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).

Textbook Multigrid Barriers

Far-Field Boundary Conditions

Grid Aligned with Characteristics

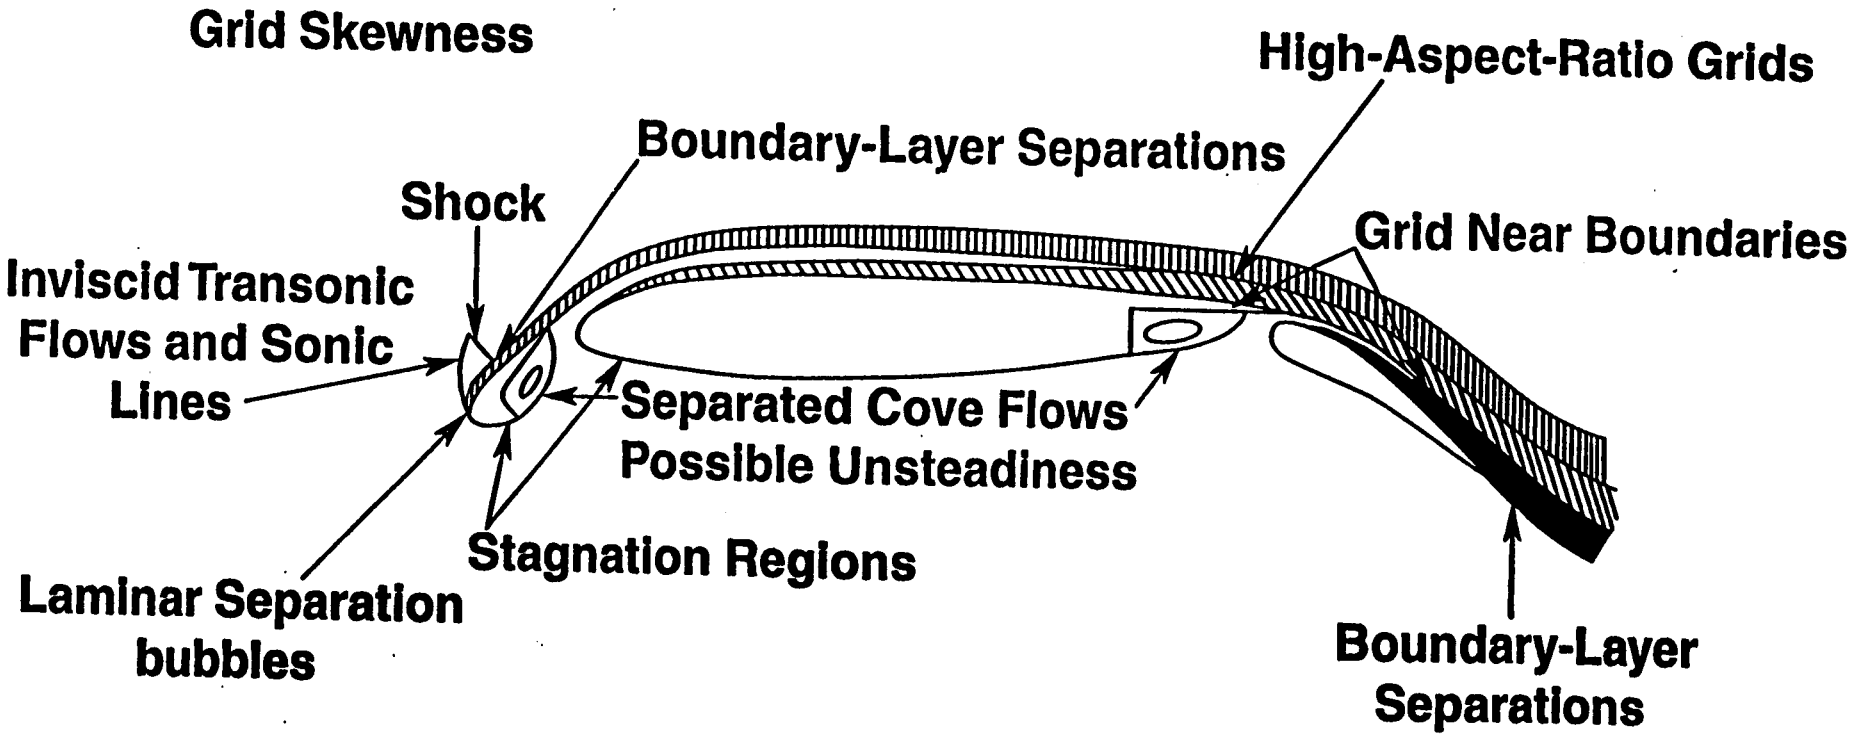

Attaining Ideal Multigrid Efficiency in CFD: Difficulties and Cures

Difficulty	***Possible Solutions***	***Status***
• Uniformly elliptic scalar equation on uniform grids in general domains	Multigrid cycles, guided by local mode analysis + FMG	TME demonstrated 1971 [B73], [B77] and rigorously proved [RLMA], [RQMA]
• Nonlinearity	(1) FAS cycles (2) Continuation processes (to obtain good initial approximations), integrated into one-shot FMG algorithm	(1) Demonstrated 1975 [South], [B77]. (2) Described in [G, §8.3.2]
• Fluid dynamics – general	See a review in [R88, §2]; at some points it is not fully up to date, but it concisely summarizes some main approaches needed for obtaining TME	

Difficulty	***Possible Solutions***	***Status***
• Non-scalar PDE systems	(1) General rules for the order of the inter-grid transfer operators are given in [G, §4.3], with some more details in [RQMA, §3.3]	
	(2) A general approach to the design of relaxation is based on the operator principal matrix L and on the factors of $\det L$ (see [G84, §§3.4, 3.7]). In this approach a distribution matrix M and a weighting (or "preconditioning") matrix P are constructed so that PLM is triangular, containing the factors of $\det L$ on the main diagonal (separated from each other as much as possible, to avoid the complication with "product operators" discussed next). This (if necessary – together with the technique described next), leads to decomposing relaxation into simple schemes for the (scalar) factors of $\det L$. Many specific examples are given below	TME demonstrated in a number of cases (see many detailsbelow). TME *proved* for uniformly elliptic systems [RLMA], [RQMA]
	(3) For systems of PDE which are of mixed type (elliptic-hyperbolic) another possibility sometimes is to introduce new unknowns in terms of which elliptic and hyperbolic parts are separated	TME demonstrated for incompressible and compresisble cases [T1]–[T5]

