SCALABLE

SPONSORED BY THE Federal Ministry of Education and Research

MINISTRY OF EDUCATION, YOUTH AND SPORTS

EuroHPC Joint Undertaking

scalable-hpc.eu @scalable_hpc

In SCALABLE, eminent industrials and academic partners team up to achieve the leap to unprecedented performance, scalability, and energy efficiency of an industrial LBM-based computational fluid dynamics (CFD) software. Lattice Boltzmann methods (LBM) have already evolved to become trustworthy alternatives to conventional CFD. In several engineering applications they are shown to be roughly an order of magnitude faster than Navier-Stokes approaches in a fair comparison and in comparable scenarios. In the context of EuroHPC, LBM is especially well suited to exploit advanced supercomputer architectures through vectorization, accelerators, and massive parallelization.

LaBS

- Industrial CFD software with capabilities at a proven high level of maturity, but potential for performance improvement

SCALABLE aims to:

- Transfer leading edge performance technology from waLBerla to LaBS
- Break silos between scientific computing & physical flow modelling worlds
- Deliver improved efficiency & scalability for LaBS in view of upcoming European Exascale systems

scalable implementation, performance,

complex boundary conditions, industrial complexity…

WaLBerla

- Public domain research code offering great performance & unlimited scalability (1+ trillion (10^{12}) lattice cells already on petascale systems)
- Uncompromising unique, architecture-specific automatic generation of optimized compute kernels
- Carefully designed parallel data structures

- However, not compliant with industrial applications due to lack of a geometry engine and user-friendliness for non-HPC experts

Performance and Energy Efficiency optimization

Performance optimization and scalability optimization is essential goal of the project. The evaluation is done using POP2 CoE methodology.
Methodology developed by the H2020 READEX project, is used to reduce energy consumption by exploiting dynamic behavior of the different phases of our applications. Below you can see a dynamic behavior of the waLBerla running Turbulent Channel testcase. In this particular case, one can save 17% of energy with minimal impact on code performance (only 2% runtime slowdown).

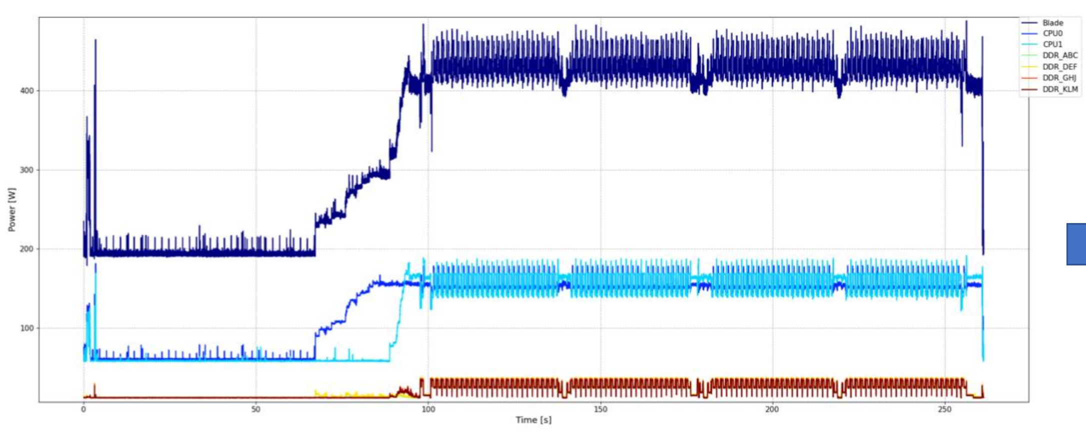

Figure 38: WaLBerla Turbulent Channel: Power consumption of the blade, CPUs and memory channels.

