ Trade-off between CPU times/ RAM consumption/stability/robustness,
  - ✓ We would probably implement less efficient algorithms.
- **Ecomical reasons**
  - ✓ Fewer codes that are not ‘core business’,
  - ✓ Collection of tremendous know-how.
- **Pragmatism and QA management**
  - ✓ Share feedback,
  - ✓ Good pratices.
- **Use a well-known package**
  - ✓ Share risks,
  - ✓ To find other partnerships,
  - ✓ To find trained applicants/users.

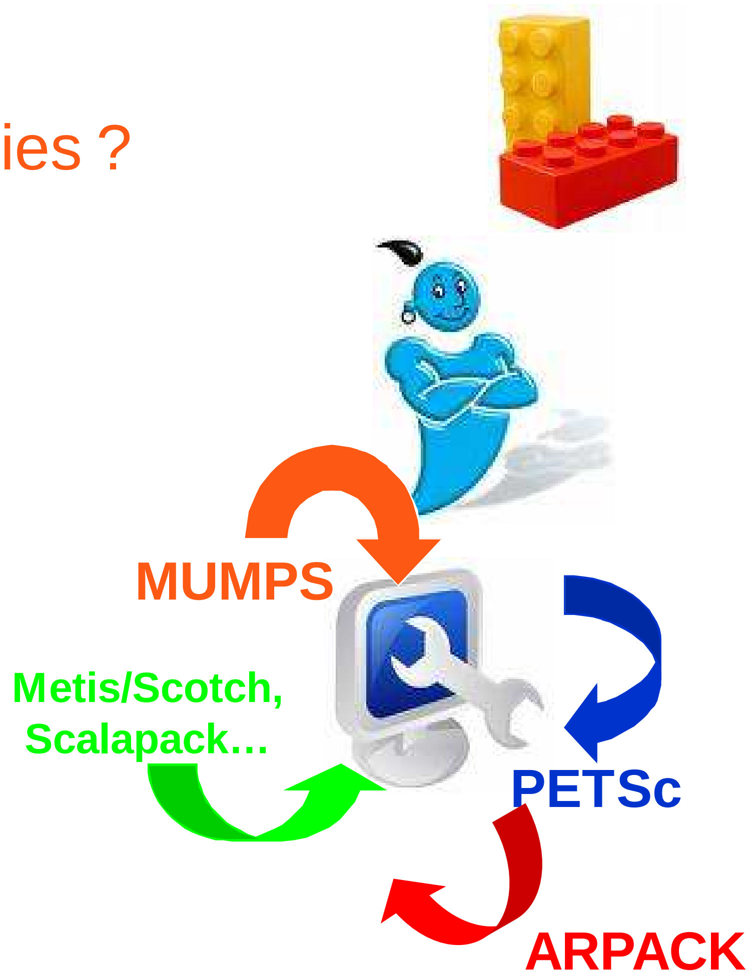

MUMPS-Consortium
A consortium of MUMPS-Solver users

CERFACS INP Toulouse Inria ENS de Lyon université de Bordeaux

2d. Our answer: yes, when it's possible ! Mainly based on MUMPS (& PETSc)

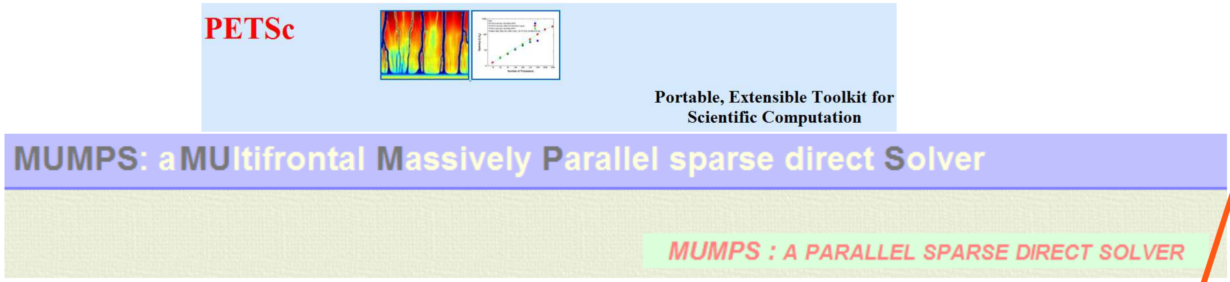

11-year fruitful and win-win partnership

EDF>MUMPS:

- Functional/numerical feedback,
- Bug report/industrial validation in our QA in-house codes.

MUMPS>EDF:

- Numerical expertise,
- Tips and tricks.

EDF supports innovation and research:

Ex: PhD about low-rank compression

Intensive industrial use

Multicore desktop computer and super-computer

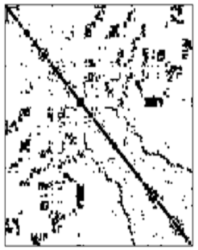

MUMPS-Consortium
A consortium of MUMPS Solver users

MUMPS User Group Meeting
Clamart
May 29th and 30th 2013

First annual Consortium Committee Meeting March 12, 2015 at EDF, Clamart, France

2e. Our answer: in order to provide a wide range of strategies

Linear system size

10^3 — 10^6 — 5.10^7 — 10^8

In-house direct solvers: CHOL, MULT_FRONT, SCHUR… DENSE…

MUMPS

Out-Of-Core

Two-level parallelism: MPI x OpenMP

Numerical pivoting/scalin

Use as a preconditioner in single precision with PETSc

BLR compression

New!

?? Hybrid algebraic linear solver ??

MaPHyS, ABCD_Solver etc.

PETSc

ML/BOOMER/GAMG

In-house preconditioners

In-house iterative solvers: GPCG, GCPG, SPLI, CROUT/JACOBI…

3a. MUMPS: versatile linear solver tool

- **Wide range of linear systems**
  - ✓ **Sparse matrices from discretization schemes,**
  - ✓ **Size: between 10^5 and 10^7 DOF and even more**,
  - ✓ Symmetric or not, mainly real,
  - ✓ No special feature or pattern, flat matrices.

