Feedback on the use of MUMPS in Safran Tech's applications:

a multithread performance evaluation of MUMPS

Ramzi MESSAHEL
Julien CORTIAL
Augustin PARRET-FREAUD
Felipe BORDEU

Safran Tech / DST / MADS / CATS

MUMPS User Days
22-23 June 2023

SAFRAN

Agenda

Chapter 1

Context

SAFRAN

Safran Tech: from research labs to design departments (and vice versa)

Industry

High TRL

Scientific research

Low TRL

Aero Boosters
Aerosystems
Aircraft Engines
Cabin
Electrical & Power
Electronics & Defense

SAFRAN

Helicopters Engines
Landing Systems
Nacelles
SA
Seats
Transmission systems

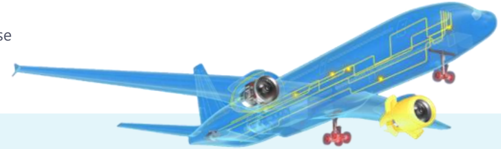

Companies goals

- Shorten development cycle
- Decrease conservatism
- Decrease numerical wasting
- Optimise products and processes
- Decarbonize aviation

Safran Tech

Dedicated research center to complement the R&T works done by Group businesses

- **link between Group and academic research**
- upstream research, long-term innovation
- three scientific departments
- ~ 400 research scientists

Digital Science & Technology department

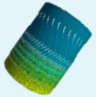

- Faster and better design of more complex parts in more demanding environments
- Shorten crisis analysis and design time
- Model and compute « right at the right cost »
- Comp. solid/fluid mechanics, applied mathematics, system engineering

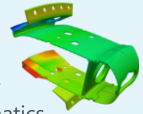

Research topics

- Fine high fidelity computations for non-linear single and multiphysics applications
  - HPC, complex boundary conditions, multiphysics, hight fidelity, non-linear behaviours, damage, cracking, turbulence, adaptative meshing…
- Simplified and mastered methods
  - model reduction, multiscale methods, structural zoom, model coupling, link with system engineering
- A cascade of model suited to real needs

