PERFORMANCE ISSUES FOR GEOSCIENCE APPLICATIONS

WORKFLOW AND NUMERICAL TOOLS

Stéphane de Chaisemartin
Soleiman Yousef

IFPEN

- Context
  - Climate change and ecological transition
- Gathering
  - **Scientific Management**
  - Business Units
    - **Energy Resources**
    - Process
    - Transport
    - Industrial Development
- Subsidiaries
  - Energy transition
  - Geoscience consulting and software
  - Alternative and renewable energies, refining, petrochemicals, gas, water
  - Training
- Energy Resources Business Unit
  - **CO_2** capture, use and **storage**
  - **Hydrogen storage**
  - Climate/soil interactions and the water cycle
  - Geothermal energy
  - Basin scale Hydrogen & Lithium exploration / CO_2 site screening

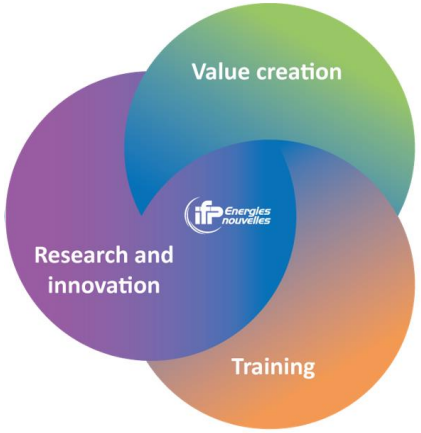

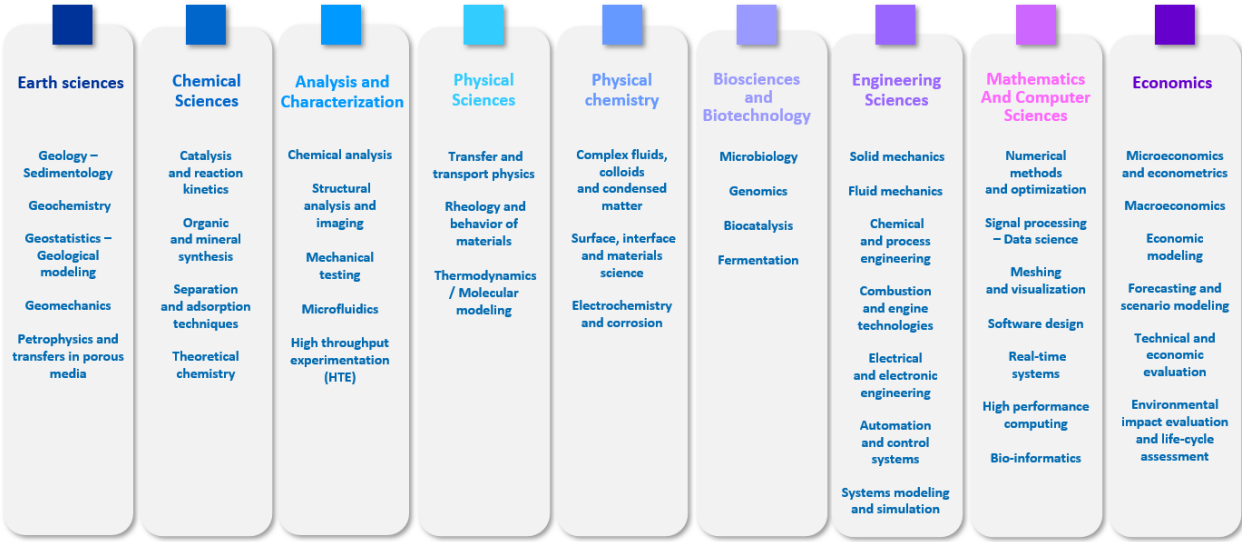

Axens
Powering integrated solutions
An IFP Group company

BeicipFranlab

TECHadvantage

IFP Training

IFPEN / NUMERICAL SIMULATION / HPC

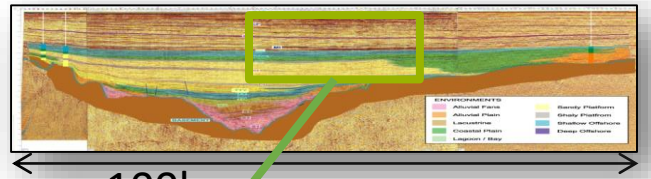

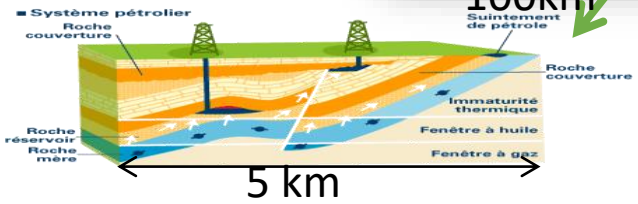

- General purpose
  - Robust and high-performance applications
    - Able to simulate more and more complexes phenomena in a reasonable time
      - To improve the quality of simulation results
      - To explore a wider range of scientific hypotheses
    - Parallel computing resources
      - European Exascale supercomputer
        - Jules Verne consortium
          - IFPEN letter of intente (https://www.genci.fr/fr/node/1236)
      - The performance of our applications is one of the major challenges for the **energy transition**
      - Contributing to high-performance and sovereign solutions
  - Be part of several PEPRs for GeoScience (physical and mathematical aspects)
    - Sous-sol, Math-vives
  - External partner
    - PEPR NumPEx

IFPEN supercomputer

Images from www.hpc.cea.fr

Images from www.genci.fr

TECHNICAL & SCIENTIFIC ISSUES

Challenges

- **Demonstrate:** Containment, Capacity and Injectivity
- **Optimize** injection
- **predict** storage behavior
- **Reduce** risks, **Assess** the environmental impact, **Ensure** the long-term future of the storage facility

- **Multi-scale modelling**
  - From pore scale to basin scale
  - From a few minutes to several thousand years
- **Multi-physics (coupled processes)**
  - Compositional multiphase flows /Phase changes
  - Mechanical deformations (variations in pore volume, fault reactivation)
  - Fluid-rock geochemical reactions
- **Addressed aspects**
  - **Numerical & Robustness**
  - **High performance computing**
  - **Shared development platform**

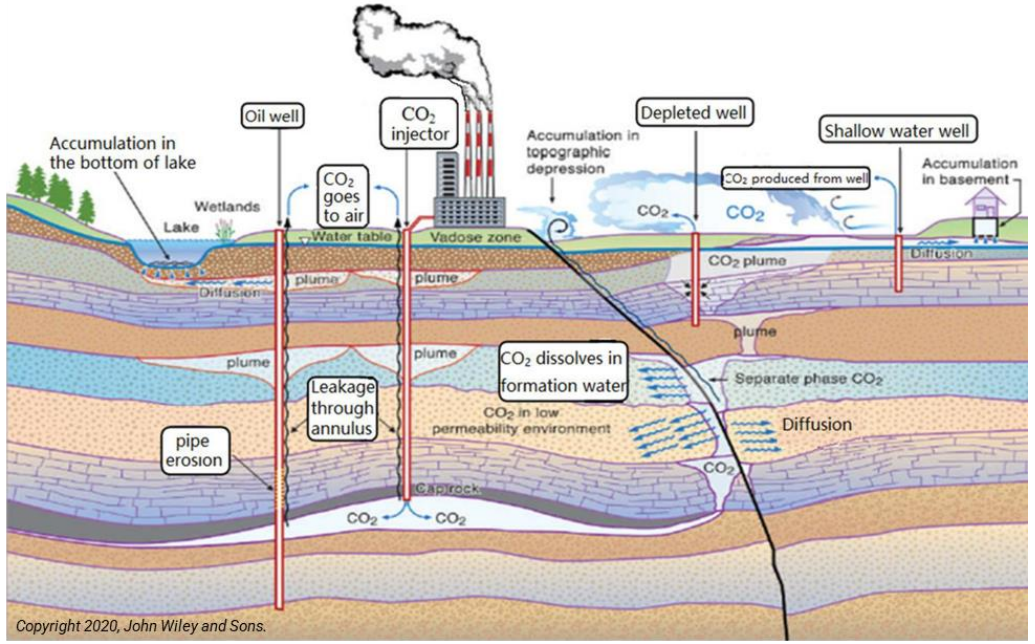

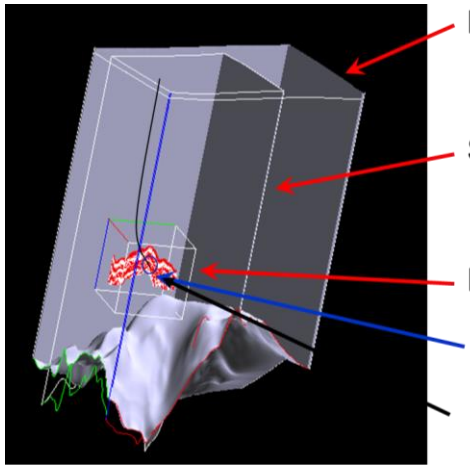

Basin scale

Interactions with other basin fluid exploitations: oil & gas production, gas storage, geothermal site

Storage Complex scale

Fluid migration and pressure dissipation

Reservoir storage scale

Anisotropy and heterogeneity

Near wellbore scale

injectivity

Pore scale

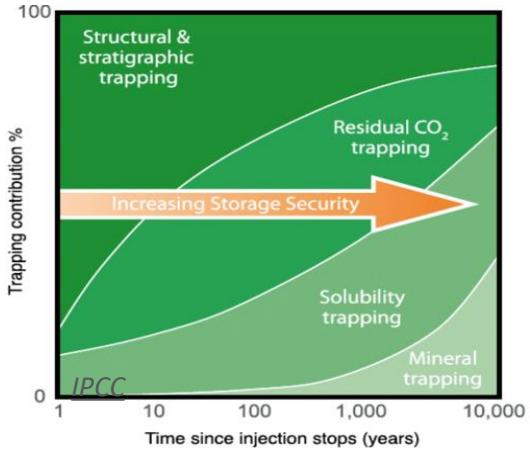

ARCGEOSIM

ArcGeoSim™

ArcGeoSim™ : environment for IFPEN new gen geosciences simulators

- *Performance of simulation codes*
- *Capitalizing on R&D*
- *Efficient development*