- **Robustness and numerical skills**
  - ✓ **Often indefinite matrix:** mixte FE, Lagrange multipliers…,
  - ✓ Detection of singular matrix: lack/excess of BC, eigenvalue problem, null space analysis…,
  - ✓ Accuracy management: error analysis, iterative refinement, pivoting,
  - ✓ **Often poorly conditioned** (>1

- **Additional features**
  - ✓ Computation of determinant and **negative pivots (eigenvalue computation)**,
  - ✓ Out-Of-Core and compression facilities,
  - ✓ Preconditioner tools **(non linear studies).**

Mixing 3D FE, shells, beams, pipes, membranes, bars…

Many Lagrange multipliers (BC, contact-friction, modelling connections…)

Specific constitutive laws, non linearities

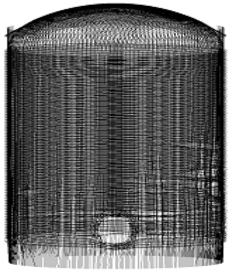

Pretension cables

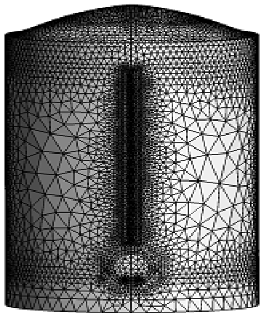

Concrete

Example from Code_Aster's studies
Reactor building model

3b. MUMPS: framework of parallelism

- **Pros**
  - ✓ Robust, all-purpose and user-friendly,
  - ✓ **Can be easily used in more ambitious frame**:
    - » Preconditioners (ex: PETSc+MUMPS).
    - » MPI+OpenMP (ex: MUMPS+BLAS),
    - » spectrum slicing (ex: ARPACK+MUMPS).
- **Cons**
  - ✓ Medium range (<1000 cores),
  - ✓ Not totally scalable (efficiency between 0.2 and 0.5),
  - ✓ **Need of memory** (RAM/disk),
  - ✓ Need of large enough problem size,

