Towards exascale simulations on GPUs for industrial turbomachinery applications

Workshop Efficient discretisation for PDE@Exascale

November 8, 2023 - CEA Saclay, Gif-sur-Yvette, France

Julien Vanharen[1], Loïc Maréchal[1], and Frédéric Alauzet[1]

[1]GammaO Inria-ONERA Team

Inria Saclay Île-De-France, Université Paris-Saclay - F-91120 Palaiseau, France

Context

- Aeronautical engines' manufacturers such as SAFRAN must **reduce their environmental footprint**.
- ACARE 2050 objectives set a **reduction** of **75%** in production of CO_2 and of **90%** of NOx.
    - Brings to the design of new and groundbreaking parts,
    - Requires more efficient and **complex numerical tools**.
- Propellers and turbines are very complex with a large number of **technological effect**.
- SAFRAN is **pushing** to develop new numerical toolchains.

Figure 1: LEAP.

Inria

Challenges

1. Handle **very complex geometries**:
    - Requires to use the automatic generation of tetrahedra meshes.
2. Predict unsteady turbulent flows with a **high-fidelity**:
    - Requires accurate and fast numerical methods on unstructured tetrahedra meshes to capture unsteady turbulent three-dimensional flows,
    - Requires anisotropic mesh adaptation.
3. And do it **efficiently**:
    - Requires a massively parallel environment which deals with anisotropic unstructured meshes.

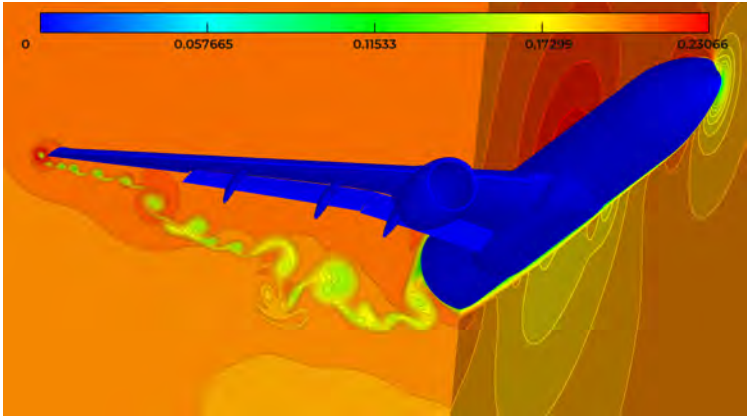

Figure 2: AIAA HLPW3 Common Research Model.

Inria

Anisotropic mesh adaptation

- Our numerical toolchain mainly composed of:
  - **Fefloa.a**: robust anisotropic local remesher,
  - **Wolf**: mixed Finite Volume Finite Element flow solver for the compressible Navier Stokes equations.
- Not limited to turbomachinery applications:
  - High-lift configuration,
  - Atmospheric reentry,
  - Ice accretion,
  - Supersonic aircraft.

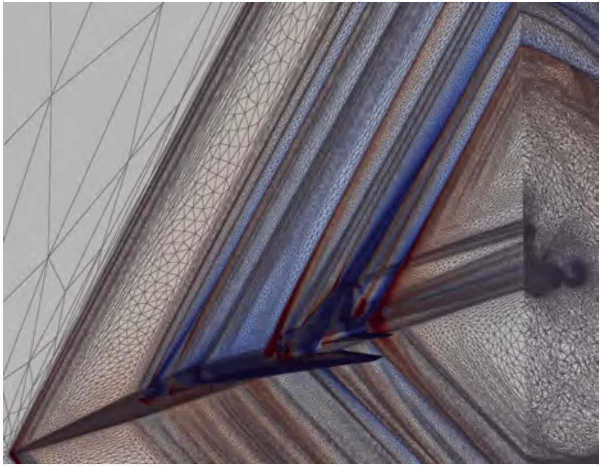

Figure 3: C608 Low-Boom Flight Demonstrator solution obtained with WOLF on the finest L2-norm adapted mesh.