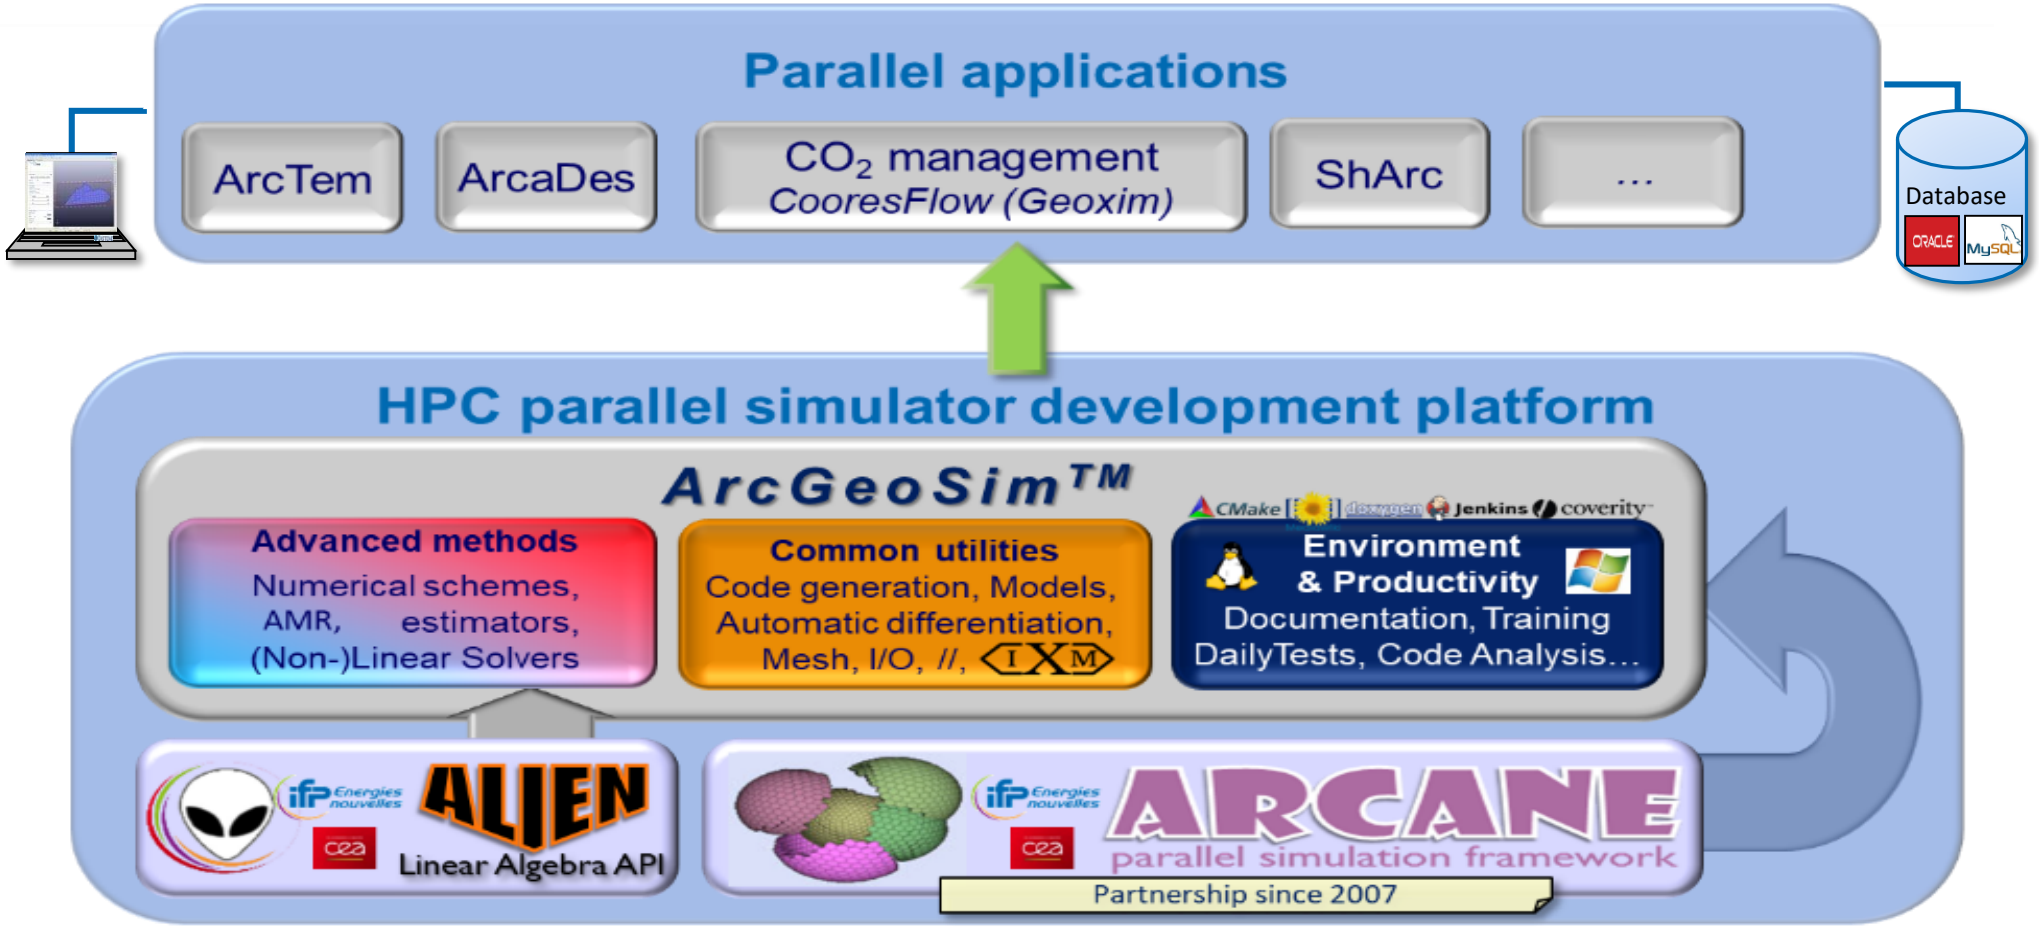

ARCANE PLATEFORM

- Started at CEA on the early 2000s, co-developed at IFPEN since 2007
- Provides the foundation to build a parallel simulator
  - Mesh and Variables,
  - Parallelism, Partitioning, Load balancing,
  - IO, Checkpoint-restart…
- Structures the application workflow into
  - Modules (independent units)
  - Services (plug-ins)
- Makes available different parallel strategies
  - Message passing (implemented via mpi, threads or mpi+threads)
  - Concurrent loops (implemented via TBB)
  - Classes for vectorization (e.g. SIMDCell, for SSE, AVX and AVX512, MIPP plug)
- An API for accelerators has been recently added (2021-2023)
  - Arcane concepts available in the device (Variables, loops on mesh, connectivity's…)
  - Dynamic choice between accelerator, threads or sequential environment
  - Two implementations currently available: Nvidia cuda & AMD rocm

ALIEN LIBRARY

AN INTERFACE TO THE LINEAR ALGEBRA WORLD

- **ALIEN** Algebra library
  - A unique interface for many linear solvers
    - For **Arcane**-based code
    - **…or not (C-Fortran API) !**
      - e.g. plugging into **Traces code - Andra** (Fortran)
  - Reduced efforts to plug solvers
    - A unique access point to the main solver libraries
    - Includes IFPEN in-house solvers
  - Shared work between CEA and IFPEN
    - Shared developments by the entities
    - ⇒Lower development and maintenance costs

