Computational Fluid Dynamics

Prof. Dr.-Ing. Siegfried Wagner

Institut für Aerodynamik und Gasdynamik, Universität Stuttgart,
Pfaffenwaldring 21, 70550 Stuttgart, Germany

The majority of papers within CFD deals with Direct Numerical Simulation (DNS) and Large Eddy Simulation (LES) or with hybrid methods, i.e. a combination of DNS or LES with RANS (Reynolds Averaged Navier-Stokes) methods because they require high performance computers (HPCs) with big capability both with respect to memory and performance.

The increasing demand of memory and performance since the last workshop at HLRS in 2009 can be demonstrated by a few highlights. The Direct Numerical Simulations of the project by Friedrich and Kloker required a memory of 0.4 Tbyte RAM since their computational grid contained 10^9 grid points. The DNS of turbulent plane Cuette flow by Garcia-Villalba, Azagro and Uhlmann needed $6.06 \cdot 10^8$ grid cells on the HP XC4000 of SCCK. The big demand on more capable computers was also demonstrated by von Terzi, Venema, Bauer and Rodi who studied unsteady heat transfer by DNS. They used 10^8 grid cells for their simulation but did foresee a demand of 10^9 grid cells in the future.

Another important point is the scaling behaviour of a code as the number of processors per platform increased steadily in the past and will do so in the future. Chan-Braun, Garcia-Villalba and Uhlmann started DNS with 64 processors on the HP XC4000 of SCCK and demonstrated a good scaling behaviour up to 256 processors. A good speed-up up to 128 processors was also experienced by Galeazzo, Habisreuther and Zazalis. It is hoped that they can preserve this behaviour up to 1000 processors and more in the future.

LES and even more DNS require HPCs with extremely high performance and big memory. But they are mandatory since many problems cannot be simulated accurately by eddy-viscosity based RANS. Von Terzi, Schneider and Bauer showed that RANS cannot inherently account for mean secondary vortices in corners of diffusers which remarkably influence the performance of these devices. Thus, the authors used LES, but needed up to $22 \cdot 10^6$ grid cells.

The development of advanced numerical methods and algorithms, i.e. the software, is equally important as the hardware. Friedrich and Kloker reached a speed-up of about 2 compared to the performance of the NEC SX-8 last year by improving communication and employing optimized FFT. They reached a performance of 1.1 TFlop/s on the NEC SX-8. Furthermore, they used for their DNS of pinpoint suction in the boundary layer of airfoils and wings sixth order finite differences to resolve the complex flow in the boundary layer. To further improve the performance of the code they replaced the fully explicit Runge-Kutta 4th order time integration by a three-stage Heune method, i.e. the standard Heune method plus an additional second corrector step. They applied a spatial filter necessary for the downstream direction since the alternating forward-backward differences did not provide enough numerical damping for high wave number modes. The computer code LESOCC2 of von Terzi, Venema, Bauer and Rodi was run on the HP XC3000 of SCCK for preliminary simulations. They could obtain a better performance on a given platform when using different implementations of the computationally most intensive algorithms. They gained this experience by testing their code on the NEC SX-8 and SX-9, SGI-Altix and on several Linux clusters besides the HP XC3000.

To simulate the industrial production of large silicon single crystals Raufeisen, Breuer, Botsch and Delgado added the explicit computation of the melt-crystal interface to the previous version of their code FASTEST and computed the dynamics of the three-phase boundary where the melt, the crystal and the surrounding atmosphere meet. They extended the code to an ALE (Arbitrary Lagrangian Eulerian) formulation and could that way explicitly compute the free surface that was necessary to more accurately simulate the process.

Bensing, Keßler and Krämer used for their simulation of flow-structure interactions of helicopter rotors also an ALE formulation. They used an unstructured code on the NEC Nehalem cluster where they reached a sustained performance of 12 per cent peak performance using 128 processors. The present mesh size of $6.09 \cdot 10^6$ will be increased up to $5 \cdot 10^8$ cells in the future.

Roidl, Meinke and Schröder used for their numerical investigation of shock wave boundary-layer interactions a zonal RANS-LES ansatz whereby the transition from RANS to LES took place in an overlapping zone using synthetic turbulence generation methods and a reconstruction of eddy viscosity for transition. Up to 31 million mesh points were used. A performance of 483.2 GFlop/s was reached on the NEC SX-9 of HLRS using 2 nodes.

A shortcoming in the reports turned out to be the usage of a relatively low number of CPUs. This situation has to be improved since some of the HPCs within GCS (Gauss Centre of Supercomputing) are already massively parallel platforms and some additional could come in the near future. Although the three centres of supercomputing in Jülich, Munich and Stuttgart do already a lot with respect to education and courses for the customers there is still a need to continue this effort and may be to even increase it. A change in the philosophy of running the job classes might also help. Roidl, Meinke and Schröder stated in their paper that they used a relatively low number of nodes to minimize the overall user time for the simulation. They took into account that a rather high number of nodes would increase the turnaround time as the scheduling system prefers jobs with fewer nodes.

http://www.springer.com/978-3-642-15747-9

High Performance Computing in Science and Engineering '10
Transactions of the High Performance Computing Center, Stuttgart (HLRS) 2010
Nagel, W.E.; Kröner, D.B.; Resch, M.M. (Eds.)
2011, XIII, 604 p., Hardcover
ISBN: 978-3-642-15747-9