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

LEVERAGING CONSUMER-GRADE GPUS FOR EFFICIENT HIGH-ORDER COMPUTATIONAL FLUID DYNAMICS

Brian C. Vermeire[1], Muhammad Saghir[1]
[1]Department of Mechanical, Industrial, and Aerospace Engineering, Concordia University, Montreal, Canada

ABSTRACT

Modern many-core hardware architectures, specifically Graphical Processing Units (GPUs), have advanced rapidly over the past decade. More recently, a clear distinction has emerged between consumer/workstation and data center GPUs. The former are characterized by low double-precision compute, low memory availability, and low cost. In contrast, the latter are characterized by high double-precision compute, high memory availability, and an order of magnitude higher cost. However, despite these differences, both consumer/workstation and data center GPUs have comparable single-precision compute and power consumption characteristics.

In this talk, we will explore implementation details for a high-order unstructured solver (HORUS) that allow it to exploit the high single-precision compute of consumer/workstation GPUs while minimizing memory footprint using strategies including array aliasing, batched computing, and time-stepping optimizations. Results from scale-resolving simulations of turbulent flows will demonstrate the ability to run simulations with up to 250 million unknowns on a single GPU, comparable accuracy to double-precision data center GPUs for turbulent flows, and significant reduction in hardware cost. Ultimately, these results demonstrate that scale-resolving simulations of turbulent flows can be performed by leveraging only the single-precision compute capability of consumer/workstation grade GPUs, at an approximately 20-40 times reduction in hardware cost.