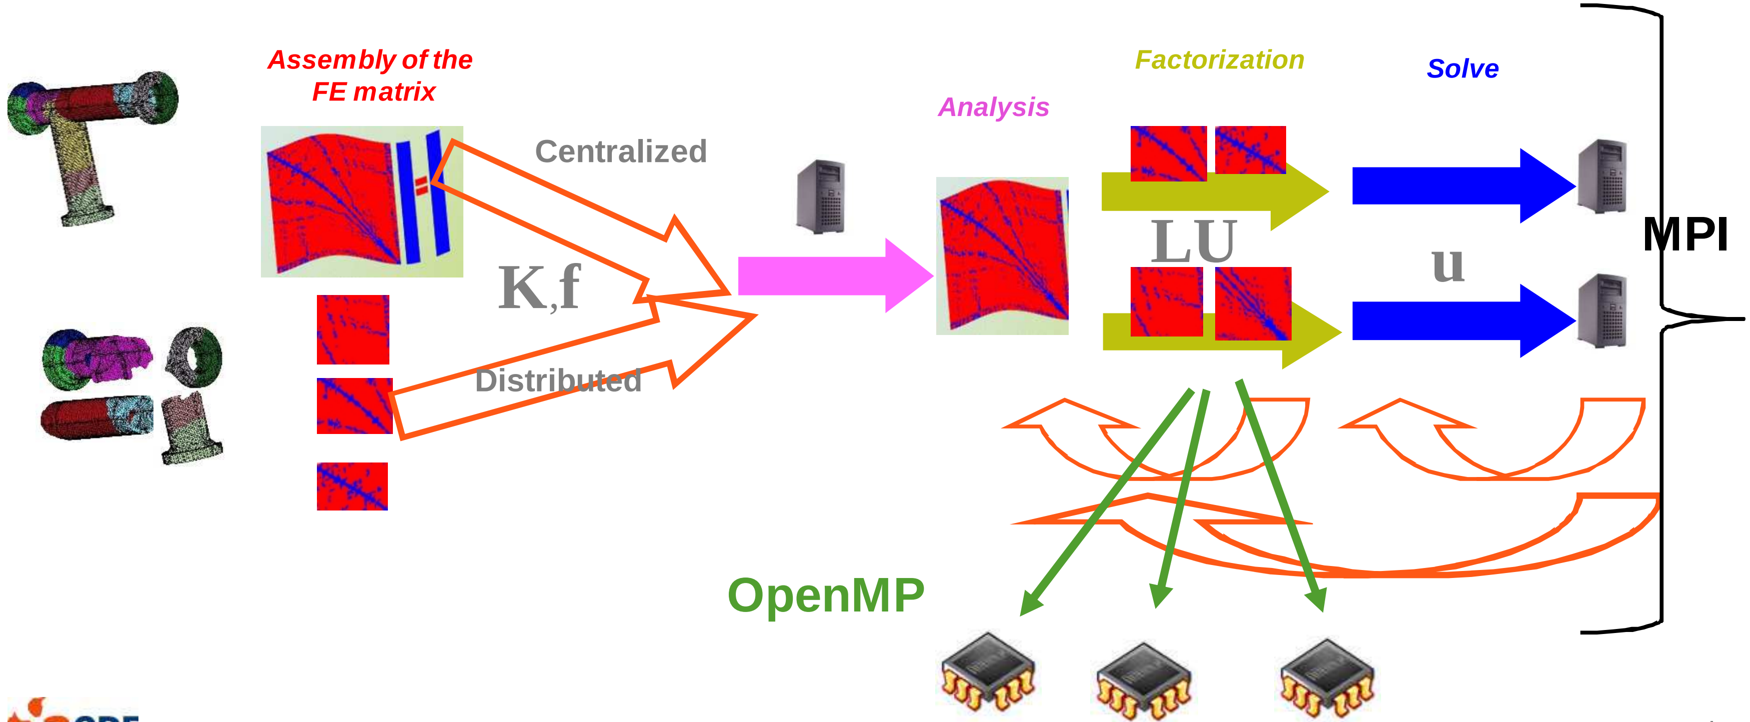

3c. MUMPS: examples of use at EDF

Real unsymmetric

Size: growing systems, <5 Millions

Multiple RHS coupled with EDO solvers: SUNDIALS/CVODES

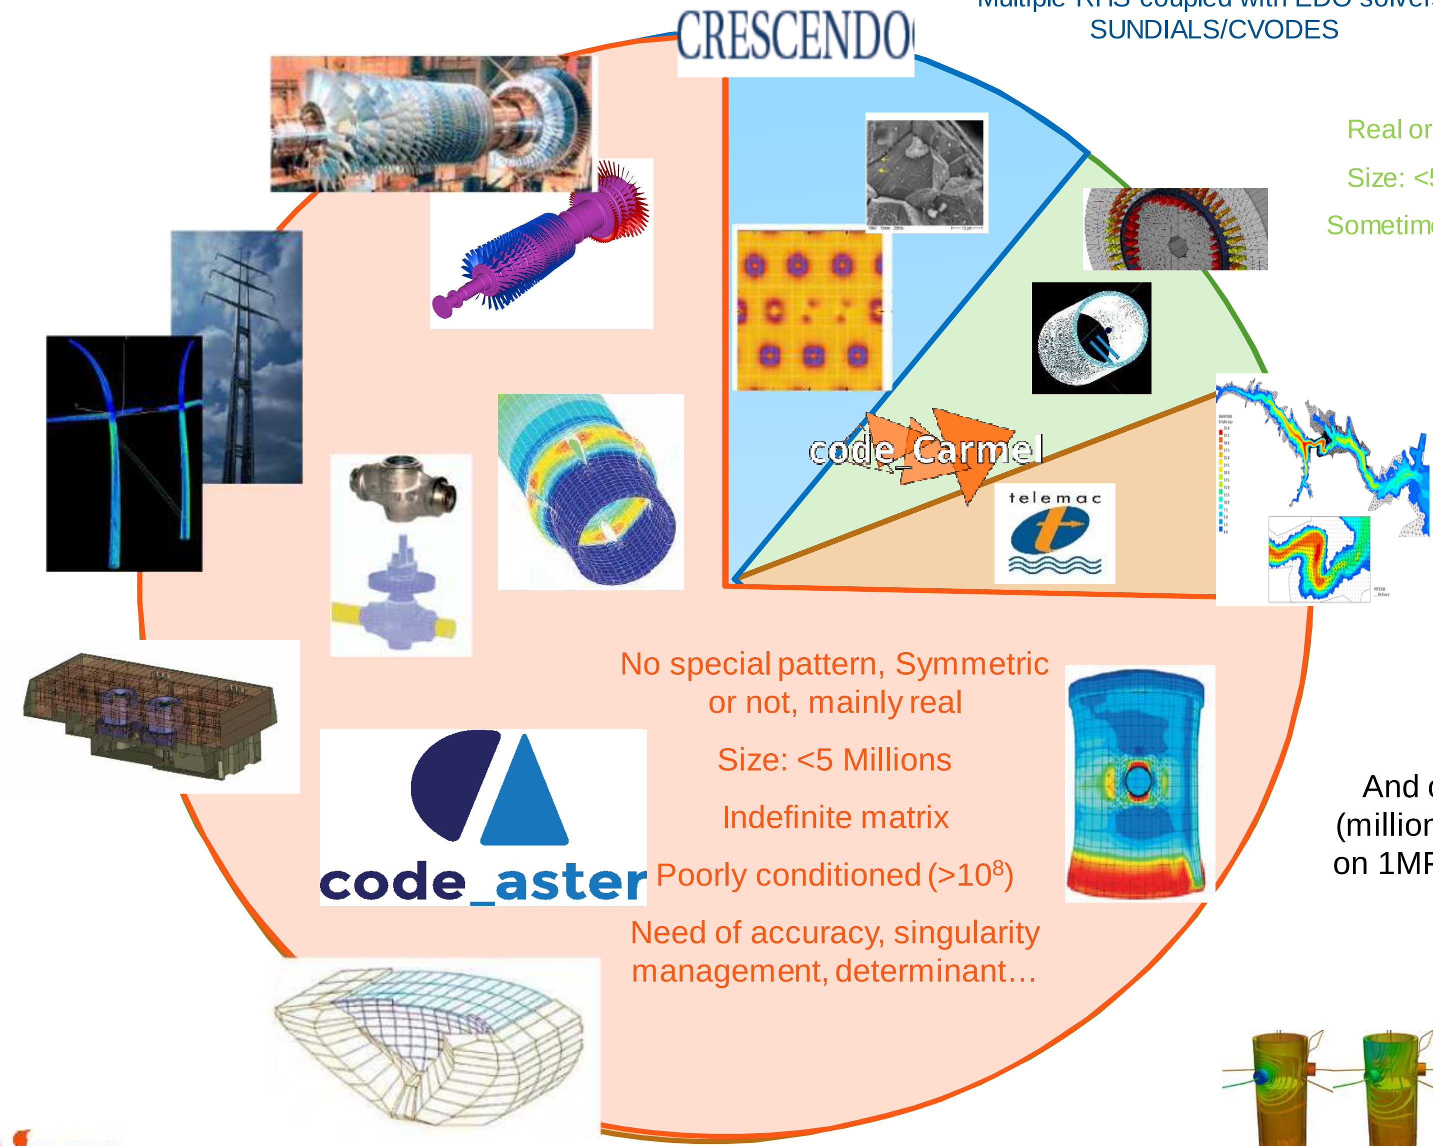

Real or complex

Size: <5 Millions

Sometimes singular

Real symmetric

Poorly conditioned

Drawback: build the assembled matrix

And others: test with EPX (millions of very small system on 1MPI), Labo Saint-Venant, STEP…

4a. Benchmark MUMPS versus MKL-PARDISO

- **LOCA studies: ‘Loss of Coolant Accident’**
  - ✓ **Lagrange multipliers:** Links between fluid and structure meshes,
  - ✓ Multitude of local and global very small linear systems,
  - ✓ **Different strategies:** local/global systems, direct (`CHOL`, `PARD`)/iterative linear solvers (`GPCG`, `SPLI`) …
- **Linear solver benchmark:**
  - ✓ `CHOL, PARD(OPT), MUMPS.`
  - ✓ MUMPS: fallback position to PARDISO.
  - ✓ Useless here: BLR, threads.
- **Speed-up of this EDF’study: X7** (`PARD_OPT/MUMPS` versus `CHOL`): **20 days**.