Safran : ~82 000 people worldwide. Everything but the plane itself…

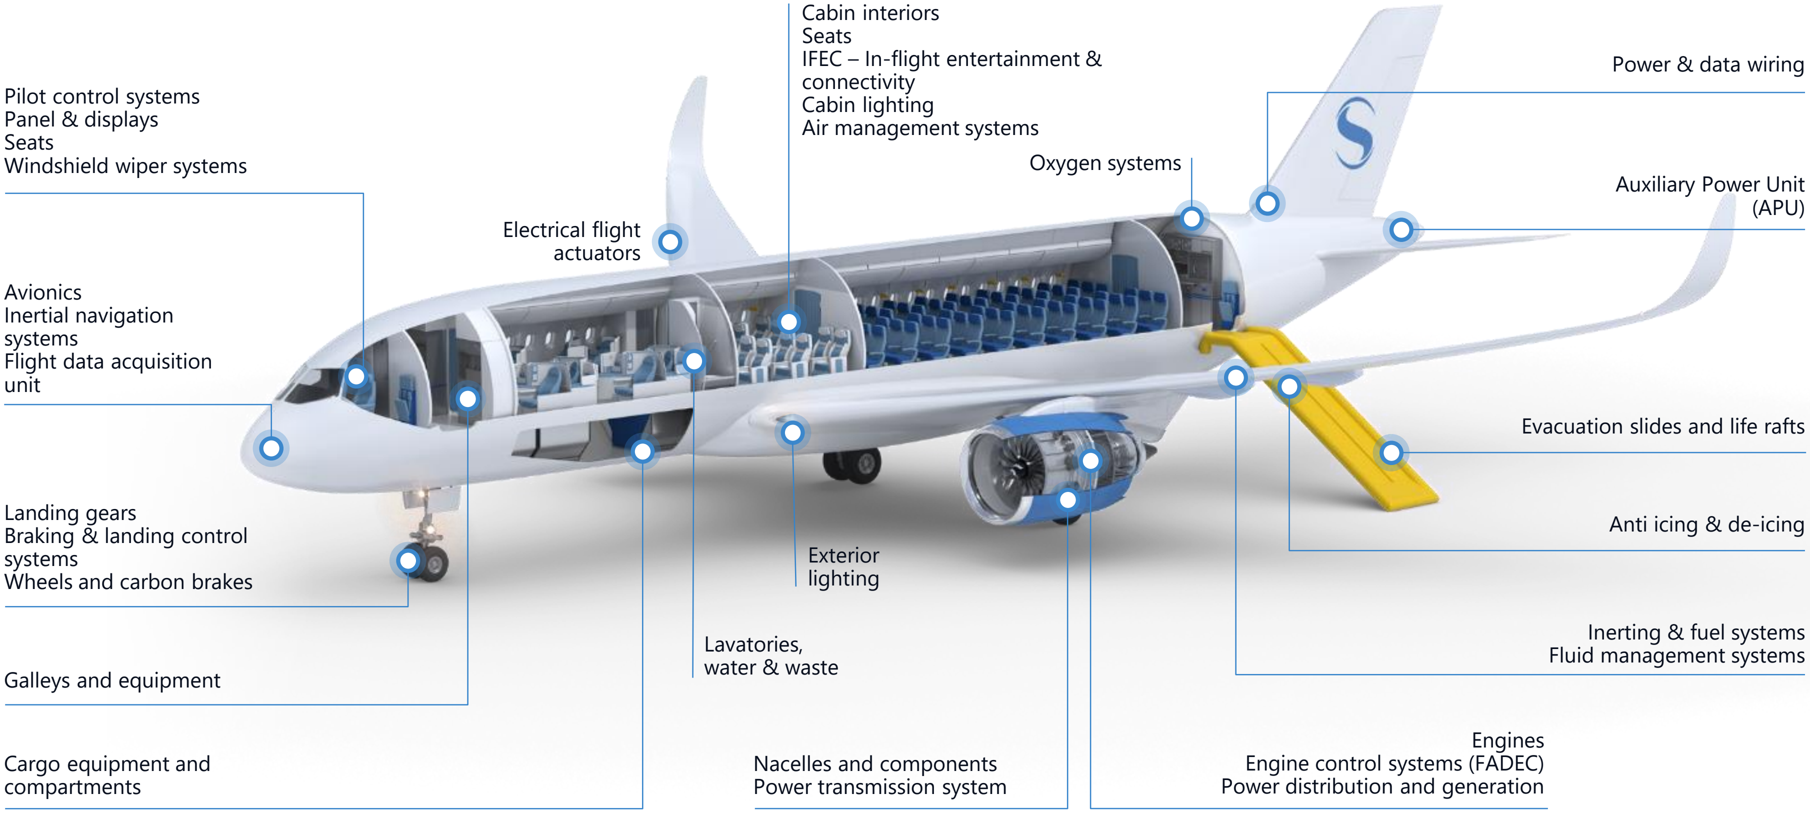

SAFRAN

Simulations at Safran – Design tools for aerospace engines

Numerical simulations are used massively in the design of Safran technologies

WHY ?

- **Master and reduce the development cycle**
- **Master the risks around manufactured parts**
- **Master the uncertainties (environnement, materials, manufacturing process …)**

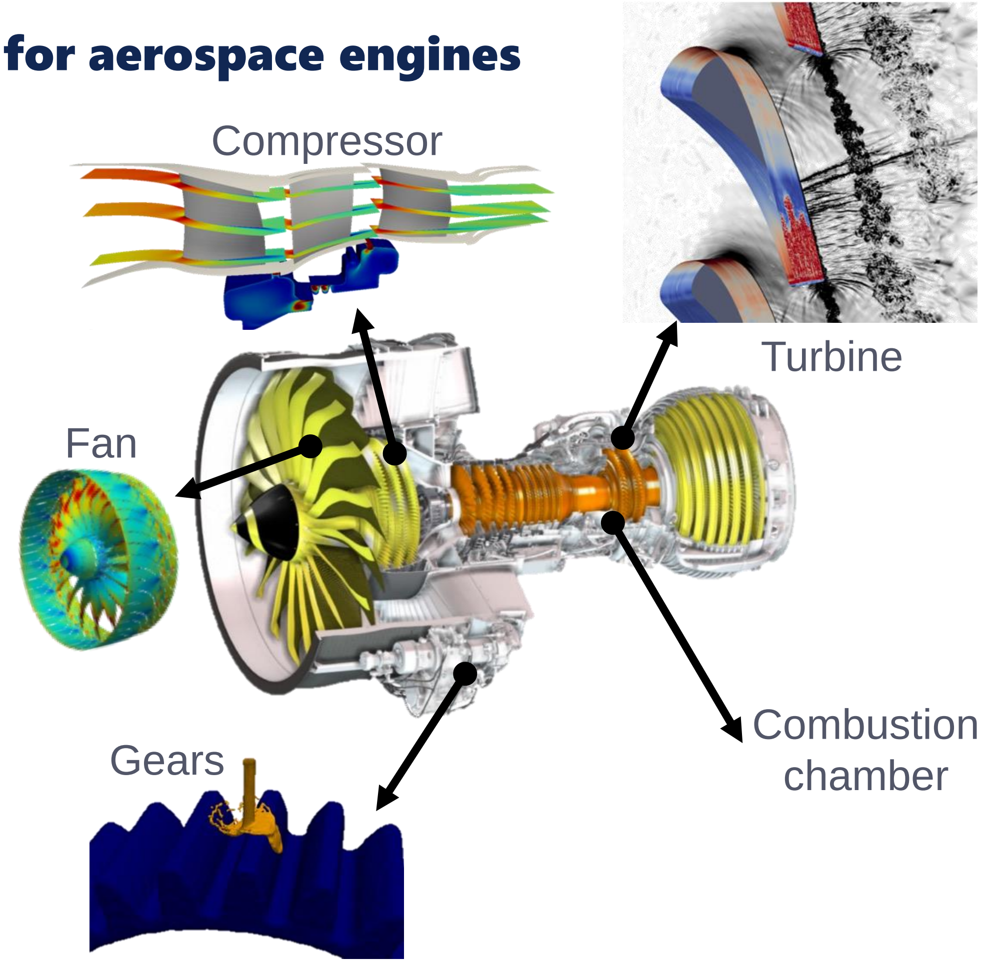

SAFRAN

Simulations at Safran – Design tools for aerospace engines

Numerical simulations are used massively in the design of Safran technologies

HOW ?

- **Capture the governing physics with the right precision at an optimal cost**
- **Both R&T and design processes need to decrease the return time of such simulations**

We rely on state of art scientific/HPC libraries and software

> **MUMPS direct HPC solver for solving large linear systems from non-linear finite element calculations**
> **MUMPS consortium member since 2016**

