Computational Fluid Dynamics

Prof. Dr.-Ing. Siegfried Wagner

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

The impact of computer simulation in engineering has been significant and continues to grow. Simulation allows the development of highly optimised designs, the investigation of hazards too dangerous to test and reduced development costs. In parallel new scientific investigations are developing understanding in areas such as turbulence and flow control necessary for future engineering concepts.

However, since the flow phenomena are usually very complex highly sophisticated numerical procedures and algorithms as well as high performance computers (HPCs) had to be developed. Because the flow processes in daily life are turbulent and may even include chemical reactions, phase changes, heat transfer and interference with structural movement the high performance computers that are available at present are still by far too small in order to simulate for instance the turbulent flow around a complete aircraft. Despite this fact the following chapter on results in the field of computational fluid dynamics (CFD) that were obtained at HLR Stuttgart and SSC Karlsruhe will demonstrate the usefulness of numerical simulation and the progress in gaining more insight into complex flow phenomena when the computer capacity is increased and the corresponding numerical methods are improved.

The highly sophisticated numerical methods necessary for simulations on HPCs include Direct Numerical Simulation (DNS), Large Eddy Simulation (LES), numerical solutions of the Reynolds Averaged Navier-Stokes (RANS) equations, Detached Eddy Simulation (DES) that is a combination of LES and RANS, Lattice Boltzmann Method (LBM), combinations of DES and LES and finally Finte Element Methods (FEM) to investigate both flow and structural problems.

The high performance computers that were applied in the present studies were mostly the vector computer NEC SX-8 and the CRAY Opteron Cluster of HLRS as well as the massively parallel computers HP XC-4000 and HP-XC-2 of SSC Karlsruhe. However, some of the authors also compared the performance of these computers with CRAY XT and IBM BLUE GENE (e.g. paper of Zeiser et al.) and with JUMP (e.g. paper of Buijssen et al.).

It is not surprising that most of the papers use highly sophisticated numerical methods that undoubtedly require HPCs. More specifically, the distribution of the papers shows that six are devoted to DNS, five use LES, two apply the Lattice Boltzmann Method, two papers investigate flow-structure interactions, one uses RANS and optimization methods, one applies Eddy Dissipation/Finte Rate Combustion Model (EDM/FRC) and finally one paper investigates the performance of FEM.

Memory requirements did not seem to be a problem so far. At NEC SX-8 memory used ranged between 15 and 1086 GB whereas at HP XC-4000 only one information was available, namely 1.04 GB. The maximum number of processors used was 128 on the NEC SX-8 and 1024 on the HP XC-2. The maximum achieved performance approached 166.4 GFLOP/s on the NEC SX-8. Despite this big performance wall clock time was 6 hours in this case whereas in other cases wall clock time approached 85.7 hours. For the HP XC-4000 no performance numbers were presented. The jobs were big. For instance one run on the HP XC-2 with 1024 processors required 80,000 CPU hours. Turn around times went up to 60 days. These numbers show that turn around times would have been unacceptable without the access to HPC. Thus, the performance seems to be the bottle neck and should be increased by the next generation of high performance computers.

At present investigations in the field of DNS and LES have to be restricted to small Reynolds numbers and simple geometries i.e. Reynolds numbers far below those of technical applications. The present situation will be demonstrated by the following example that stems from flow-structure interactions in helicopter aeromechanics.

The CPU requirements for future simulations of helicopter aeromechanics are not yet known exactly but are bound to be substantial. Since little experience is available for full-scale/full-helicopter simulations researchers are still in the process of discovering the required grid resolution. An accurate simulation of an isolated rotor on today's supercomputer systems (3 GHz processor) would roughly require 5 million grid points per blade, 2000 time steps for each turn of the rotor, around 5 rotor revolutions until a trimmed state is reached and around 30000 CPU hours. This may well increase as the fidelity of the data required by the researchers and engineers increases. Accurate drag and sectional pitching-moment predictions require viscous flow simulations and this can easily triple the above estimate. Adding a fuselage and a tail rotor will increase this requirement to around 150000 CPU hours. This is due to the bluff body aerodynamics of the fuselage with shed vortices and massive flow separation but also due to the complex interaction between the fuselage and the rotors. Several calculations that are an order of magnitude larger would allow an assessment of issues such as grid dependency and establish confidence levels for smaller simulations on local facilities. The incorporation of flow control devices will potentially boost the computational requirements significantly. Finally, the helicopter system requires a delicate balance between the aeromechanics of the aircraft, the flight control system and the pilot. Simulation of manoeuvres, in contrast to design conditions, becomes an inherently multi-disciplinary problem with contributing modules requiring different treatments when it comes to parallel computing. The aerodynamic analysis is to have the lion's share but also a large effort will be required to simulate the effect of the pilot's action and of the control systems on the flight mechanics. This creates a challenging project where different solvers and modelling techniques must come together in a single parallel environment and require a PFlop/s sustained performance.

The presentation shows that a vector computer like the NEC SX-8 would be advantageous for many applications. The next generation, the NEC SX-9, will provide both several times higher performance and new cache-like memory concepts. These new concepts would probably require much more effort to optimally adapt existing vectorized codes to this new design than by the movement from NEC SX-4 to NEC SX-8. On the other hand increasing the performance of parallel systems by increasing the number of CPUs up to 100,000 and more will also be a big task in software development. The maximum number of CPUs in the papers presented was only 1024.

http://www.springer.com/978-3-540-88301-2

High Performance Computing in Science and Engineering ' 08
Transactions of the High Performance Computing Center, Stuttgart (HLRS) 2008
Nagel, W.E. (Ed.)
2009, XII, 613 p. 308 illus., 214 illus. in color., Hardcover
ISBN: 978-3-540-88301-2