Inria

Flow solver

- Compressible turbulent Navier-Stokes equations

$$\frac{\partial W}{\partial t} + \nabla \cdot \mathcal{F}(W) = \mathcal{Q}(W) + \mathcal{S}(W) \tag{1}$$

- Spalart-Allmaras turbulence model
- Mixed Element Volume method (MEV)
  - Convective and source terms by Finite Volume method
  - Diffusive terms by Finite Element method
- Vertex-centered using median cells
- Implicit time integration
  - SGS iterative solver
  - All terms are fully differentiated except for convective terms due to memory considerations
  - Strong implicit solver

Metric-based local remesher

- Based on a unique cavity operator
- Remesh the surface and the volume at the same time based on geometry surface definition to **insert points on the surface**

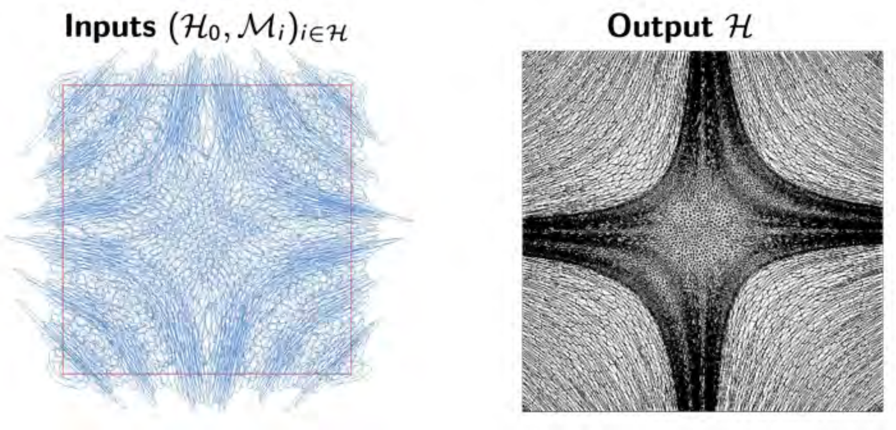

Figure 4: Metric-based local remeshing.

Recent results - NASA Rotor 37

- NASA Rotor 37 cold geometry
  - Accurate turbulent flow predictions was obtained with unstructured meshes composed only of tetrahedra,
  - Result independent of the initial mesh,
  - Mesh-converged solutions are achieved in 3D thanks to anisotropic mesh adaptation,
  - The stall on the transonic NASA Rotor 37 comes from tip vortices that increases with the pressure ratio inducing large flow separation at high pressure ratio.

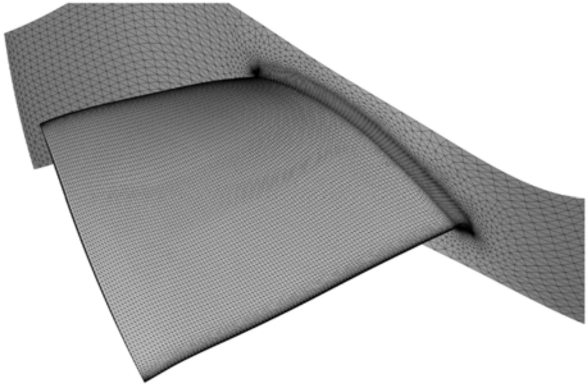

Figure 5: NASA Rotor 37.

Inria

Recent results - NASA Rotor 37

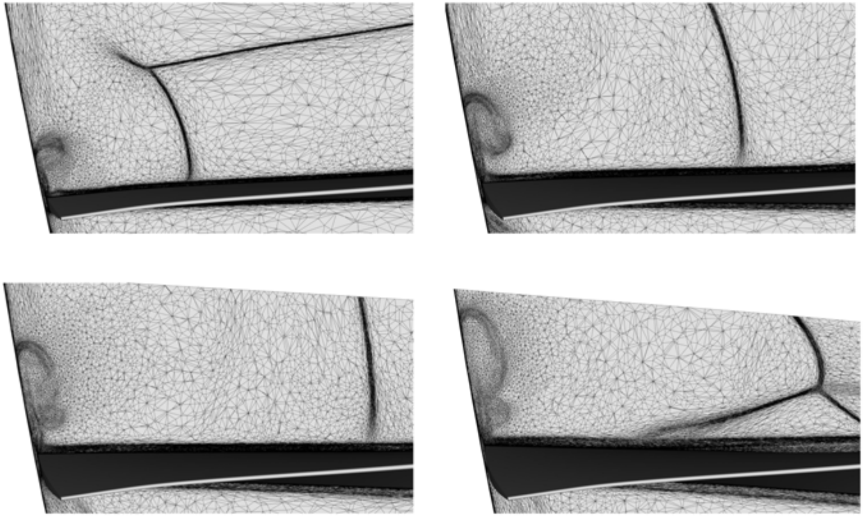

Figure 6: *NASA Rotor 37 adaptation to the tip vortices.*

Inria

Recent results - TATEF

- TATEF geometry:
  - Three rows of shaped holes on the suction side,
  - Four rows of cylindrical holes inclined by 45° in the opposite direction of the coolant feeding,
  - Three rows of shaped holes on the pressure side,
  - Cooling holes are fed through a plenum inside the blade.

