code_saturne neptune_cfd

Visualisation à distance avec ParaView client/serveur sur le supercalculateur Olympe de CALMIP

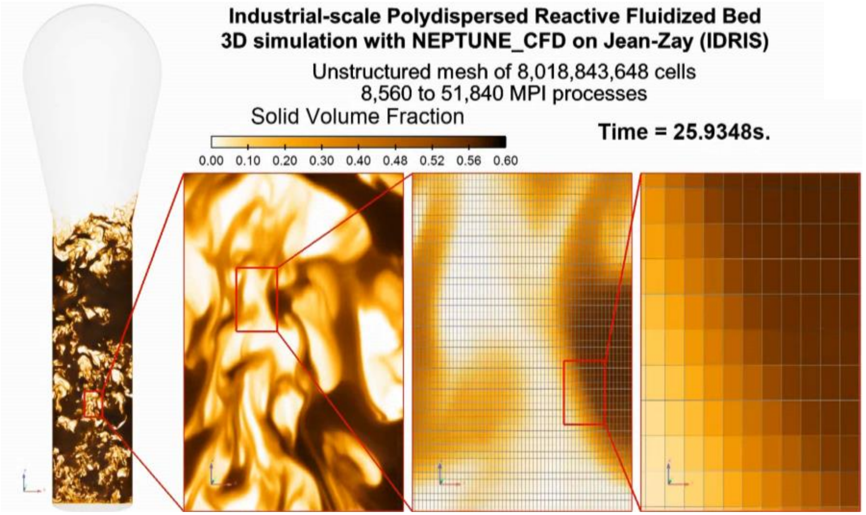

Maxime PIGOU[1,2,4]

Hervé NEAU[1,2,4]

Nicolas RENON[3,4]

1 IMFT Institut de Mécanique des Fluides de Toulouse
2 CNRS
3 CALMIP
4 Toulouse INP, Université de Toulouse

2018-2022: neptune_cfd Meso- and Grands-Challenges from Tier2 to Tier0

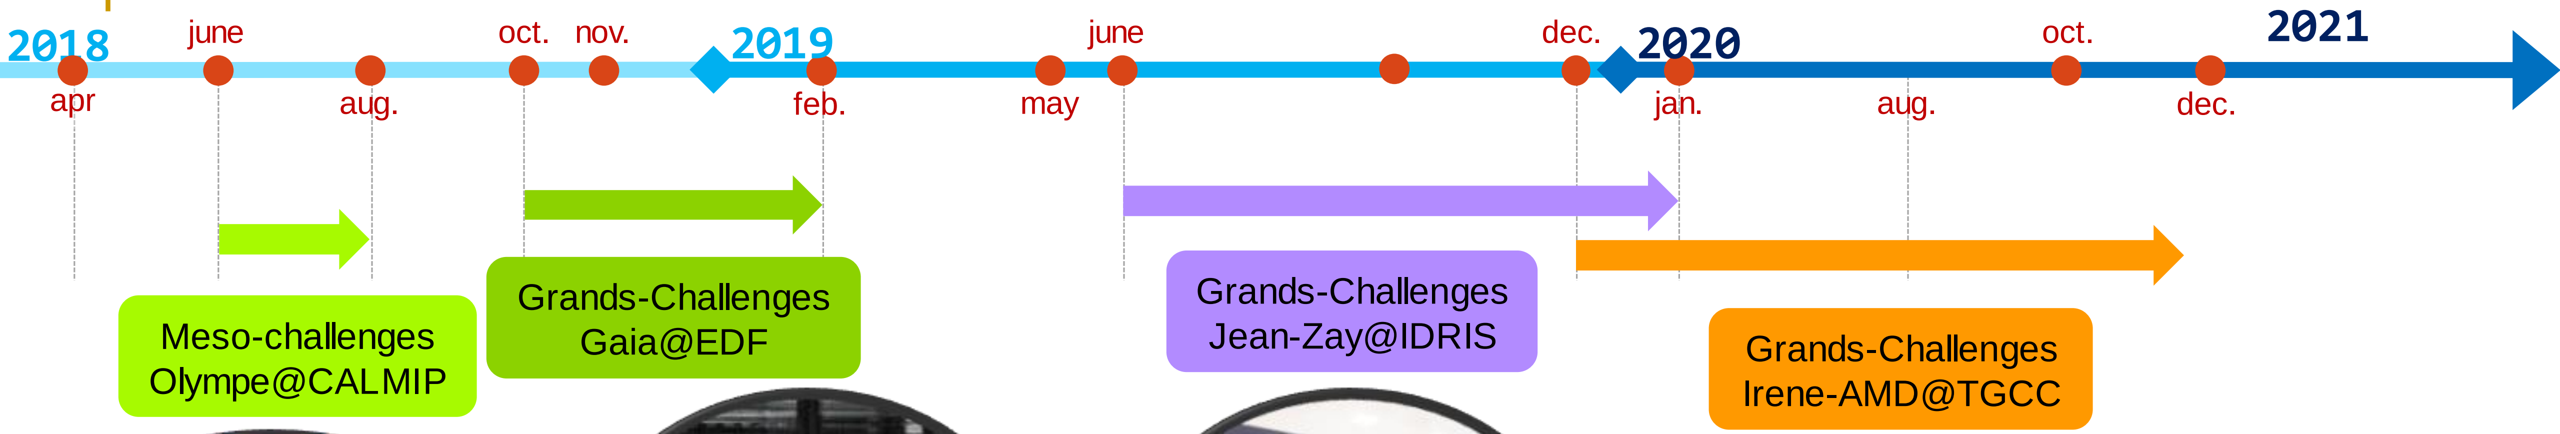

Olympe at CALMIP (2018)
Tier2 / ATOS Sequana
Perf. Peak: 1.37 Pflop/s
13,392 cores Intel (2.3 GHz)
360 nodes / 192 GB/node

Gaïa at EDF R&D (2018)
Tier1 / ATOS
Perf. Peak: 3.05 Pflop/s
42,912 cores Intel (2.3 GHz)
1,192 CPU nodes/192 GB/node