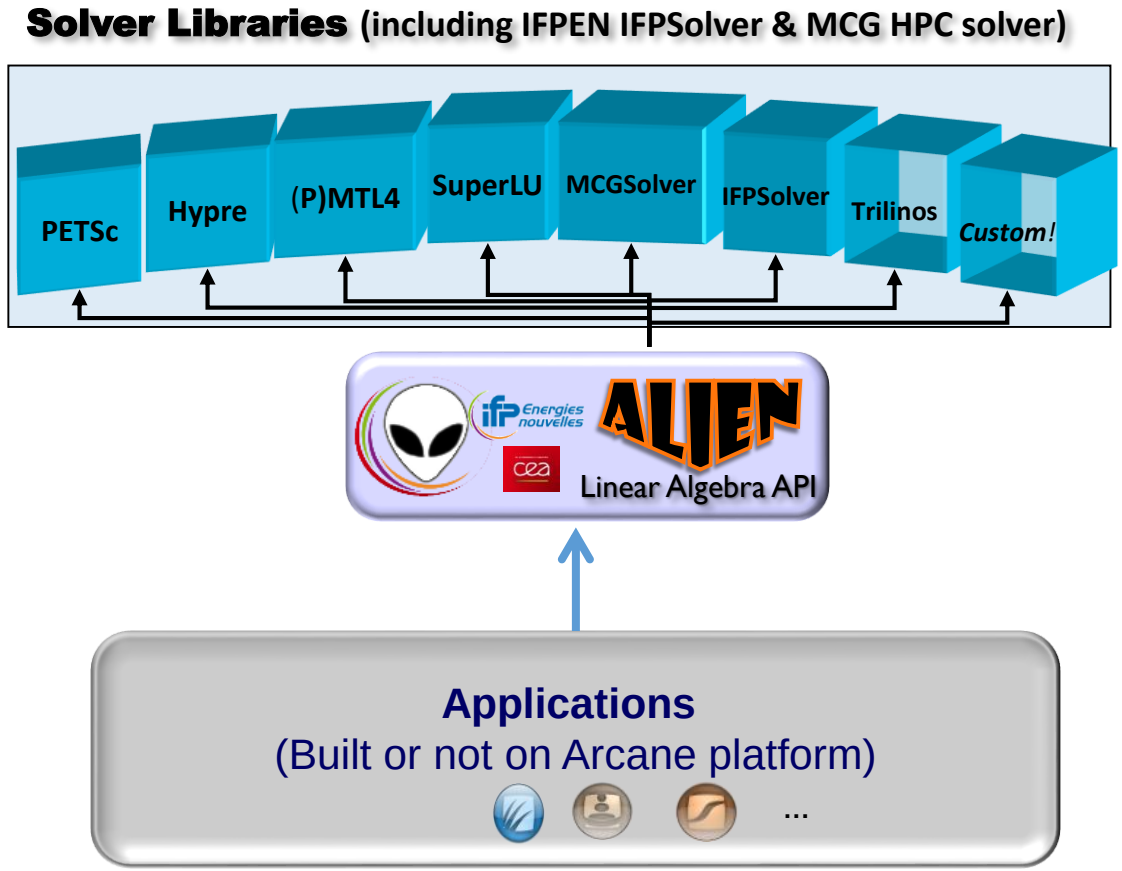

OPEN SOURCE STRATEGY

PLATFORMS MADE OPEN SOURCE

- Towards Open Source building blocks
  - Easier collaboration between CEA & IFPEN and with academics
  - Aims at better visibility for IFPEN in scientific computing community
  - May benefit from Open Source community (user community, new use cases…)
- Alien made Open Source end 2020
  - Used in Andra, code TRACES : internship in 2021, plugged in production version in 2023
- Arcane made Open Source end 2021
  - Open Source efforts defined in CEA-IFPEN cooperation agreement
- Join Arcane and Alien teams on GitHub/Arcaneframework !

THE APPLICATIONS

COORESFLOW
TOOL FOR SIMULATING FLUID-ROCK INTERACTIONS

- A model used for **Hydrogen & Co2 storage** and also for **geothermal energy** to **address carbon footprint reduction challenges**

- **R&D simulator** : Two calculators and one interface
  - **CooresFlow** Interface
  - **ArximCpp**
    - Chemistry / 1D Transport
    - Detailed modelling of water-rock interactions
  - **Geoxim** multi-D simulator
    - Multiphase flow / Transport
    - Chemistry / Thermal / Mechanics

- Involved in international collaborations (e.g. Petrobras)

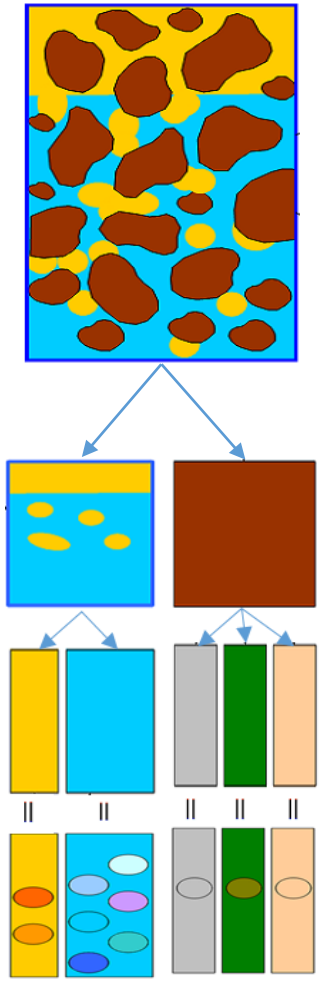

gas – water – minerals

GEOXIM

- Thermal-hydro-chemical-mechanical modelling
  - Porous medium: « water - gas/liquid – rock » system, multiphase flows
  - Chemical reactions, particularly fluid-rock: precipitation/dissolution
  - Partial coupling with mechanical equilibrium equations
- Multiple “industrial" models
  - **Compositional multiphase flows**: relative permeability, capillary pressures with hysteresis, thermodynamic models, dispersion, diffusion, etc.
  - **Heat transfer**: Conduction, convection, etc.
  - **Fluid-rock interactions**: Chemical species database, Equilibrium/Kinetics
  - **Mechanical deformation**
- Advanced object design for model flexibility and development reliability
  - **Generic** description of a reactive multiphase system and its data model
  - Management of behavior laws
    - **Flexible** dependencies and forms: physical/analytical/tabulated/ML
    - Easily extendable library
  - **Automatic differentiation** tool for easy, reliable assembly of the Jacobian

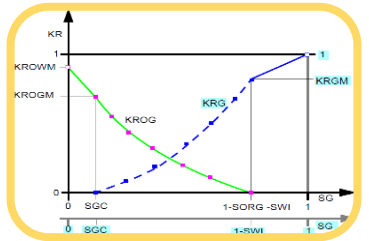

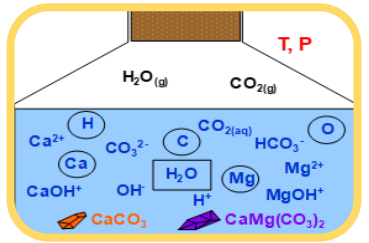

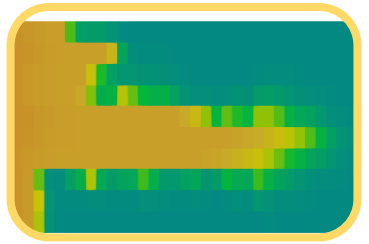

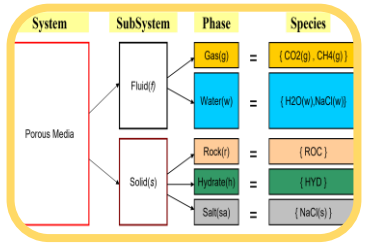

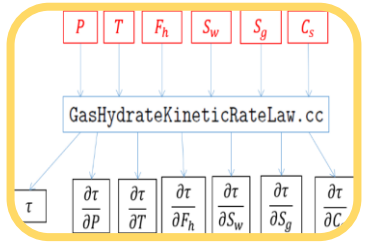

INDUSTRIAL REALIZATIONS: MESHES

ADVANCED CAPACITY FOR GEOSCIENCE MESHES

- Several types of eligible mesh
  - Enriched unstructured Mesh for Geoscience
    - Contains all the degeneracies of the hexahedron
  - Cartesian mesh
  - Polyhedral mesh (under development)

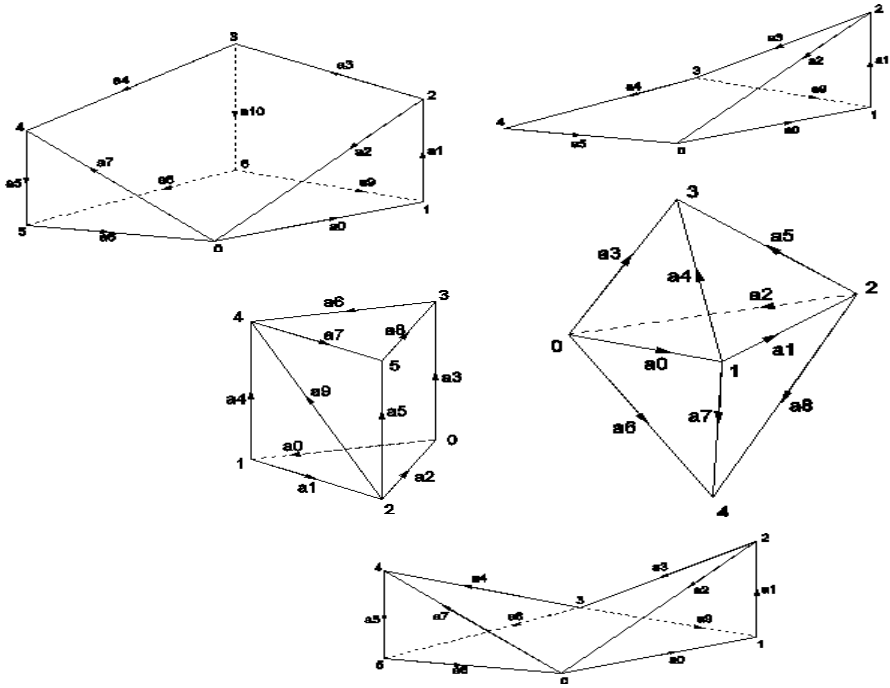

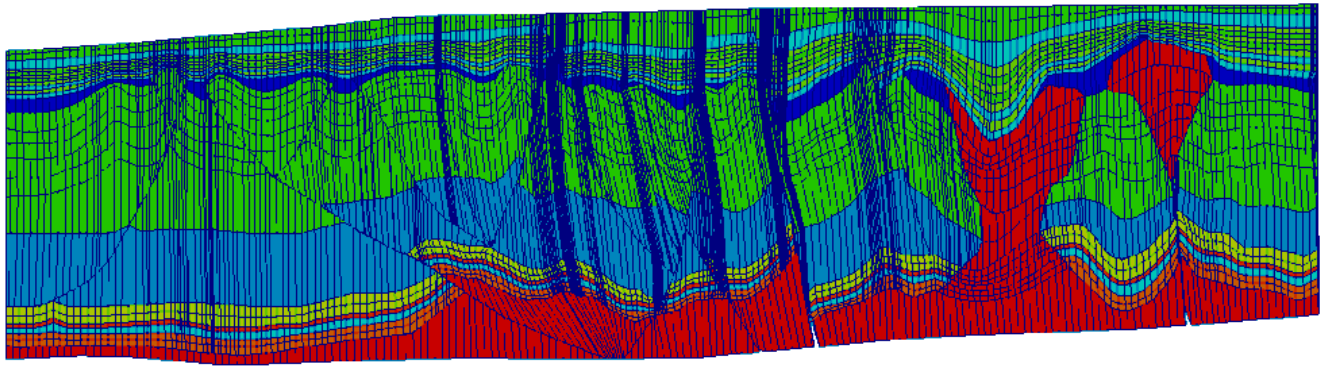

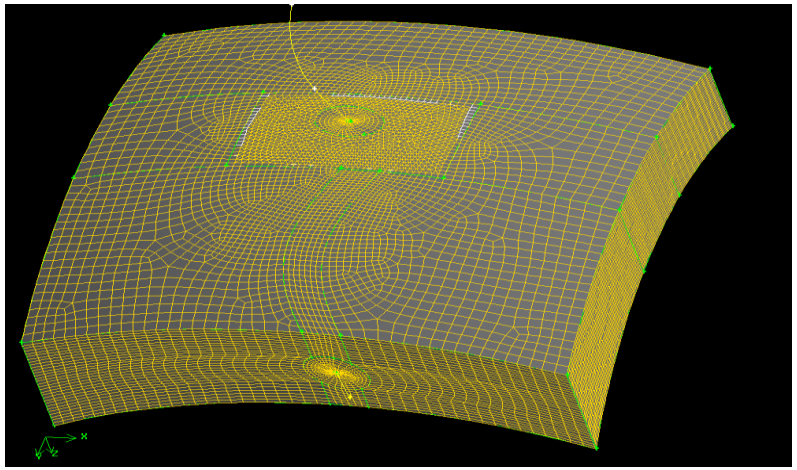

INDUSTRIAL REALIZATIONS: MESHES

ADVANCED CAPACITY FOR GEOSCIENCE MESHES

- Numerous tools for meshing
  - Hierarchical AMR
  - Co-raffinement for calculation across faults
  - Evolving mesh for basin simulation