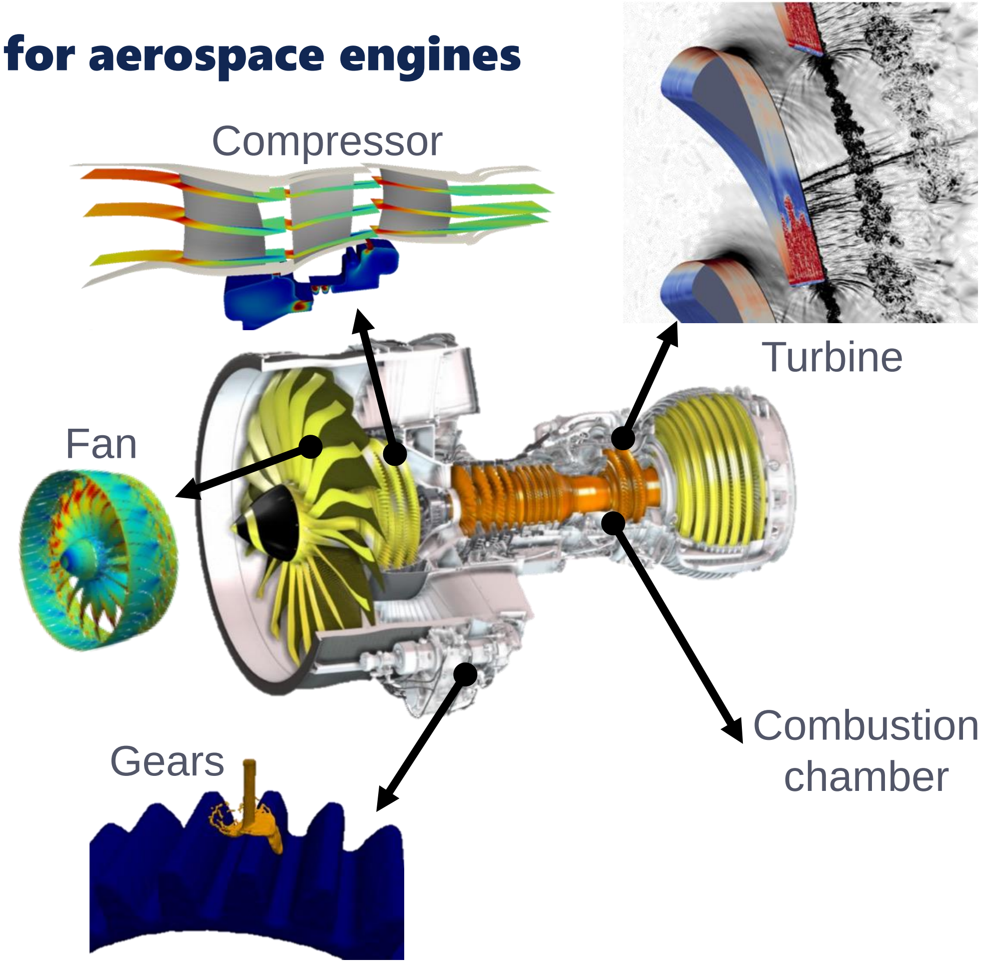

SAFRAN

HPC for structural mechanics simulation

- ## Context
  - **Growing need** for **high-fidelity** simulation during design cycles
    - Extreme loads, non-linearities, new architectural materials
  - **Availability of HPC computing resources** but under-used in structural design
- ## Challenges
  - **Achieve the right level of fidelity** in an **"acceptable"** simulation time
  - **Reduce simulation time** to increase the use of simulation in design cycles
  - **Controlling the accuracy** of resolutions to increase confidence in predictions
- ## Strategy
  - Development and maturity of **HPC solver bricks**
  - Design of **mesh adaptation** methods
  - **Increasing readiness level** of academic developments
  - **Targeted simulation platform:** Z-set finite element code (**MUMPS direct solver**)
  - **Collaborators**: ONERA / Mines ParisTech / Transvalor SA / ENS Paris-Saclay / **MUMPS Consortium**

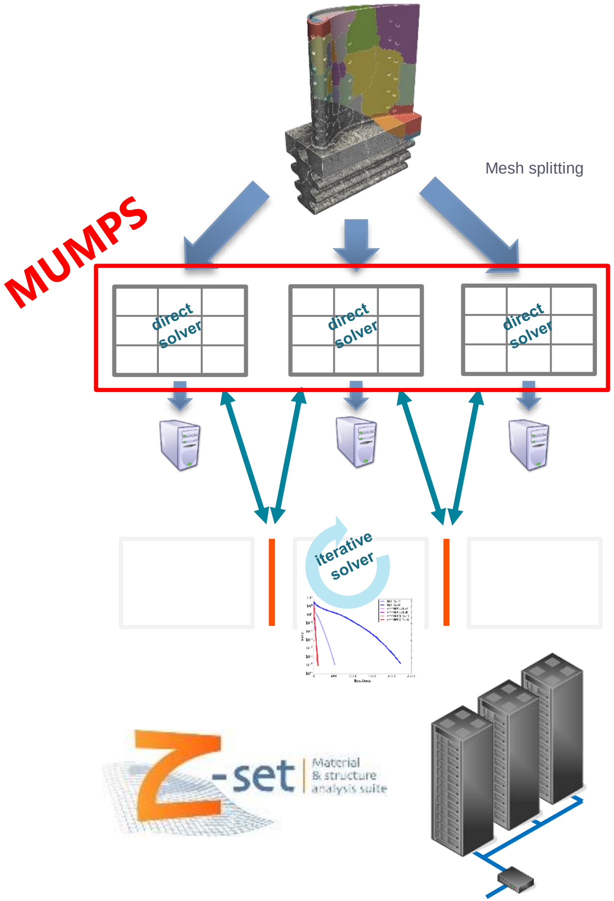

SAFRAN

OpenPisco

- Open source modular software for topology optimization: https://openpisco.irt-systemx.fr/
- OpenPisco structure design process supports: grids; unstructured meshes and body-fitted meshes
- User interfaces: User-friendly **GUI** (OpenPisco), **command line** application (OpenPiscoCL) and **Python APIs**
- Interfaced to state of art scientific libraries and open source

  FE softwares : Code Aster, FreeFem++, **Basic Tools**, *etc.*
- R&D framework → modularity, extensibility and interoperability.
- Conda package: https://anaconda.org/openpisco