Jean-Zay at IDRIS (2019)
Tier1 / HPE SGI 8200
Perf. Peak: 16 Pflop/s
61,120 cores intel (2.5 GHz)
1,528 CPU nodes/192Go/nodes

Joliot-Curie Irene-AMD at TGCC (2020)
Tier0 / ATOS Sequana
Perf. Peak: 11.75 Pflop/s
293,376 cores AMD (2.6 GHz)
2,292 CPU nodes - 256 GB /node

2018-2022: neptune_cfd Meso- and Grands-Challenges from Tier2 to Tier0

2018 — 2019 — 2020 — 2021

CALMIP/EDF: Worldwide Premiere with 10^9 hexahedra unstructured mesh

Meso-challenges Olympe@CALMIP

Grands-Challenges Gaia@EDF

25s simulated ⇨ physical and statistical analysis

⇨ ~**10 TB** saving 16 variables (EnSight Gold binary) for 660 time steps on 12 selected planes, cylinders and external surface

⇨ **Large data set post-processed after transfer on standard workstation**

Application case: Olefin polymerization in industrial gas-solid fluidized bed reactor

⇒ Polydispersed multi-scale turbulent and reactive flow

⇒ **JCAD'18, JCAD'20**

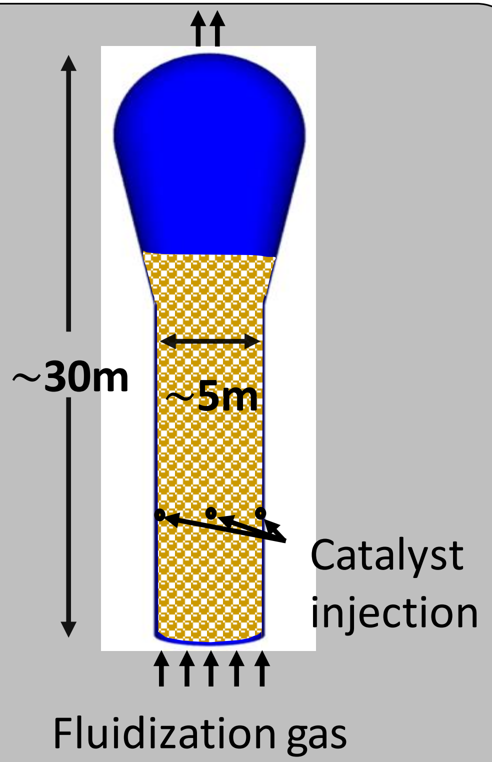

Industrial Scale Bidispersed Reactive Fluidized Bed Reactor

100 tonnes of particles - D~5m - H~30m - Unstructured Mesh: 1,002,355,456 cells

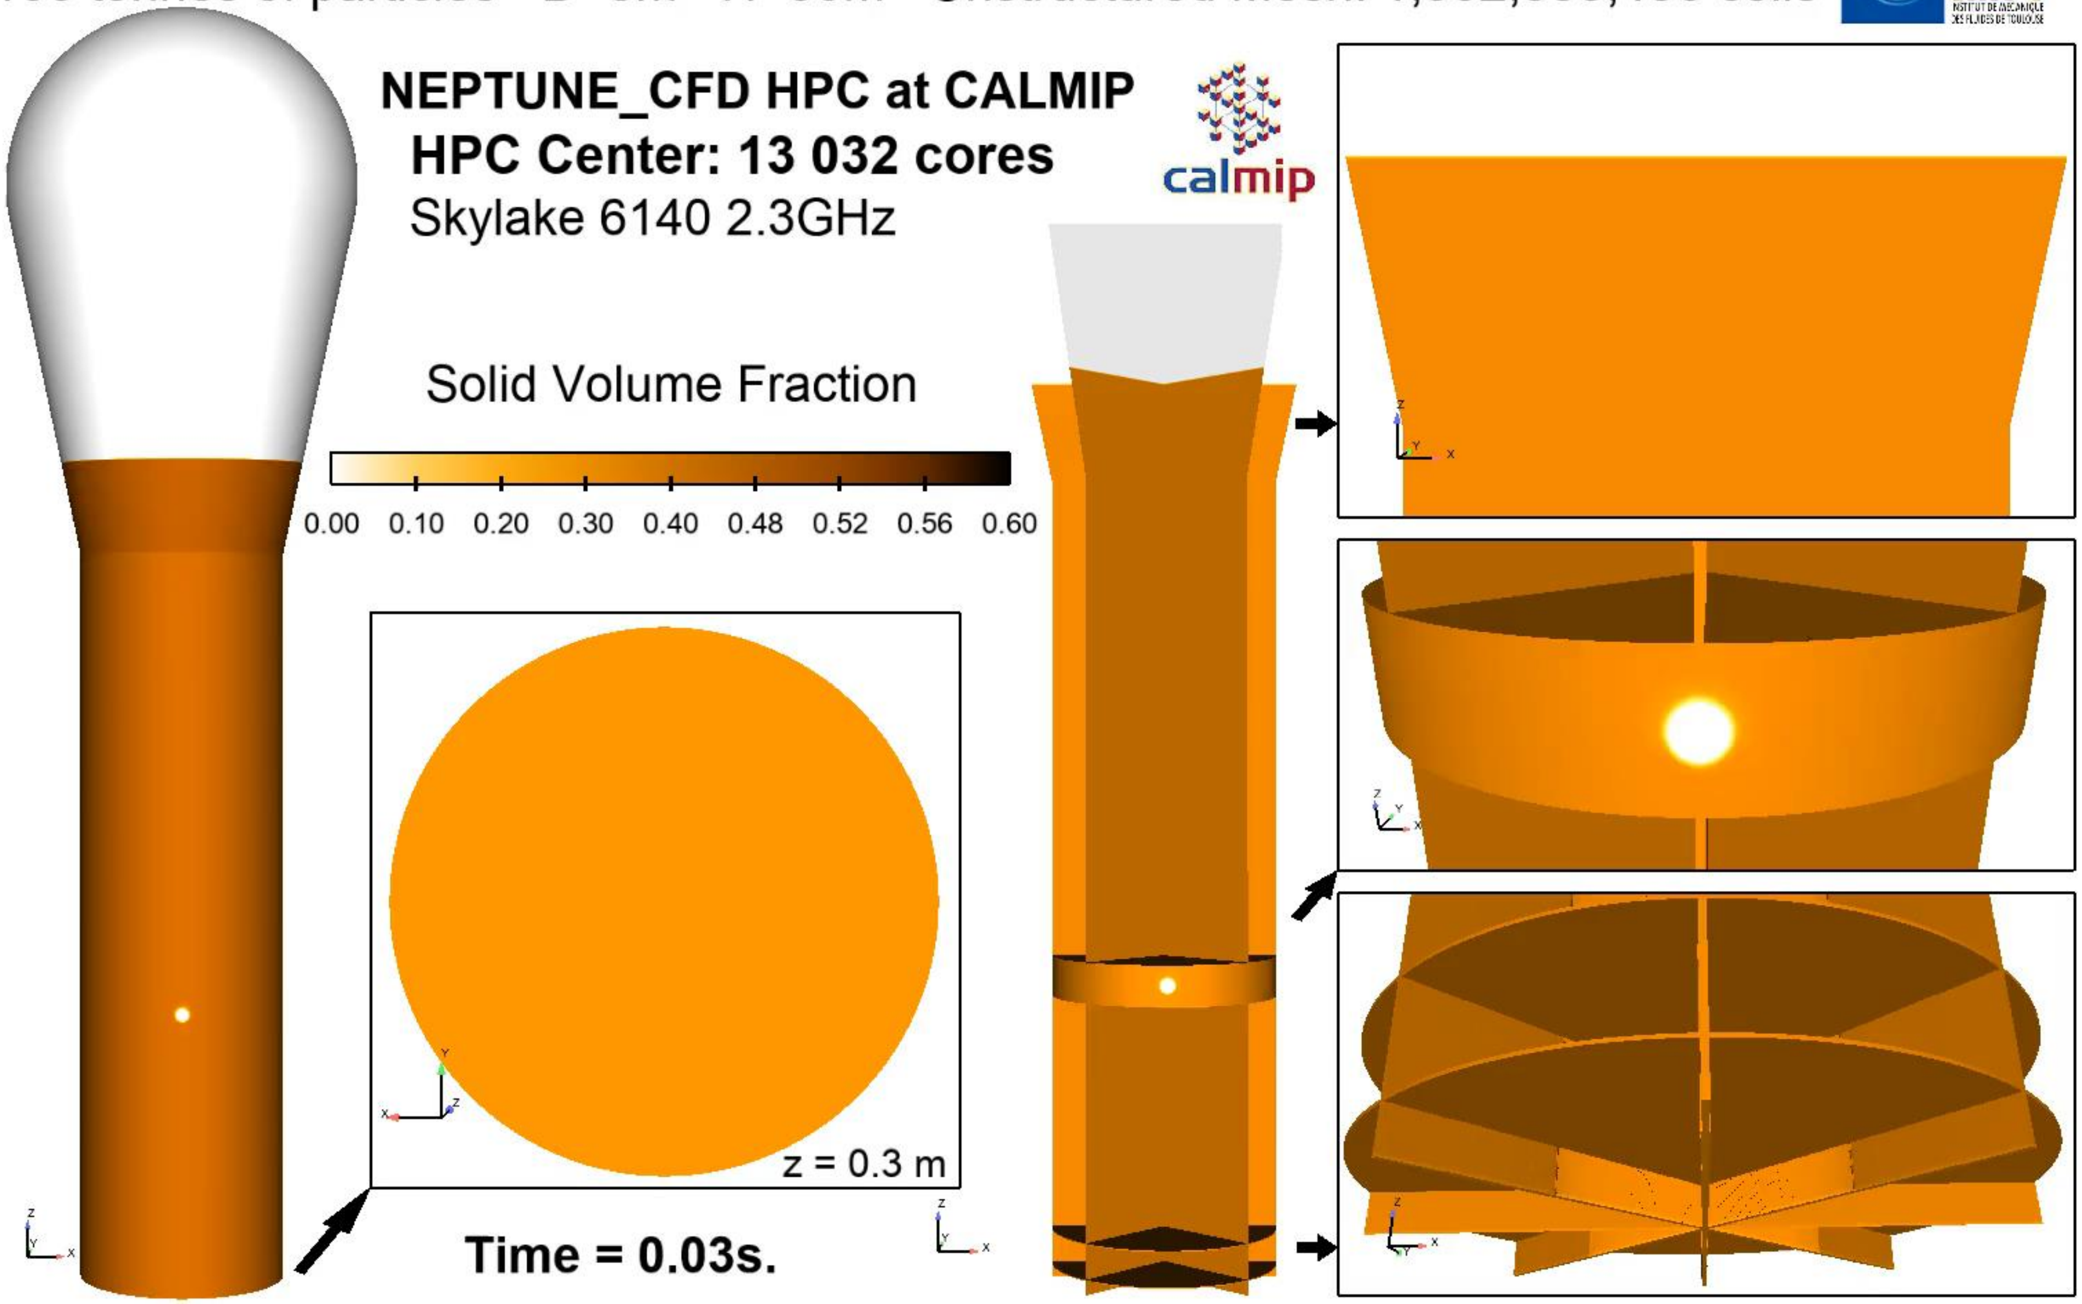

2018-2022: neptune_cfd Meso- and Grands-Challenges from Tier2 to Tier0

2018 — 2019 — 2020 — 2021

CALMIP/EDF: Worldwide Premiere with 10^9 cells unstructured mesh → **IDRIS: 8 times bigger mesh*
8 billions cells**