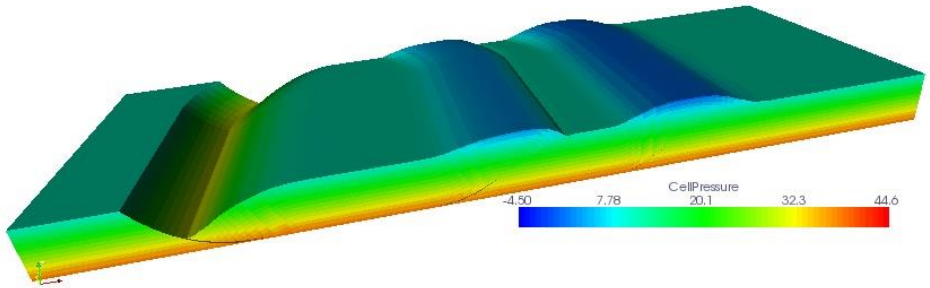

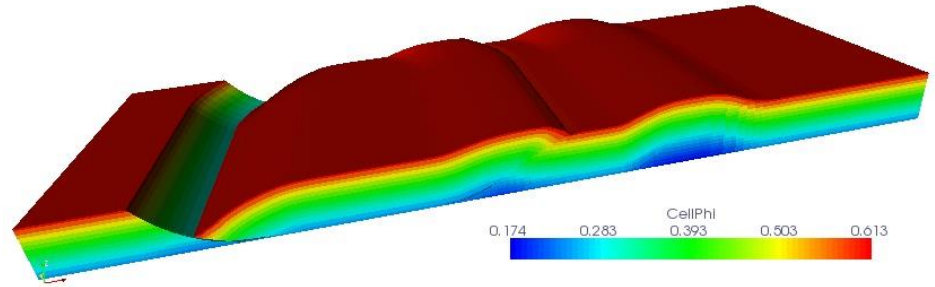

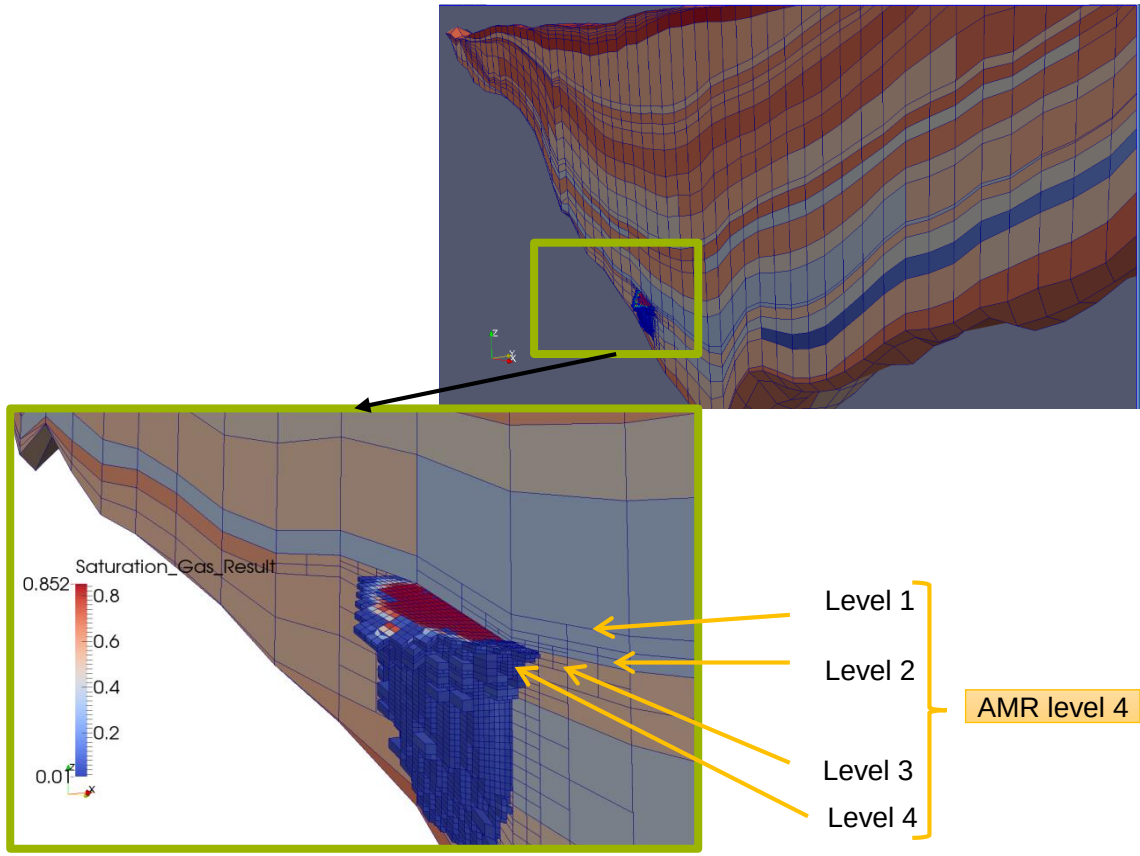

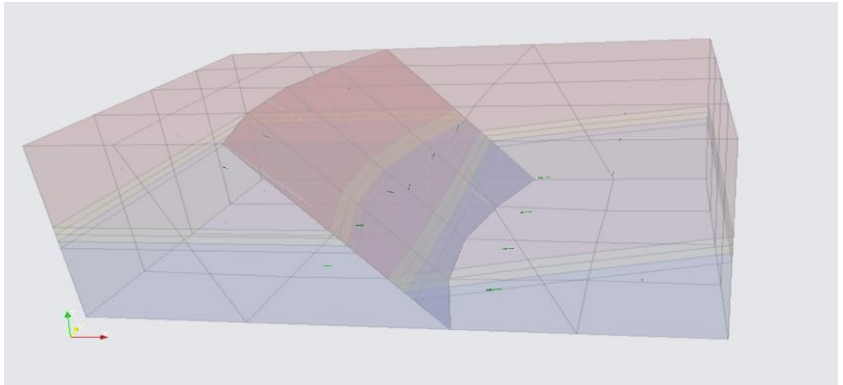

DISCRETIZATION METHOD

- Two-Point / Multi-Point flux approximation **(TPFA / MPFA)**
- Hybrid finite methods **HFV**

- Virtual element method **(VEM)** for **geomechanics**
  - Research work / integration ArcGeoSim
  - Integration (one-way mechanical-flow coupling) in Geoxim (collaboration with **Petrobras**)
- Non-linear two-point scheme **(NLTPFA)**
  - Research work / integration ArcGeoSim
- Coupling (VEM-NLTPFA)
  - Preliminary tests using the ArcGeoSim platform
  - Tests in progress on polyhedral meshes

- Hybrid High-Order methods **(HHO)**
  - Research work – fracture propagation

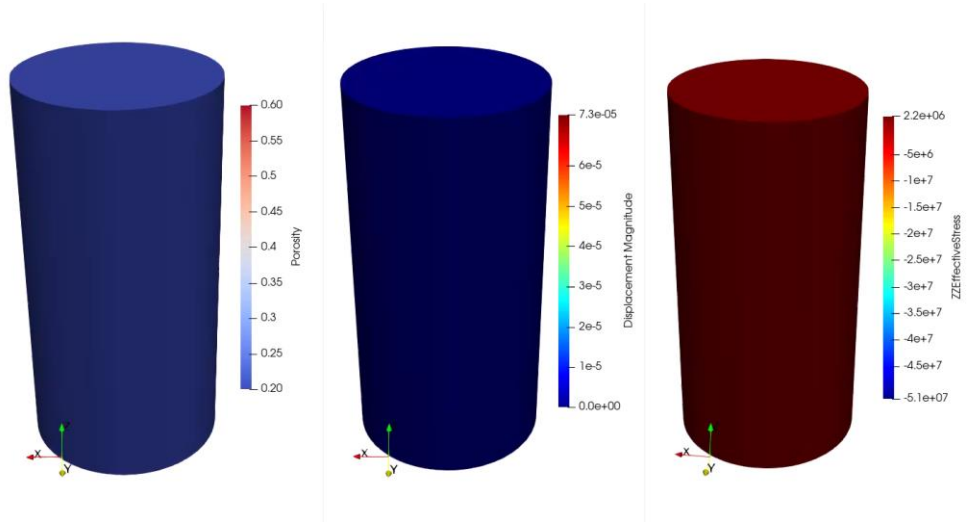

Mechanical coupling and reactive flow on a rock plug (Geoxim-video)

