RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

Inria

NUMPEX GT: UNSTRUCTURED MESH ROADMAP

Julien Vanharen

GammaO, Centre Inria de Saclay,
Université Paris-Saclay,
91120 Palaiseau, France

Objectives & current status

- Meeting every month since September 2024
- Workshop git repository: https://gitlab.inria.fr/numpex-pc5/wp2-co-design/Working-Groups-Playground
- Math2Product conference in June 2025 (Valencia, Spain)

- Guix packaging
- Ongoing integration of PT-Scotch into the Ouranos library
- Ongoing integration between Ouranos and ShArc

Objectives & current status

- **Generate unstructured meshes containing tens of billions of elements**
- **Provide a set of services to discretize partial differential operators**
- **Reproduce the results obtained on cartesian generated meshes presented in the IFPEN paper (https://comptes-rendus.academie-sciences.fr/mecanique/articles/10.5802/crmeca.336/), but this time in the context of NumPEx Exa-DI with real geoscience unstructured meshes available in the ExaMesh-Gallery**
- **Realistic meshes with around 800 million cells obtained by refining coarse meshes available in the ExaMesh-Gallery**
- **Technological bottlenecks:**
  - Parallel reading,
  - Scalable re/partitioning under constraints.

Roadmap

- **Guix packaging:**
  - *guix pull -C channels.scm*
  - *guix install sharc*
  - Integration to the guix-science/guix-hpc channels
- **PT-scotch integration**
- **KalpaTARU integration**
- **Zoltan integration**
- **Evaluation of EZTrace/Pallas within Ouranos**, in collaboration with Exa-SofT
- **Integration of the Ouranos reader within ShArc:**
  - Mapping Ouranos data structures to ShArc's data structures

Roadmap

- **Validation of parallel reading, partitioning and refining algorithms on extremely large mesh sizes (> 800 million cells):**
  - Develop tests to validate mesh reading and partitioning,
  - Very few efficient parallel mesh readers exist at this scale,
  - Few partitioning tools capable of scaling up to 32,000 cores.
- **Benchmarking of parallel tools for reading, refinement, and partitioning:**
  - Expected types of partitions:
    - Up to 32,000 CPU cores with 10,000 to 50,000 cells per MPI process,
    - Up to 8,000 GPUs with 50,000 to 100,000 cells per GPU.
  - Compare read times with existing parallel readers such as binary MSH readers,
  - Compare performance with existing tools such as ParMETIS, PT-Scotch, KalpaTARU, and Zoltan,
  - Establish a reproducible benchmarking procedure on GENCI machines (using Guix and tools for processing benchmarking results),
  - Write the benchmark report/publication.

Perspectives

- **Customized task decomposition with a dependency graph to enable the use of the StarPU runtime:**
  - Renumbering,
  - Coloring,
  - Sub-partitioning,
  - Leveraging the Kokkos backend (KOKTAILS project).
- **Topological operators:**
  - Iterating over elements and retrieving vertex data,
  - Iterating over vertices and identifying all elements sharing a given vertex,
  - Iterating over faces or edges and retrieving neighboring elements,
  - Performing reductions on data associated with vertices or elements.
- **Real-world simulation on Alice Recoque to solve a true geoscience problem.**

Thank you for your attention!