Meso-challenges
Olympe@CALMIP

25s simulated

Grands-Challenges
Gaia@EDF

Grands-Challenges
Jean-Zay@IDRIS

1.7s simulated ⇨ limited physical analysis
⇨ 57 TB of visualization data

⇨ **Simulation time: x20 slower**

***Automated mesh refinement** (split by 2 in each direction using code_saturne features) **coupled with the interpolation of latest time step results onto the new refined mesh**

10^9 cells mesh (25s)
⇨ *interpolation onto a 8.10^9 cells mesh (26.7s)*
⇨ *2^{nd} interpolation up to 64.10^9 cells*

Teratec 2022

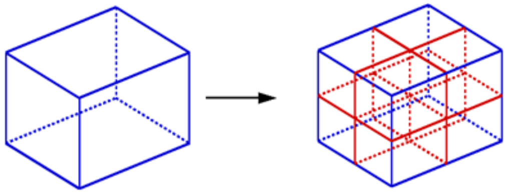

2018-2022: NEPTUNE_CFD Meso- and Grands-Challenges from Tier2 to Tier0

2018 — 2019 — 2020 — 2021

CALMIP/EDF: Worldwide Premiere with 10^9 cells unstructured mesh → **IDRIS: 8 times bigger mesh 8 billions cells**

Meso-challenges Olympe@CALMIP

Grands-Challenges Gaia@EDF

Grands-Challenges Jean-Zay@IDRIS

Huge data volume to analyze and visualize

Moving data from IDRIS to lab: slow, insufficient space
⇨ transfer directly to CALMIP

Post-processing of heavy data requires significant RAM and CPU resources

25s simulated

1.7s simulated ⇨ limited physical analysis
⇨ 57 TB of visualization data

- ⇨ **Simulation time: x20 slower**
- ⇨ **Reaching post-processing limits: storage of 53 TB of data, data transfer limitations, limited toolset for visualization, …**

Industrial-scale Polydispersed Reactive Fluidized Bed 3D simulation with NEPTUNE_CFD on Jean-Zay (IDRIS)

Unstructured mesh of 8,018,843,648 cells
8,560 to 51,840 MPI processes

Time = 25.0006s.