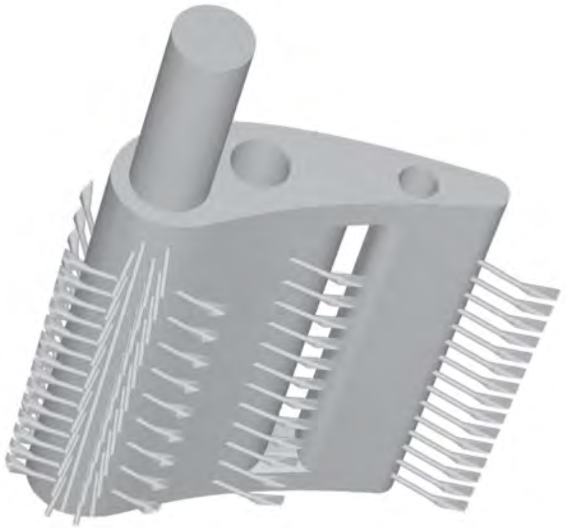

Figure 7: NASA Rotor 37.

Recent results - TATEF

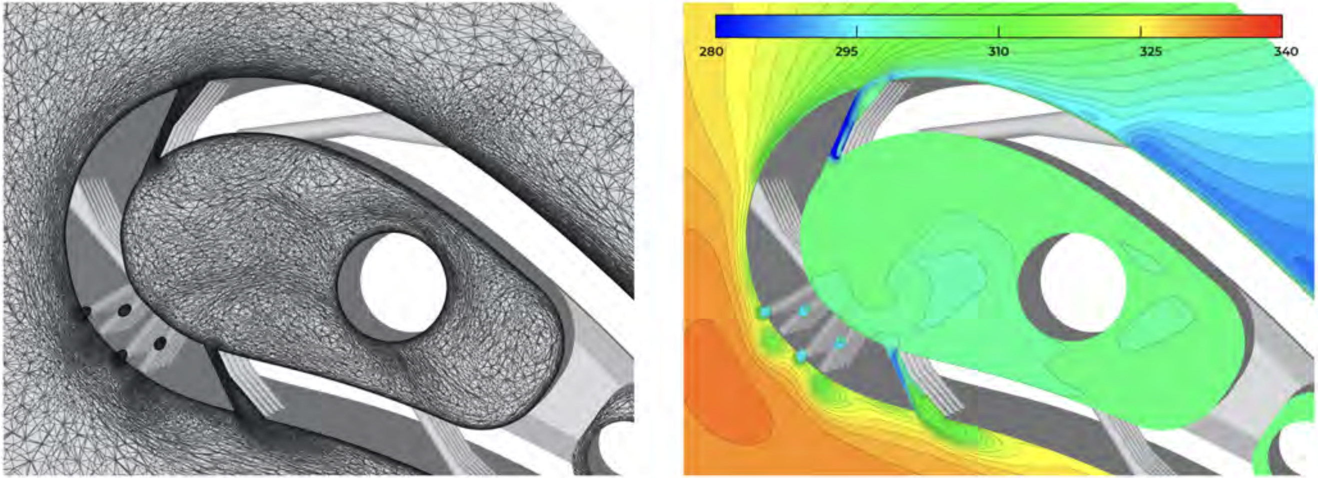

Figure 8: *TATEF cut plane. Adapted anisotropic mesh (left) and temperature field (right).*

Application demonstrator - CREATE

- CREATE compressor available at Ecole Centrale de Lyon in the Laboratory of Fluid Mechanics and Acoustics (LMFA):
    - Four stator rows and three rotors, for a total of **592** blades,
    - A very coarse initial mesh using 300,000 elements per blade leads to a final mesh composed of **177.6** million elements,
    - High-fidelity simulation of the single blade of the NASA Rotor 37 compressor requires **66** million elements with a computational time of **3200** CPU hours for an accurate prediction close to the stall,
    - A very simple extrapolation leads to an adapted mesh of **39.072** billion elements with a computational time of **1,894,400** CPU hours.