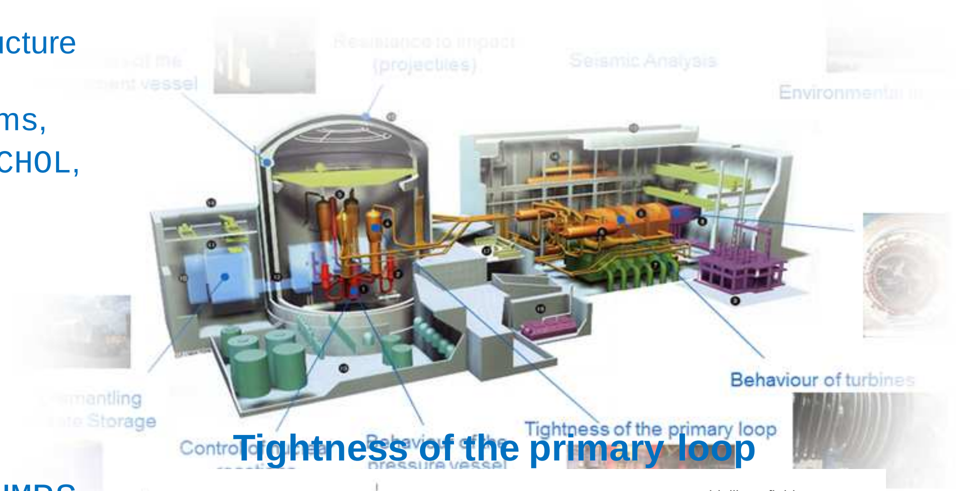

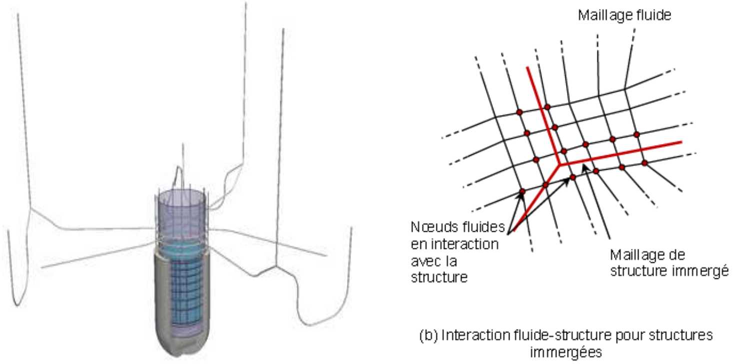

(b) Interaction fluide-structure pour structures immergées

Portion of LOCA study

(#dof=0.6M, #link=25000)

Linear system solve time

#MPI=12

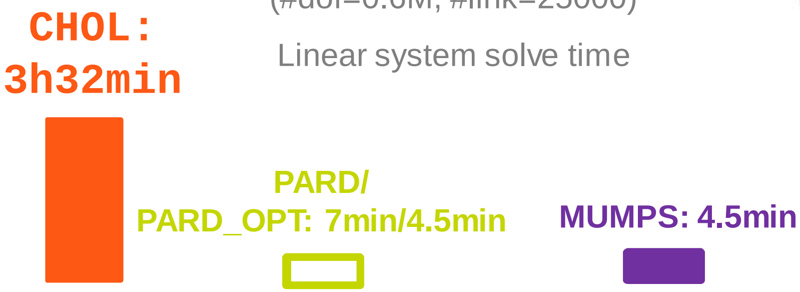

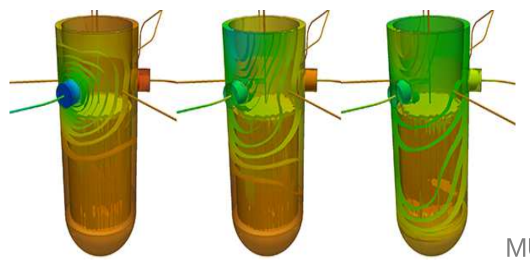

4b. MUMPS vs (F)GMRES_PETSc+MUMPS

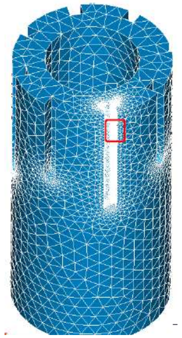

Model of graphite reactor core

Non-linear study (#dof=2.7M)

5 days 17 hours

<1 day

#MPI=48

MUMPS

PETSc+MUMPS

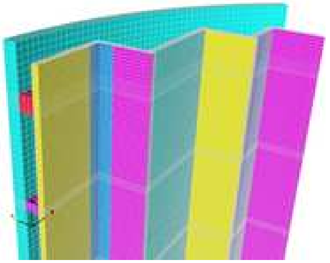

Inner panels of the reactor vessel

Big non-linear study (#dof=6.7M) ,

Some Newton steps

#MPI=1

3 hours 20 mins

1 hour

#MPI=4

Strategy PETSc+MUMPS doesn't converge (contact-friction) !