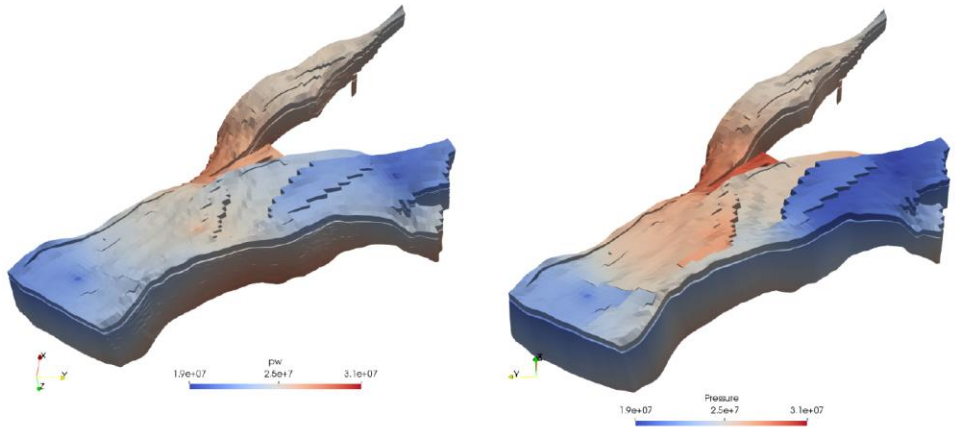

(a) NLFV TPFA (44644 unknowns) (b) HFV (388978 unknowns)

ADAPTIVITY

- Enhance the simulator run-time without affecting the accuracy.
- Use **a posteriori error estimate** to efficiently drive adaptive algorithms.
- A simple a posteriori estimate on general polytopal meshes
  - Research work / integration **ArcGeoSim**
  - Goal-oriented a posteriori error control
  - Time stepping adaptivity
  - Adaptively regularized inexact semismooth Newton method
  - Adaptive linear solvers

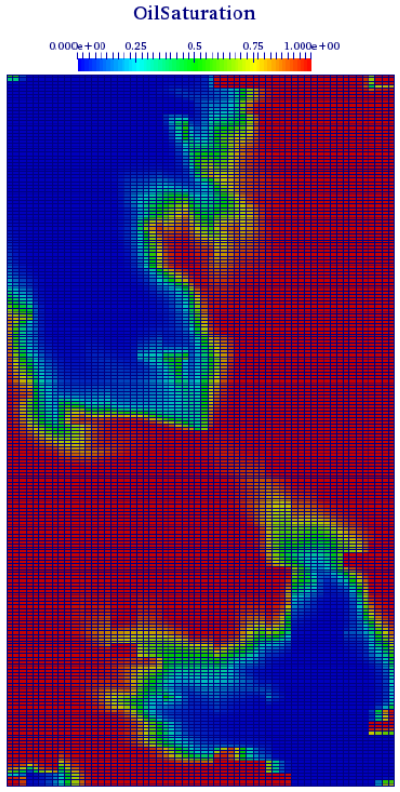

Fine mesh

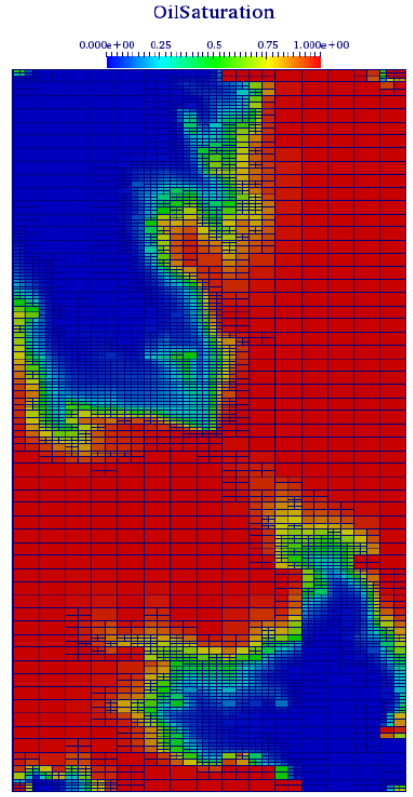

AMR (a posteriori)

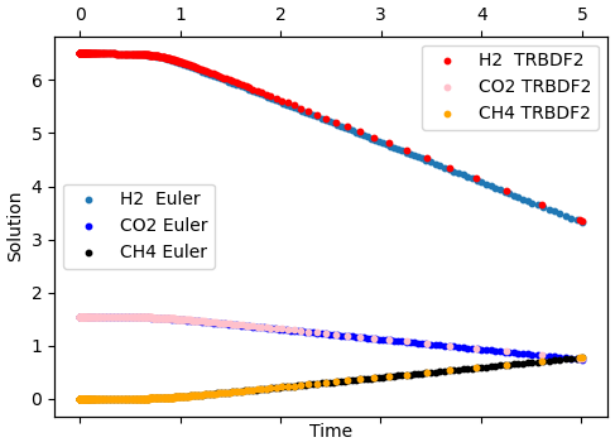

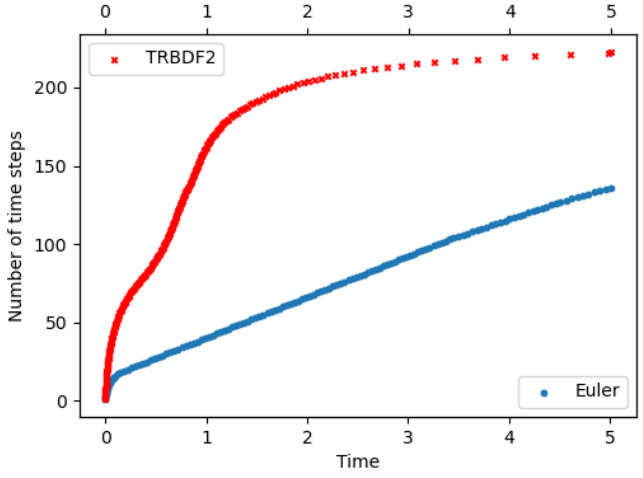

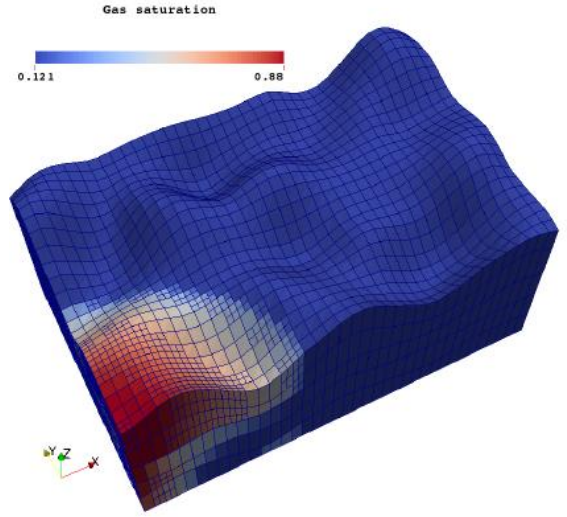

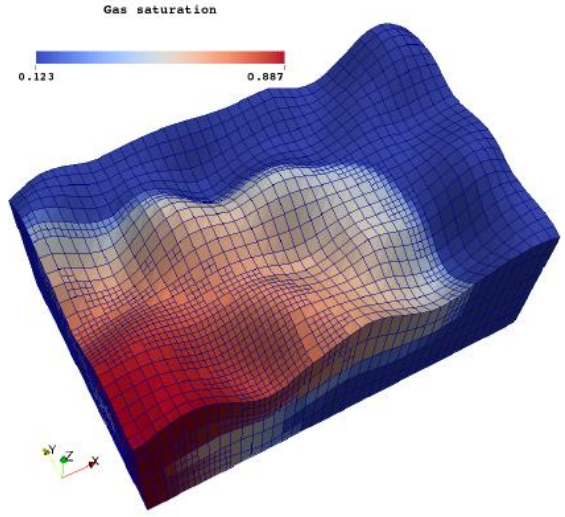

ADAPTIVITY / AMR

AMR enables the zone of interest to be dynamically refined

CooresFlow / Geoxim - Ultimate CO2 European project : Paris Basin with refined well head

