Exascale Oriented GPGPU Algorithm for Efficient Sparse Matrix Assembly in FEM

POSTER # 181

SCA/HPC Asia 2026
Jan 26-29, 2026, Osaka, Japan

Mohd Afeef Badri[1], ✉ , Gilles Grospellier[2], Evelyne Foerster[3]
✉ mohd-afeef.badri@cea.fr

[1]Université Paris-Saclay, CEA, SGLS, 91191, Gif-sur-Yvette, France

[2]CEA, DAM, Division Île-de-France, 91297, Arpajon, France

[3]Université Paris-Saclay, CEA, DM2S, 91191, Gif-sur-Yvette, France

Motivation & Context

The Exascale Challenge

- France's first exascale supercomputer **Alice Recoque** in 2027 at CEA
- Over **80%** of computational power derived from **GPU acceleration**
- **Paradigm shift**: from CPU-centric petascale to GPU-dominated exascale computing

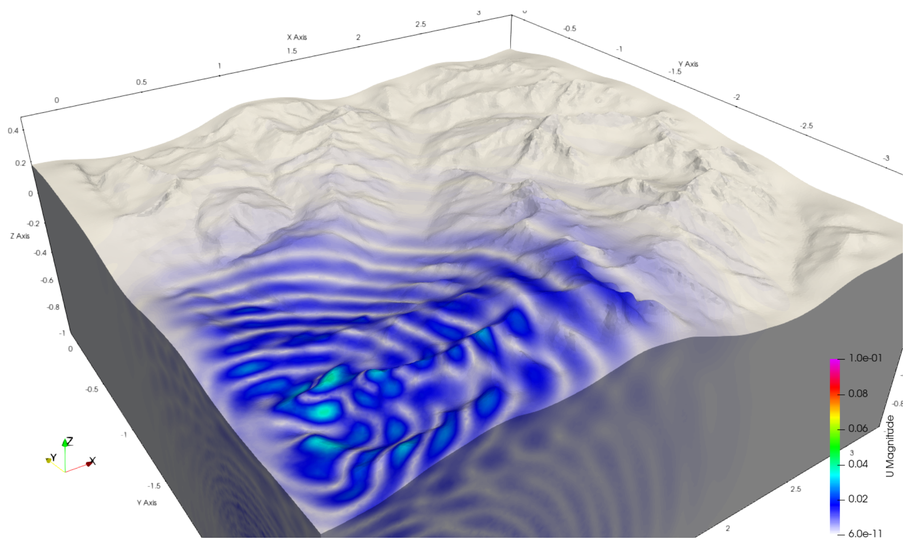

The Scientific Challenge:

- Complex seismic wave propagation requires high-fidelity modeling
- Finite Element Method (FEM) with nonlinear, hyperbolic PDEs
- Need to **rethink FEM algorithms** optimized for GPUs

ArcaneFEM: GPU-Accelerated FEM Framework

Core Features:

- Built on **Arcane**: C++ based HPC scientific computing framework
- **Fully multi-GPU/multi-CPU** solver with minimized CPU-GPU data transfers
- **GPU-architecture agnostic**: NVIDIA (V100, A100, H100) & AMD (MI250X, MI300)
- State-of-the-art **DOF-wise assembly algorithms** (atomic-free, deterministic)
- Performance comparable to leading GPU FEM solvers (e.g., MFEM)