4c. MUMPS vs (F)GMRES_PETSc+MUMPS

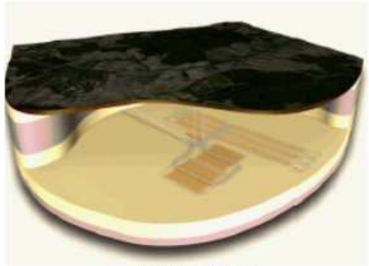
Image Andra

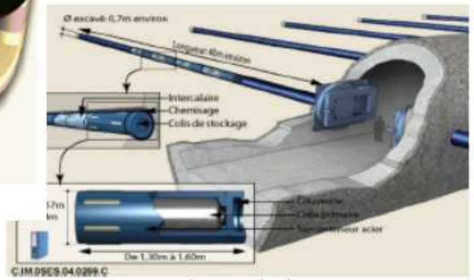
Image Andra

Waste storage

Non linear thermohydromechanics

Large excavation model

#TETRA10=1.3M

#dof=5.2M unsymmetric

#NNZ=673M

#time steps=108

#MPI=24
#threads=7
#nodes_EOLE=6

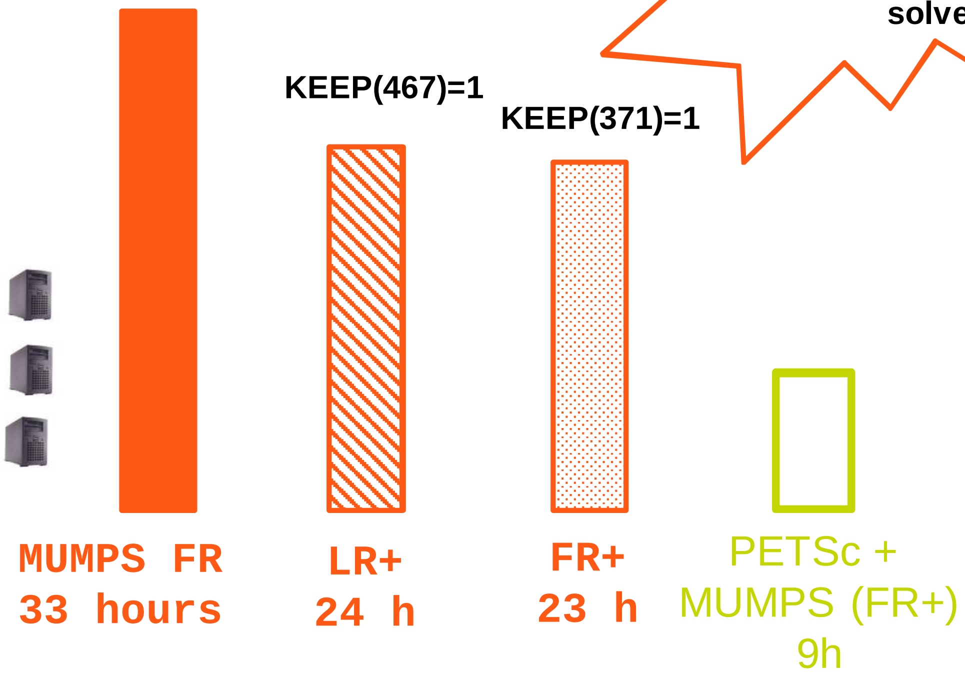

Breakthrough
Very large THM model
#dof=40M
#NNZ=5000M
#nodes_EOLE=24

5h per linear system solve !

4d. Parallelism in MUMPS

Strategy PETSc+MUMPS doesn't converge (condition number > 10^{15}) !