Difficulty	***Possible Solutions***	***Status***
• *Product operator*: an equation $LU = f$, where $L = L_2L_1$. Assume a relaxation process for L_j is given, with the amplification factor $\mu_j(\theta)$ and the smoothing factor $\overline{\mu}_j$, $(j = 1, 2)$	Two alternative approaches: (1) Introduce an explicit new unknown function V, replacing the equation with the pair of equations $L_1U - V = 0$ and $L_2V = f$, throughout the MG solution process (including, e.g., transferring residuals of both equations to coarse grids and correcting both u and v by interpolations from the corresponding coarse-grid values). The smoothing factor for this process is $\overline{\mu} = \max(\overline{\mu}_1, \overline{\mu}_2)$ (2) Use V only as an auxiliary function in relaxation. That is: starting with $v = L_1u$, where u is the current approximation to U, perform ν_2 sweeps on the equation $L_2V = f$, yielding a new value $\widetilde{v}$. Then perform ν_1 sweeps on the equation $L_1u = \widetilde{v}$. The resulting amplification factor is $\mu(\theta) = \mu_1(\theta)^{\nu_1} + \left[1 - \mu_1(\theta)^{\nu_1}\right]\widehat{L}_1(\theta)^{-1}\mu_2(\theta)^{\nu_2}\widehat{L}_1(\theta)$, where the Fourier symbols are defined by $\widehat{L}_j(\theta) = e^{-i\theta\cdot x/h}L_je^{i\theta\cdot x/h}$. Hence in scalar cases $\overline{\mu} < \overline{\mu}_1^{\nu_1} + \overline{\mu_2}^{\nu_2}$	(1) TME demonstrated for $L = \Delta^2$ [Lin], [Dym] (2) Not tried

Difficulty	***Possible Solutions***	***Status***
• **Smoothing for special CFD systems**	M = distribution operator P = preconditioner (see discussion above)	
• Cauchy Riemann on staggered grid $L = \begin{pmatrix} \partial_x & \partial_y \\ \partial_y & -\partial_x \end{pmatrix}$	Two alternatives: (1) $M = L, \quad P = I$ (2) $P = L, \quad M = I$	(1) TME demonstrated [BD], [Dinar] (2) TME validated [T6]
• Stokes on staggered grid $L = \begin{pmatrix} -\Delta & 0 & \partial_x \\ 0 & -\Delta & \partial_y \\ \partial_x & \partial_y & 0 \end{pmatrix}$	(1) $M = \begin{pmatrix} 1 & 0 & -\partial_x \\ 0 & 1 & -\partial_y \\ 0 & 0 & -\Delta \end{pmatrix}, \quad P = I$	(1) TME demonstrated [BD], [Dinar]
	(2) $P = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \partial_x & \partial_y & -\Delta \end{pmatrix}, \quad M = I$	(2) TME validated [T6]
• Stokes, non-staggered (1) Quasi-elliptic discretization $L = \begin{pmatrix} -\Delta & 0 & \partial_x^{2h} \\ 0 & -\Delta & \partial_y^{2h} \\ \partial_x^{2h} & \partial_y^{2h} & 0 \end{pmatrix}$ with averaging of the resulting pressure	Analogous to the staggered case; e.g., $M = \begin{pmatrix} 1 & 0 & -\partial_x^{2h} \\ 0 & 1 & -\partial_y^{2h} \\ 0 & 0 & -\Delta \end{pmatrix}$	(1) In a quasi-elliptic approach, TME demonstrated [G84, §18.6], [quasi]
(2) h-elliptic discretization, e.g. $L = \begin{pmatrix} -\Delta & 0 & \partial_x^{2h} \\ 0 & -\Delta & \partial_y^{2h} \\ \partial_x^{2h} & \partial_y^{2h} & -\alpha h^2 \Delta \end{pmatrix}$	No modifications of the FMG algorithm is required, even in the quasi-elliptic case (as explained in [G84, §18.6]). In generalization to NS, pressure averaging is required of coarse-level results before their interpolation to the next finer level (whenever the coarse-level employs the quasi-elliptic discretization)	(2) TME demonstrated? (perhaps by J. Linden)

Difficulty	***Possible Solutions***	***Status***
• Non-conservative incompressible Euler, whose principal operator in 2D is $L = \begin{pmatrix} \underline{u}\cdot\underline{\nabla} & 0 & \partial_x \\ 0 & \underline{u}\cdot\underline{\nabla} & \partial_y \\ \partial_x & \partial_y & 0 \end{pmatrix}$ (similarly 3D), on staggered grid, second (or higher) order discretization	(1) Employ cycle index $\gamma = 2^p$, where p is the order of discretization, with $M = \begin{pmatrix} 1 & 0 & -\partial_x \\ 0 & 1 & -\partial_y \\ 0 & 0 & \underline{u}\cdot\underline{\nabla} \end{pmatrix}$ (2) With the same M, for each of the momentum equations employ a relaxation scheme which is fast *converging* for the advection operator $\underline{u}\cdot\underline{\nabla}$ (i.e., converging fast not only for h–f, but also for smooth characteristic components; see discussion of advection below) (3) Use canonical variable (u, v, P) on staggered grid, where $P = (u^2 + v^2)/2 + p$. Upwind only P, use central discretization for (u, v). Relaxation is marching for P, and weighted (preconditioning) for (u, v)	(1) TME for first-order discretization using W cycles shown in [BD], [Dinar] (2) TME demonstrated for 2D entering flows with second-order discretization [BY2] and for recirculating flows with first-order discretization [BY3] (3) TME in [T1–T3]
• Low-Reynolds Incompressibble Navier-Stokes, staggered or not	Fully analogous to Stokes solvers: just replace Δ in L by $Q = -R^{-1}\Delta + \underline{u}\cdot\underline{\nabla}$	TME demonstrated 1978 [BD], [Dinar]
• High-Reynolds Incompressible Navier-Stokes, staggered or not	Fully analogous to Incompressible Euler (outside boundary layers: see discussion on such layers below): just replace $\underline{u}\cdot\underline{\nabla}$ everywhere with Q	TME demonstrated for first-order discretization on staggered ([BD], [Dinar]) and non-staggered grids [G84, §19.5], and for second-order staggered discretization[BY2]