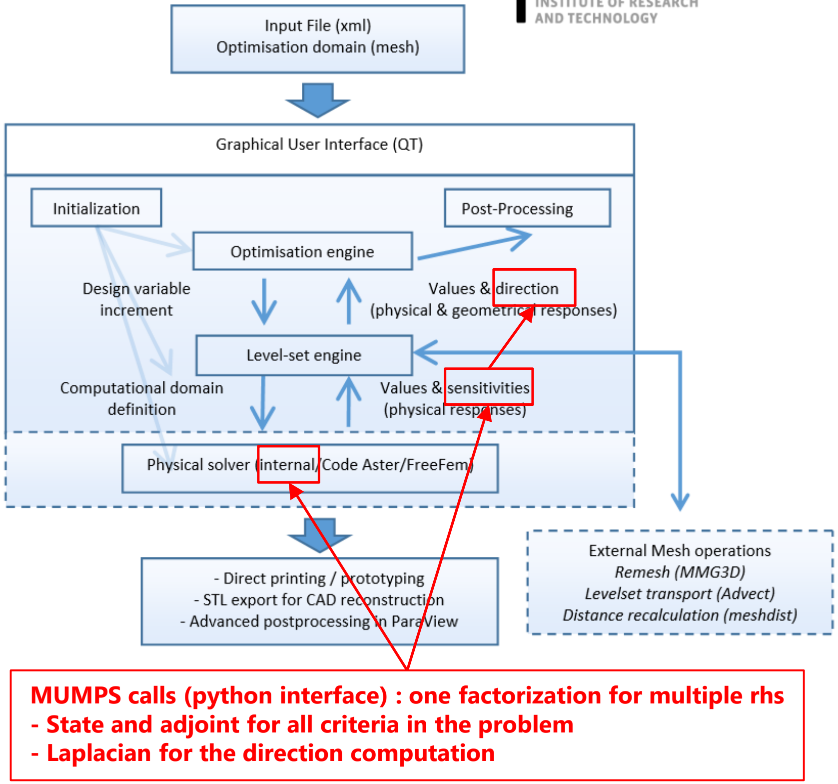

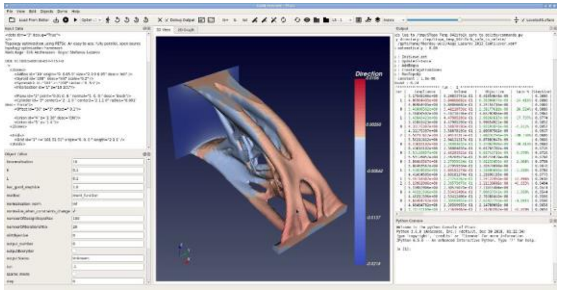

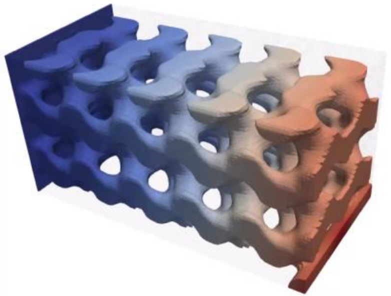

Nardoni, C and Danan, D and Mang, C and Bordeu, F and Cortial, J. A R&D software platform for shape and topology optimization using body-fitted meshes, Mesh Generation and Adaptation, 2022

Chapter 2

A multithread performance evaluation of MUMPS

SAFRAN

Standard non-linear FE solver architecture

$$K(u(t),t)u(t) = f(t)$$

Euler + Newton ➔

$$r(u_n) = \tilde{f}_n - \tilde{K}(u_n)\, u_n$$
$$T(u_n)\Delta u_n = r(u_n)$$
$$u_{n+1} = u_n - \Delta u_n$$

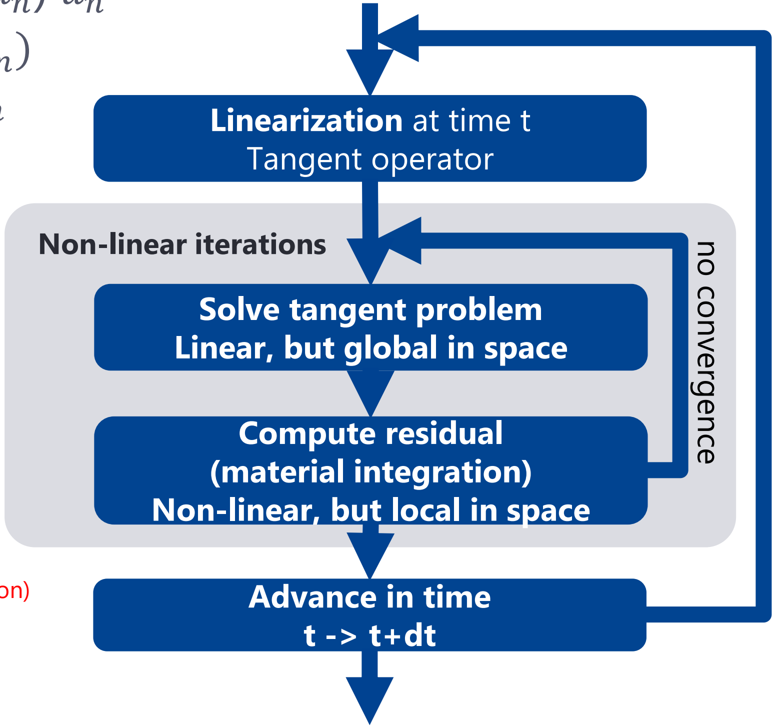

- **Two main steps**
  - Linearization and residual computation
    - 1 Non-linear integ. scheme / integ. point

    ➔ **embarassingly parallel (SMP)**
  - Resolution of tangent problem
    - Global linear system Ax=b ➔ Large sparse operator (sym. or non-sym.)

    ➔ **not embarassingly parallel !**
      - Computation of elemental matrices is embarassingly parallel
      - Assembly in the global matrix is not embarassingly parallel (race condition)

        ➔ **Task delegated to MUMPS by providing (i,j,value) triplets !**
- **Successive resolution of large sparse systems**

SAFRAN

Numerical setup

- **Physical properties**
  - Quasi-static simulation
  - Thermo linear elasticity analysis
  - Multipoint constraints (normal displacement) at the blade root
  - Centrifugal and thermal loading
- **Numerical properties**
  - MUMPS solver in full OpenMP mode (#MPI =1)
  - Shared memory computations on Intel-Haswell (2x12cores) and Intel-IceLake (2x24cores) nodes

Quadratic mesh: 103 108 elements & 1 123 050 DOFs

Type	C3d10 (10%)	C3d13 (20%)	C3d15 (3%)	C3d20 (67%)
#elements	10 654	20 823	2 695	68 936
#integ points per element	4	10	18	27

MUMPS’ setup

- **Configuration and installation using SPACK ‡**
  - GCC v11.3.0, OpenMPI v4.1.5
  - OpenBLAS v0.3.23/Netlib-ScaLAPACK v2.2.0 and Intel MKL v2020.4.304
  - Metis v5.1.0/Parmetis v4.0.3 and SCOTCH v7.0.3
- **Matrix properties**
  - Full Matrix: N ≈ 1E+6 / NNZ ≈ 3E+8 with **44% of duplicate entries**
- **MUMPS versions**
  - 5.6.0 / 5.6.0c (Activation of L0 thread based multithreading)
- **#MPI = 1 and #OMP = 24/48 (Haswell/Ice Lake)**
- **Factorization**
  - LDL^T / LU
- **Ordering**
  - METIS / SCOTCH

