

NASA/TM-2004-213037



Textbook Multigrid Efficiency for Leading Edge Stagnation

Boris Diskin
National Institute of Aerospace, Hampton, Virginia

James L. Thomas and Raymond E. Mineck
Langley Research Center, Hampton, Virginia

May 2004

The NASA STI Program Office ... in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) Program Office plays a key part in helping NASA maintain this important role.

The NASA STI Program Office is operated by Langley Research Center, the lead center for NASA's scientific and technical information. The NASA STI Program Office provides access to the NASA STI Database, the largest collection of aeronautical and space science STI in the world. The Program Office is also NASA's institutional mechanism for disseminating the results of its research and development activities. These results are published by NASA in the NASA STI Report Series, which includes the following report types:

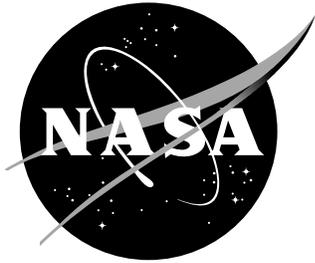
- **TECHNICAL PUBLICATION.** Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.
- **TECHNICAL MEMORANDUM.** Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- **CONTRACTOR REPORT.** Scientific and technical findings by NASA-sponsored contractors and grantees.
- **CONFERENCE PUBLICATION.** Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.
- **SPECIAL PUBLICATION.** Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.
- **TECHNICAL TRANSLATION.** English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services that complement the STI Program Office's diverse offerings include creating custom thesauri, building customized databases, organizing and publishing research results ... even providing videos.

For more information about the NASA STI Program Office, see the following:

- Access the NASA STI Program Home Page at <http://www.sti.nasa.gov>
- E-mail your question via the Internet to help@sti.nasa.gov
- Fax your question to the NASA STI Help Desk at (301) 621-0134
- Phone the NASA STI Help Desk at (301) 621-0390
- Write to:
NASA STI Help Desk
NASA Center for AeroSpace Information
7121 Standard Drive
Hanover, MD 21076-1320

NASA/TM-2004-213037



Textbook Multigrid Efficiency for Leading Edge Stagnation

Boris Diskin
National Institute of Aerospace, Hampton, Virginia

James L. Thomas and Raymond E. Mineck
Langley Research Center, Hampton, Virginia

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

May 2004

The use of trademarks or names of manufacturers in this report is for accurate reporting and does not constitute an official endorsement, either expressed or implied, of such products or manufacturers by the National Aeronautics and Space Administration.

Available from:

NASA Center for AeroSpace Information (CASI)
7121 Standard Drive
Hanover, MD 21076-1320
(301) 621-0390

National Technical Information Service (NTIS)
5285 Port Royal Road
Springfield, VA 22161-2171
(703) 605-6000

Abstract

A multigrid solver is defined as having textbook multigrid efficiency (TME) if the solutions to the governing system of equations are attained in a computational work which is a small (less than 10) multiple of the operation count in evaluating the discrete residuals. TME in solving the incompressible inviscid fluid equations is demonstrated for leading-edge stagnation flows. The contributions of this paper include (1) a special formulation of the boundary conditions near stagnation allowing convergence of the Newton iterations on coarse grids, (2) the boundary relaxation technique to facilitate relaxation and residual restriction near the boundaries, (3) a modified relaxation scheme to prevent initial error amplification, and (4) new general analysis techniques for multigrid solvers. Convergence of algebraic errors below the level of discretization errors is attained by a full multigrid (FMG) solver with one full approximation scheme (FAS) cycle per grid. Asymptotic convergence rates of the FAS cycles for the full system of flow equations are very fast, approaching those for scalar elliptic equations.