Difficulty	***Possible Solutions***	***Status***
• *Compressible Euler*, non-conservative, on staggered grid: The subprincipal operator on $(u_1, u_2, u_3, \rho, \varepsilon, p)$ is $L = \begin{pmatrix} \rho\underline{u}\cdot\underline{\nabla} & 0 & 0 & 0 & 0 & \partial_1 \\ 0 & \rho\underline{u}\cdot\underline{\nabla} & 0 & 0 & 0 & \partial_2 \\ 0 & 0 & \rho\underline{u}\cdot\underline{\nabla} & 0 & 0 & \partial_3 \\ \rho^2\partial_1 & \rho^2\partial_2 & \rho^2\partial_3 & \rho\underline{u}\cdot\underline{\nabla} & 0 & 0 \\ p\partial_1 & p\partial_2 & p\partial_3 & 0 & \rho\underline{u}\cdot\underline{\nabla} & 0 \\ 0 & 0 & 0 & -\partial p/\partial\rho & -\partial p/\partial\varepsilon & 1 \end{pmatrix}$ $\det L = \rho^5(\underline{u}\cdot\underline{\nabla})^3\big((\underline{u}\cdot\underline{\nabla})^2 - a^2\Delta\big)$, $a = \left(\dfrac{\partial p}{\partial\rho} + \dfrac{p}{\rho^2}\dfrac{\partial p}{\partial\varepsilon}\right)^{1/2}$ is the sound speed, ρ, ε, p defined at cell centers, u_i – at center of cell faces perpendicular to the i-th coordinate	$M = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & -\rho(\underline{u}\cdot\underline{\nabla})\partial_1 \\ 0 & 1 & 0 & 0 & 0 & -\rho(\underline{u}\cdot\underline{\nabla})\partial_2 \\ 0 & 0 & 1 & 0 & 0 & -\rho(\underline{u}\cdot\underline{\nabla})\partial_3 \\ 0 & 0 & 0 & 1 & 0 & -\rho^2\Delta \\ 0 & 0 & 0 & 0 & 1 & -p\Delta \\ 0 & 0 & 0 & 0 & 0 & \rho^2(\underline{u}\cdot\underline{\nabla})^2 \end{pmatrix}$ The advection and full-potential operators are each relaxed by one of the approaches described for them below (in the chapter on non-elliptic operators. The *semi*-coarsening described there would then be used as an *inner* multigrid cycle for *relaxing* one factor of the determinant, to be distinguished from the *outer* multigrid cycle, which can use *full* coarsening.)	Not tried
• 2D Compressible Euler, nonconservative and conservative, staggered grid, using canonical variables (u, v, S, H). Structured and unstructured grids	Use (u, v) at cell edges, H at middle of cell, S at vertices. Upwind only S at momentum equations. Relax S, H by marching. (u, v) by a weighting relaxation. Crocco's form is used here to define relaxation	TME in [T2–T5]
• 2D/3D incompressible and compressible Euler: Canonical variables in which velocities are replaced by vector potential representation. Nonstaggered structured and unstructured grid	All variables at cell nodes. Relax hyperbolic quantities using marching. Relax vector potential using point Gauss-Seidel	TME acheived (unpublished) for interior and exterior flows in 2D, interior in 3D

Difficulty	***Possible Solutions***	***Status***

- *Compressible Navier-Stokes*, non-conservative.
  The subprincipal operator on $(u_1, u_2, u_3, \rho, \varepsilon, p)$ is

$$L_s = \begin{pmatrix} Q_\mu - \overline{\lambda}\partial_{11} & -\overline{\lambda}\partial_{12} & -\overline{\lambda}\partial_{13} & 0 & 0 & \partial_1 \\ -\overline{\lambda}\partial_{21} & Q_\mu - \overline{\lambda}\partial_{22} & -\overline{\lambda}\partial_{23} & 0 & 0 & \partial_2 \\ -\overline{\lambda}\partial_{31} & -\overline{\lambda}\partial_{32} & Q_\mu - \overline{\lambda}\partial_{33} & 0 & 0 & \partial_3 \\ \rho^2\partial_1 & \rho^2\partial_2 & \rho^2\partial_3 & Q_0 & 0 & 0 \\ p\partial_1 & p\partial_2 & p\partial_3 & 0 & Q_\kappa & 0 \\ 0 & 0 & 0 & -\partial p/\partial\rho & -\partial p/\partial\varepsilon & 1 \end{pmatrix}$$

  where $Q_\alpha = -\alpha\Delta + \rho\underline{u}\cdot\underline{\nabla}$, $\overline{\lambda} = \lambda + \mu$, $\lambda = \frac{2}{3}\mu$,
  $\kappa = k/c_v$ (coefficient of thermal conductivity divided by the specific heat at constant volume),
  $\det L_s = Q_\mu^2 \det L_c$, where L_c is the "core operator"

$$L_c = \begin{pmatrix} Q_0 & 0 & -\rho^2\Delta \\ 0 & Q_\kappa & -p\Delta \\ -\partial p/\partial\rho & -\partial p/\partial\varepsilon & Q_{\mu+\overline{\lambda}} \end{pmatrix}$$

  At standard conditions of laminar air flow the Prandtl number $\gamma\mu/\kappa \approx 0.72$; for turbulence $\gamma\mu/\kappa \approx 0.9$, with $\gamma = c_p/c_v = 1.4$

(1) Where $\overline{\lambda}, \mu, \kappa \ll \rho h|\underline{u}|$ relax as in Euler above — Not tried
(2) Otherwise use

$$M = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & -\partial_1 \\ 0 & 1 & 0 & 0 & 0 & -\partial_2 \\ 0 & 0 & 1 & 0 & 0 & -\partial_3 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ \overline{\lambda}\partial_1 & \overline{\lambda}\partial_2 & \overline{\lambda}\partial_3 & 0 & 0 & Q_{\mu+\overline{\lambda}} \end{pmatrix}$$

relaxing each Q_μ by one of the approaches described for the advection-diffusion below, and L_c by procedures discussed for it below (in the chapter on non-elliptic operators)

Difficulty	***Possible Solutions***	***Status***
• Non-conservative *not staggered* Euler and NS	(1) Probably similar to the staggered (cf. transition from staggered to non-staggered in Stokes) (2) In the 2D incompressible case: Premultiply L by a projection operator P, obtaining a Poisson equation for the pressure. Solve pressure equation with multigrid and the advection equation by marching downstream.	TME demonstrated for 2D incompressible Euler [RSS] in the cases of channel (with bump) and airfoil flows
• *Conservative* discretization of any of the above systems	Apply a prefactor P such that PL has principally the above non-conservative form. See, however, the difficulty associated with FDA *factorizability* (discussed in the chapter on non-elliptic operators), which may arise with such PL operators	Mentioned in [G, §3.4], but not tested