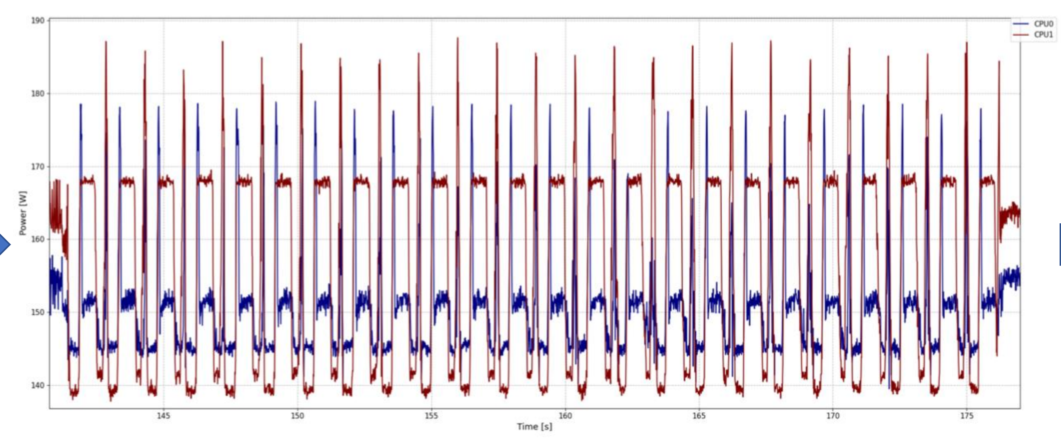

Figure 39: WaLBerla Turbulent Channel: CPUs power consumption during the iterative solver execution.

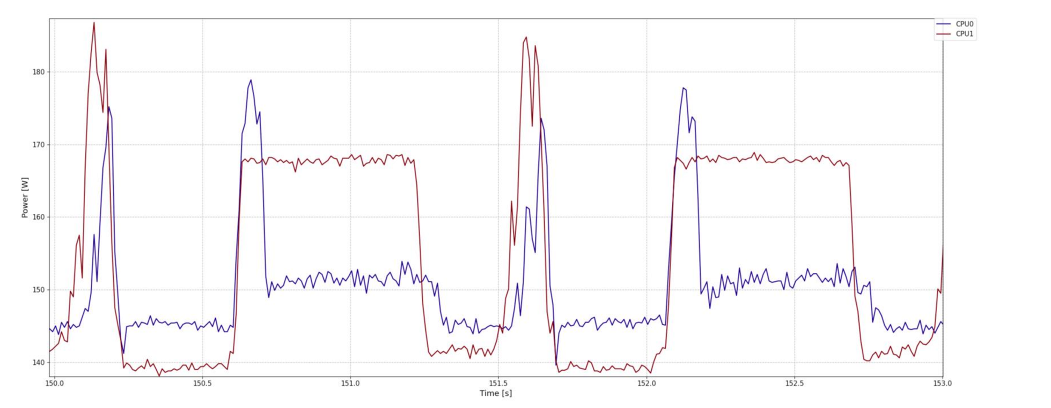

Figure 40: WaLBerla Turbulent Channel: Power consumption of the CPUs, detail at one iteration of the iterative solver.

