Barriers to Achieving Textbook Multigrid Efficiency (TME) in CFD

Achi Brandt†

Abstract

"Textbook multigrid efficiency" (TME) means solving a discrete PDE problem in a computational work which is only a small (less than 10) multiple of the operation count in the discretized system of equations itself. As a guide to attaining this optimal performance for general CFD problems, the table below lists every foreseen kind of computational difficulty for achieving that goal, together with the possible ways for resolving that difficulty, their current state of development, and references.

Included in the table are staggered and nonstaggered, conservative and nonconservative discretizations of viscous and inviscid, incompressible and compressible flows at various Mach numbers, as well as a simple (algebraic) turbulence model and comments on chemically reacting flows. The listing of associated computational barriers involves: non-alignment of streamlines or sonic characteristics with the grids; recirculating flows; stagnation points; discretization and relaxation on and near shocks and boundaries; far-field artificial boundary conditions; small-scale singularities (meaning important features, such as the complete airplane, which are not visible on some of the coarse grids); large grid aspect ratios; boundary layer resolution; and grid adaption.

Introduction (by James L. Thomas, NASA LaRC)

Computational fluid dynamics (CFD) is becoming a more important part of the complete aircraft design cycle because of the availability of faster computers with more memory and improved numerical algorithms. As an example, all of the external cruise-surface shapes of the new Boeing 777 wide-body subsonic transport were designed with CFD [R1]. The cruise shape of such a vehicle is designed to minimize viscous and shock wave losses at transonic speeds and can be analyzed with potential flow methods coupled with interacting boundary layers. Off-design

† Department of Applied Mathematics and Computer Science, The Weizmann Institute of Science, Rehovot 76100, Israel. mabrandt@weizmann.weizmann.ac.il This research was supported by the National Aeronautics and Space Administration under NASA Grant No. NAG1-1823 and by the Israeli Ministry of Science and Technology under Grant 9680-1-097. It was also suported by NASA Contract Nos. NAS1-19480 and NAS1-97046 while the author was in residence at the Institute for Computer Applications in Science and Engineering (ICASE), Mail Stop 403, NASA Langley Research Center, Hampton, VA 23681-2199.

performance associated with maximum lift, buffet, and flutter and the determination of stability and control derivatives, involving unsteady separated and vortical flows with stronger shock waves, are determined largely by experimental methods. Computational simulations of these flowfields require the use of Reynolds-averaged Navier-Stokes (RANS) methods; these computations for high-Reynolds flows over complex geometries are very expensive, the turnaround time is too long to impact the design cycle, and the turbulence models for separated flows have a high degree of variability. Thus in these areas experiments, rather than computations, are preferred for reasons of cost and uncertainty

Inroads are being made into these off-design areas with RANS methods. A major lesson learned from industrial use of RANS methods is that both the numerics and the physics must be improved substantially for a new procedure to replace an older procedure. Also, there is a synergistic interplay between the speed of the simulation and the fidelity of the turbulence model, since a larger parameter variation and/or model formulation can be explored on fine enough grids with a faster simulation. For example, the TLNS3D Navier-Stokes code [R2] found its way into use because it was the first three-dimensional Navier-Stokes code to show true multigrid performance, in which the cost scales linearly with the number of unknowns, and it incorporated a better turbulence model than the algebraic models then in use. Solutions with 1 million grid points could be converged in approximately 1 hr of Cray-2 time, which allowed spatial convergence studies to be conducted to ensure that the level of truncation error is sufficiently low, and the prediction of the angle of attack to attain a desired lift coefficient was improved over interacted potential methods [R3]. The faster turnaround of the multigrid procedure enabled the extension and calibration of the original two-dimensional turbulence model to three-dimensions, thus allowing a more accurate prediction of the transonic shock/boundary-layer interaction.

