EuroHPC Summit
2023 GÖTEBORG
European Supercomputing Excellence in the Exascale Era

SCALABLE

In SCALABLE, eminent industrials and academic partners will 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

The industrial CFD software LaBS already has such industrial capabilities at a proven high level of maturity, but it still has performance worthy of improvement. Therefore, SCALABLE will transfer the leading-edge performance technology from the waLBerla to LaBS, thus breaking the silos between the scientific computing world and physical flow modelling world to deliver improved efficiency and scalability for LaBS to be prepared for the upcoming European Exascale systems.

scalable implementation, performance,

complex boundary conditions, industrial complexity, …

waLBerla

In the public domain research code waLBerla, superb performance and unlimited scalability has been demonstrated, reaching more than a trillion (10^12) lattice cells already on Petascale systems.

waLBerla performance excels because of its uncompromising unique, architecture-specific automatic generation of optimized compute kernels, together with carefully designed parallel data structures. waLBerla, however, is not compliant with industrial applications due to of a geometry engine and user friendliness for non-HPC experts.

Test cases

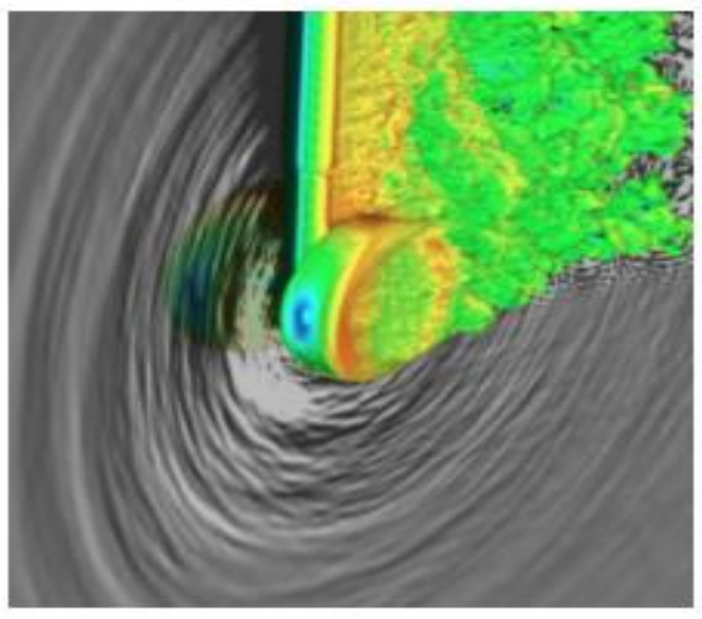

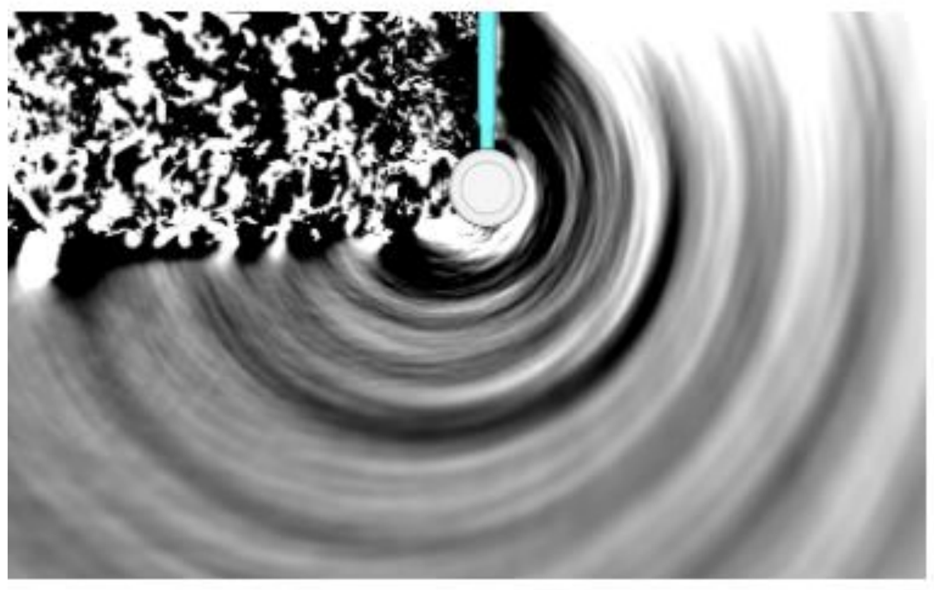

The Lagoon case is a simplified landing gear. During the landing phase, landing gear is an important source of noise and produce a non-trivial turbulent flow. It is convenient to bench CFD codes since the geometry is simple and a large experiment database is available for both aerodynamic and acoustic.

