GPU-ACCELERATED COMPUTING FOR COMPUTATIONAL PLASTICITY

ALI R. HASHEMI*, NARGES DIALAMI†

Centre Internacional de Mètodes Numèrics en Enginyeria (CIMNE), C/Gran Capità s/n,
Campus Nord UPC, 08034, Barcelona, Spain
Escola Tècnica Superior d'Enginyers de Camins, Canals i Ports, Universitat Politècnica de Catalunya – BarcelonaTech (UPC), C/Jordi Girona 1,
Campus Nord UPC, 08034, Barcelona, Spain

* alireza.hashemi@upc.edu † narges.dialami@upc.edu

ABSTRACT

Simulating elastoplastic, crystal-plasticity, and other strongly nonlinear behaviour increasingly involves large-scale, computationally demanding problems — finite-strain kinematics, history-dependent constitutive updates, contact, and adaptive discretization — arising in both finite-element (FE) and mesh-free methods such as Smoothed Particle Hydrodynamics (SPH), and challenging conventional CPU implementations. Graphics Processing Units (GPUs), built on the Single Instruction, Multiple Threads (SIMT) paradigm, offer substantial acceleration, yet raise open questions on parallelizing constitutive integration, point-location search, and nonlinear solvers, as well as on exploiting automatic differentiation for inverse characterization and digital-twin applications.

Recent advances span multi-GPU material point methods for extreme deformation, differentiable GPU-accelerated finite-element and crystal-plasticity solvers achieving substantial speedups, and GPU-resident mesh and point-location algorithms for particle tracking and contact detection. This rapidly maturing landscape — spanning computational mechanics, high-performance computing, and machine learning — underscores GPU computing's promise not only for faster forward simulations but also for optimization, uncertainty quantification, and digital-twin workflows increasingly reliant on GPU-accelerated reduced-order and surrogate models for real-time simulation.

This Invited Session gathers researchers developing GPU-accelerated methods for computational plasticity to share strategies, benchmarks, and open challenges, such as: matrix-free/assembly-free GPU FE solvers, GPU-parallel constitutive integration, GPU-accelerated SPH and particle methods for large-deformation and impact plasticity, differentiable and AI-augmented frameworks, GPU-accelerated contact and mesh-search algorithms, reduced-order and surrogate modeling for real-time simulation, and next-generation HPC strategies, while comparing scalability across frameworks and identifying transferable algorithmic strategies for memory management, data locality, and heterogeneous CPU-GPU architectures. The session aims to clarify best practices, identify common bottlenecks, and foster collaborations advancing robust computational methods for nonlinear solid mechanics.

REFERENCES

[1] Hashemi A.R., Dialami N., Chiumenti M., GPU-Resident Point Location on Adaptive Tetrahedral Meshes via a Mesh-Aligned Multi-Level Octree, Comput. Methods Appl. Mech. Eng., under review.

[2] Hu F. et al., Efficient GPU-computing simulation platform JAX-CPFEM, npj Computational Materials, Vol. 11, 2025.