
Michael J Brazell
Scalable solvers for the RANS equations

Dept 3295 1000 E University
Laramie
WY 82071
mbrazell@uwyo.edu
Dimitri J. Mavriplis
Behzad R. Ahrabi

The Reynolds-Averaged Navier-Stokes (RANS) equations are the engineering standard for solving aerodynamic flows and will be for the next 20 plus years. For some time finite-volume discretizations have been solving the RANS equations and have been successful. For example, NSU3D has incorporated an agglomerated non-linear multigrid strategy and has shown strong scaling up to 64k cores. More recently, the addition of a Newton-Krylov GMRES method preconditioned with linear multigrid has demonstrated full newton convergence. Machine-zero residuals are obtained for difficult aerodynamic flows which are not typical for finite-volume solvers. These simulations are reaching the order of 10^8 degrees of freedom per solution variable and are pushing the limits of mesh generators.

More recently, high-order methods such as discontinuous Galerkin (DG) and Streamline-upwind Petrov-Galerkin (SUPG) have been attempting to solve the RANS equations for aerodynamic flows. High-order methods have the advantage of obtaining the same accuracy with less degrees of freedom compared to low-order methods. However, these methods struggle when there are steep gradients or sharp geometrical features. Shock capturing and strong solvers have done a lot to obtain solver robustness, but high-order methods are still slower than traditional finite-volume solvers. This opens the question will high-order methods ever overtake low-order finite volume solvers in the solution to the RANS equations.

Two areas of research will determine this outcome, the first is *hp*-adaption and the second is the extension of these solvers on emerging computer architectures. High-order methods have the ability to both adapt the mesh size (*h*) and polynomial degree (*p*) without extending the stencil size beyond neighboring elements. *h*-refinement can be used where there are sharp gradients and *p*-refinement can be used in smooth regions. *hp*-adaptation gives the ability for high-order methods to obtain solutions with much fewer degrees of freedom and possibly faster time to solution compared finite-volume solvers. All of this hinges on the ability to adapt meshes in parallel with good load balancing and the ability to maintain solver performance. With the emergence of new computer architectures

high-order methods may have an advantage due to their dense computational kernels. Which in turn could mean that high-order methods although slower on smaller problems might have an advantage on larger problems due to the ability to scale more efficiently.

In this work we will investigate different solver technology between discretizations and attempt to draw conclusions about the future of high-order methods.