Scan for more details:

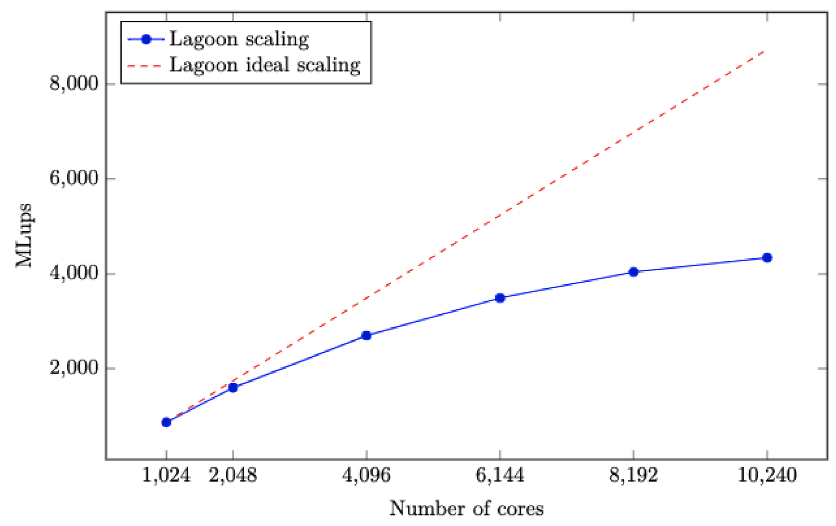

Strong scaling for LaBS on the Lagoon case using the KAROLINA CPU supercomputer up to 10K using a 276 999 093 fluid nodes. At 30 000 nodes per core performance is acceptable at 10 240 cores

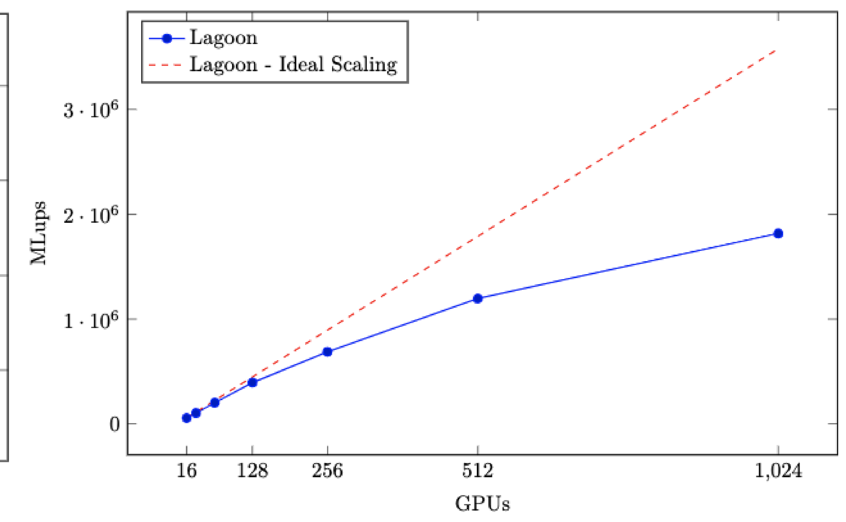

Strong scaling for walBerla on the Lagoon case using the GPU system JUWELS up to 1024 GPUs. These test were performed with 2 097 152 per A100 GPU. This larger mesh allows for better GPU occupancy and scaling.

Performance and energy efficiency optimization

Optimization of energy consumption of the simulation is essential goal of the project. Methodology developed by the H2020 READEX project, is used for non-accelerated executions to reduce energy consumption by exploiting dynamic behavior of the different phases of our applications. Bellow you can see a dynamic behavior of the LaBS running the Lagoon testcase.

Using a MERIC runtime system to scale CPU core and uncore frequencies (Intel Xeon Cascade Lake 6240) to fit the executed phase of the application let to major energy savings (overall blade) with negligible performance penalty.