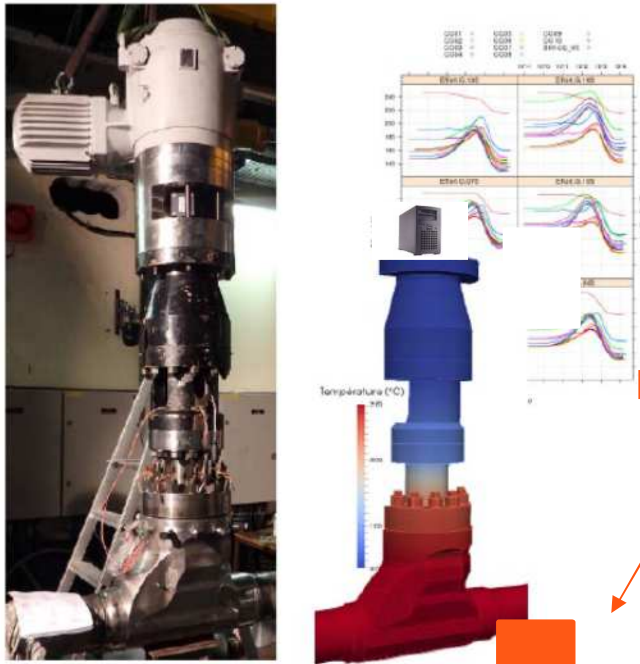

Thermo-mechanical behavior of valves

Thermal shocks,

Links between dof

Plasticity, Contact

Few time steps

#dof=1M

Factorization elapsed time

MUMPS direct solver

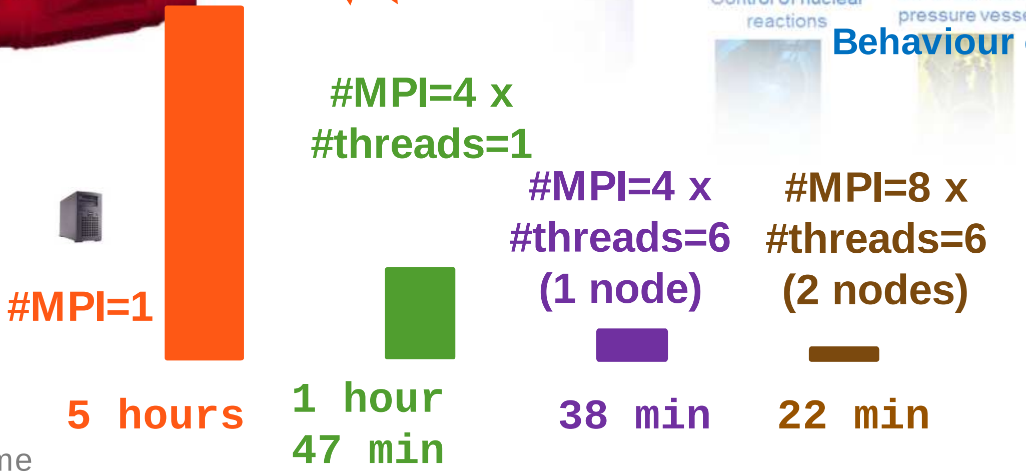

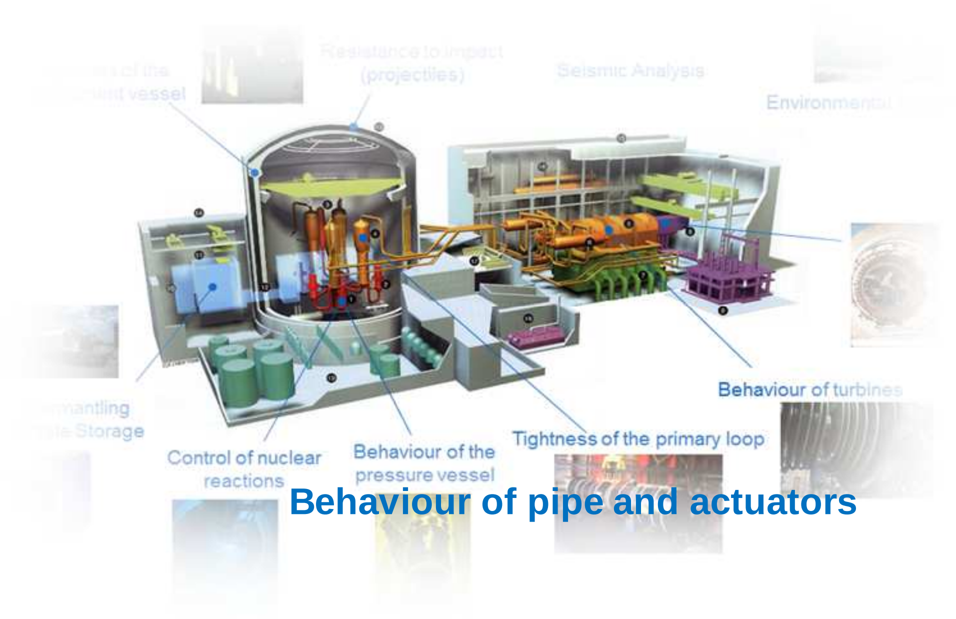

Behaviour of pipe and actuators

BLR strategy useless here
Matrix too small

4e. BLR(+) compression in MUMPS

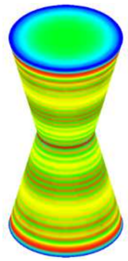

Perf008d (#dof=2M)

Elastic computation of a full diabolo subjected to a thermal loading

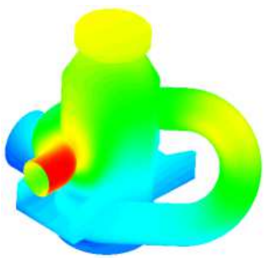

Perf009d (#dof=5.4M)