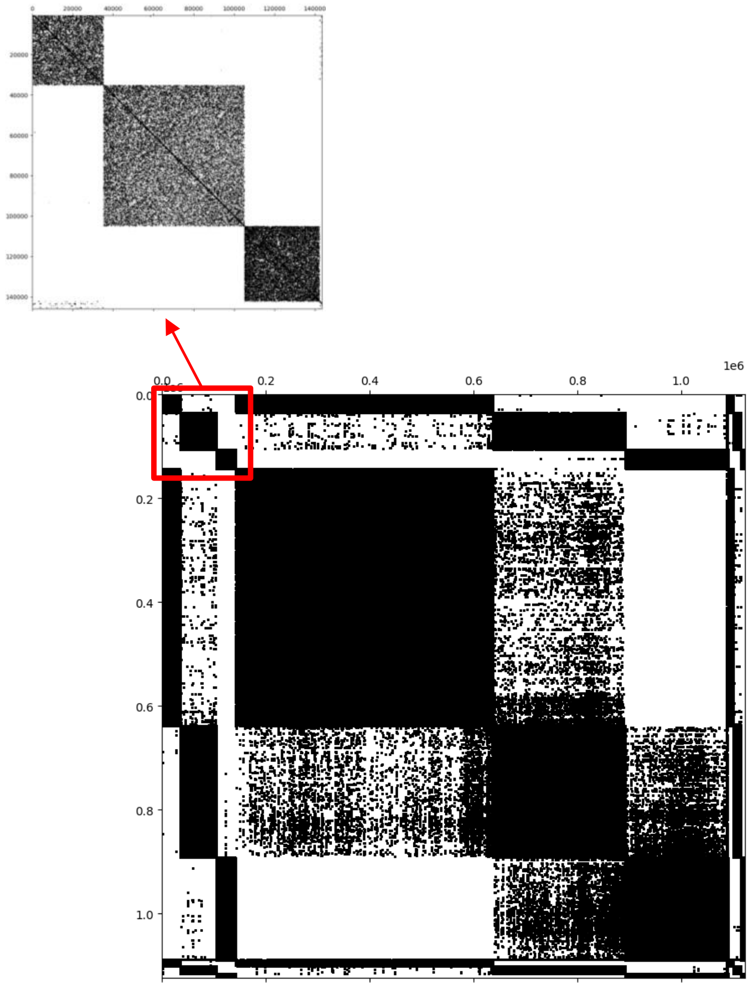

‡ https://spack.readthedocs.io/en/latest/index.html

Duplicate entries impact evaluation

with duplicate entries

```
Entering DMUMPS 5.6.0c from C interface with JOB, N, NNZ =
1  1123050 293278152
      executing #MPI =       1 and #OMP =     48
...

L U Solver for unsymmetric matrices
Type of parallelism: Working host

 ****** ANALYSIS STEP ********
...
 Elapsed time in analysis driver         =      11.3603
...
****** FACTORIZATION STEP ********
...
** Memory effectively used, total in Mbytes (INFOG(22 )): 50285
...
 Elapsed time in factorization driver    =      40.1097
...
 ****** SOLVE & CHECK STEP ********
...
 Elapsed time in solve driver=      0.9155
...
```

without duplicate entries

```
Entering DMUMPS 5.6.0c from C interface with JOB, N, NNZ =
1  1123050 170281056
      executing #MPI =       1 and #OMP =     48
...

L U Solver for unsymmetric matrices
Type of parallelism: Working host

 ****** ANALYSIS STEP ********
...
 Elapsed time in analysis driver         =      10.4813
...
****** FACTORIZATION STEP ********
...
** Memory effectively used, total in Mbytes (INFOG(22)): 49167
...
 Elapsed time in factorization driver    =      39.6549
...
 ****** SOLVE & CHECK STEP ********
...
 Elapsed time in solve driver=      0.9062
...
```


SAFRAN

Stiffness matrix assembly and LU factorization performance

- **Configuration**
  - **MUMPS version 5.6.0c / OpenBLAS v0.3.23 / SCOTCH v7.0.3**

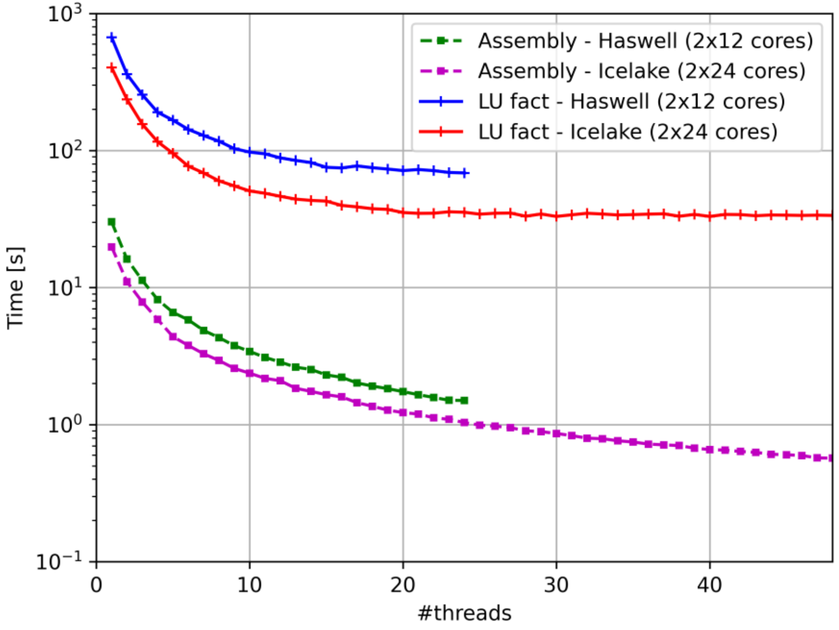

- **Memory effectively used in factorization**
  - **Haswell: ~50 Gbytes**
  - **Ice Lake: ~50 Gbytes**