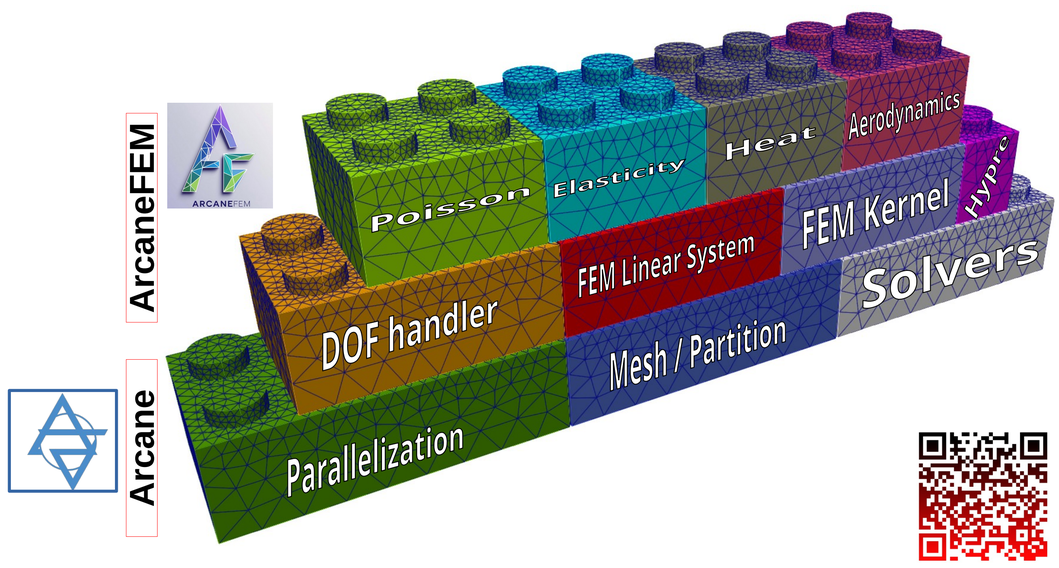

Domain-Decomposition for Multi GPUs

Mono GPU

- Single mesh
- 1 GPU sequential
- No MPI needed

Multi-GPU

- $N_{\text{GPU}} = N_{\text{CPU}}$
- 4 mesh partitions
- 4 MPI procs. for 4 GPU

Multi-GPU

- $N_{\text{GPU}} < N_{\text{CPU}}$
- 8 mesh partitions
- 8 MPI procs. for 4 GPU

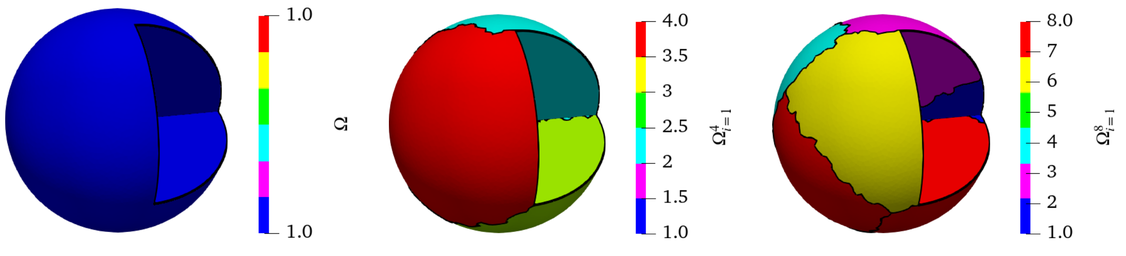

Tailored GPGPU Algorithms for FEM

Novel DOF-Wise Assembly Algorithm:

Domain change principle: Reformulate FEM assembly to match GPU architecture

$$\int_{\Omega_h} \nabla u \cdot \nabla v \rightarrow \int_{P_i \in \Omega_h} (\cdot) + \int_{E_i \in \Omega_h} (\cdot) \rightarrow \mathbf{A}$$

- **Traverse mesh by DOF's** instead of elements
- **Eliminates race conditions** without atomic operations
- Fully reproducible and numerically consistent: **deterministic**