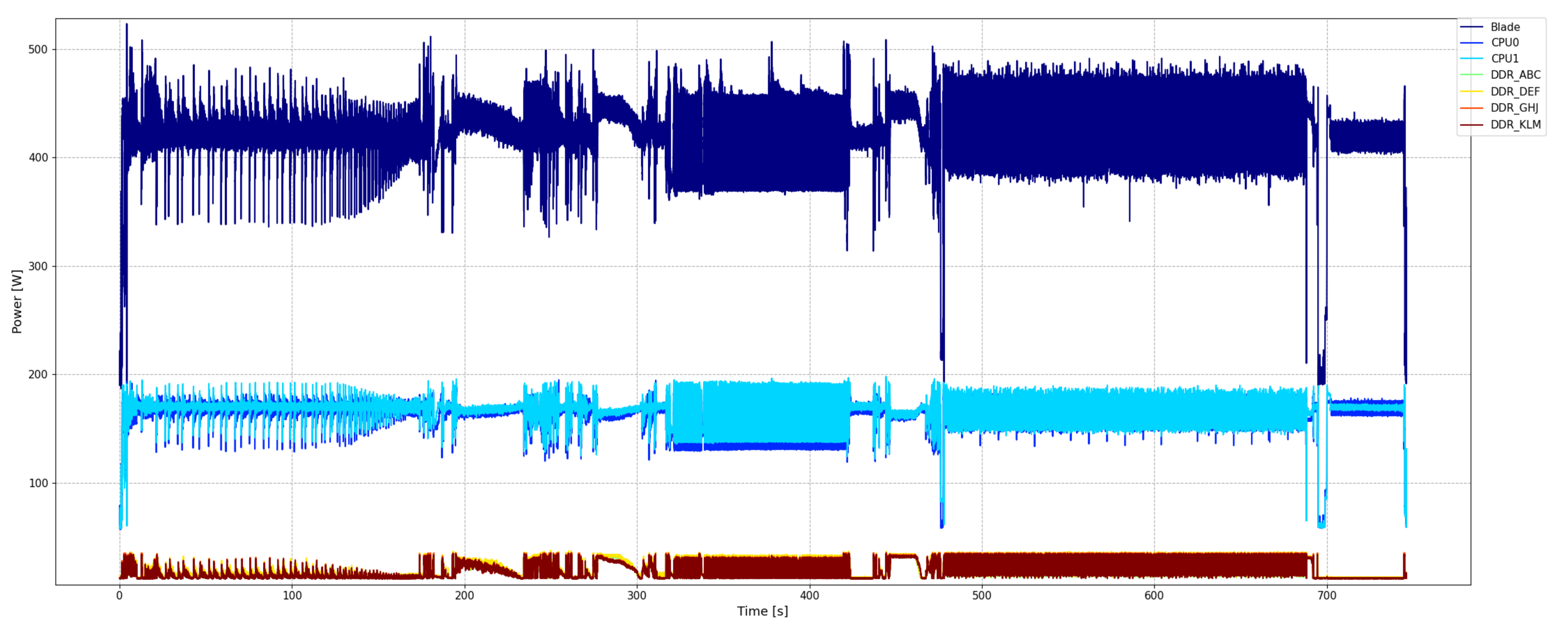

	Default	Static tuning	Dynamic tuning constant runtime	Dynamic tuning
Runtime [s]	1797.9	1942.73	1807.13	1871.14
Energy consumption [kJ]	3102.3	1942.73	2726.7	2496.71
Solver energy-efficiency [MLups/W]	0.054	0.059	0.056	0.056
Runtime extension [%]	-	8.1	0.5	4.1
Energy savings [%]	-	15.1	12.1	19.5

waLBerla has been analyzed in both non- and accelerated configuration. The GPU accelerated run of the Lagoon use case was optimized by specifying a static GPU streaming multiprocessor frequency.