Difficulty	***Possible Solutions***	***Status***
• **Non-elliptic operators**, or more precisely: small ellipticity measures at some (e.g., large) scales. The main operators of interest here are (1) The advection operator (or, similarly, the convection-diffusion operator at large Reynolds numbers). (2) The near-sonic full-potential operator or more generally the core operator L_c. (3) The *supersonic* full-potential operator or L_c. (See below a separate discussion of anisotropies caused by the discretization)	The DGS relaxation of the full flow equations allows a specific individual treatment for each of these cases, taking into account its particular set of characteristic, as detailed below	
• Grid *aligned* with the characteristics: Pointwise relaxation has only *semi*-smoothing capability	Block (e.g., line or plane) relaxation schemes and/or semi-coarsening, possibly in alternating directions, guided by mode analyses [B77], [Stages]; or ILU relaxation [K2], [S9]	TME demonstrated in many cases
• Distinguishing different regimes (open vs. closed characteristics)	Running separately the relaxation subroutine of a given non-elliptic factor can (1) Separately check its convergence properties (2) Produce a scalar $\sigma \approx 1$ at regions of open characteristics and $\sigma \ll 1$ on closed characteristics (as in separated flow zones)	

Difficulty	***Possible Solutions***	***Status***
• *Non-aligned* grids, with *open* characteristics (e.g., entering flow): The main difficulty is the shorter distance (along the characteristics) for which a coarser grid still approximates some *smooth* solution components (characteristic components with intermediate cross-characteristic smoothness) [NESP], [BY1]	Three possible approaches, all guided by *half-space* two-level FMG mode analysis, using for simplicity the First Differential Approximation (FDA) to the discrete operator [NESP], [G, §7.5]:	
	(1) Downstream-ordered relaxation marching [R88, §2.3]: Suitable only for the advection factor; sometimes still requires W cycles, and not very good for massively parallel processing. In the case of an $O(h^p)$ discretization which is not purely upstreamed, relaxation should involve a predictor-corrector downstream marching. If the predictor order is q, the corrector should be applied at least p/q times.	(1) TME demonstrated in [BY2], and in Ruge's recent calculations, both for the advection operator by itself and as part of the incompressible Euler system
	(2) Similarly, with downstream-ordered ILU relaxation: Suitable for the advection operator (in 2D and 3D) *and* for the near-sonic full-potential operator in 2D (*not* in 3D)	
	(3) Semi-coarsening with controlled artificial dissipation at coarse levels (to match the target-grid numerical dissipation): Suitable for all operators in 2D and 3D, and for massive parallel processing	(3) TME has been shown for the sonic full-potential operator [D1], [D2] [D3]
	(4) Cycle index $= 2^{p/m}$, where p is the order of discretization and m is the order of the differential factor. (Suitable actually only for the advection operator, for which $m = 1$; especially attractive for $p = 2$ in 3D; not requiring ordered relaxation, but still disadvantageous for massively parallel processing because of the high cycle index)	(4) For $p = 1$, TME has been shown on various occasions. For $p = 2$, should be tried

Difficulty	***Possible Solutions***	***Status***
• The mixed *convection-diffusion* operator with order p approximation, having natural viscosity ν and artificial viscosity αh^p	Treatment as elliptic operator on levels where $\nu \gtrsim (2^p \cdot 4 - 5)\alpha h^p$ and as the non-elliptic advection operator otherwise	Not precisely tried
• *Closed* characteristics (recirculating flows). Here *uniformity* of viscosity (including numerical viscosity) is important for accuracy, while the viscosity *size* is less important (except at resolved boundary layers, discussed below). In fact, a uniform $O(h)$ artificial viscosity can yield higher order approximations. Full convergence may also be less important here (since in reality, too, steady state may take exceedingly long to establish)	Using the above-mentioned scalar σ, form a σ-dependent convergence test, to tell between slowness of open and closed characteristics (and possibly ignore the latter). Also based on σ, at recirculation regions use uniform (explicit) $O(h)$ numerical viscosity, with continuation from large to small viscosity integrated into the FMG algorithm. The cycles can employ one of the following 3 options. (1) DCW method (using Defect Corrections within W cycles), with suitable overweighting of residuals [BY3]. Suitable only for $O(h)$ discretizations. (2) Effectively downstream relaxation ordering (using alternate-direction sweeps) and doubling of transferred residuals (for $O(h)$ discretization) [YVB]. (3) Semicoarsening, generally similar to [D1]	TME cycles by methods (1) and (2) were shown in [BY3] and [YVB] respectively. Method (3), which should be best for massive parallelization, has not been implemented
• Full-potential operator $(\underline{u} \cdot \underline{\nabla})^2 - a^2\Delta$, $M_0 = \lvert\underline{u}\rvert/a \lesssim .7$ (uniformly elliptic)	Any classical algorithm is suitable, but the algorithm of the next case is also adequate	TME well established

Difficulty	***Possible Solutions***	***Status***
• Full potential $.7 \lesssim M_0 \lesssim 1.4$	Relaxation marching downstream (for transition to the supersonic case below) together with semicoarsening in the characteristic (cross-stream) direction	TME shown for the case $M = 1$ [D1]. Other cases have not yet been implemented
• Full potential $1.4 \lesssim M_0$ (uniformly hyperbolic, with the stream as the time-like direction, and with $O(1)$ "Courant number".)	Marching in the stream direction, possibly with a predictor-corrector procedure. For full massive parallelization, however, wave methods (extending standing wave methods [Ira]) should be used	Not yet tried?