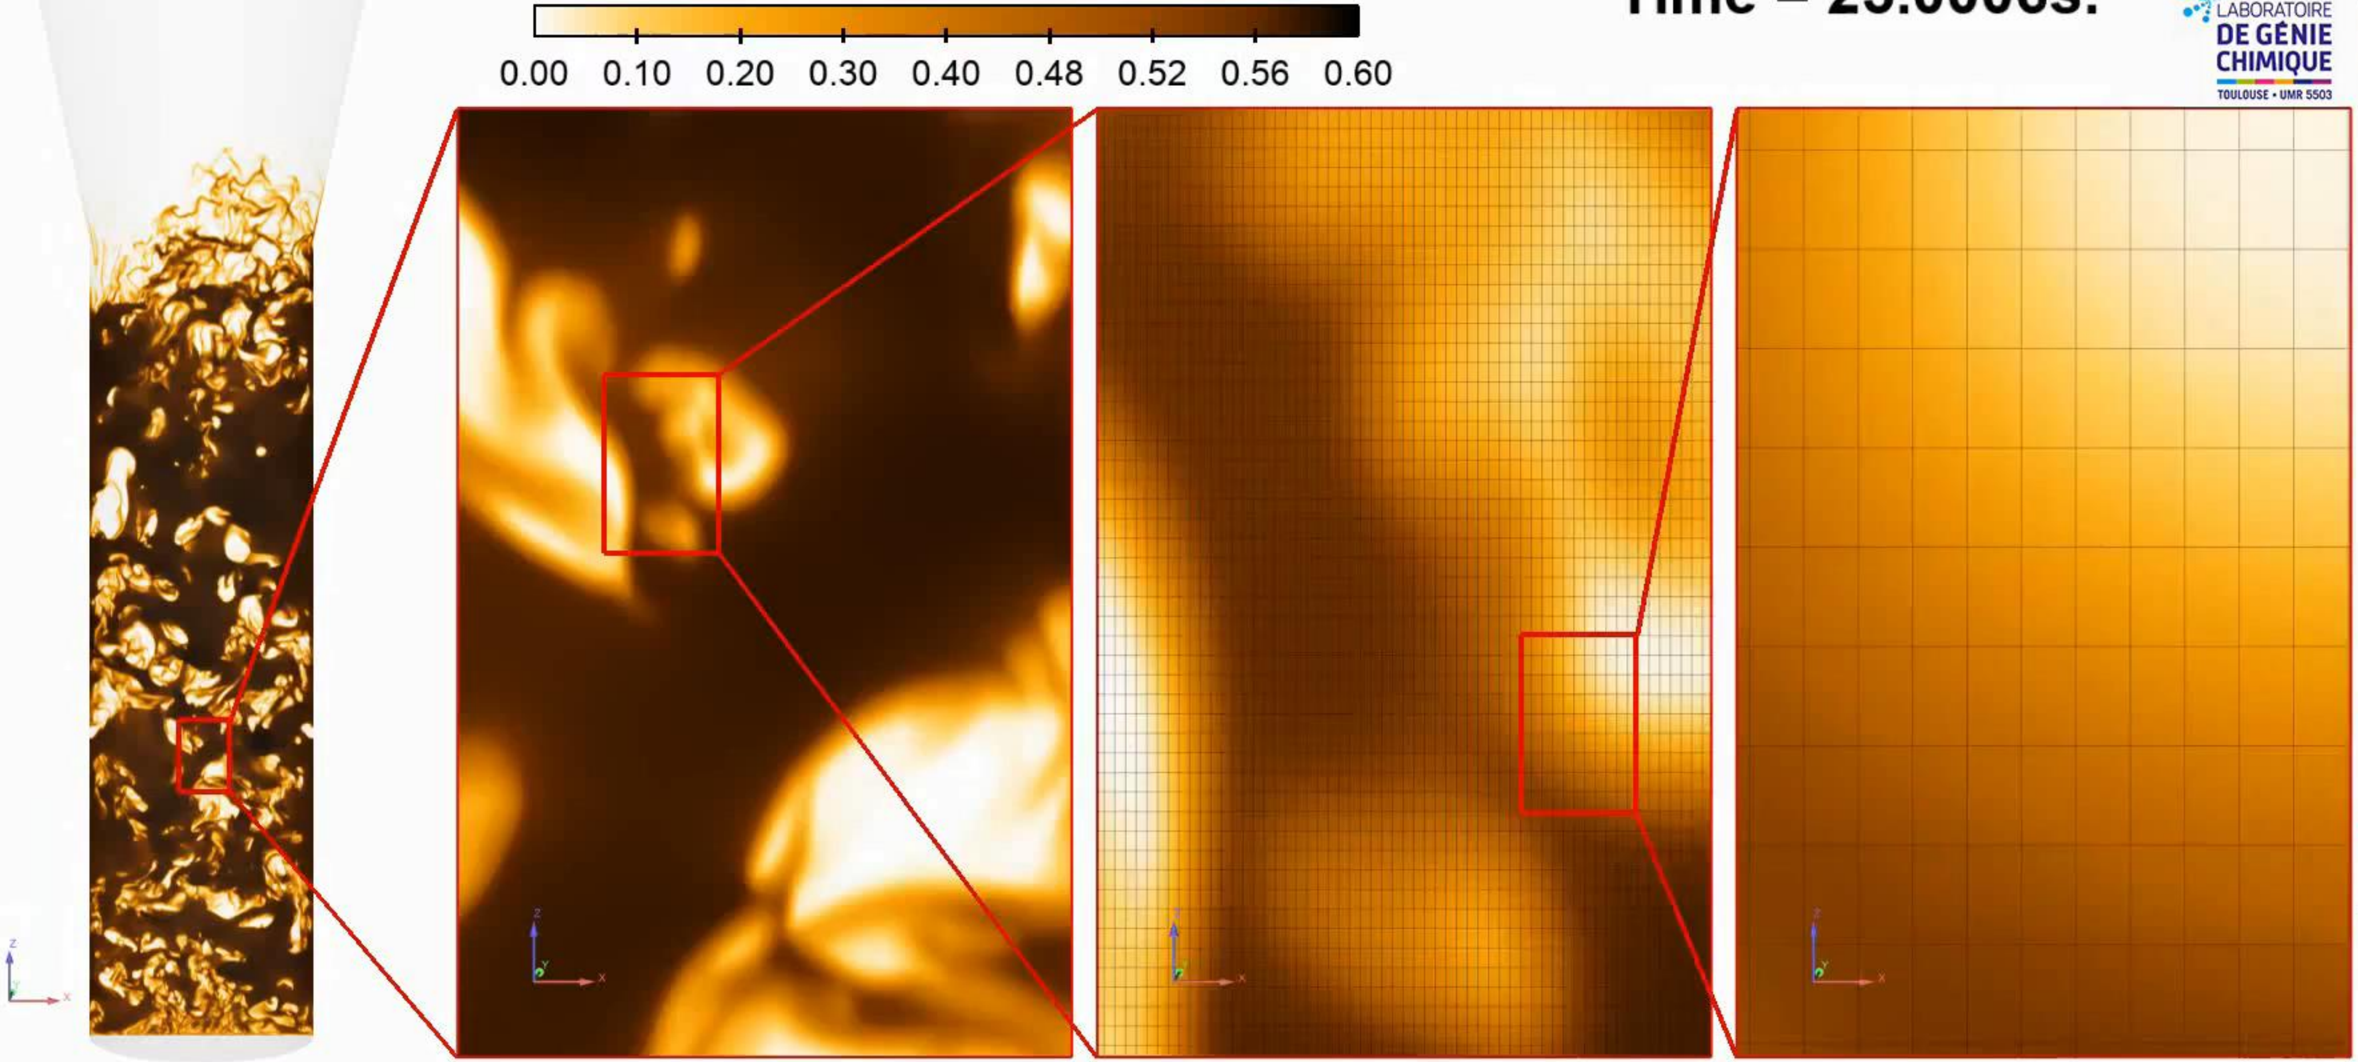

H. Neau
M. Pigou

JCAD 2020: From HPC computations to HPC post-processing: ParaView client/server at CALMIP from home

Remote display solutions considered:

⇨ **Turbo VNC** solution as suggested by CALMIP
https://www.calmip.univ-toulouse.fr/espace-utilisateurs/doc-technique-olympe/se-connecter-olympe/visualisation-graphique-distance

↳ 1 node Volta max (4 cores max), 50GB max of RAM and variable display quality and latency

⇨ **ParaView Client/server:**
A solution presented by Jean Favre (CSCS)
in 2017 at Toulouse (CUTIS, Groupe Calcul CNRS)

⇨ **Volume rendering of a cylinder close to injectors on 8 billion cells mesh data results using 10 nodes of Olympe CALMIP**