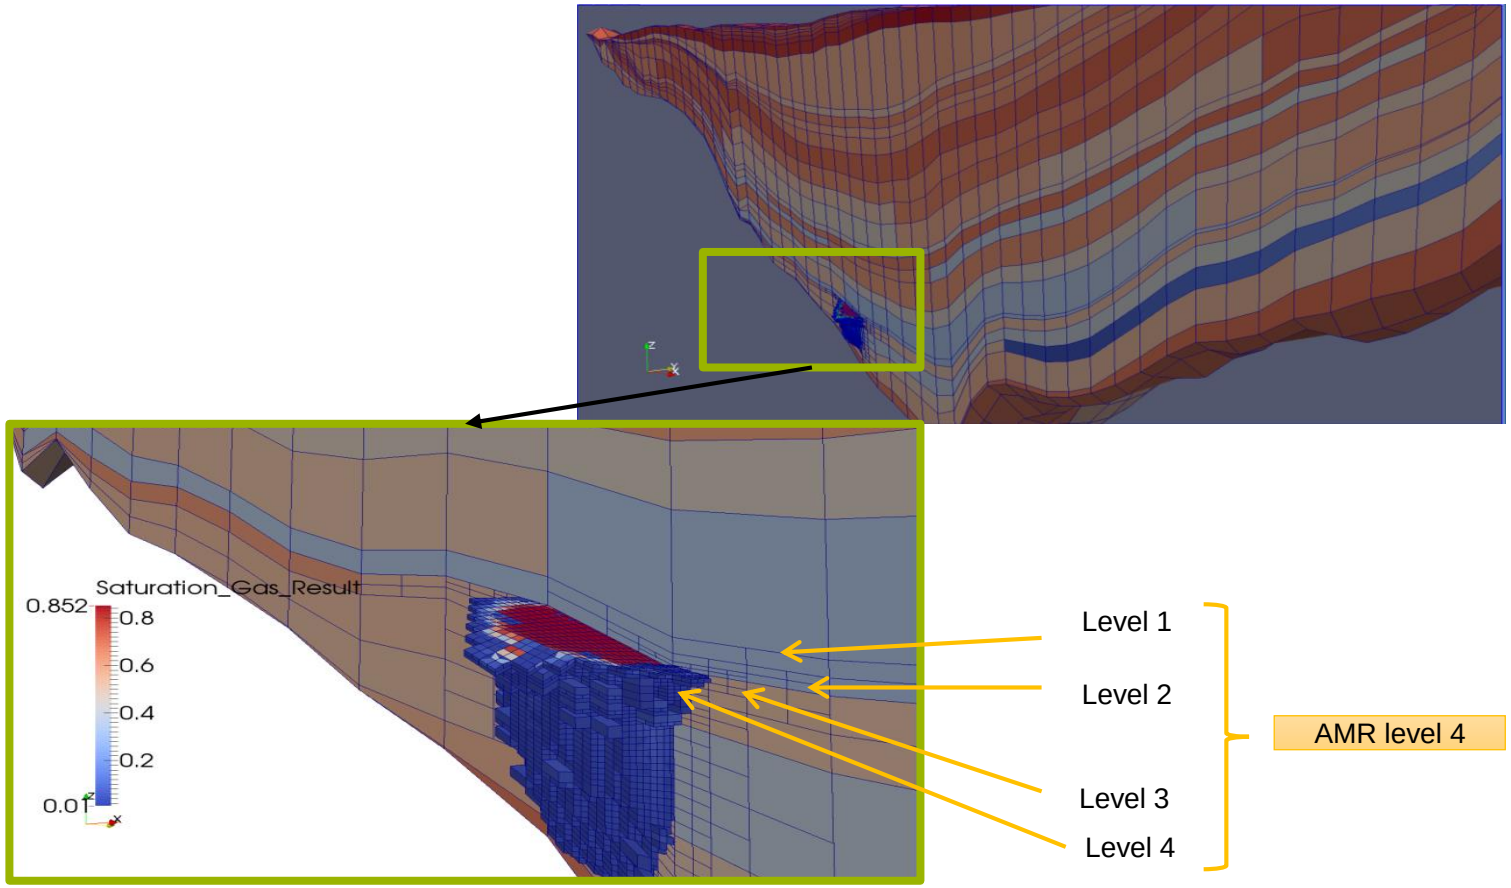

Not addressed issues :

- Load balancing for AMR

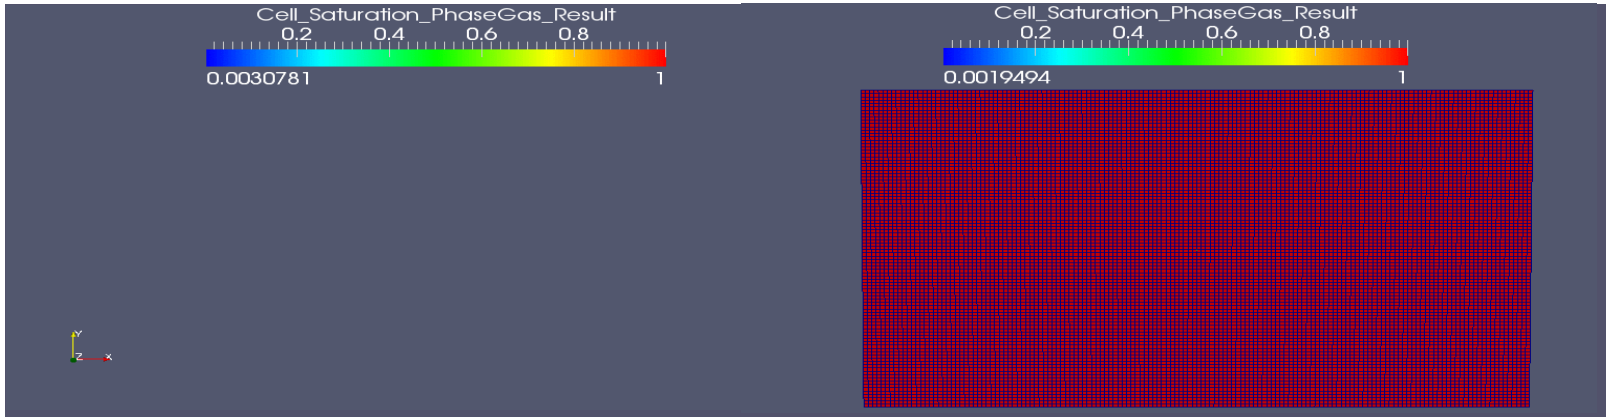

Refined Mesh (AMR)

Fine Mesh Reference

Mesh refinement improves performance

CooresFlow / Geoxim – diphasic flow : equivalent details 3.5 faster with AMR refinement.

OPEN SOURCE MINI APPLICATION (**ShArc**)
HPC performance / Grand Challenge

SHOWING ARCANE

TOWARDS OPEN SOURCE APPLICATIONS

- ***ShArc*** demonstrator (*Showing Arcane*)
  - mid-2022 on GitHub
- What is ShArc ?
  - An application for simulating flow in a two-phase porous medium
  - Solved using conventional numerical tools :
    - Fully-implicit discretization
    - Two-point finite volume schemes
    - Newton solver (ArcGeoSim)
    - Alien for linear solvers
- What tools does it use ?
  - Gump : a tool for generating data models of a physical system
  - Law : a generation tool for evaluating physical laws
  - Audi :an automatic differentiation tool

```
const Law::Contribution grad_kl = T[iface] * (P[cell_k] - P[cell_l]);
ENUMERATE_PHASE(iphase, fluid.phases()) {
  const Law::Contribution mobility_k = rho[iphase][cell_k] * kr[iphase][cell_k] / mu[iphase][cell_k];
  const Law::Contribution mobility_l = rho[iphase][cell_l] * kr[iphase][cell_l] / mu[iphase][cell_l];

  const Law::Contribution flux_kl =
        ( audi::one(grad_kl >= 0) * mobility_k
        + audi::one(grad_kl <  0) * mobility_l ) * grad_kl;
```

Flux computation (Gump + Lois + Audi)

```
const Arcane::Integer iequation = iphase.index();
if (cell_k.isOwn()) {
  residual[iequation][cell_k]          += flux_kl;
  jacobian[iequation][cell_k][stencil] += flux_kl;
}
if (cell_l.isOwn()) {
  residual[iequation][cell_l]          -= flux_kl;
  jacobian[iequation][cell_l][stencil] -= flux_kl;
}
```

Adding the flux to the residual and the Jacobian

PERFORMANCES SHARC

Processor info : **AMD Rome**

- CPUs : 2x64
- Frequency : 2.6Ghz
- cores/Node : 128
- Sockets/Node : 8
- Cores/Socket : 16
- Threads/ Cores : 2