The current RANS solvers with multigrid require on the order of 1500 residual evaluations to converge the lift and drag to one percent of their final values for wing-body geometries near transonic cruise conditions. Complex geometry and complex physics simulations require many more residual evaluations to converge, if indeed convergence can even be attained. It is well-known for elliptic problems that solutions can be attained using full multigrid (FMG) processes in far fewer, on the order of 3–6, residual evaluations; this efficiency is known as textbook multigrid efficiency (TME). Thus, there is a potential gain of two orders of magnitude in operation count reduction if TME could be attained for the RANS equation sets. This possible two order of magnitude improvement in convergence represents an algorithmic floor since it is unlikely that faster convergence for these nonlinear equations could be attained. This algorithmic speed-up, however, coupled with further increases in computational speed can open up avenues and accelerate progress in many areas, including: the application of steady and time-dependent simulations in the high-lift, off-design, and stability and control areas; the usage of RANS solvers in the aerodynamic and multidisciplinary design areas; and the development of improved turbulence models.

The RANS equation sets are a system of coupled nonlinear equations which are not, even for subsonic Mach numbers, fully elliptic, but contain hyperbolic factors. The theory of multigrid for hyperbolic and mixed-type equations is much less developed than that for purely elliptic equations. Resolution of complex geometries and the thin boundary layers at high Reynolds number cause the grid to be highly irregular and stretched, leading to a slowdown in convergence. Discontinuities, such as shocks and slip surfaces, introduce additional difficulties. These difficulties are illustrated in the sketch in Fig. 1 for a typical multi-element section of a three-dimensional wing with the flaps deployed at takeoff and landing conditions. Overcoming these difficulties poses a formidable challenge, especially because in order to attain optimal and robust convergence rates for the applications of interest in aircraft design, they must all be overcome.

Brandt, in 1984 [G84], summarized the state of the art for attaining multigrid performance for fluid dynamics. Since that time, there has been considerable progress in the field, although optimal results have only been shown for inviscid flows, viscous flows at low Reynolds number, and simple geometries. The methodology and theory that Brandt and others have developed is applicable to the RANS equations and can lead to optimal convergence rates; however, a rational and systematic attack on the barriers which stand in the way needs to be mounted. The purpose of this paper is to delineate clearly the barriers which exist to attaining optimal convergence rates for solutions to the fluid dynamic equations for complex geometries. The following sections identify the barriers, possible solutions, and current status of the problem. The paper is intended as a guide to attaining the optimal convergence goal and is written for the most part in a tabular form so that new solutions and updates to the current status can be made. When completed, the document is intended to list every type of computational difficulty encountered on the road to attaining TME for RANS and the solution paths taken. The insights, lessons learned, and methodologies gained from aerodynamic applications should be applicable to other areas such as acoustics, electromagnetics, hypersonic propulsion, and aerothermodynamics.

Preliminary comments

The table below does *not* refer to a vast literature on multigrid methods in CFD (see for example [AJ]), in which enormous improvements over previous (single-grid) techniques have been achieved, but without adopting the systematic TME approach. This approach insists on obtaining basically the same ideal efficiency to *every* problem, by a very systematic study of each type of difficulty, through a carefully chosen sequence of model problems. Several fundamental techniques are typically absent in the multigrid codes that have not adopted the TME strategy. Most important, those codes fail to decompose the solution process into separate treatments of each factor of the PDE principal determinant, and therefore do not identify, let alone treat, the separate obstacles associated with each such factor. Indeed, depending on flow conditions, each of those factors may have different ellipticity measures (some are uniformly elliptic, others are non elliptic at some or all of the relevant scales) and/or different set of characteristic surfaces, requiring different combinations of relaxation and coarsening procedures.

The table deals only with *steady-state* flows and their *direct* multigrid solvers, i.e., not through pseudo-time marching. Time-accurate solvers for genuine *time-dependent* flow problems are in principle simpler to develop than their steady-state counterparts. Using semi implicit or fully implicit discretizations, large and adaptable time steps can be used, and parallel processing across space *and time* is feasible [R88]. The resulting system of equations (i.e., the system to be solved at each time step) is much easier than the steady-state system because it has better ellipticity measures (due to the time term), it does not involve the difficulties associated with recirculations, and it comes with a good first approximation (from the previous time step). A simple multigrid "F cycle" at each time step can solve the equations much below the discretization errors of that step [Par]. It is thus believed that fully efficient multigrid methods for the steady-state equations will also yield fully efficient and highly parallelizable methods for time-accurate integrations.

Acknowledgement

Contributions to the table were made by David Sidilkover (ICASE), Jerry C. South (NASA Langley Research Center, retired), R. Charles Swanson (NASA Langley Research Center), and Shlomo Ta'asan (Carnegie Mellon University).