Proceedings of the Canadian Society for Mechanical Engineering International Congress
32nd Annual Conference of the Computational Fluid Dynamics Society of Canada
Canadian Society of Rheology Symposium
CSME-CFDSC-CSR 2025
May 25–28, 2025, Montréal, Québec, Canada

A GPU-Based High-Order Discontinuous-Galerkin-Hancock Method with Adaptive Mesh Refinement

Osman El-Ghotmi and James McDonald
Department of Mechanical Engineering, University of Ottawa, Ottawa, Canada

ABSTRACT

The discontinuous Galerkin Hancock (DGH) scheme, originally developed by Suzuki and Van Leer, is a high-order numerical technique used to efficiently solve systems of hyperbolic-relaxation partial differential equations (PDEs) with stiff local source terms. These types of equations are commonly used to model various scientific problems, such as non-equilibrium gas flows, multiphase flows, reactive flows, radiation transport prediction, and ionized plasma flows. In fully explicit schemes, stiff local processes, such as drag and ionization, can severely restrict timestep sizes to maintain stable and accurate solutions. The DGH scheme addresses this issue by treating the local source term implicitly, while treating global transport terms explicitly. Using only linear basis functions, the DGH scheme achieves third-order accuracy in both space and time. In order to maintain stability, slope limiting techniques are used to discard higher-order information in specific regions of high gradients and discontinuities. Slope limiting techniques are effective but highly expensive, because at every timestep they require each cell to check all neighbouring cell values, which degrades the parallel efficiency of the scheme. This presentation demonstrates how we extended our existing GPU-based implementation of the DGH method to work with adaptive mesh refinement (AMR) and a load-balancing algorithm. Given a certain sensitivity criteria, AMR can dynamically refine or coarsen a mesh to adapt the accuracy of the solution in specific regions of the mesh. This efficient approach provides a method to focus computational resources on areas of the mesh that require greater precision. The load-balancing algorithm ensures, after each refinement step, that the global problem is evenly distributed across all available GPUs to ensure an efficient use of all available computational resources. We also explore the use of entropy-stability in DGH, hoping to replace the need for computationally costly slope limiting techniques and further optimizing the numerics of the scheme for parallel platforms. The presentation demonstrates how we incorporate a solution-dictated adaptive mesh refinement scheme into our GPU-based implementation of DGH and how load-balancing was used to efficiently distribute the global problem across multiple GPUs on large clusters of compute nodes. Numerous test cases for the solution of hyperbolic relaxation PDEs resulting from kinetic equations are also presented, illustrating massive computational speedups with much greater accuracy.