Contents

1	Introduction	2
2	Ingredients of Textbook Multigrid Efficiency	3
3	Incompressible Euler Equations	4
3.1	Primitive Formulation and Solutions	4
3.2	Pressure-Equation Formulation	7
3.3	Discrete Boundary Conditions	8
3.3.1	Boundary Conditions at Inflow Boundary	8
3.3.2	Boundary Conditions at Outflow Boundary	8
3.3.3	Boundary Conditions at Body Surface	8
4	Relaxation Scheme	9
4.1	Interior Relaxation	9
4.2	Boundary Relaxation	12
5	Analysis	13
5.1	Limitations of Local-Mode Fourier (LMF) Analysis	13
5.2	Ideal Relaxation (IR) and Ideal Coarse-Grid (ICG) Iterations	17
5.3	Examples of Analyses	19
5.3.1	Example 1: Laplace Equation	19
5.3.2	Example 2: Constant Coefficient Convection Equation	21
5.3.3	Example 3: System of Constant-Coefficient Partial Differential Equations	24
6	Numerical Tests	26
6.1	Convergence of Newton Iterations	26
6.2	Amplification of Algebraic Errors in Stagnation Region	27
6.3	Multigrid Description	30
6.4	Isotropic Grids	31
6.4.1	Convergence Rates	31
6.4.2	FMG Performance	33
6.5	Stretched Grids	34
7	Conclusions	35

1 Introduction

This paper addresses one aspect – multigrid treatment of leading-edge stagnation flows – towards the goal of attaining textbook multigrid efficiency (TME) in solving general computational fluid dynamics (CFD) equations. A multigrid method is defined [4, 7] as having TME if the solutions to the governing system of differential equations are attained in a total computational work of less than 10 minimal work units; a *minimal work unit* is defined as the operation count in evaluating the residuals of the discretized equations. Such efficiencies have been demonstrated for elliptic problems [12, 28, 30]. The compressible Euler and Navier-Stokes equations are sets of coupled nonlinear equations of mixed type that admit discontinuous solutions (shocks, slip lines, etc.) Flow- or grid-induced singularities can occur. A list of envisioned difficulties and possible solutions in attaining TME for CFD simulations is discussed elsewhere [7, 10, 27]. Compared to current solvers, there is a potential gain of several orders of magnitude in operation count reduction if TME could be attained for CFD simulations.

The equations of steady, viscous, incompressible flow in two dimensions are governed by uniformly elliptic and convection-diffusion factors. In multigrid solvers, it is often possible to separate these factors; corrections associated with the elliptic factors can be computed efficiently with multigrid while corrections associated with the convection-dominated convection-diffusion factors can be computed by downstream marching. The convergence rate of the solution is then determined by the convergence of the slower factors, which are the elliptic factors in this particular scenario. This separation, or factorization, can be achieved using several approaches, such as distributed relaxation [11] or equation reformulation [23, 26].

One such approach, the pressure-equation formulation [23], has been extended to general coordinates and implemented for lifting airfoils in inviscid flow [20–22] and viscous flow [25]. Grid-independent convergence rates were demonstrated in both situations; the convergence rates, although quite fast, were not as fast as the optimal multigrid rates for the underlying elliptic factor. Difficulties were reported near stagnation, especially for inviscid flows. Other (unpublished) simulations showed that on some coarse grids with a poor leading-edge resolution, no solutions could be attained. The research reported in this paper evolved from attempts to analyze and understand the erratic behavior of multigrid solvers for stagnation flows.

During the course of the work the following four principal difficulties associated with achieving TME for stagnation-flow problems have been identified and overcome:

1. The accuracy of the discrete solution is very sensitive to boundary condition formulations in the vicinity of stagnation point. With a poor choice of discrete boundary conditions near stagnation, the solutions may not exist on the coarse grids, even if the second order accuracy is observed on fine grids. This sensitivity is illustrated in Section 6.1. In the course of this work, we have carefully experimented with different formulations of boundary conditions at stagnation. Compact numerical closure conditions leading to sufficient solution accuracy on all uniform, including anisotropic, grids have been found and applied. Even with this formulation, convergence on stretched grids with large stretching ratios cannot be guaranteed.

