9th European Thermal Sciences Conference
EUROTHERM 2024
10-–13 June 2024, Bled, Slovenia

ENSEMBLE AVERAGING PARALLEL-IN-TIME APPROACH FOR INDUSTRIAL LES

Josep Plana-Riu [1], F.Xavier Trias [1], Àdel Alsalti-Baldellou[1,2], Guillem Colomer [1], Asensio Oliva [1]

[1] Heat and Mass Transfer Technological Centre, Technical University of Catalonia, ESEIAAT, Carrer de Colom 11, 08222 Terrassa (Barcelona), Spain
[2] Termo Fluids S.L., Carrer de Magí Colet 8, 08204 Sabadell (Barcelona), Spain

Real-world computational fluid dynamics (CFD) simulations demand not only robust and stable numerical methods, but also to be completed in a limited time. For this reason, industrial applications usually rely on RANS modeling, which have some remarkable limitations. Alternatively, using Large Eddy Simulations (LES) leads to higher fidelity simulations, while having much higher computational costs.

CFD simulations are virtually always carried out with spatial parallelization, which, if exploited to a bigger and bigger number of cores will eventually lead to a saturated system and thus, using more cores will not provide a further speed-up. In order to deal with this, some concepts such as the Multigrid Reduction In Time (MGRIT) have been proposed as a way to extend the parallelization to both space and time domains.

On the other hand, most CFD simulations are based on sparse matrix-vector products (SpMV) or SpMV-like stencil-based kernels, which have a low arithmetic intensity, i.e. the ratio of floating-point operators is low compared to the data traffic. This sets the bottleneck of CFD simulations in data transfer instead of in computational power. In order to deal with this, ensemble averaging parallel-in-time approaches consist on running a set of shorter simulations. Afterwards, the average result among this set of simulations is computed instead of running a single long simulation. In this work, this set of simulations is proposed to be run at once on the same device. This allows an increase of the arithmetic intensity, as the operation requires less data transfer, and it eventually generates speed-up.

This will be tested in canonical flow and heat transfer cases targetting concentrated solar power (CSP) simulations, on both CPU and GPU-accelerated nodes, in order to prove that this method can lead to improvements in the efficiency of simulations and thus moving towards reliable industrial overnight LES simulations.