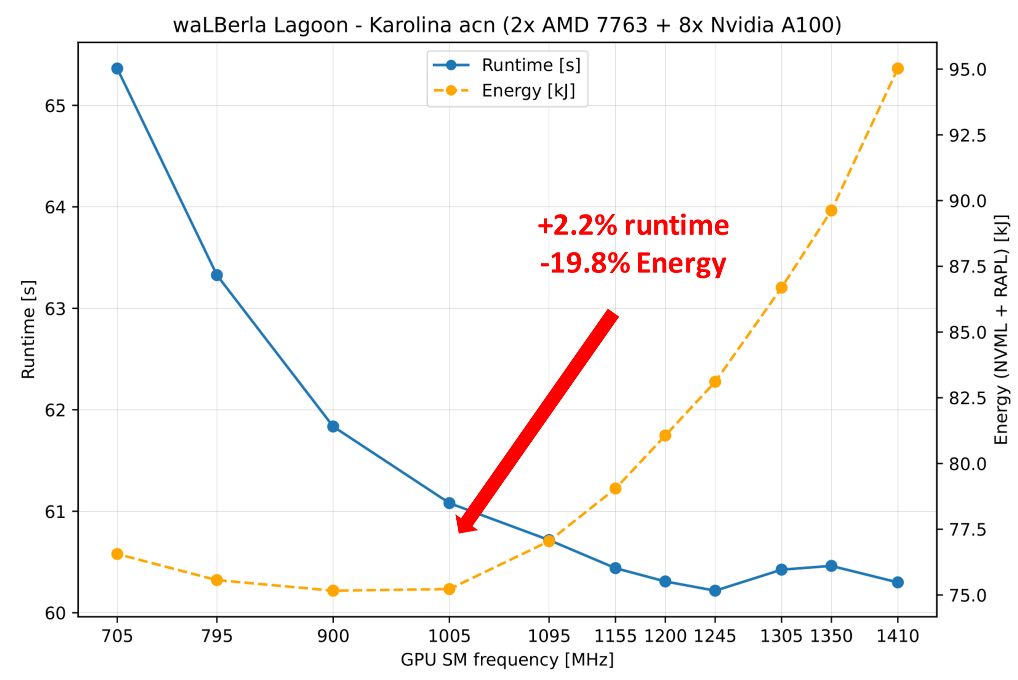

Scaling of the CPU core frequency does not bring any additional energy savings, because the CPUs are used to control the GPUs only. Most of the time the CPU is idle, and its frequency does not have impact on power consumption.

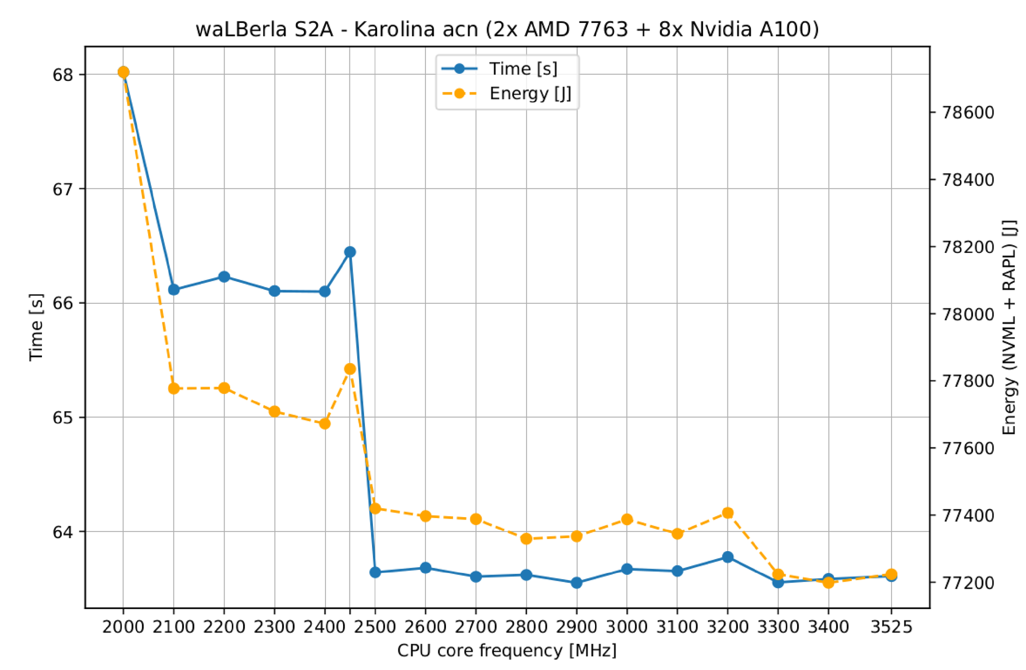

EuroHPC Joint Undertaking
The SCALABLE project has received funding from the European High-Performance Computing Joint Undertaking Joint Undertaking (JU) under grant agreement No 956000. The JU receives support from the European Union's Horizon 2020 research and innovation program and France, Germany, Czech Republic.