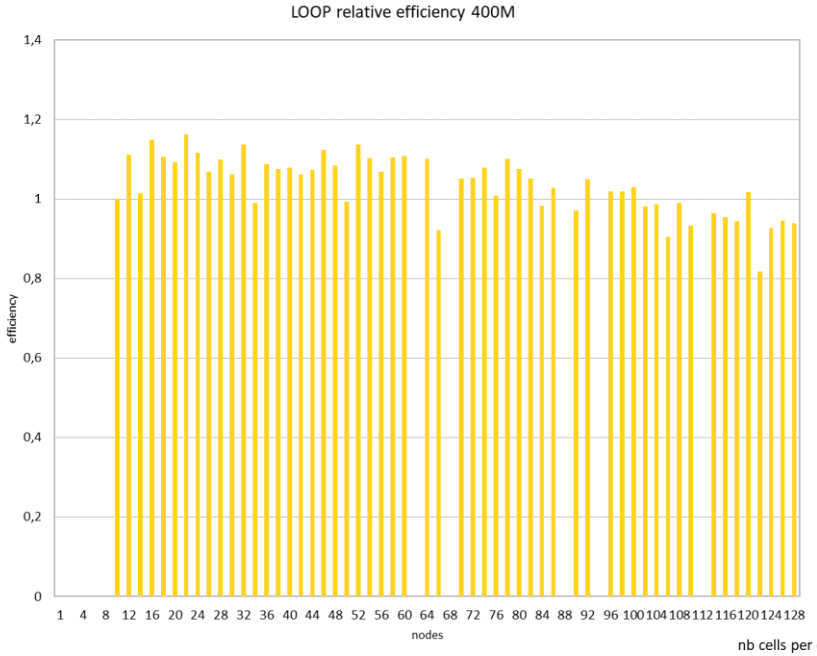

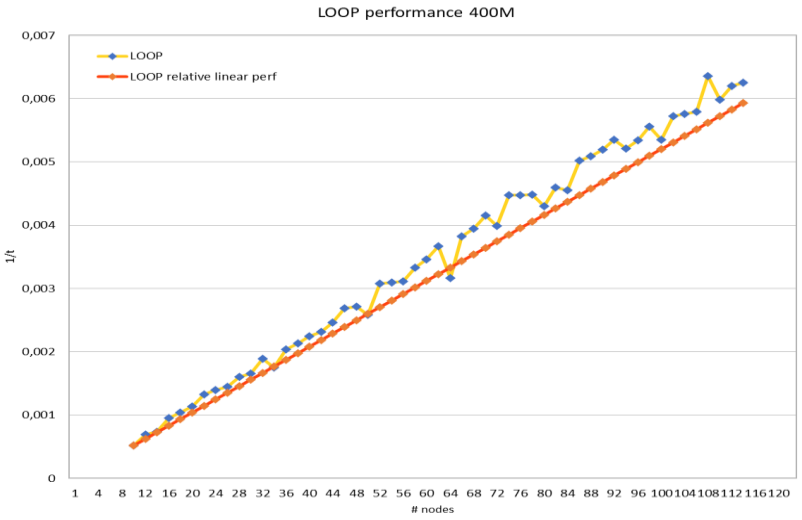

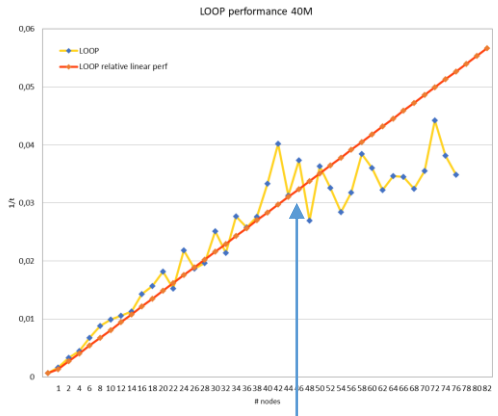

44 nodes

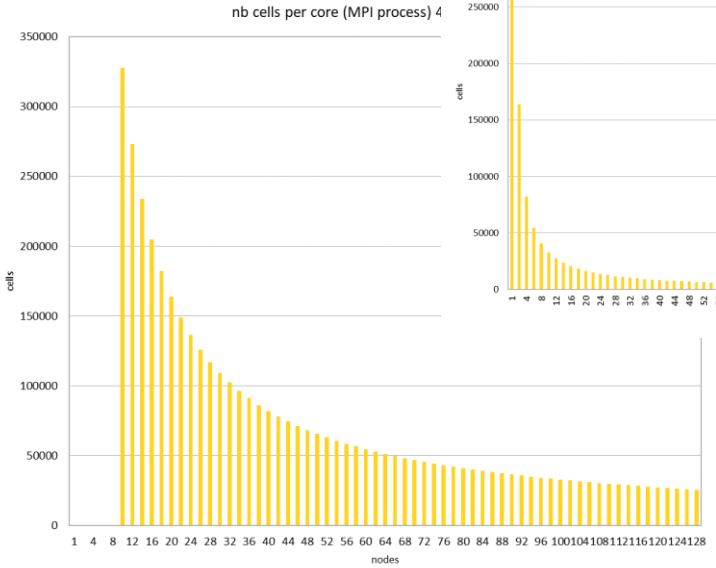

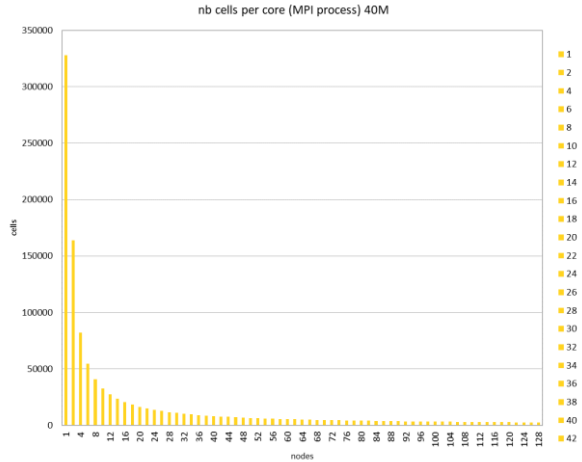

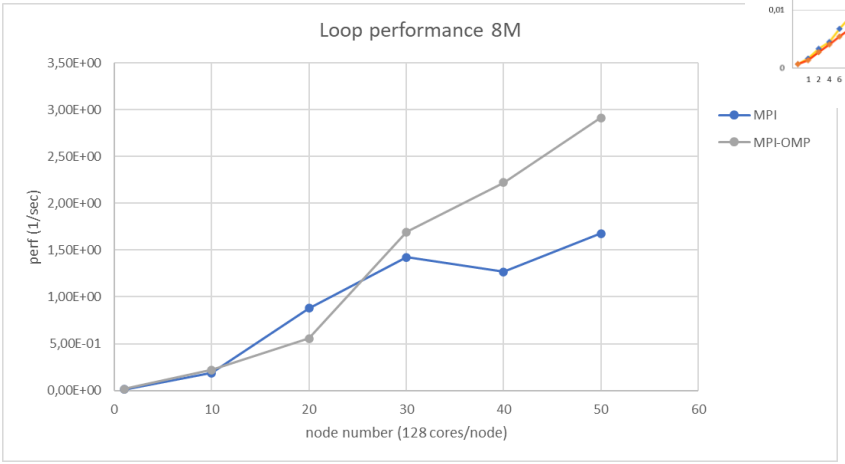

Test case: 400 000 000 cells, 3 unknowns per cell

- Good linear solver scalability (LOOP)
  - MCGSolver (Non convergence for Petsc)
- 12 000 cells / cores.

PERFORMANCES SHARC

HPC overview for codes@Arcane

- CEA
  - **Unstructured** code : **50 million cells + 4.1 billion particles on 32 000 cores**
  - **Structured** code + **AMR** : **20 billion cells on 130 000 cores**
- IFPEN
  - **Structured** code : up to **400 million cells on 16 000 cores**
  - Good speedup, down to **12 000 cells per core**
  - Requires iterative inhouse solver (MCG)

Going further

- Decrease minimum cell number per core : **hybrid parallelism**
- Use **accelerator** : API available in Arcane but not used yet.

CONCLUSION

- Issues and barriers need to be more investigated
  - Exascale Algorithms: Redesign algorithms to improve scalability (e.g., reduce communication, avoid/hide synchronization) and computational efficiency on accelerators.
  - Intensive compute at node level, using appropriate algorithms
  - Scaling up and handling highly heterogenous computing resources & data representation
  - Load balancing for AMR using homogeneous and heterogenous systems

- Some references
  - **The Arcane development framework,** Grospellier et al, *Proceedings of the 8th workshop on Parallel/High-Performance Object-Oriented Scientific Computing*, 2009
  - **Alien: a Flexible Wrapper API on Linear Solvers,** Chevalier et al, *10th International Workshop on Parallel Matrix Algorithms and Applications*, 2018
  - **Leveraging Single Node Linear Algebra GPU Solvers from a Multi-Nodes User Code with Alien**, Chevalier et al, *Conference on Parallel Processing for Scientific Computing*, 2022
  - **Evaluation of the performance portability layer of different linear solver packages with Alien**, Gratien et al, *Proceedings of ECCOMAS*, 2022

Innover les énergies

Retrouvez-nous sur :

www.ifpenergiesnouvelles.fr

@IFPENinnovation

- Proposed Joint ONERA/IFPEN PhD thesis
  - Neural Linear Solvers and Preconditioners for General Sparse Matrices