Elastic computation of a RIS pump subjected to internal pressure

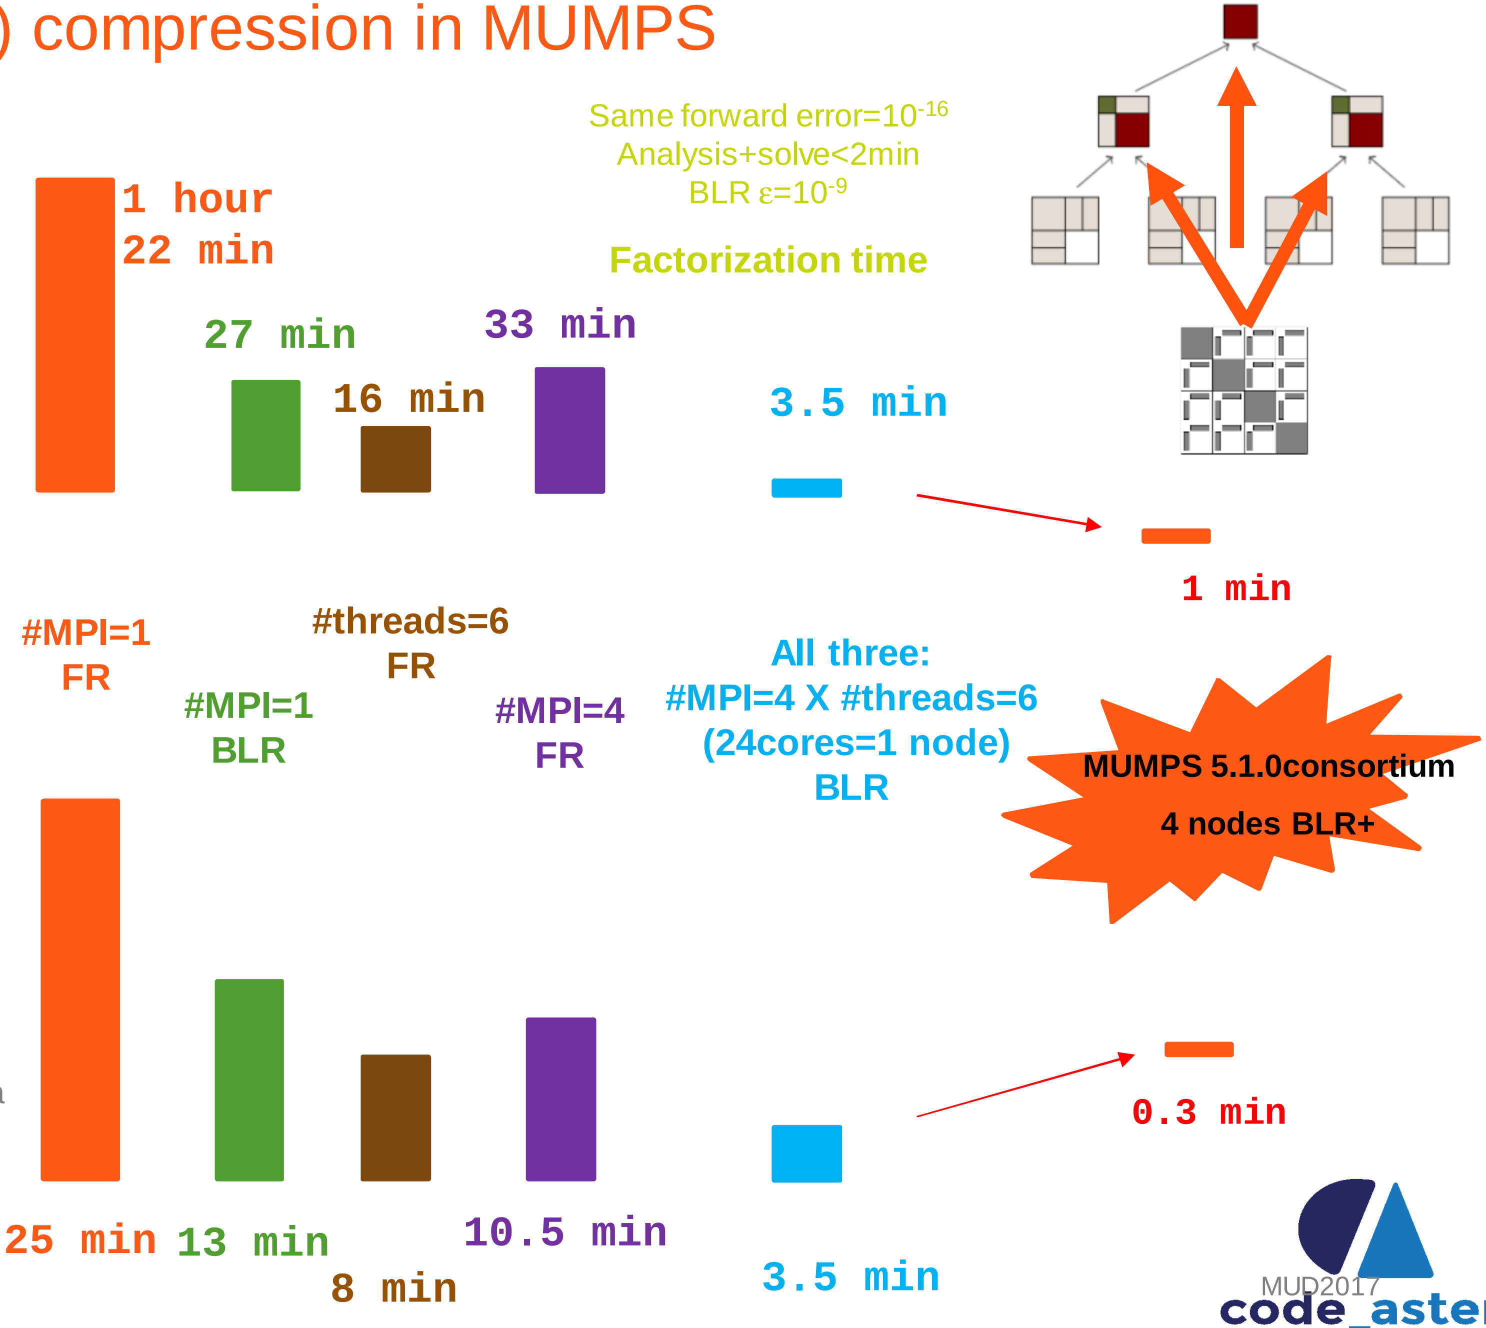

4f. Very large models: selective 64-bits integers

- Performance test-cases

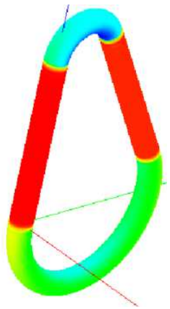

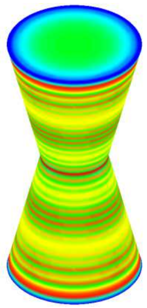

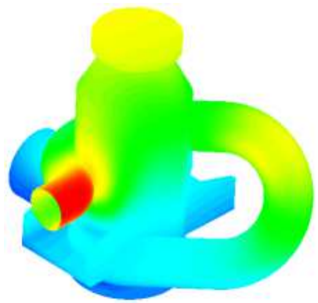

Perf002c (**#dof=58M**)	Perf008a (**#dof=31M**)	Perf008a (**#dof=39M**)
FR_parmetis=51min	FR_metis=2h9min	FR_parmetis=14min
LR_parmetis=29min	LR_metis=47min	LR_parmetis=18min

24 nodes
EOLE
256/512Go
One full MUMPS solve

- Industrial study stuck until now

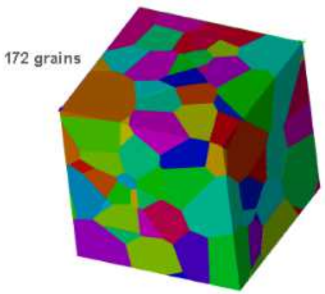

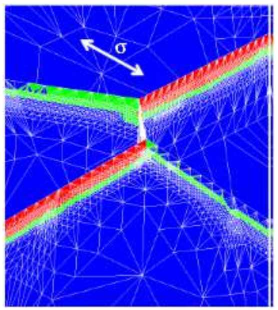

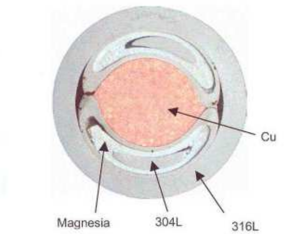

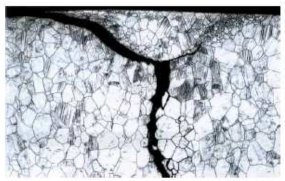

Steel subjected to corrosion (pressuriser)

Polycristal model

#dof=10M (unsymmetric)

#NNZ=2011M

FR_parmetis=29min

LR_parmetis=23min

4g. Large eigenvalue computation: ARPACK+MUMPS

- **Fluid-structure interactions in water tank of nuclear power plant**
  - ✓ Thin structure > extended eigenvector basis,
  - ✓ Computation of 6100 eigenvalues/vectors (#dof=50000).

- **Three level parallelism scheme**
  - ✓ Spectrum slicing with MPI and ARPACK,
  - ✓ MPI in MUMPS,
  - ✓ Threads in BLAS.