Difficulty	***Possible Solutions***	***Status***
• The "core operator" $L_c = \begin{bmatrix} Q_0 & 0 & -\rho^2\Delta \\ 0 & Q_\kappa & -p\Delta \\ -\partial p/\partial\rho & -\partial p/\partial\varepsilon & Q_{\mu+\overline{\lambda}} \end{bmatrix}$ should be relaxed as part of relaxing the compressible NS system, in the case that $\rho\|u\|h \lesssim \max(\overline{\lambda}, \mu, \kappa)$. In the case of alignment between the grid and the flow, with meshsize h_1 and h_2 in the stream and cross-stream directions, respectively, and $h_2 \le h_1$ (e.g., in boundary layers), the case where L_c need be relaxed is when $\rho\|u\|h_2^2 \lesssim h_1 \max(\overline{\lambda}, \mu, \kappa)$. In aerodynamics, $\overline{\lambda}$, μ and κ are comparable, so the case of interest is $\|u\|h_2^2 \le \nu h_1$, where $\nu = \mu/\rho$	Best relaxation scheme depends on the flow parameters. For example: (1) If $\kappa \ll \rho\|u\|h$, then $Q_\kappa \approx Q_0$ (in principal terms) and one can use DGS with $M = \begin{pmatrix} 1 & 0 & \rho^2\Delta \\ 0 & 1 & p\Delta \\ 0 & 0 & Q_0 \end{pmatrix}$ resulting in the need to relax the first two equations each on an advection operator (see methods above), and the third equation on the operator $Q_0 Q_{\mu+\overline{\lambda}} - \rho^2 a^2 \Delta$. In the case of interest the principal part of the latter is $\left[(\mu+\overline{\lambda})Q_0 + \rho^2 a^2\right]\Delta$, so it can be relaxed by the general method for relaxing a product opertor (see $L = L_2 L_1$ above). (2) In the aerodynamics and aligned case of interest, the term $Q_{\mu+\overline{\lambda}}$ in L_c is not principal. Therefore relaxation can easily be conducted with the weighting (preconditioning) matrix $P = \begin{pmatrix} 1 & 0 & 0 \\ -p & \rho^2 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ and the distribution matrix $M = \begin{pmatrix} 0 & 1 & 0 \\ 0 & -p_\rho/p_\varepsilon & 1 \\ 1 & 0 & 0 \end{pmatrix}$ yielding PLM whose principal part is its main diagonal, on which separately appear the Laplace operator Δ, the convection-diffusion operator $Q_{\overline{\kappa}}$ where $\overline{\kappa} = p_\rho \rho^2/(2pp_\varepsilon)$ $= 1.25\kappa$ (for air), and a free function	Not tried Not tried

Difficulty	***Possible Solutions***	***Status***
• *FDA factorizability question*: The decomposition of a system relaxation into its scalar factors depends on the equality of the different occurrences of the advection-diffusion operator Q (or $Q_{\mu+\overline{\lambda}}$) appearing in PL, the prefactoring by P of a conservative discretization L. However, for smooth-characteristic convergence in relaxing a non-elliptic discrete operator, important is not only the differential operator it approximates, but also its First Differential Approximation (FDA) terms in non-characteristic directions; e.g., the cross-stream *numerical* viscosity of Q. This may not be the same in the different occurrences of Q, putting the factorization into question	(1) Examining several examples of conservative discretization of transonic flows, the FDA terms in various occurrences of $Q_{\mu+\overline{\lambda}}$ turn out sufficiently close to each other (e.g., only 4% discrepancy) to allow full efficiency of the proposed relaxation schemes.	(1) Further examination is needed
	(2) Conservative schemes may be designed so that the various FDAs of $Q_{\mu+\overline{\lambda}}$ are identical, or at least so that the scheme is still factorizable.	(2) Some "genuinely multidimensional upwind" schemes turn out to yield factorizable schemes, e.g., in the subsonic case in the control-volume structured-grid context [DS2]. Further studies are in progress
	(3) A general practical approach is a defect-correction relaxation: the residuals are calculated by the given PL system and fed into a DGS relaxation scheme whose driving factors may have different discretizations (as long as their numerical viscosities are not larger than those in the PL system)	(3) Not tried

Difficulty	***Possible Solutions***	***Status***
• High order discretization (away from shocks)	(1) "*Double discretization*" *schemes*: Use high-order only in calculating residuals transferred to the coarse grid, *not* in relaxation (unless the high order scheme is preferable also for h–f modes). (2) However, in relaxing *non-elliptic* factors (e.g., downstream relaxation marching for convection operator) the high order must be used (e.g., by a predictor-corrector downstream relaxation)	Introduced 1978 [BD]. Successfully implemented in various *elliptic* cases (see description and refs in [G, §10.2]). Methods for non-elliptic have not been tested beyond second order. *Comment*: High order approximations on unstructured grids are very expensive

Difficulty	*Possible Solutions*	*Status*
• **Albegraic turbulence models** employ the (compressible or incompressible) Navier-Stokes equations, adding to the laminar viscosity μ_ℓ (and similarly to λ_ℓ and κ_ℓ) a turbulent viscosity μ_t (similarly λ_t, κ_t), which is defined in terms of geometric functions (such as the distance to the wall), flow-dependent boundary-layer-wide (BLW) parameters (such as the boundary layer thickness, the maximum and minimum total velocity across the layer, and the flow wall friction) and in terms of the local total vorticity $\omega = \lvert \operatorname{curl} \underline{u} \rvert$. For example, the two-layer Baldwin-Lomax model [BL], is defined in two regions as follows: (1) *Outer layer*. Here μ_t is defined only in terms of distance from the boundary and BLW parameters. (2) *Inner layer*. Closer to the wall, $\mu_t = \rho\ell^2\omega$, where ℓ depends on the distance to the wall and on BLW parameters. In the 2D incompressible case, and neglecting the laminar viscosity, the resulting principal operator, on the vector of unknowns (ω, u, v, p), is $L = \begin{pmatrix} -\omega & (u_y - v_x)\partial_y & -(u_y - v_x)\partial_x & 0 \\ -A & Q_\mu & 0 & \partial_x \\ -B & 0 & Q_\mu & \partial_y \\ 0 & \partial_x & \partial_y & 0 \end{pmatrix}$ where $A = \rho\ell^2[2u_x\partial_x + (u_y + v_x)\partial_y]$, $B = \rho\ell^2[2v_y\partial_y + (u_y + v_x)\partial_x]$, so that $\det L = \Delta\Big\{\omega\partial_s - 2\rho\ell^2(u_y - v_x) \times \big[v_x\partial_{xx} - u_y\partial_{yy} + (v_y - u_x)\partial_{xy}\big]\Big\}$	*In the outer layer* the principal operator, hence also relaxation, are exactly as for the laminar case, with $\mu = \mu_t + \mu_\ell$. The BLW parameters are held unchanged during relaxation at scales finer than the boundary-layer width. Only at a suitable coarser level, where the cross-layer meshsize approaches the layer width, the dependence of the BLW parameters on the flow is relaxed together with the flow equations themselves, by applying box relaxation near the boundary (cf. the section on boundary relaxation). *In the inner layer*, suppose for example that the coordinate along the wall is x, and $u_y \gg \max\{\lvert u_x\rvert, \lvert v_x\rvert, \lvert v_y\rvert\}$. Then the principal operator takes the form $L = \begin{pmatrix} -\mu & \mu\partial_y & -\mu\partial_x & 0 \\ -\mu\partial_y & Q_\mu & 0 & \partial_x \\ -\mu\partial_x & 0 & Q_\mu & \partial_y \\ 0 & \partial_x & \partial_y & 0 \end{pmatrix}$ A suitable distribution operator then is $M = \begin{pmatrix} 1 & \partial_y & -\partial_x & 0 \\ 0 & 1 & 0 & -\partial_x \\ 0 & 0 & 1 & -\partial_y \\ 0 & \mu\partial_x & -\mu\partial_y & Q_\mu \end{pmatrix}$ yielding LM with principal terms only on the main diagonal, where there appear the operators Δ (for the continuity equation ghost function) and $\underline{u} \cdot \underline{\nabla} - 2\mu\partial_{yy}$ (for each of the momentum equations). The latter is nonelliptic, and its characteristics would often be aligned with the grid (cf. the section on nonelliptic operators)	TME should first be demonstrated for a simple turbulence model, such as the one described here