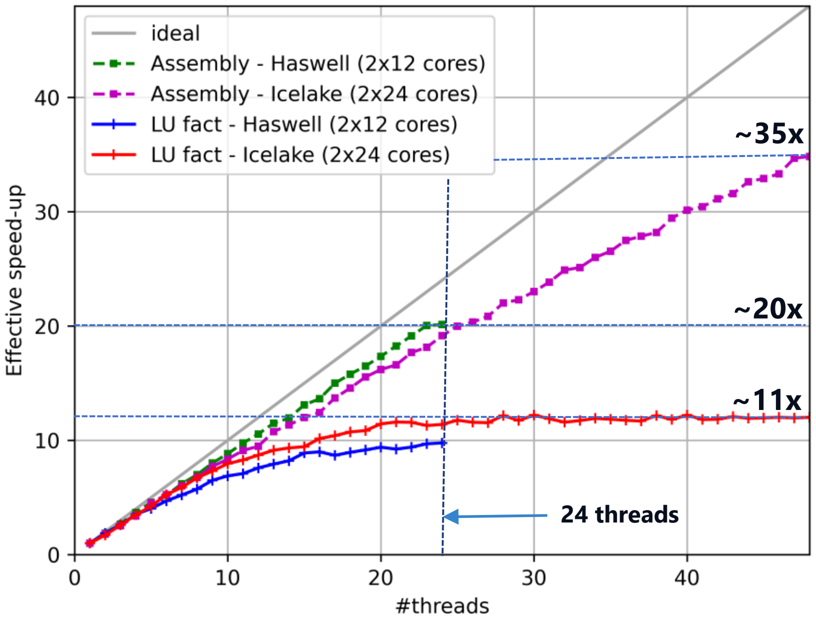

MUMPS versions: 5.6.0 vs 5.6.0c with L0 based multithreading

- **Configuration**
  - **LU factorization / OpenBLAS v0.3.23 / SCOTCH v7.0.3 / Intel-Ice Lake procs (2x24cores)**

- **Memory effectively used in factorization**
  - **v5.6.0: ~50 Gbytes**
  - **v5.6.0c: ~50 Gbytes**

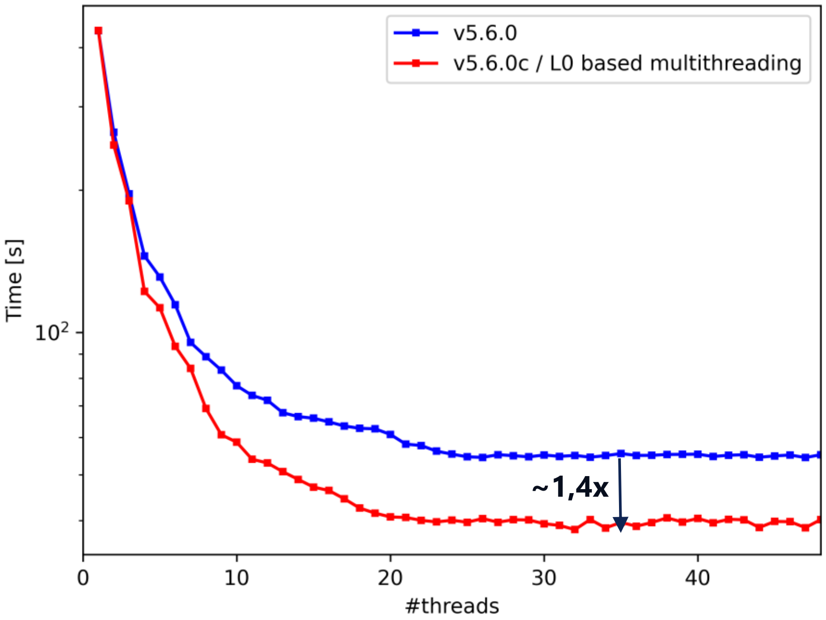

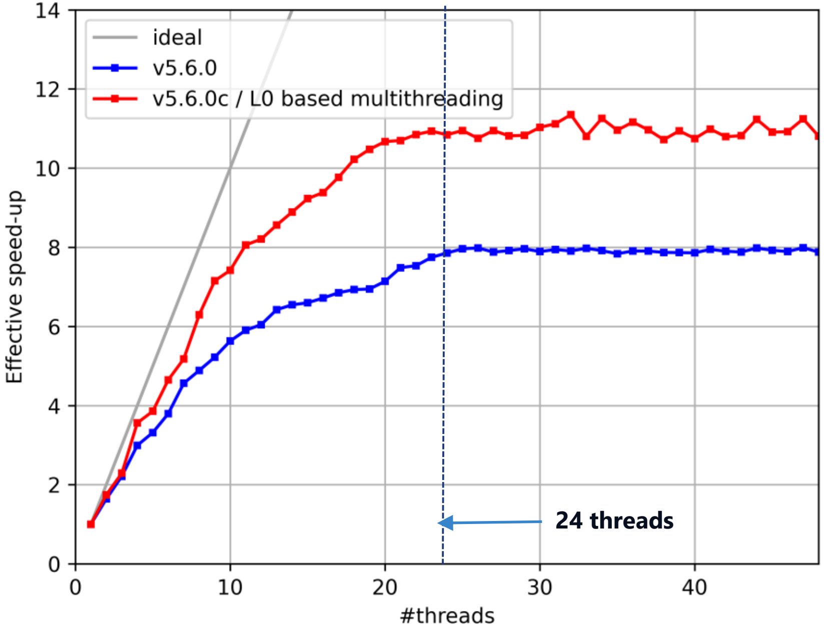

SAFRAN

Factorization: LDL^T & LU

- **Configuration**
  - **MUMPS v5.6.0c / OpenBLAS v0.3.23 / SCOTCH v7.0.3 / Intel-Ice Lake procs (2x24cores)**

- **Memory effectively used in factorization**
  - **LDLT: ~28 Gbytes**
  - **LU: ~50 Gbytes**