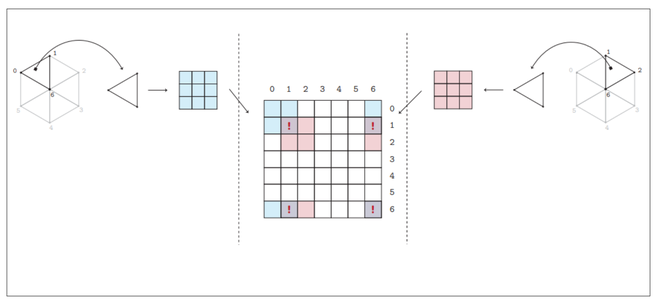

Classical Approach

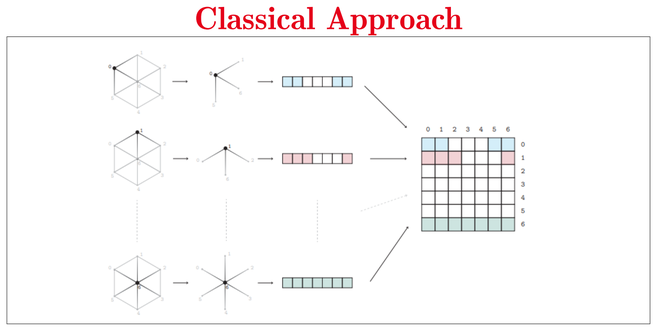

Novel DOF-Wise Approach

Applicable on Sparse Matrix Formats:

- 20-30× speedup vs CPU for matrix assembly

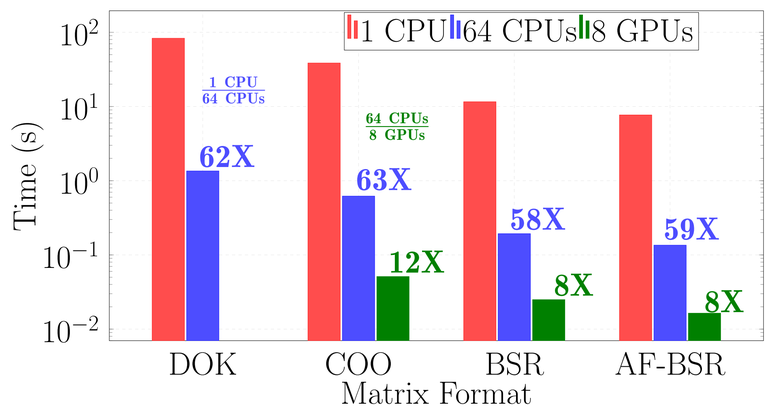

Assembly time comparison DOK/COO/BSR/AF-BSR

Performance & Scalability

Strong Scaling Results:

- Near-optimal scaling up to **800 GPUs** (AMD MI250X)
- Problem size: 300 million DOFs for Poisson equation
- Successfully validated on meshes exceeding **1 billion elements**
- Tested on: Adastra (France), Topaze (CEA), multiple architectures

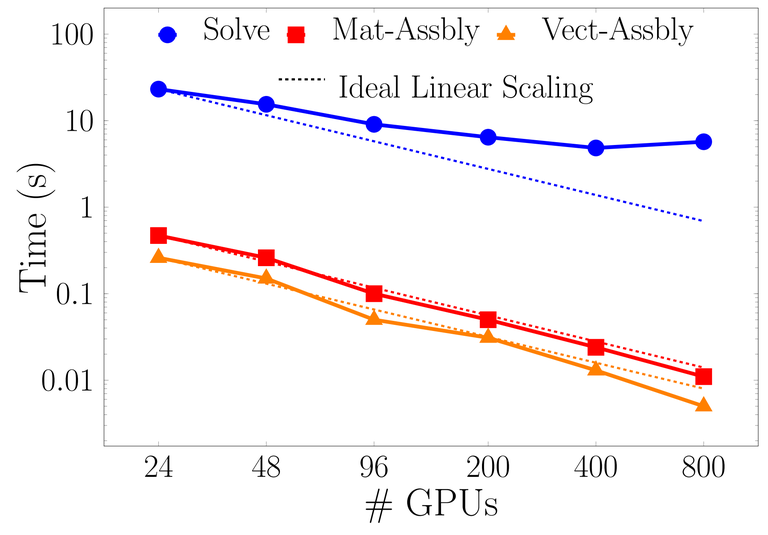

Strong scaling plot for 300 Million Poisson Problem - 24 to 800 GPUs