GEOXIM – PETROBRAS TEST CASE

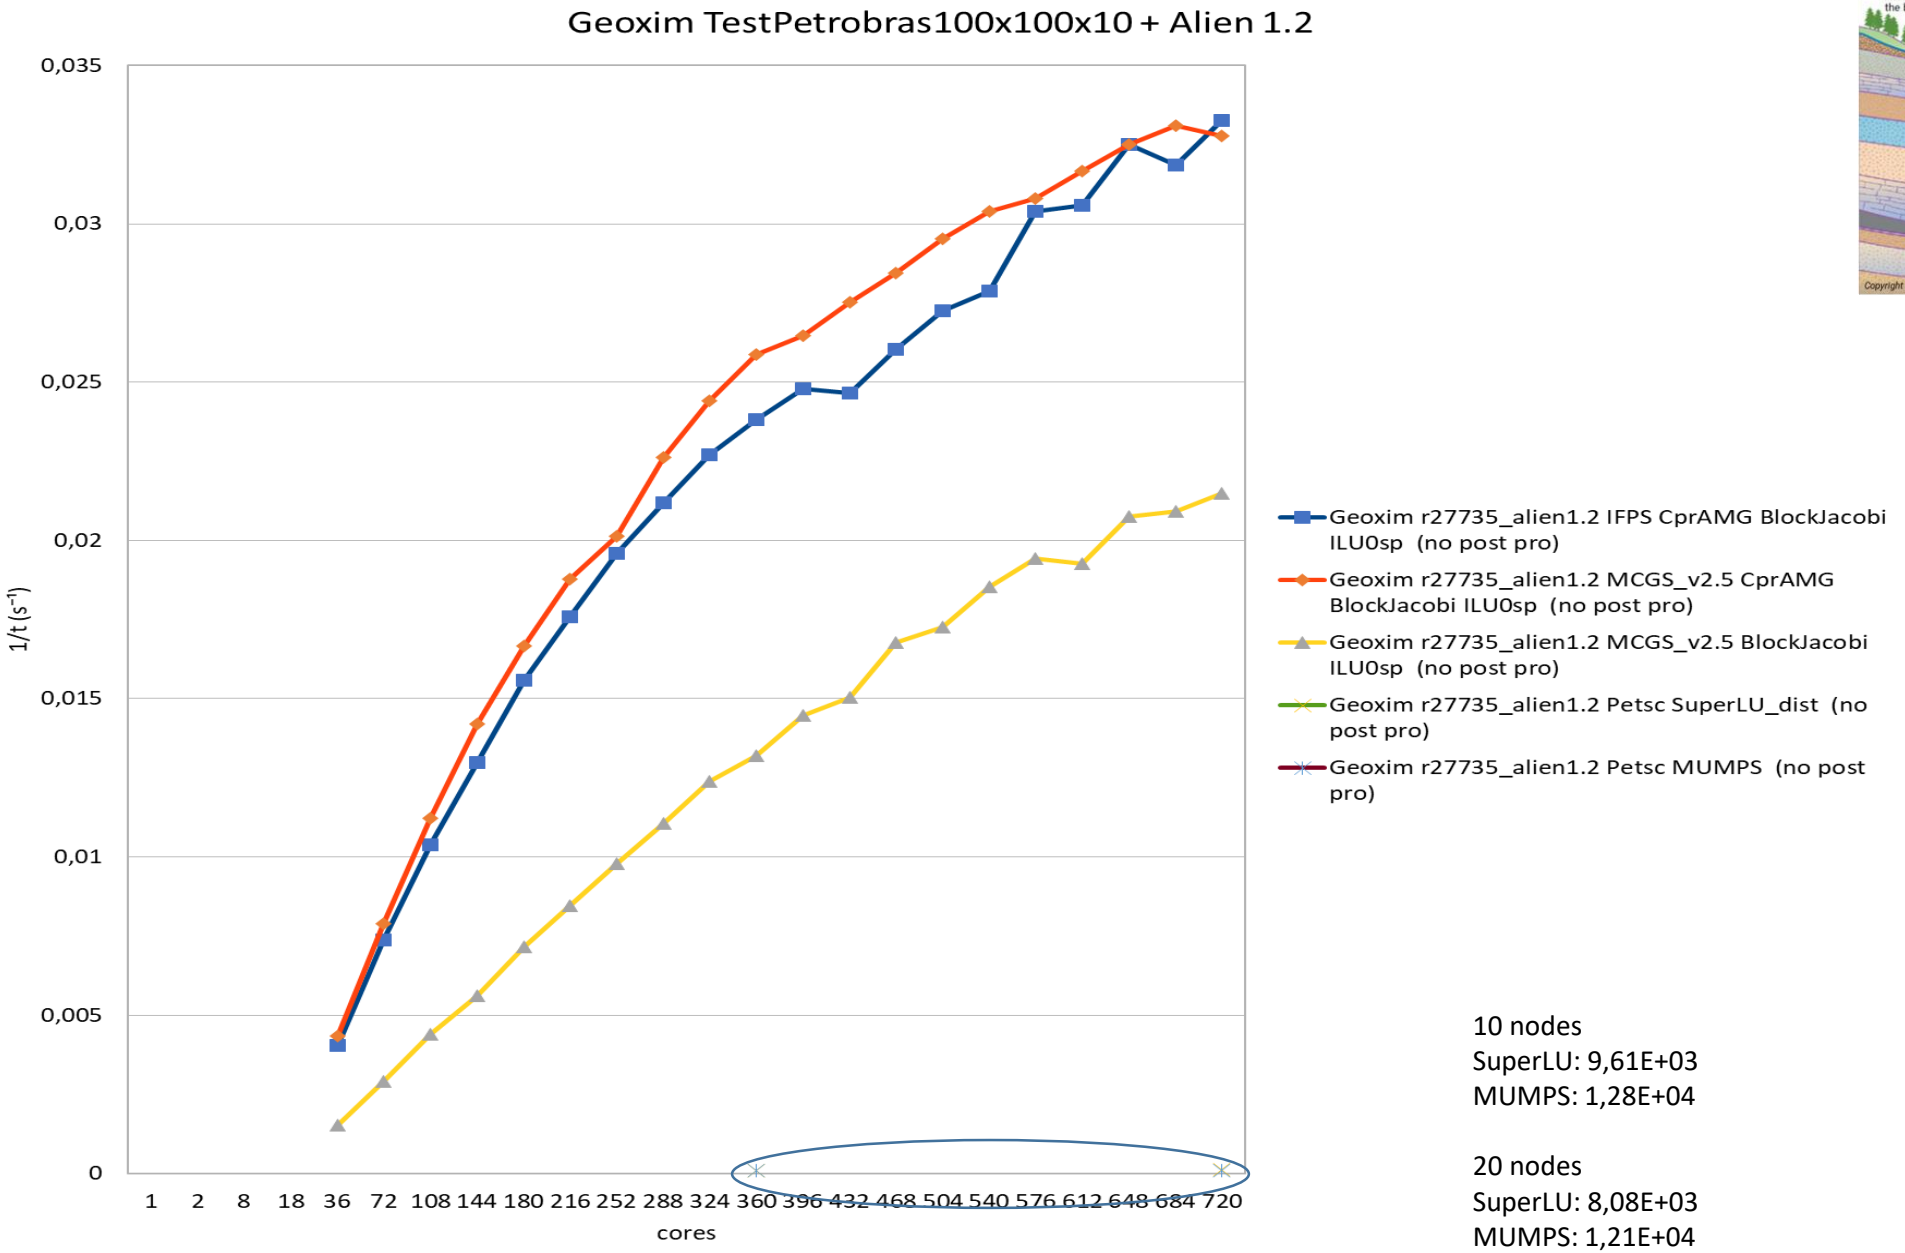

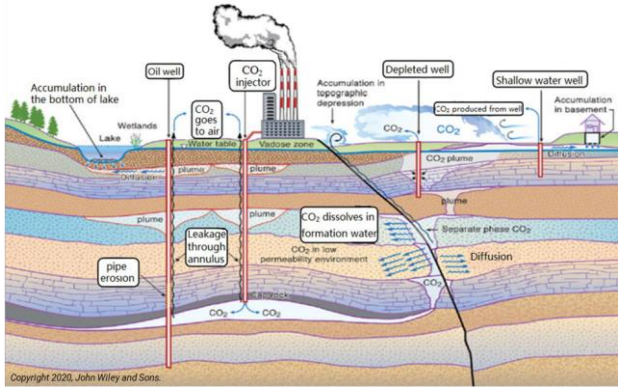

IFPEN supercomputer

10 nodes
SuperLU: 9,61E+03
MUMPS: 1,28E+04

20 nodes
SuperLU: 8,08E+03
MUMPS: 1,21E+04

LES SERVICES **ARCGEOSIM**

DES SERVICES BASES SUR ARCANE

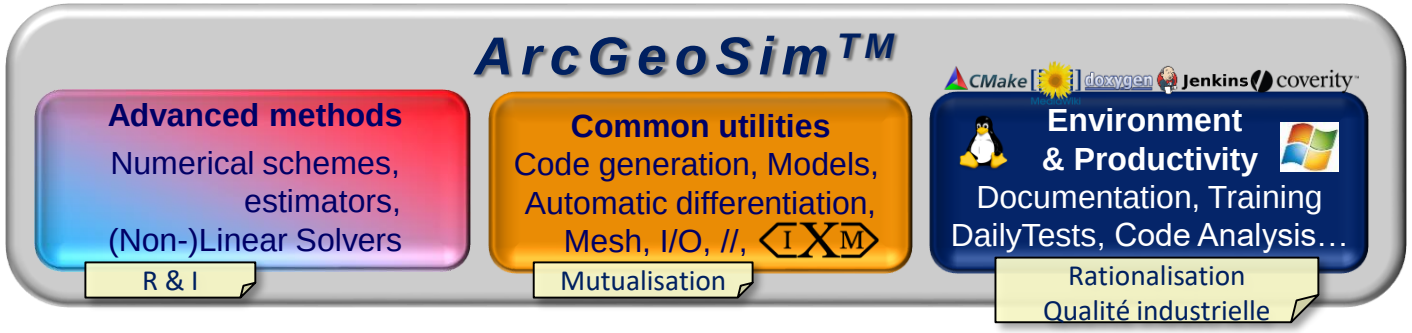

- Une intégration de la R&I calcul scientifique
  - Schémas numériques
  - Solveurs (non) linéaires
- Des utilitaires communs
  - Outils numériques (lois, différentiation automatique…)
  - Formats de maillage spécifiques(ixm)
  - Outils de post-processing
- Un environnement de développement et de production de code industriel :
  - Kit de développement (multi-plateforme)
  - Intégration continue et gestion de versions
  - Production de versions automatisée

LA SUITE…

COLLABORATION ARCANE, OPEN SOURCE, DIFFUSION

- Préparation du prochain avenant collaboration Arcane (2025-2028) avec le CEA
- Open Source : poursuite des efforts et instruction
  - De nouveaux modules dans ShArc
  - Instruction pour de nouvelles ouvertures : ArcGeoSim ? Une application complète…
- Rencontres Arcane : journée pour les développeurs d'applications IFPEN et CEA
  - Précédentes éditions : 17 Avril 2023
    - Participation du CEA-DES, de l'ONERA, de l'Andra et de l'INRIA (STORM)…
  - Prochaine édition peut-être dès 2024

- Quelques références
  - **The Arcane development framework,** Grospellier et al, *Proceedings of the 8th workshop on Parallel/High-Performance Object-Oriented Scientific Computing*, 2009
  - **Alien: a Flexible Wrapper API on Linear Solvers,** Chevalier et al, *10th International Workshop on Parallel Matrix Algorithms and Applications*, 2018
  - **Leveraging Single Node Linear Algebra GPU Solvers from a Multi-Nodes User Code with Alien**, Chevalier et al, *Conference on Parallel Processing for Scientific Computing*, 2022
  - **Evaluation of the performance portability layer of different linear solver packages with Alien**, Gratien et al, *Proceedings of ECCOMAS*, 2022

RÉALISATIONS INDUSTRIELLES: HPC

ETUDE PERFORMANCE SOLVEURS LINÉAIRES POUR LA GÉOSCIENCE

- Évaluation de la portabilité de la performance des solveurs utilisés
- Comparaison des solveurs sur des supports d'exécution différents :
  - Parallélisme par décomposition de domaine MPI
  - Parallélisme de tâches OpenMP
  - Accélérateur GPU (CUDA)
- Étude sur notre cluster IFPEN
  - 240 dual-socket Nodes with Intel Skylake G-6140 processors at 2.3 GHz. (18 cores per socket)
  - GP-GPU NVidia P100 (16 Go)
- Les performances sont portables sur les différents supports

Preconditioner	Cheb-HTS	Cheb-Trilinos	SSOR-Trilinos
ExecSpace	MPI/OpenMP/SIMD	MPI/OpenMP/Cuda	MPI/OpenMP/Cuda
Nb iter	577/584/531	649/528/581	519/523/600
Exec time	2.80/3.84/3.59	3.74/4.94/2.11	2.71/4.65/2.02

A VENIR

COLLABORATION CEA-IFPEN, OPEN SOURCE, DIFFUSION

- Poursuite avenant 2021-2024 et préparation du prochain 2025-2028
- Dépôt GitHub arcane-benchs
  - Mini-applications pour des tests HPC
- A confirmer : nouvel arrivant pour l'Open Source ShArc
  - Mini-application Arcane/Alien
  - Outillage pour le numérique (Gump, Audi, Laws)
- Rencontres Arcane : journée pour les développeurs d'applications IFPEN et CEA
  - Précédentes éditions : 2014, 2020
  - Prochaine édition prévue : 17 Avril 2023
    - Formule plus ouverte : participation du CEA-DES, de l'ONERA, de l'INRIA et de l'Andra…