More information at:
www.ceec-coe.eu

CEEC

Centre of Excellence in Exascale CFD: Finding Solutions to Grand Problems at the Frontier of CFD

Abstract

CEEC is dedicated to advancing cutting-edge computational fluid dynamics (CFD) algorithms and models to the exascale level in order to address the specific needs of real-world applications crucial for both scientific research and industrial use. To achieve this, CEEC focuses on six highly challenging CFD problems, referred to as light-house cases (LHCs), which are large and detailed enough to require the capabilities of next-generation high-performance computing (HPC) systems up to exascale performance. HPC has finally entered the exascale era with GPU-enabled heterogeneous cluster systems being first to breach the one exaFLOP barrier. While these systems offer massive computing resources, their stream-based architecture renders branching and data transfer to the GPU host comparatively expensive. As such, the design and optimization of CFD solvers must inherently be adapted to these changes in hardware layout to benefit from their potential performance gains and power savings. CEEC has selected five individual demonstrator frameworks to adapt for exascale readiness: FLEXI/GALÆXI, Alya and SOD2D, waLBerla, Neko, and NekRS/Nek5000. Our interim progress is described below and detailed on our website.

Lighthouse Case Codes

FLEXI/GALÆXI

FLEXI/GALÆXI is a high-order accurate HPC framework based on the discontinuous Galerkin (DG) spectral element method for hyperbolic-parabolic conservation laws with special focus on compressible flows (at higher Mack numbers). FLEXI has demonstrated excellent scaling on up to five hundred thousand CPUs, and GALÆXI scales well up to thousands of GPUs utilizing MPI-aware CUDA and HIP. The strong scaling results presented below were conducted on various EuroHPC machines, including LUMI-C (AMD EPYC 7763), MareNostrum5 GPP (Intel Sapphire Rapids 8480+), JEDI (GH200), Leonardo Booster (A100), as well as HAWK (AMD EPYC 7742) and Juwels Booster (A100).

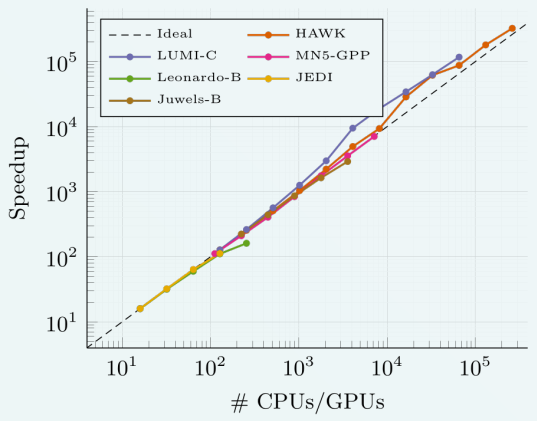

Alya/SOD2D

Alya/SOD2D is an HPC simulation framework designed for multiphysics problems. It explicitly addresses structural mechanics and compressible fluid flows, with a special focus on fluid-structure interaction (FSI). The framework integrates two finite element codes which are optimized to harness the capabilities of modern HPC infrastructures, supporting both CPU-only and GPU-accelerated computing nodes. Alya has shown an excellent speed-up on MareNostrum5-GPP, and SOD2D demonstrates outstanding performance on various GPU-accelerated EuroHPC machines, including MareNostrum5-ACC (H100), Leonardo Booster (A100), and Karolina (A100) using CUDA-aware MPI and OpenACC.

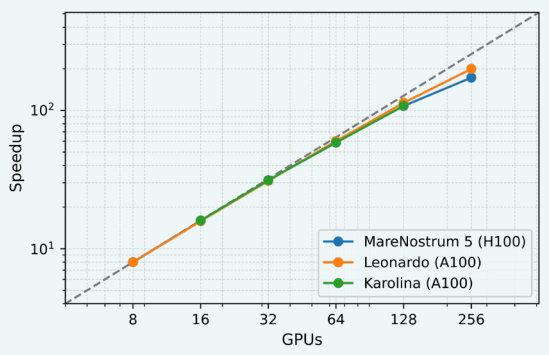

waLBerla

waLBerla is a modern highly scalable open source multi-physics simulation software framework with a focus on CFD applications such as three-dimensional fully-resolved coupled fluid-particle simulations. The framework employs the lattice Boltzmann method for the fluid phase, and the discrete element method for the granular soil. waLBerla has shown perfect scalability on conventional CPU architectures, and on up to thousands of GPUs utilizing MPI-aware CUDA and HIP. The strong scaling results presented below were obtained on a variety of EuroHPC machines, including LUMI-C (AMD EPYC 7763), LUMI-G (AMD MI250X), MareNostrum5-GPP (Intel Sapphire Rapids 8480+), and MareNostrum5-ACC (H100).