Difficulty	*Possible Solutions*	*Status*
• **Chemically reacting flows** feature three types of difficulty (1) A set of N continuity equations, where N, the number of chemical species, may be quite large. (2) The nonlinear source terms in these equations may be very stiff. (3) Some densities at some (few) points may become negative upon the coarse-to-fine FAS interpolation	At any grid point where any source term is *principal* (meaning: its production rate of species i per unit volume is large compared with $\max\big(h^{-2}D_i, h^{-1}\rho_i\lvert u\rvert\big)$, where D_i is the local diffusion coefficient of species i) it should be included in the principal matrix operator L. As a result, at each gridpiont the weighted-distributive relaxation step (local inversion of the principal terms of PLM) may involve the inversion of a matrix of size upto $N \times N$. This would correspond to the (relaxation part of the) point-implicit method [BM]. Fortunately, this will usually happen only on some coarser multigrid levels and/or at some restricted zones, thus consuming only relatively small work. Nonlinearity is treated by an FAS in which, instead of the fine-to-coarse transfers of densities $\rho_1, \ldots, \rho_N$ and the coarse-to-fine interpolation of the changes $\delta\rho_1, \ldots, \delta\rho_N$, transferred are the functions $f_1(\rho_1), \ldots, f_N(\rho_N)$ and interpolated are $\delta f_1(\rho_1), \ldots, \delta f_N(\rho_N)$, where $f_i(\rho_i)$ are properly chosen functions; e.g. $f_i(\rho_i) = \log \rho_i$, so that after interpolation $\rho_i = \exp\big(f_i(\rho_i^{\text{OLD}}) + \delta f_i(\rho_i)\big) > 0$. Furthermore, the continuity-equation residual restriction should be conservative (strictly full weighting)	TME should first be demonstrated for a simple model case; e.g. a 2D incompressible inviscid flow with two reacting species

Difficulty	*Possible Solutions*	*Status*
• **Shocks**		
• *Shock displacement question*: A (small) displacement Should result from global solution changes that occur on coarse levels of the cycle. How can one obtain an accurate displacement, when those levels are too coarse to resolve it?	An accurate shock displacement is obtained if the fine-to-coarse residual transfer is conservative (e.g., "full weighting") and the coarse-to-fine correction interpolation is followed by local relaxation passes near the shock	Full efficiency shown [DS]
• Relaxation near strong shocks	Add extra relaxation passes, using general robust schemes (e.g., box Kacmarz), until all local residuals drop to their level away from chocks	Not tried
• Poor h-ellipticity of high-resolution schemes	Construction of new, genuinely multidimensional upwind schemes	Developed in the context of unstructured triangular grids [DS1]

Difficulty	*Possible Solutions*	*Status*
• **Boundary related difficulties**		
• *Discretization* near boundaries	For best multigrid efficiency, use Cartesian coordinates throughout the domain, with boundary-fitted local grid patches, regarded as finer multigrid local levels [B77], [G]. Only crude (e.g., first order) discretization is then needed near the boundaries on the Cartesian grids	Ruge & Brandt devised a near-general-boundary discretization for incompressible Euler on staggered Cartesian grid
• *Relaxation at and near boundaries*: Difficulties: (1) There is no smoothing analysis in case the boundaries are not aligned with the grid. (2) The fine-to-coarse residual weighting near boundaries is generally very imprecise, hence the residuals should be reduced there *more* than in the interior. (3) Larger residuals are created near boundaties upon coarse-to-fine interpolations (of solution or corrections)	A general-type robust relaxation scheme, e.g., *box Kacmarz*, throughout several-meshsize-wide zone near the boundary. Box size in each direction should be several meshsizes and the boxes should have substantial overlap. One can afford several passes of such a relaxation per each full interior sweep since the zone width is $O(h^{1-\varepsilon})$, with $0 \leq \varepsilon < 1$. In particular, add near-boundary relaxation passes after the FMG interpolation (allowing the latter to be of lower order near the boundary). The local relaxation passes should continue at least until all local residuals have dropped well below their global average magnitude	For uniformly elliptic equations it has been proved [RLMA], [RQMA] and demonstrated computationally (for cases of reentrant corners [Bai]) that the interior efficiency as predicted by mode analysis (implying TME) can always be obtained. TME demonstrated (by Ruge & Brandt) for incompressible Euler on staggered cartesian grids

Difficulty	***Possible Solutions***	***Status***
• *Boundary layers* (in the case that they need be resolved). also grid adaptation below.)	Resolved by boundary-fitted local grid (Seepat-(ches, with local semi refinements: finer levels, in narrower layers near the boundary, have smaller cross-layer meshsizes, allowing the physical cross-stream viscosity to dominate over the numerical one. Additional terms in the governing equations (NS instead of Euler, or turbulent modelling etc.) may be used in these patches. Downstream marching relaxation and cross-stream semi coarsening should feature in the multigrid cycles. A "λ-FMG" kind of algorithm [G, §9.6], should be employed, so that coarse FMG stages already include local semi-refinements at the boundary, thus effectively incorporating into the FMG stages a process of continuation in R_e	Description in [R88, §2.4]; not implemented. The local refinement techniques for Poisson equation, with TME, are demonstrated in [Bai]
• *Far-field* artificial boundary conditions: requiring in some cases non-local *absorbing boundary conditions* (ABC) for some wave factor	Increasingly coarser grids covering increasingly larger domains. The size of each domain is based on accuracy-to-work optimization criteria (similar to those in [B77, §8], [G, §9.5]), implying also a natural criterion for the largest needed domain. On *interior* boundaries (boundaries of a grid residing in the interior of the next coarser grid) the solution is interpolated from the coarser grid. On such boundaries, if ABC is at all needed, only high-frequency components need be absorbed, for which the ABC are *local*, and can be enforced as part of the relaxation process (of the corresponding wave factor)	Details of the algorithm have been worked out, and TME (or its equivalent accuracy-to-work relation) has recently been demonstrated for the 2D Poisson equation in the unbounded plane (cf. [ETNA, §4]). Techniques for non-elliptic or indefinite cases have not been systematically studied

