The Continuing Need for Method Development in Computational Fluid Dynamics

Philip Roe
Department of Aerospace Engineering, University of Michigan

Since its inception about 1950, CFD has progressed on two main fronts, the development of bigger, faster computers, and the development of subtler, more efficient algorithms. Until about 1990, it seems that both aspects were of similar importance in making CFD more effective. Since then, the major gains have come from exploiting hardware improvements. In particular, the use of workstation clusters has made large scale, massively parallel computing into an affordable option for almost all CFD users. By comparison, the gains from continuing investment into algorithmic improvements have seemed less certain, and harder to implement in practice. Although such improvements have continued to be made in the mathematical community, there have been few, if any, successful engineering implementations. Almost all current CFD is performed using methods developed during the 1980s or even earlier.

Moreover, the bigger machines that are now employed are not generally exploited to improve the accuracy of the calculations by using finer computational meshes, but instead to carry out bigger and more ambitious calculations that are no better resolved. There has been a huge increase in the quantity of the calculations, but little increase in the quality. Indeed, the more ambitious applications can reveal previously unsuspected defects. It is my opinion that the time is ripe for a resurgence of improved methods. In this talk I will discuss methods that are available for exploitations, some of the barriers to exploiting them, and some issues that remain to be addressed. Mostly, I will restrict myself to taking an Aerospace perspective.

Complex Problems

With available resources, we can now tackle problems that are governed by equations more complicated than those of Euler or Navier-Stokes, often incorporating additional unknowns that represent chemistry, electromagnetic effects, or non-equilibrium. Solving a Riemann problem in these circumstances to determine an interface flux becomes expensive. Approximate methods are worth trying in these situations.

High-order methods

First-order methods are seldom accurate enough for quantitative purposes, and virtually all practical CFD is performed using second-order methods. To get better accuracy than this requires using more information, either in the form of increased stencil size, or bestowing each computational element with additional degrees of freedom. The former strategy is typified by the ENO/WENO methods. The large stencils involved imply large overhead costs in communication, especially on unstructured meshes, and hence poor parallel performance for practical geometries. It also implies difficulty in maintaining accuracy adjacent to non-trivial boundaries. The second strategy is typified by the Discontinuous Galerkin method. An attraction of this method is that a rather automatic methodology is available to raise the order to any arbitrary level, but the computational cost rises rapidly with the order and with the number of dimensions. There are other ways to equip the elements with additional degrees of freedom, for example with Fourier expansions, as in the Spectral Element method of , for example, Gottlieb, or with subcell averages, as in the Spectral Volume method of Z J Wang.

Multidimensional Methods

Most compressible flow codes achieve robustness through some form of upwinding, almost always applied by a pseuodo-one-dimensional argument that lacks any proper mathematical justification. It can be shown that such codes are, in many situations, overly diffusive, but the problem is how to reduce the diffusion while retaining the robustness. Methods exist that do this based on multidimensional pattern recognition. They are less sensitive to imperfections in the grid, and lead to greatly reduced anomalies such as numerical entropy generation. They are based on the intensive processing of local information, which makes them inherently parallel. High-order `spectral' versions are under development. A crucial test is underway, to use these methods to improve the prediction of heat transfer at hypersonic speeds.

Preconditioning

The application of one-dimensional arguments to CFD methodology leads directly to inappropriate scaling of the diffusion terms in the incompressible limit. Preconditioning is a technique that restores proper scaling, leading to both improved convergence and better accuracy. Interestingly, this is one mathematical development from the past decade that has been quite widely implemented in practice, becausea version of it can be applied as a `fix' to existing code structures.

The Adjoint Revolution

Adjoint analysis is a borrowing from Control Theory that deals with the sensitivity of some predicted quantity, such as the drag of a wing, to all aspects of the computation, such as the boundary conditions, and the local error. Studying the sensitivity to the boundary conditions allows the shape to be optimized, and studying the sensitivity to the grid allows the computation itself to be done on an optimal grid. A traditional approach to such a study involves varying one thing at a time, and measuring the response, so as many computations have to be performed as there are parameters to be varied. With adjoint analysis, all of the sensitivities are available from just one computation, whose cost is comparable to that of one direct computation. CFD-based design could be potentially revolutionized, both in terms of the speed to reach design decisions, and in terms of confidence in the results.

The Practical Problem

Because of the focus on parallelization of existing codes, CFD engineers have lost touch with the kind of mathematical developments of which they used to take alert advantage. They no longer have the time, or the inclination to implement such developments themselves. Mathematicians and other researchers, however, lack the resources to give the kind of large-scale practical demonstrations that engineers would be convinced by. Theoretical and practical CFD, which enjoyed very close ties in the 1970s and 1980s, are beginning to grow apart. The best hope for a reunification is the development of programming tools that enable practical application of theoretical concepts in a shorter time than is currently feasible.