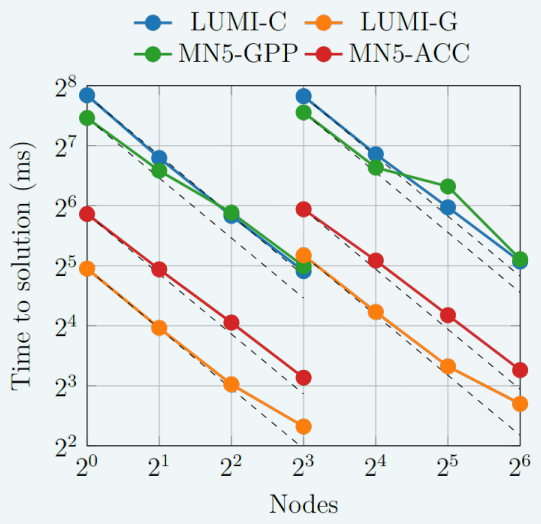

Neko

Neko is a portable framework for high-order spectral element based simulations, focusing on the incompressible regime. The framework adopts an object-oriented approach, allowing for multi-tier abstractions of the solver stack and facilitating various hardware backends, ranging from general-purpose processors, accelerators, and vector processors, to limited support of FPGA. Neko demonstrates excellent scaling on the EuroHPC machines LUMI-G (AMD MI250X) and Leonardo Booster (A100).

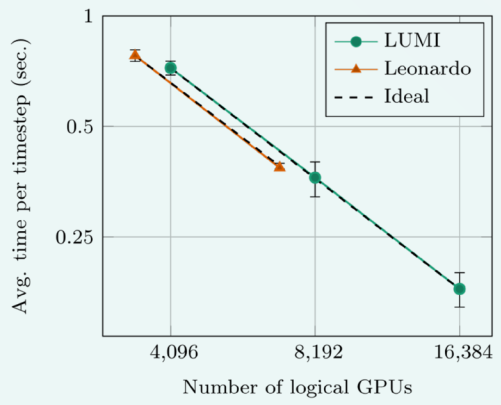

NekRS/Nek5000

Nek5000/NekRS is a high-order spectral element framework for thermal-fluids simulations, solving the incompressible Navier-Stokes and thermal transport equations. The GPU-accelerated version NekRS is designed for high-performance on exascale platforms, leveraging performance-portable frameworks like OCCA and optimized kernels from libParanumal. NekRS demonstrates efficient scaling on Summit (V100) and JUWELS Booster (A100) up to thousands of GPUs.

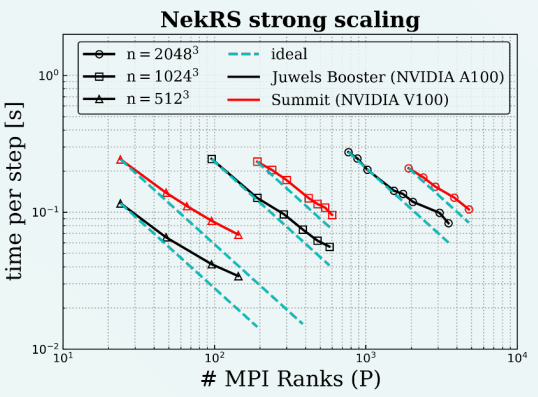

Code repositories

Interested in learning more about the codes and frameworks we're working with to run our lighthouse cases?

Visit our repos:
https://ceec-coe.eu/our-codes

This work has been supported by the CEEC project, which has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101093393. The JU receives support from the European Union's Horizon Europe research and innovation programme and Sweden, Germany, Spain, Greece, and Denmark.

EuroHPC Joint Undertaking

Funded by the European Union