Test cases

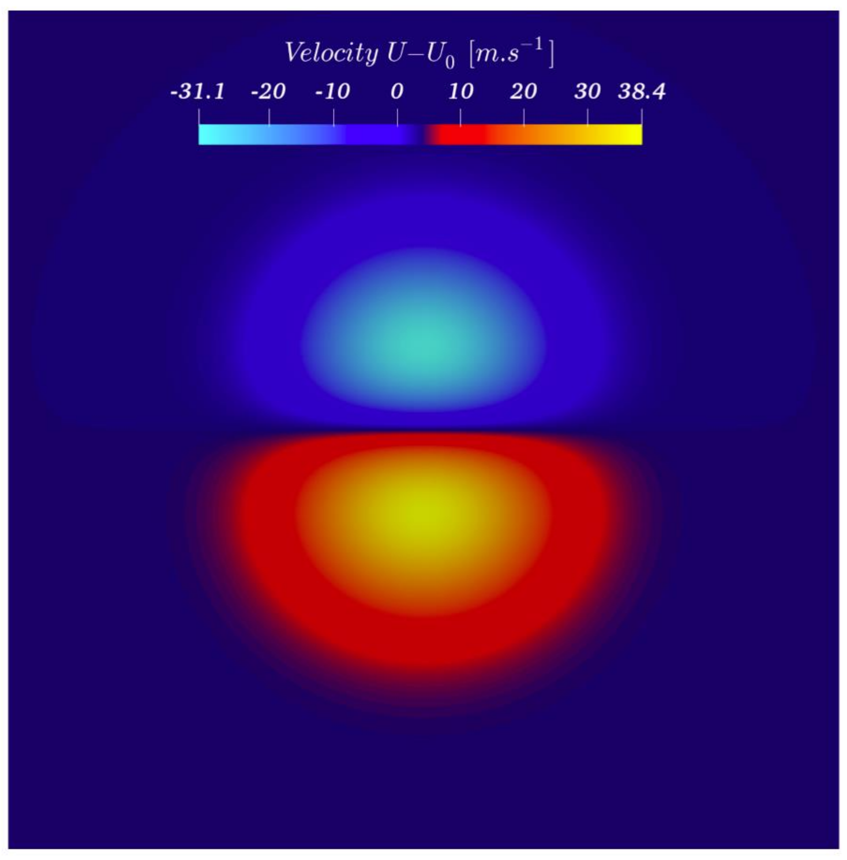

#1 CERFACS CO-VO, in which an isentropic vortex is simply moving to the east of a periodic square grid with time. The flow is supposed inviscid Euler Equations.

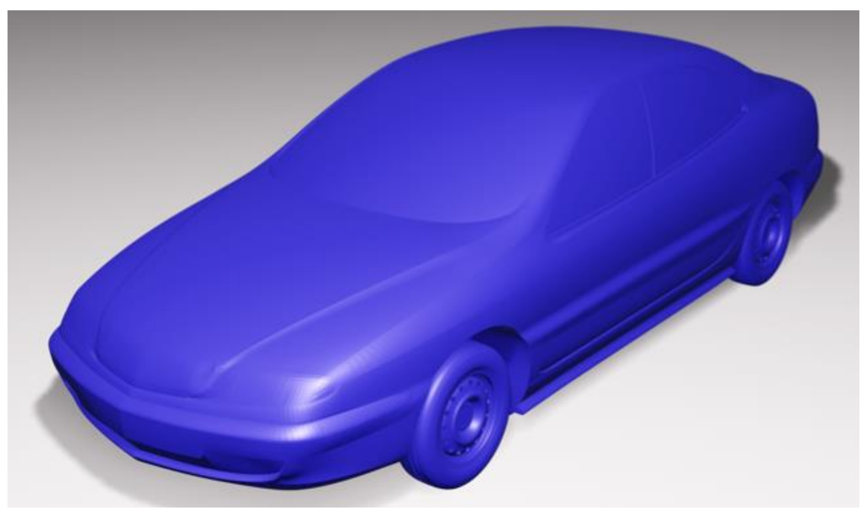

#2 S2A flow around a modern sedan car placed inside a wind tunnel.

#3 Turbulence modeling must be carefully assessed, starting with cases where the analytical theory allows to be quantitative. The first test case is a turbulent channel, the turbulent flow between two parallel plates.