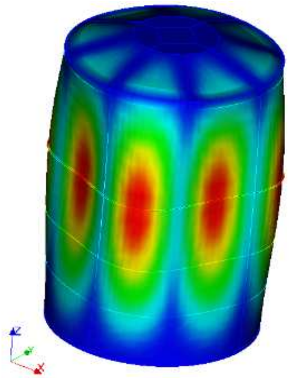

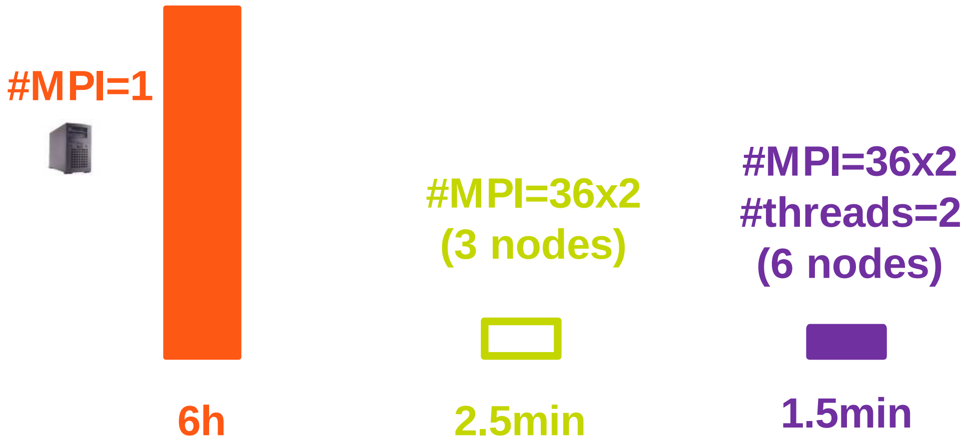

36 homogeneous frequency intervals in [0, 100Hz]

Generalized Eigenvalue Problem

4h. Large electromagnetic computation

- **3D non linear study of a turbo-alternator**
  - ✓ Newton non linear algorithm,
  - ✓ 17M Tetrahedra,
  - ✓ #dof=8M,
  - ✓ Thousands of linear systems to solve
  - ✓ Each MUMPS computation: 10min (4.5min for analysis)

- **In test**
  - ✓ BLR compression,
  - ✓ MUMPS as a preconditioner,
  - ✓ Mutualization of the tangent stiffness matrix between Newton steps.

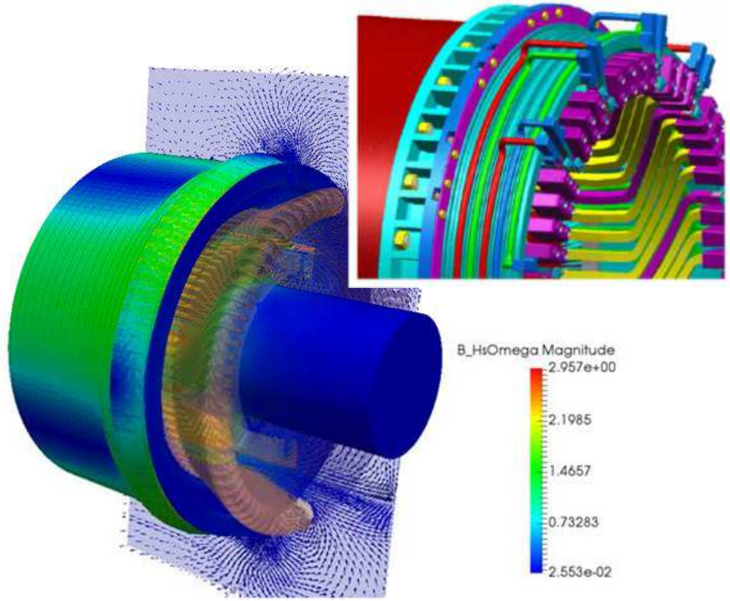

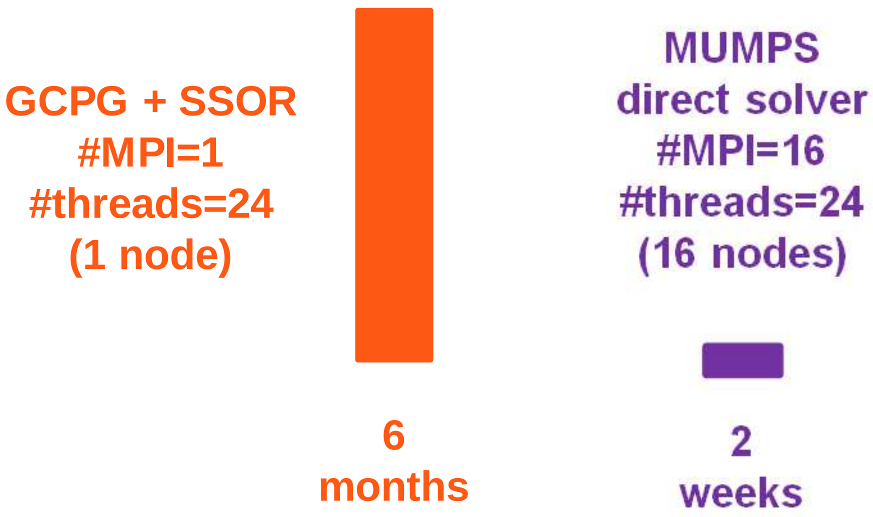

5. Conclusion and perspective

- **Daily use of MUMPS/PETSc in EDF's *in-house* codes ('best-in-class' tools).**
- **11-year fruitful and win-win partnership EDF-MUMPS,**
- **In EDF codes context, to achieve quicker and bigger computations:**
  - ✓ **Direct solver:** BLR variants, coupled with MPI-OpenMP and OOC, avoid of refactoring, management of Lagrange multipliers…
  - ✓ **Hybrid algebraic linear solvers:** MaPHyS, ABCD_Solver…
  - ✓ **Better use of preconditioner**: BLR double precision with strong compression, multigrid, DD…

- **Links « in-house code » - « linear algebra package » needs steady adjustments. Often, questioning about external libraries induces improvment in the caller code.**

Thank you for you attention !