Difficulty	***Possible Solutions***	***Status***
• *Small-scale singularities* invisible on the next coarser grid, such as small "islands" or "holes" in the domain (e.g., an airplane smaller than the meshsize of some coarser grid) or small BC features (e.g., small regions of Neumann BC in otherwise Dirichlet BC)	Local relaxation passes around the singularities after return from the next coarser grid, together with either one of the following three devices: (a) Enlarging the singularity on the coarser grid. (b) Modifying the interior coarse-grid equation near the singularity. (c) If the coarse grid equations are not modified, then the convergence is slow, but slow to converge are just few very special components. Hence slowness can be eliminated by recombining iterants (see below)	TME shown in elliptic cases [Rec]

Difficulty	*Possible Solutions*	*Status*
• **Grid-induced slow convergence**	One can avoid many of the following maladies by using suitable multi-grid structures (described below under "grid adaptation")	
• Large aspect ratios	Either one of the following: (1) Block (part-line or part-plane) relaxation, analyzed by mode analysis [B77]. (2) Semi coarsening [Arl], [Stages, §3.2] (often natural, since the large aspect ratio is created in the first place by semi refinements) with relaxation "semi smoothing" analysis [Stages, §2.1], [G, §3.3]. (3) Combinations of block relaxation in some directions and semi coarsening in others	TME has been shown in a variety of elliptic cases
• Expanding grids	Relaxation marching in the direction of increasing meshsize; or distributive relaxation [Njm, §6]	
• Grid adaptation	Use local patches of multigrid levels in creating any desired local refinement, aspect ratio, boundary fitting or even flow fitting (see [R88, §2.7]). Base refinement criteria on the fine-to-coarse multigrid correction (τ). Adaptation can be integrated into the λ-FMG algorithm together with proper (e.g., Reynolds-number) continuations	Introduced in [B77] and [G], but tried only for Poisson equation near singularities [Bai]

Difficulty	*Possible Solutions*	*Status*
• **Stagnation point** causing an instability in the coarse-grid corrections	Coarse-grid numerical viscosity depending on the *average* (e.g., "full-weighting") of the fine-grid numerical viscosity (not on its injected value) [BY3, §4.5]	TME shown in an example [BY3]
• **A small number of slowly-converging components** may arise in many situations, especially when the "Possible Solutions" described in many of the sections above are not *fully* implemented. That "small number" would often slowly but unboundedly incrase with decreasing meshsize	A general method to expel few slow components is by *recombining iterants*, or equivalently, using the multigrid cycle as a preconditioner for *Krylov subspace acceleration* [WO]. To inexpensively expel a larger number of slow components (without executing many mg cycles and storing many fine-grid iterants), iterants may also be recombined *at various coarse levels* of the multigrid cycle [Rec], [WO]	

References

[AJ] Jameson, A., Solution of the Euler equations for two dimensional transonic flow by a multigrid method, *Appl. Math. Comp.* **13** (1983) 327–356.

[AMGT] Brandt, A., Algebraic multigrid theory: The symmetric case, in *Preliminary Proc. Int. Multigrid Conf.*, Copper Mountain, Colorado, April 6–8, 1983; *Appl. Math. Comp.* **19** (1986) 23–56.

[Arl] Arlinger, B., Multigrid technique applied to lifting transonic flow using full potential equation, SAAB-SCANIA Rep. L-0-1 B439, 1978.

[Bai] Bai, D. and Brandt, A., Local mesh refinement multilevel techniques, *SIAM J. Sci. Stat. Comp.* **8** (1987) 109–134.

[BD] Brandt, A. and Dinar, N., Multi-grid solutions to elliptic flow problems, in *Numerical Methods for Partial Differential Equations* (Parter, S., ed.), Academic Press, New York, 1979, pp. 53–147.

[BL] Baldwin, B. and Lomax, H., Thin-layer approximation and algebraic model for separated turbulent flows, AIAA Paper 78–257, Jan. 1978.

[BM] Bussing, T.R.A. and Murman, E.M., Finite-volume method for the calculation of compressible chemically reacting flows, *AIAA J.* **26** (1988) 1070–1078.

[BY1] Brandt, A. and Yavneh, I., Inadequacy of first-order upwind difference scheme for some recirculating flows, *J. Comp. Phys.* **93** (1991) 128–143.

[BY2] Brandt, A. and Yavneh, I., On multigrid solution of high-Reynolds incompressible entering flows, *J. Comp. Phys.* **101** (1992) 151–164.

[BY3] Brandt, A. and Yavneh, I., Accelerated multigrid convergence and high-Reynolds recirculating flows, *SIAM J. Sci. Comp.* **14** (1993) 607–626.

[B73] Brandt, A., Multi-level adaptive technique (MLAT) for fast numercal solutions to boundary value problems, in *Proc. 3^{rd} Int. Conf. on Numerical Methods in Fluid Mechanics* (Cabannes, H. and Temam, R., eds.), Lecture Notes in Physics 18, Springer-Verlag, 1973, pp. 82–89.

[B77] Brandt, A., Multi-level adaptive solutions to boundary value problems, *Math. Comp.* **31** (1977) 333–390.

[D1] Brandt, A. and Diskin, B., Multigrid solvers for the non-aligned sonic flow: the constant coefficient case. To appear in *Computer and Fluids*.

[D2] Brandt, A. and Diskin, B., Multigrid for non-aligned sonic flows.

[D3] Brandt, A. and Diskin, B., Efficient multigrid solvers for the linearized transonic full-potential operator.

[Dinar] Dinar, N., Fast Methods for the Numerical Solutions of Boundary Value Problems, Ph.D. Thesis, Weizmann Institute of Science, Rehovot, 1979.

[DS] Sidilkover, D., Higher order accurate method for recovering shock location, portion of Ph.D. Thesis, Weizmann Institute of Science, Rehovot, 1989.

[DS1] Sidilkover, D., A genuinely multidimensional upwind scheme and efficient multigrid solver for the compressible Euler equations, ICASE Report 97–84, 1994.

[DS2] Sidilkover, D., Some approaches towards constructing optimally efficient multigrid solvers for the inviscid flow equations, ICASE Report 97–39. Submitted to *Comptuer and Fluids*.