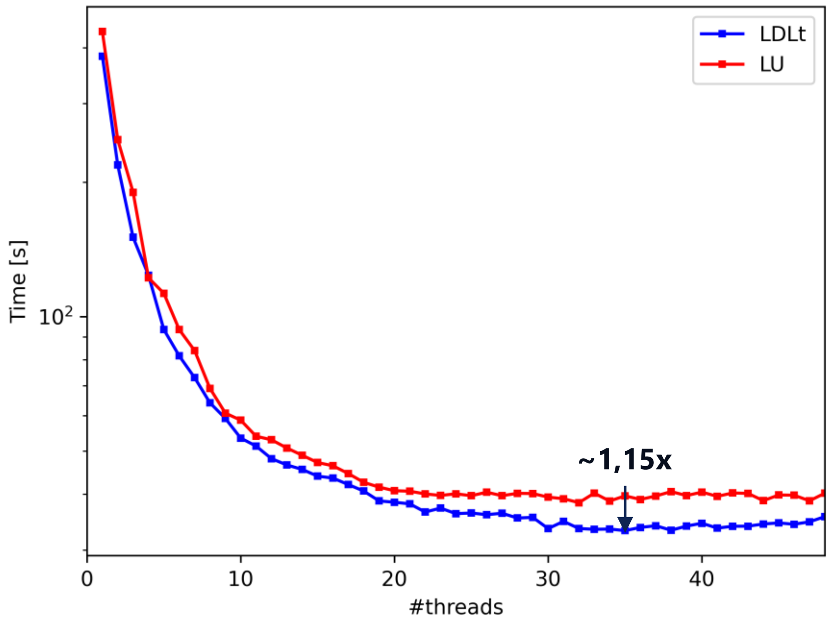

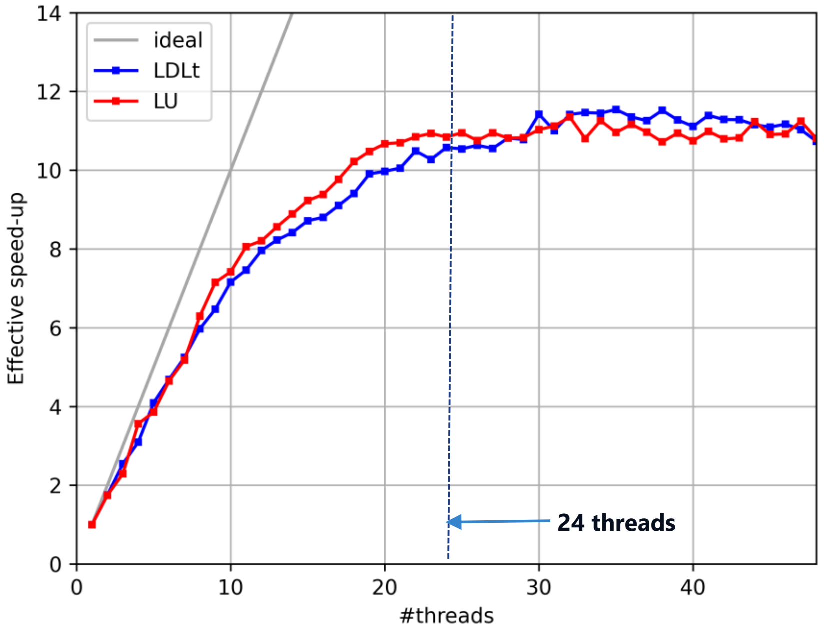

Metis v5.1.0 vs SCOTCH v7.0.3

- **Configuration**
  - **MUMPS v5.6.0c / LU factorization / OpenBLAS v0.3.23 / Intel-Ice Lake procs (2x24cores)**

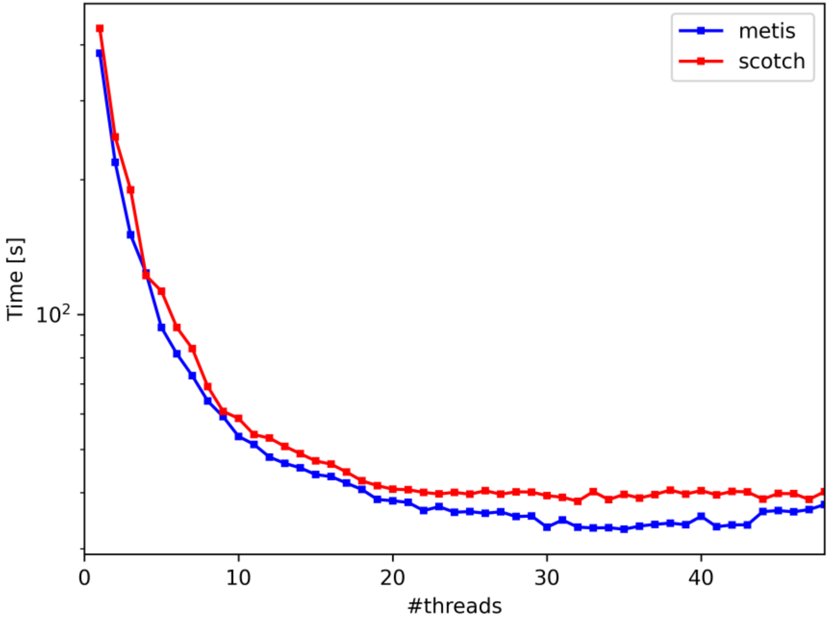

- **Elapsed time in reordering/analysis driver**
  - **METIS: ~13s / ~20s**
  - **SCOTCH (seq): ~13s / ~21s & SCOTCH (MT): ~6s / ~10s**
- **Memory effectively used in factorization**
  - **METIS: ~48 Gbytes**
  - **SCOTCH: ~50 Gbytes**