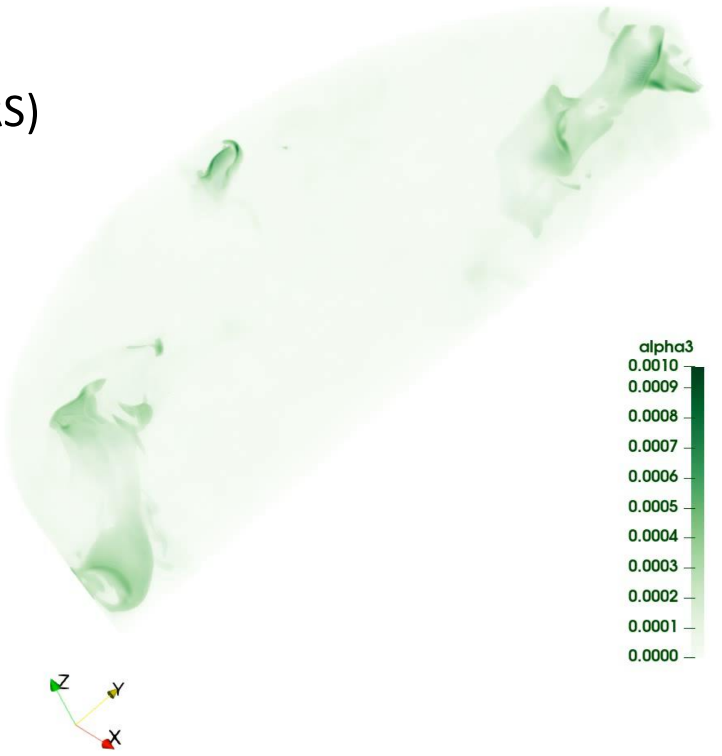

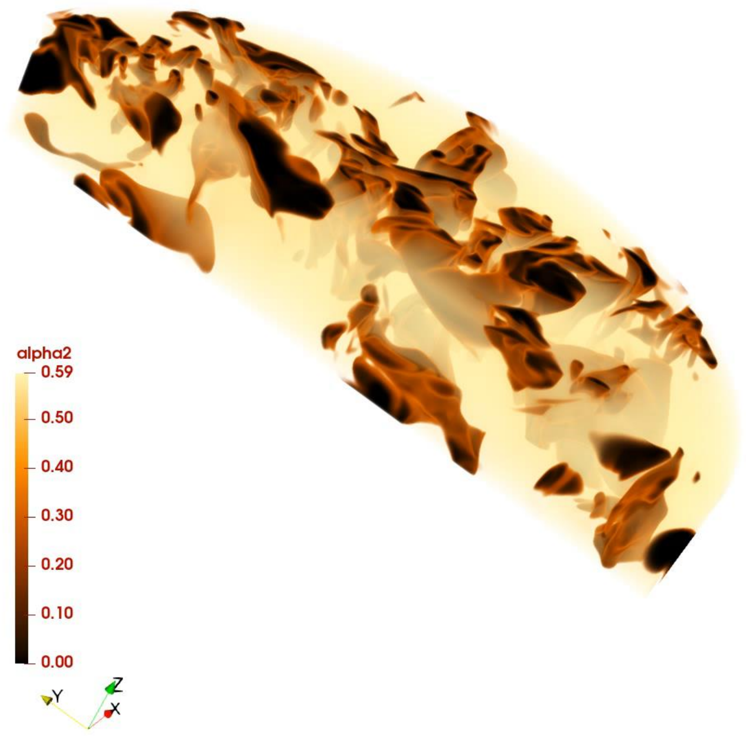

ParaView

![ParaView logo]

ParaView client/server

"ParaView is designed to work well in client/server mode. In this way, users can have the full advantage of using a shared remote high-performance rendering cluster" https://www.paraview.org/Wiki/Setting_up_a_ParaView_Server

Basic principle: post-process and visualize data using ParaView server on the computing center (where were produced) and remote only display on ParaView Client on a classical workstation or laptop

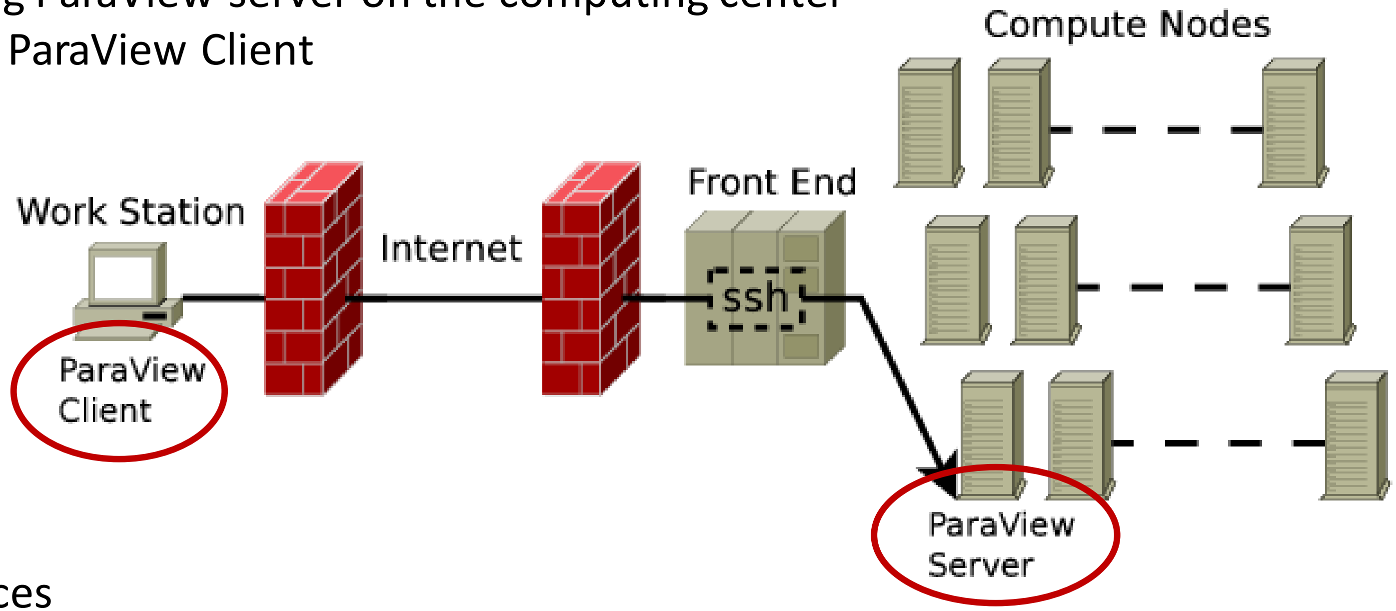

https://www.paraview.org/Wiki/File:Two-hop-tunnel.png

Main interests:
- Better display quality and latency
- Possibility to visualize huge data using HPC resources
- Works directly with pre-compiled ParaView binaries downloaded from ParaView website for both server *(osmesa MPI for CPU)* and client
- ParaView server span over multi-nodes: many cores, RAM, GPU (ray tracing)
- Solution compatible with co-processing and visualization in situ (catalyst)

