C3.3: Transitional flow over a SD7003 wing

Cenaero

1 Introduction

For this computation we target DNS; in order to provide sufficient resolution, we refine until continuity of the vorticity field has been obtained. We have done a similar computations on the eppler E387 airfoil, and are doing this kind of study for LP turbine blades ($Re \sim 80.000$).

Due to circumstances, we have not been able to do the computations yet, since we did not have access to the necessary computational time (competition with consultancy activities from our colleagues and other DNS computations). However we should be able to do the computations before the deadline, or the end of november at the latest.

About the time we were made aware of the workshop, we had started some preliminary computations on the SD7003 airfoil since it had received quite some interest from the high-order community. The specifications for the domain were based on our previous computations of the Eppler E387 airfoil at $Re = 60000$. We did not follow up on these computations since we needed to redefine the setup to comply with the specifications of the workshop. What follows is the description of these preliminary computations.

2 Code description

Discretization. Argo is based on the discontinuous Galerkin / interior penalty method. It has been implemented for hybrid two and threedimensional curved grids, featuring triangles, quadrangles, tetrahedra, prisms and hexahedra. The maximum order of interpolation, both in terms of geometry and solution, is 4.

Iterative methods and time integration. On meshes featuring large differences in stable time steps (eg. due to boundary layers), we use second order backward differencing time-integration, with Jacobi preconditioned Newton-GMRES iterations as the implicit solver. For isotropic meshes, explicit Runge-Kutta iterations are used.

Parallellization. Argo features a hybrid shared / distributed memory parallellisation. The shared memory parallellisation uses an internal domain partitioning per element type and operation. Distributed memory parallellisation is based on ghost elements. Weak and strong scaling tests have been done up to 512 cores on an intel based cluster, and up to 32.000 cores on a BlueGene/P machine.

Postprocessing. All types of monitors (forces, fluxes, volume integrals) are embedded in the code. For these integration rules of order 2p+1 are used (where p is the order of interpolation), for any quantity. The high-order polynomial output of solution, state functions and derivative quantities, is done in Gmsh 2.0 format; the views are then recursively refined up to 4 times with Gmsh.

3 Case summary

Machine. Each of the nodes of the cluster is composed of 2 quad-core processors (Intel Xeon L5420). The cores are clocked at 2.5GHz, and have 6MB of level 2 cache. The nodes are organised following the *SMP (symmetric multiprocessing)* paradigm, meaning that all of the cores access the shared 16GB RAM memory through a common *(FSB) front side bus*. Hence the performance is much impacted by the memory bandwidth, in function of the number of active jobs, and in particular on their memory access requirements. This cluster will be superceded this month by a more recent architecture, this time featuring *NUMA (non-uniform memory access)* nodes (Intel Nehalem), thereby greatly alleviating memory bandwidth issues.

Performance. The taubench statistics are as a consequence rather dependent on the load of the node. During 3 tests, results ranged from 11.5 to 15.3 seconds (each time averaged over 4 runs). In all of these cases all 7 other cores were executing other jobs. As the queueing system does not permit to choose the nodes in function of their load, nor to exclusively reserve a node, it is difficult to really cover all possibilities, in particular the extreme cases (no jobs vs very memory-intensive jobs on the other cores). For these cases jobs are moreover rather light in memory usage (20% of RAM), as the benchmark itself (10%). One would expect the timings to degrade further, as the memory transfer load on the node increases.

Parallel issues. In practice, it difficult to foresee let alone control the load that will be encountered when running a parallel computation. For the moment, the queuing system does not yet allow for hybrid parallel jobs, so all computations are run in pure MPI mode.

Case setup. The convergence criterion per timestep is 4 orders of magnitude, for a timestep value of $dt = 0.002c/u$, with c the chord and u the upstream velocity. The computation was run as a compressible flow at Mach number 0.1. Adiabatic wall conditions have been imposed at the airfoil surface. The distance between the periodic boundaries is $0.25c$.

4 Meshes

The mesh is generated by extruding a hybrid two-dimensional mesh defined on the symmetry plane. The 2D mesh is generated by a general purpose unstructured meshes; it is composed of a quadrilateral boundary layer combined to triangular elements in the farfield. The main characteristics are:

- 25 layers are used to cover a total span of $s = 0.25c$.
- the height of the first boundary layer elements is $10^{-3}c$;
- the freestream boundary is at 5c upstream and 10c elsewhere;
- second order Lagrangian interpolation is used for the element mappings;
- third order Lagrangian interpolation is used for the solution.
- the total mesh contains a total of 200k wedge-shaped and 87k hexahedral elements.