[ETNA] Brandt, A., The Gauss Center Research in Scientific Computation, *Electronic Trans. Num. An.* **6** (1997).

[Dym] Brandt, A. and Dym, J., Effective boundary treatment for the biharmonic Dirichlet problem, in *Proc.* 7^{th} *Copper Mountain Conference on Multigrid Methods* (N.D. Melson *et al.*, eds.) NASA Conference Publication 3339, 1996, pp. 97–108.

[G] Either [G82] or [G84].

[G82] Brandt, A., Guide to multigrid development, in *Multigrid Methods* (Hackbusch, W. and Trottenberg, U., eds.), Springer-Verlag, 1982, pp. 220–312.

[G84] Brandt, A., *Multigrid Techniques*: *1984 Guide, with Applications to Fluid Dynamics*, 1984, 191 pages, ISBN-3-88457-081-1; GMD-Studien Nr. 85; Available from GMD-AIW, Postfach 1316, D-53731, St. Augustin 1, Germany, 1984.

[Ira] Livshits, I., Multigrid Solvers for Wave Equations, Ph.D. Thesis, Bar-Ilan University, 1995.

[K2] Kettler, R., Analysis and comparison of relaxation schemes in robust multigrid and preconditioned conjugte gradient methods, in *Multigrid Methods* (Hackbusch, W. and Trottenberg, U., eds.), Lecture Notes in Math. 960, Springer-Verlag, 1982.

[Lin] Linden, J., A multigrid method for solving the biharmonic equation on rectangular domains, Arbeitpapiere der GMD, St. Augustin, 1985.

[NESP] Brandt, A., Multi-grid solvers for non-elliptic and singular-perturbation steady-state problems, Weizmann Institute of Science, Rehovot, December 1981.

[Njm] Brandt, A., Multi-level adaptive techniques (MLAT) for singular perturbation problems, in *Numerical Analysis of Singular Perturbation Problems* (Hemker, P.W. and Miller, J.J.H., eds.), Academic Press, New York, 1979, 53–142.

[Par] Brandt, A. and Greenwald, J., Parabolic multigrid revisited, in *Multigrid Methods, III* (Hackbusch, W. and Trottenberg, U., eds.), Birkhäuser Verlag, Basel, 1991, pp. 143–154.

[Quasi] Brandt, A. and Ta'asan, S., Multigrid solutions to quasi-elliptic schemes, in *Progress and Supercomputing in Computational Fluid Dynamics* (Murman, E.M. and Abarbanel, S.S., eds.), Birkhäuser, Boston, 1985, pp. 235–255.

[Rec] Brandt, A. and Mikulinsky, V., Recombining iterants in multigrid algorithms and problems with small islands, *SIAM J. Sci. Comp.* **16** (1995) 20–28.

[RLMA] Brandt, A., Rigorous local mode analysis of multigrid, in *Preliminary Proc. 4th Copper Mountain Conf. on Multigrid Methods*, Copper Mountain, Colorado, April 1989.

[RQMA] Brandt, A., Rigorous quantitative analysis of multigrid: I. Constant coefficients two level cycle with L_2 norm, *SIAM J. Num. Anal.* **31** (1994) 1695–1730.

[RS] Ruge, U. and Stueben, K., Algebraic multigrid, in *Multigrid Methods* (McCormick, S.F., ed.), SIAM, 1987.

[RSS] Roberts, T.W., Sidilkover, D. and Swanson, R.C., Textbook multigrid efficiency for the steady Euler equations, AIAA Paper 97–1949, in 13th *Computational Fluid Dynamics Conf.*, June 1997.

[R1] Rubbert, P.E., CFD and the changing world of airplane design, AIAA Wright Brothers Lecture, Anaheim, CA, 1994.

[R2] Vatsa, V.N. and Wedan, B.W., Development of a flexible and efficient multigrid code for 3-D Navier-Stokes equations and its application to a grid-refinement study, *Computers and Fluids* **18**(4) (1990) 391–403.

[R3] Garner, P., Meredith, P. and Stoner, R., Areas of future CFD development as illustrated by transport aircraft applications, AIAA Paper 91–1528, 1991.

[R88] Brandt, A., The Weizmann Institute Research in Multilevel Computation: 1988 Report, in *Proc. 4th Copper Mountain Conf. on Multigrid Methods* (Mandel, J. et al, eds.), SIAM, 1989, pp. 13–53.

[South] South, J.C. and Brandt, A., Application of a multi-level grid method to transonic flow calculations, in *Transonic Flow Problems in Turbo Machinery* (Adam, T.C. and Platzer, M.F., eds.), Hemisphere, Washington, 1977, pp. 180–207.

[Stages] Brandt, A., Stages in developing multigrid solutions, in *Proc. 2nd Int. Congr. on Numerical Methods for Engineers* (Absi, E., Glowinski, R., Lascaux, P. and Veysseyre, H., eds.), Dunod, Paris, 1980, pp. 23–43.

[S9] Sonneveld, P., Wesseling, P. and de Zeeuw, P.M., Multigrid and conjugate gradient methods as convergence acceleration techniques, in *Multigrid Methods for Integral and Differential Equations* (Paddon, D.J. and Holstein, H., eds.), Clarendon Press, 1985.

[T1] Ta'asan, S., Canonical forms of multidimensional steady state inviscid flows, ICASE Report 93–34, 1993.

[T2] Ta'asan, S., Optimal multigrid method for inviscid flows, *Proc. European Multigrid Conf.*, EMG93, Amsterdam, July 1993.

[T3] Ta'asan, S., Canonical-variables multigrid method for steady-state Euler equations, *Proc. 14^{th} Int. Conf. on Numerical Methods for Fluid Dynamics*, Bangalore, India, 1994.

[T4] Ta'asan, S., Essentially optimal multigrid method for steady state Euler equations, *33^{rd} Aerospace Sciences Meeting and Exhibit*, AIAA Report 95–0209, Reno, NV, January 1995.

[T5] Ta'asan, S., Essentially optimal multigrid method for steady state Euler equations, *33^{rd} Aerospace Sciences Meeting and Exhibit*, AIAA Report 95–0209, Reno, NV, January 1995.

[T6] Ta'asan, S., Lectures on numerical solution of PDE (unpublished).

[WO] Washio, T. and Oosterlee, C.W., Krylov subspace acceleration for nonlinear multigrid schemes, *Elec. Trans. Num. Anal.* (1998).

[YVB] Yavneh, I., Venner, C.H. and Brandt, A., Fast multigrid solution of the advection problem with closed characteristics, *J. Comp. Phys.*, submitted.