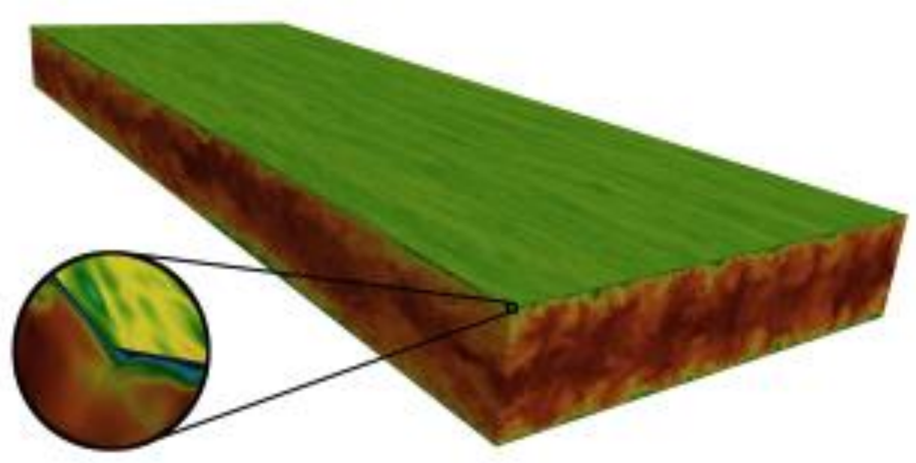

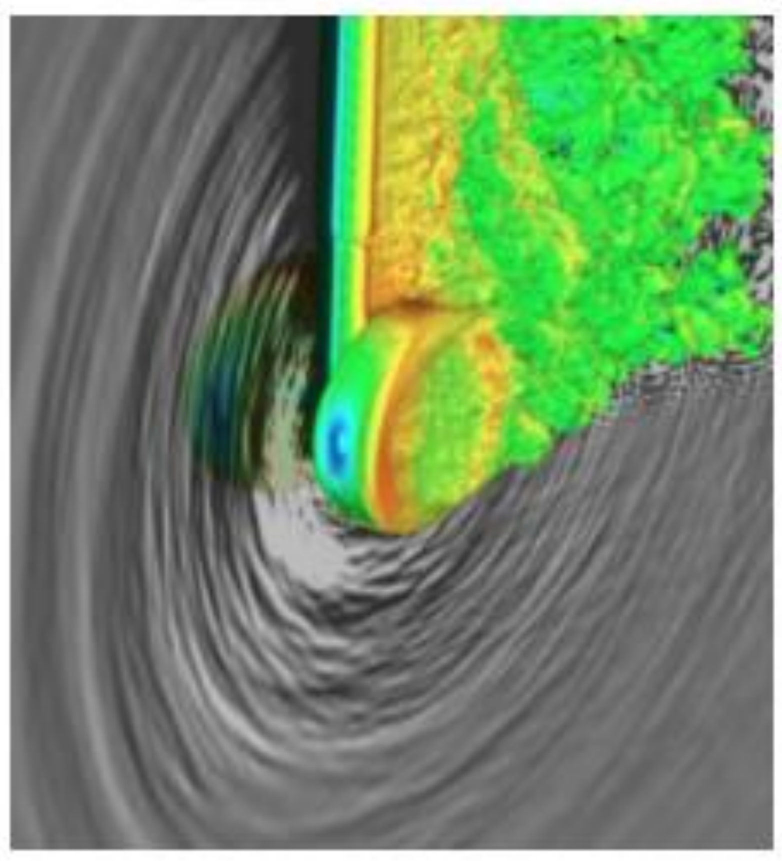

#4 Lagoon programme, assessing existing techniques for landing gear noise simulation.

Consortium

CS GROUP · AIRBUS · CERFACS · FAU FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG

GROUPE RENAULT · JÜLICH Forschungszentrum · VSB TECHNICAL UNIVERSITY OF OSTRAVA IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER · NEOVIA INNOVATION

The SCALABLE project has received funding from the European High-Performance Computing Joint Undertaking Joint Undertaking (JU) under grant agreement No 956000. The JU receives support from the European Union's Horizon 2020 research and innovation program and France, Germany, Czech Republic. This poster reflects only the author's view, the JU is not responsible for any use that may be made of the information it contains.