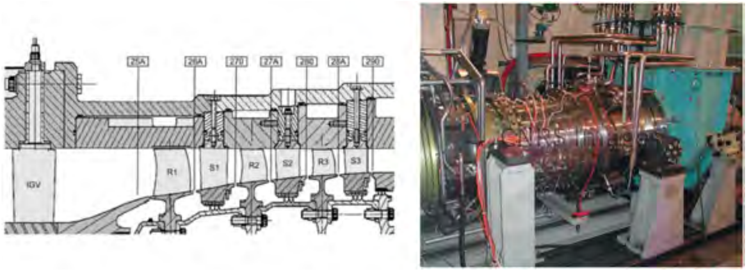

Figure 9: Left, meridian view of the CREATE compressor. Right, picture of the CREATE compressor (courtesy of CNRS).

Application demonstrator - CREATE

- SAFRAN expects to simulate the **whole pressure ratio characteristic** from choke flow to stall for certification purpose and fuel consumption reduction:
  - About **100** points are necessary to accurately predict such a characteristic,
  - Leading to a computational time of **189,440,000** CPU hours,
  - Or **1,372,754** GPU hours.

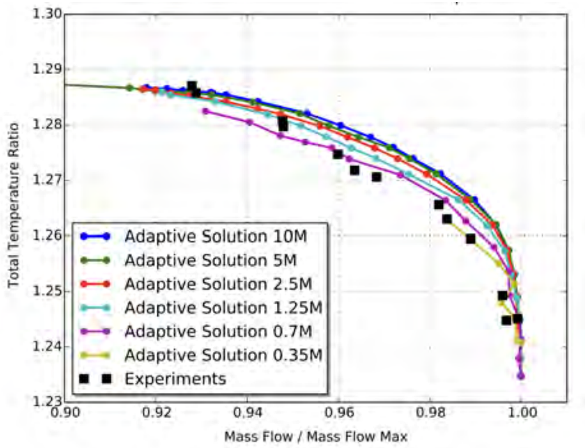

Figure 10: NASA Rotor 37. Mesh convergence of the temperature ratio characteristic.

Inria

Application demonstrator - CREATE

- Three levels of parallelism can be distinguished:
  - 1 GPU and 10 TFLOPS: steady RANS solution of several stages made of single blade coupled by mixing planes,
  - Several GPUs and 10 PFLOPS: steady RANS solutions of the 2π/16 CREATE compressor,
  - 1 EFLOPS: simultaneously running 100 RANS/URANS simulations of the 2π CREATE compressor to predict the characteristic.
- The ultimate aim is to carry out hundreds, even thousands of such simulations to design and optimize aircraft engines to reduce fuel burn and pollutant emission, which is out of reach today.
- To achieve this, **MPI** & **GPU** libraries to handle efficiently anisotropic unstructured meshes.

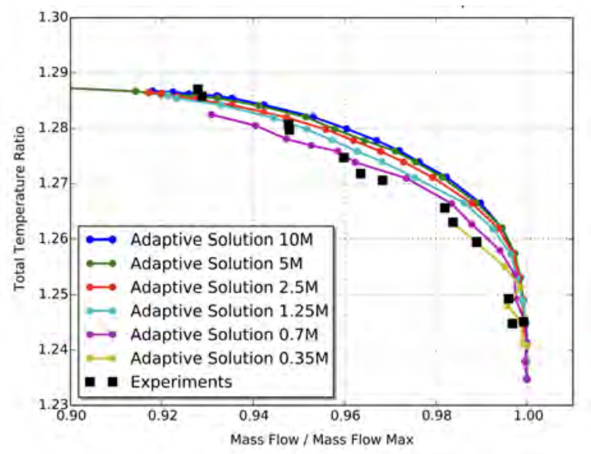

Figure 11: *NASA Rotor 37. Mesh convergence of the temperature ratio characteristic.*

Thank you for your attention!
Any questions?

E-mail: julien.vanharen@inria.fr
Homepage: https://jvanhare.github.io

Towards exascale simulations on GPUs for industrial turbomachinery applications - Workshop Efficient discretisation for PDE@Exascale © 2023 by Julien Vanharen is licensed under CC BY-ND 4.0

This is a human-readable summary of (and not a substitute for) the license.

You are free to:

Share - copy and redistribute the material in any medium or format for any purpose, even commercially.

Under the following terms:

Attribution - You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

NoDerivatives - If you remix, transform, or build upon the material, you may not distribute the modified material.