Full packaged and secured script to use ParaView from home (Linux) using Client/Server mode on several CALMIP compute nodes

```
#!/bin/bash
# This script should be used by users willing to run ParaView in a client/server
# mode with the server part being hosted on Olympe (CALMIP) compute nodes.
# Script developpe par Herve Neau et Maxime Pigou en 2020
# Service CoSiNus - Institut de Mecanique des Fluides de Toulouse (IMFT)

####################################
# STEP I - GET USER INFORMATION #
####################################
# Ia - Get user info through interactive prompt
calmip_username="mpigou"          # User login for connecting to calmip
job_nnode="10"                     # Number of node requested for allocation
job_ntask_per_node="36"           # Number of tasks per node
job_time="0:20:00"                # Duration of the allocation
job_start_timeout="600"           # How long (in sec.) to wait for the job to start
local_paraview="/home/jcad2022/DemoParaviewClientServer/ParaView/bin/paraview"

# Ib - Define script tuning variables
calmip_scname="Olympe" # How the supercomputer should be referenced
calmip_hostname="olympe.calmip.univ-toulouse.fr" # Hostname of the supercomputer
calmip_paraview="/tmpdir/mpigou/202209_JCAD_pvserver_demo/ParaView" # Path to ParaView on

# Ic - Define script constants
SSH_TIMEOUT=30 # Duration to wait for a SSH connexion to be dropped

PVSERVER_STARTUP_TIME=20 # How long to wait before the job start and the connect attempt
PVCLIENT_STARTUP_TIME=20 # How long to wait between client start and the delete of its pyt

```

<u>Prerequisites:</u>

⇨ ssh public-key authentication

⇨ to ensure data confidentiality patch the library libvtkCommonSystem.so
(by default pvserver can be accessed by anybody logged on supercomputer)

<u>To use the full packaged script:</u>

⇨ Only 2 fields to adapt

⇨ Fully automatized

Script used at IMFT to access ParaView at CALMIP

Contact your computer center to access similar setup

```
jcad2022@demojcad2022: ~/DemoParaviewClientServer

jcad2022@demojcad2022:~/DemoParaviewClientServer$ ./start_paraview_on_olympe.sh
Trying to connect to Olympe using pubkey authentication...
Successfully connected to olympe
Job name: pvserv-38ABFF4E
Submitted batch job 950638
Job submitted, job id: 950638
Waiting for job to start running.
Job has started!
Master node detected: olympecomp207
Waiting 20s for pvserver to be initialized. Done.

Creating a tunnel to the master node: olympecomp207
Tunnel opened.
Starting ParaView with auto-connect.
Waiting for the session to end properly
Warning: Ignoring XDG_SESSION_TYPE=wayland on Gnome. Use QT_QPA_PLATFORM=wayland
to run on Wayland anyway.
Warning: Permanently added 'olympecomp207' (ED25519) to the list of known hosts.
Connection to olympecomp207 closed by remote host.
Removing temporary files created on Olympe.
Done!
jcad2022@demojcad2022:~/DemoParaviewClientServer$
```

ParaView Client/Server on Windows

<u>Same principle but more manual:</u>

- Allocate Resources on computing center
- ParaView Server (pvserver) and create SSH Tunnel (Putty, mobaxterm … using Windows)
- Connect Desktop/Laptop ParaView Client to server

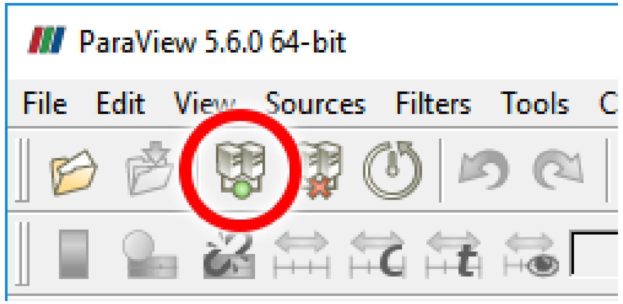

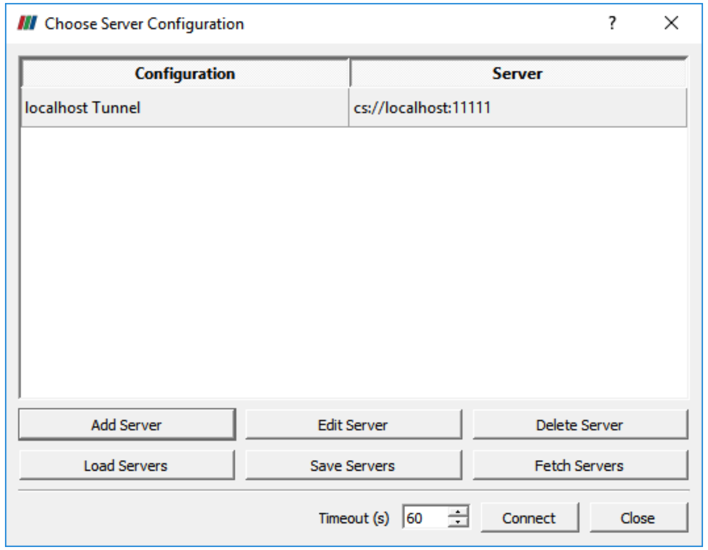

Many web sites to find information:

https://hpc.llnl.gov/running-paraview-client-server-mode

https://user.cscs.ch/computing/visualisation/paraview/

https://ciarc.mines.edu/visualization-home/paraview-connection-guide/

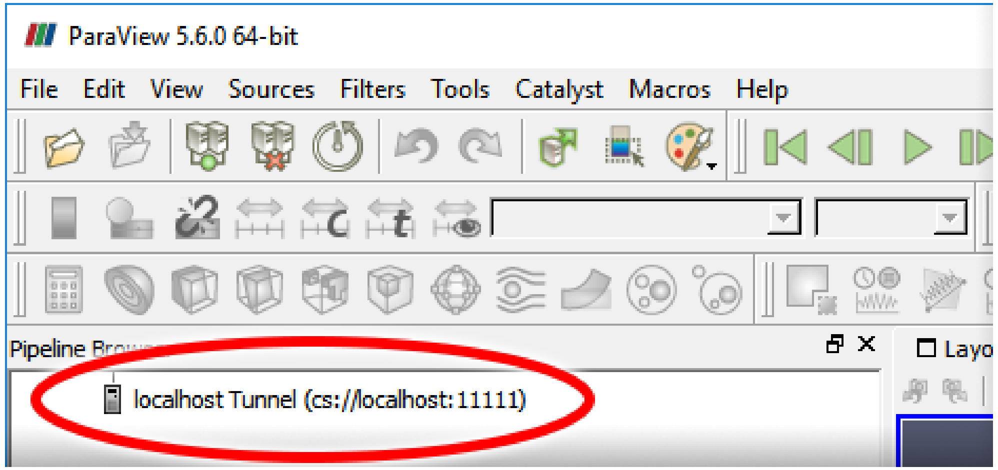

Now, the demo

A efficient solution…

Remote solutions such as ParaView client/server required to post-process results considering the exponential growth of simulated case size

Solution well deployed at IMFT especially thanks to COVID lockdowns

<u>Next steps:</u>

⇨ Evaluate co-processing during computation

⇨ Try specific compilation to optimize GPU ray tracing and produce photorealistic visualizations

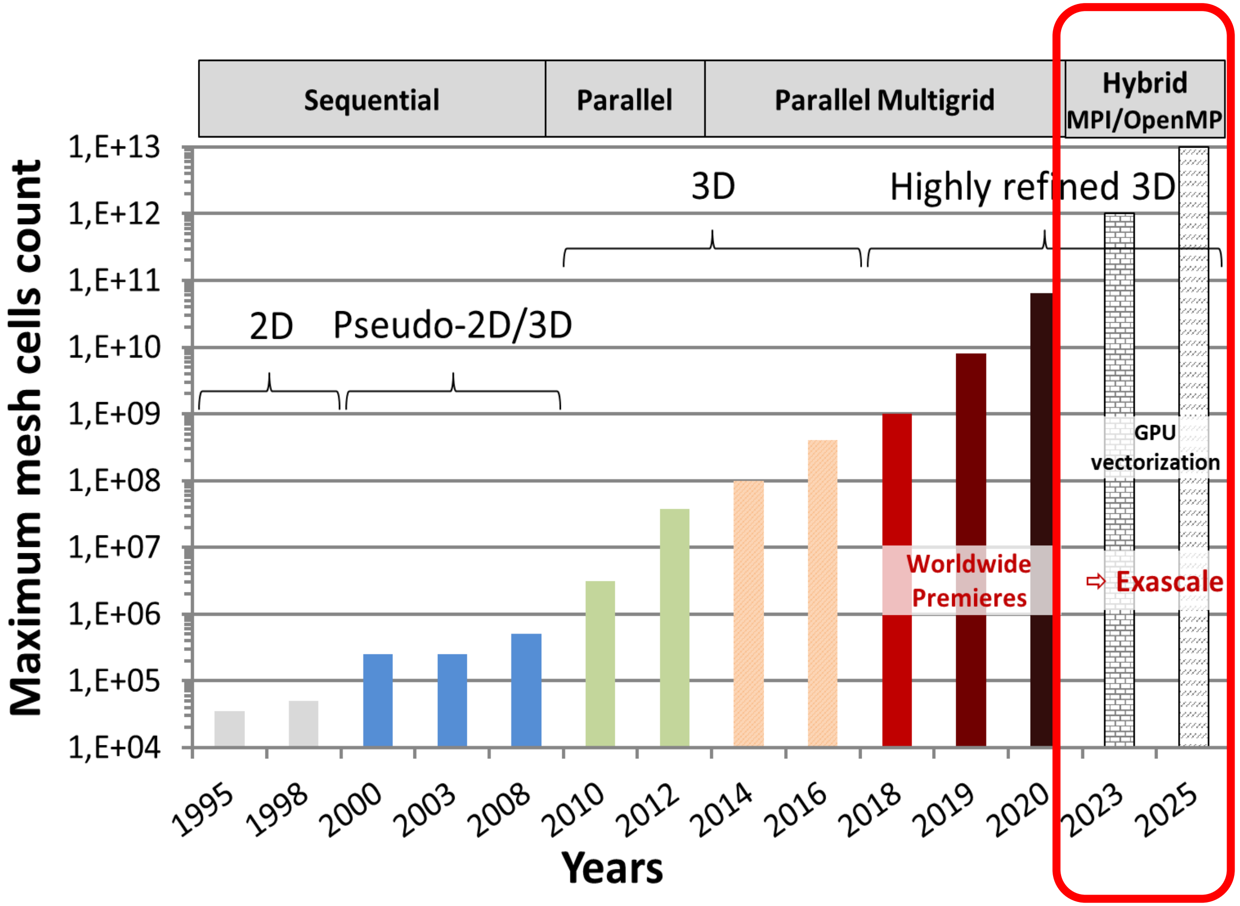

2018-2022: NEPTUNE_CFD Meso- and Grands-Challenges from Tier2 to Tier0

2018 — 2019 — 2020 — **2021**

CALMIP/EDF: Worldwide Premiere with 10^9 cells unstructured mesh → IDRIS: New Worldwide Premiere with 8 times bigger mesh : 8 billions cells → **TGCC: 64.10^9 cells mesh**

Meso-challenges Olympe@CALMIP

Grands-Challenges Gaia@EDF

Grands-Challenges Jean-Zay@IDRIS

Grands-Challenges Irene-AMD@TGCC

25s simulated

1.7s simulated

Only few iterations ⇨ sensitivity studies, profiling

⇨ **Reaching limits of both solver, MPI libraries and supercomputers**

⇨ **Failure when attempting to generate 512 billion cells mesh due to these limitations**