# Parallel Efficiency of Multi-Block CFD Solver (COSA) Improved by HECToR dCSE Team

*M. Sergio Campobasso, Lancaster University,*
*Adrian Jackson, Edinburgh Parallel Computing Centre (EPCC),*
*HECToR CSE Team, Numerical Algorithms Group Ltd (NAG)*

An HPC expert from EPCC, working under NAG's Computational Science and Engineering (CSE) support service for HECToR, the UK's national academic supercomputing facility, has successfully improved the parallel efficiency of the COSA multi-block solver by optimising both the MPI and hybrid OpenMP/MPI implementation. Halo buffers were used to reduce the number of MPI point-to-point communication messages, and the hybrid OpenMP/MPI code was improved by ensuring cache coherency in shared data, and by utilizing single parallel regions to reduce overhead while placing MPI operations within OpenMP regions to reduce synchronization overhead.

COSA is a novel CFD code which is based on the compressible Navier-Stokes model for unsteady aerodynamics and aeroelasticity of fixed structures, rotary wings and turbomachinery blades. COSA includes a steady, a time domain, and a harmonic balance flow solver. All three solvers use a finite volume scheme with structured multi-block grids. All solvers can also use Low-Speed Preconditioning, which enables their use also for low-speed flows, such as those associated with horizontal axis wind turbines.

COSA has primarily been parallelised using MPI, but there is also a hybrid parallelisation that adds OpenMP functionality to the MPI parallelisation to enable larger numbers of cores to be utilised for a given simulation as the MPI parallelisation was limited to the number of geometric partitions (or blocks) in the simulation. This project was to optimise these two parallelisation strategies, improving the efficiency of both and therefore reducing the computational time required to perform simulations.

Commenting on the dCSE project success, *Dr M. Sergio Campobasso of the Department of Engineering at Lancaster University said: "The research and development work carried out in this project has been extremely valuable in many ways. By greatly improving the parallel performance of the MPI and MPI/OpenMP Harmonic Balance solver and the MPI time-domain solver, it has not only greatly reduced runtimes of our production analyses (sometimes by nearly 50%) on most clusters on which we use COSA, but it has also substantially consolidated the foundation (the COSA HB solver) for continuing the successful development and application of this software for challenging Renewable Energy problems."*

The combined MPI and general optimisations have resulted in improved scaling on HECToR. For some simulations, the overall runtime can now be reduced to half that of the original code. E.g. for a large test case (2048 block grid with 4,194,304 cells and 17 real harmonics) the new code is 30% faster for 2048 processes.

## HECToR

HECToR is managed by EPSRC on behalf of the participating Research Councils with a mission to support capability science and engineering in UK academia. The Cray XE6 supercomputer, located at the University of Edinburgh, is managed by UoE HPCx Ltd. The CSE Support Service is provided by NAG Ltd and ensures users have access to appropriate HPC expertise to effectively exploit advanced supercomputers for their science. A critical feature of the CSE Support Service is the distributed CSE (dCSE) programme which, through lightweight peer review, delivers dedicated performance and scalability projects on specific codes in response to proposals from users. The dCSE programme now consists of over 70 focused projects complementing the traditional HPC user applications support and training also provided by NAG.

The dCSE projects completed so far have delivered outstanding examples of the cost savings and new science that can be enabled through dedicated CSE effort. The COSA project reported here adds to these success stories with a successful performance improvement.

## Project Background

Whilst the different parallelisations of COSA enabled large simulations to be performed correctly, and in a much shorter time than the serial code, the efficiency of the different parallel implementations was not optimal. The OpenMP parallelisation allows simulations to be run on local computers such as desktop machines and small shared-memory clusters; however it cannot be used on large scale resources. Furthermore, the OpenMP parallelisation did not scale well even on small numbers of cores. The MPI code was generally efficient, although there were some areas that could be improved. The MPI code can only utilise a maximum number of cores equal to the maximum number of grid blocks in the simulation and the hybrid code had been developed to address this restriction, enabling a parallelisation that could use more cores than the number of blocks in the simulation. The hybrid parallelisation, however, was less efficient than the pure MPI parallelisation as it used the inefficient OpenMP functionality as well as the reasonably efficient MPI functionality.

Dr M. Sergio Campobasso of the Department of Engineering at Lancaster University was the Principal Investigator for the project. Adrian Jackson of EPCC carried out the 6 person-month project, in close collaboration with the NAG CSE team.

## Project Results

For the MPI work, messages were combined into buffers, and a significant speedup was demonstrated by up to 100x. Further work to improve the existing use of MPI I/O for the output files was also performed. For a representative grid of 262,144 cells and 31 real harmonics, with 512 processes, the original run-time of 701 seconds is now 547 seconds. For a larger test case (2048 block grid with 4,194,304 cells and 17 real harmonics) the new code is 30% faster for 2048 processes and 40% faster with 512 processes.

The combined MPI and general optimisations have resulted in improved scaling on HECToR, and the overall runtime can now be reduced to half that of the original code. These developments have been incorporated back into the main COSA source code. COSA has also been installed as a central module on HECToR.

---

A full technical report on this work can be found at http://www.hector.ac.uk/cse/distributedcse/reports/cosa/

For more information contact: HECToR CSE Team
The Numerical Algorithms Group Ltd, Wilkinson House, Jordan Hill Road, Oxford, OX2 8DR, United Kingdom
Telephone: 01865 511 245 Email: hector-cse@nag.co.uk Web: http://www.hector.ac.uk/cse/