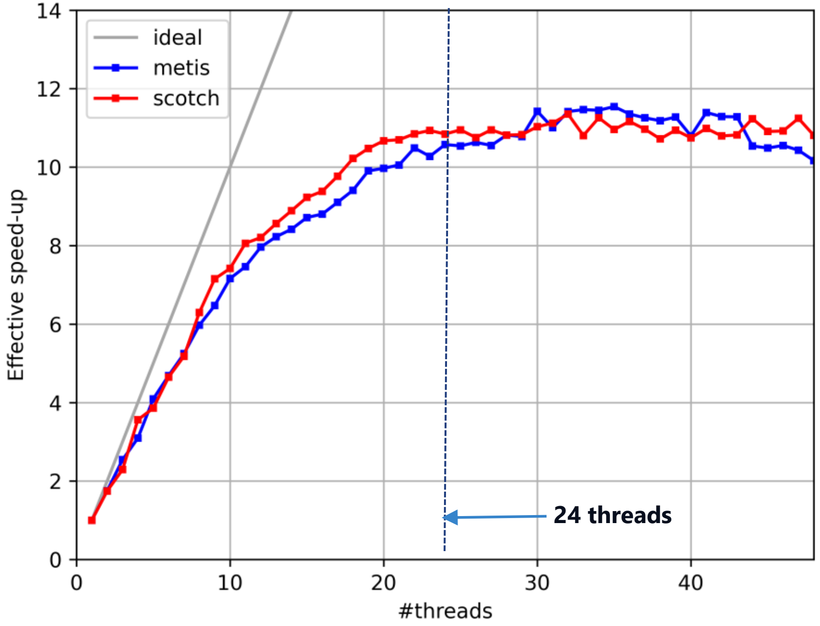

SAFRAN

OpenBLAS v0.3.23 vs Intel MKL v2020.4.304

- **Configuration**
  - **MUMPS v5.6.0c / LU factorization / SCOTCH v7.0.3 / Intel-Ice Lake procs (2x24cores)**

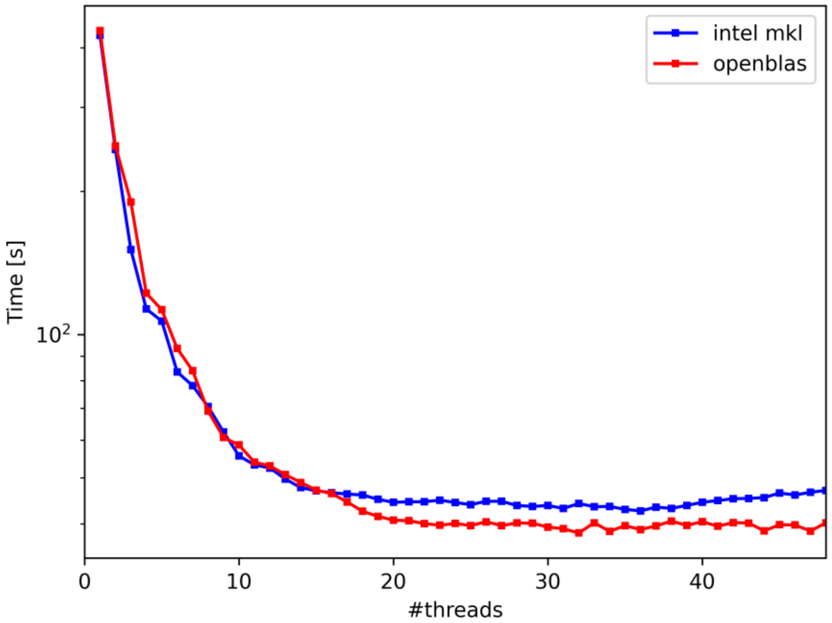

- **Memory effectively used in factorization**
  - **OpenBLAS: ~50 Gbytes**
  - **Intel MKL: ~50 Gbytes**

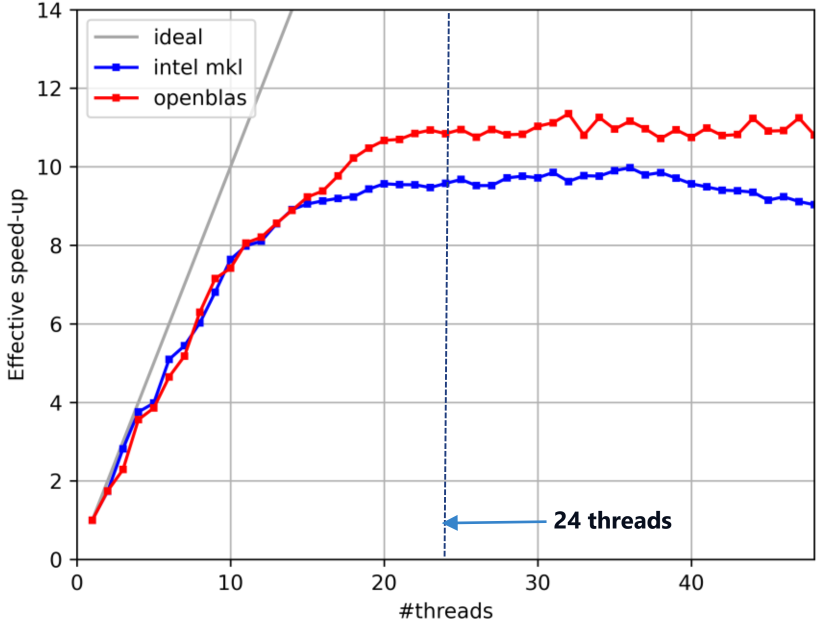

Conclusion & Prospects

- **Conclusion**
  - Easy and efficient configuration and installation of MUMPS using SPACK
  - Delegating matrix duplicate entries summation to MUMPS was free of cost !

    ➔ Efficient matrix assembly processing
  - Activation of L0 thread based multithreading provides a 1,4x speedup
  - MUMPS’ full OpenMP factorization performance on targeted application: 11x speedup on 48 cores (Intel Ice Lake)

- **Prospects**
  - Domain decomposition ➔ Massively parallel computations (hybrid MPI/OpenMP & Full MPI)
  - Benchmark with EVP constitutive law and validation
  - Investigate MUMPS’ performance on “AMD EPYC Milan 7763” processors ➔ CCRT-Topaze supercomputer‡
  - Investigate activation of BLR and iterative refinement vs full-rank approaches

‡ https://www-ccrt.cea.fr/fr/moyen_de_calcul/index.htm

POWERED BY TRUST

SAFRAN