Similar mesh specifications will be used during the computations for the workshop.

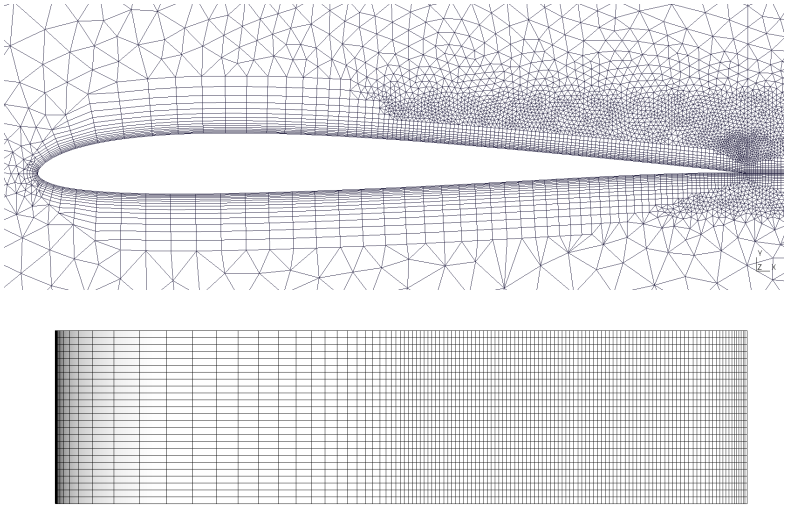

Figure 1: Extrusion mesh - view on the symmetry plane and the airfoil suction side

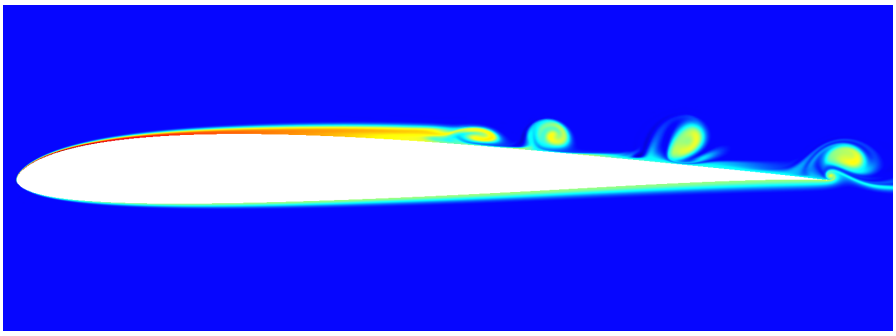

Figure 2: Entropy contours after 5 convective times

5 Results

The computation has been run for 3000 timesteps, totalling to 192 hours CPU on 512 processors. The flow field is illustrated in figures 2 and 3, showing the onset of three-dimensional instabilities. Unlike the computations we previously did for the Eppler airfoil, the natural transition is only now starting to show up, after 5 convective times, starting from a steady solution. We will introduce an initial perturbation to speed-up the transition process.

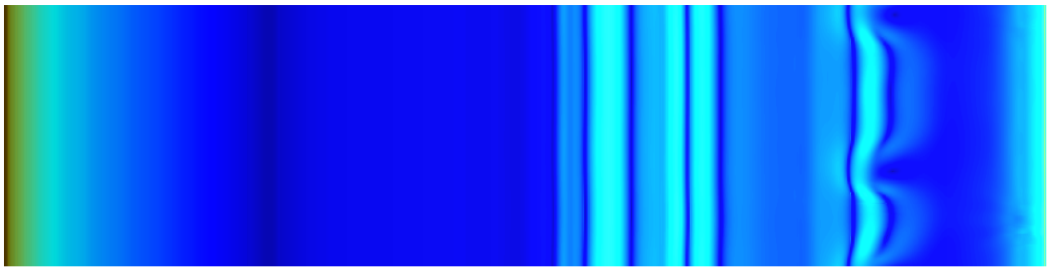

Figure 3: Skin friction contours after 5 convective times