2. Strong coupling between the equations near the boundaries and associated difficulties with restricting fine-grid residuals to coarser grids led to a serious slowdown of multigrid convergence rates. This difficulty has been solved by implementing the concept of boundary relaxation (Section 4.2).

3. Initial amplification of the algebraic errors in relaxation was an obstacle to achieving the TME goal of solution to the discretization accuracy in less than 10 minimal work units. A special downstream marching scheme has been developed (Section 6.2) to avoid amplification of the specific errors arising in the full multigrid (FMG) solver.

4. Lack of reliable quantitative analytical tools hampered our abilities to optimize multigrid solvers. The analysis of stagnation flows is a difficult task because the simplest model problem representative of stagnation is a two-dimensional system of variable-coefficient equations. Simplified (one-dimensional or constant-coefficient) analyses proved misleading, predicting instability of the formulation and divergence of multigrid iterations that contradicted the existing practical experience. The limitations of the classical local-mode Fourier (LMF) analysis in application to stagnation-flow problems have been understood; certain restrictions have been applied (Section 5.1). With the imposed restrictions, several important areas in the stagnation flow-field cannot be addressed through the LMF analysis. New, more general, analysis methods for multigrid solutions of stagnation flows have been developed (Section 5.2). These new analysis methods replace one of the complimentary parts of a multigrid cycle, relaxation or coarse-grid correction, with an idealized substitute. The methods allow multigrid performance to be probed, similarly to that using the LMF analysis, but with far fewer limitations.

Following this introduction, the paper is organized as follows: Ingredients of TME are discussed in Section 2. The governing equations, discretization, and boundary conditions for the computational simulations are presented in Section 3. The relaxation scheme, including the interior and boundary parts, is discussed in Section 4. In Section 5, analysis methods, including the LMF analysis and the two new methods developed during the course of this work, are described with examples in application to pertinent model problems. In Section 6, numerical results confirming attainment of TME as well as fast asymptotic convergence rates for stagnation flows in different regimes are shown. Conclusions are given in Section 7. Examples of simplified, one-dimensional and constant-coefficient, analyses are discussed in Appendices A, B, and C.

2 Ingredients of Textbook Multigrid Efficiency

The basic framework for nonlinear TME solvers is full multigrid (FMG) algorithms [1, 12, 28]. An FMG- k algorithm starts the solution process on a very coarse grid, where the computational cost of solution is negligible. The coarse-grid solution is then interpolated to the next finer grid to form an initial approximation. On the current grid, k multigrid full approximation scheme (FAS) cycles [1, 12, 28] (ideally $k = 1$) are performed to obtain an improved solution. This process of obtaining an initial approximate solution from a coarser grid followed by FAS cycling continues to finer grids until the solution on the target finest grid is achieved.

The objective of FMG- k algorithms (and TME methods in particular) is an accurate discrete approximation to the solution of the differential equations. In this context, the solution to a discrete problem is considered accurate if its algebraic error is below the level of the discretization error. The algebraic error is defined as the difference between the exact and approximate solutions of the discrete problem. The discretization error is defined as the difference between the exact solutions of the discrete and differential problems.

A major part of a multigrid cycle is relaxation; the role of relaxation in a multigrid cycle is to reduce the error components that cannot be represented in the coarser grids. If the target discretization is strongly *h-elliptic* (or *semi-h-elliptic*) one can design a local (or block-wise) relaxation procedure efficiently reducing all high-frequency error components. By definition [2–4, 8, 28], a discrete constant-coefficient scalar (not necessarily elliptic) operator $L[u]$ possesses a good measure of *h-ellipticity* if the absolute value of its symbol

$$|L(\bar{\theta})| = |e^{-i(\bar{\theta} \cdot \mathbf{j})} L[e^{i(\bar{\theta} \cdot \mathbf{j})}]| \quad (1)$$

is well separated from zero for all high-frequency Fourier modes. Here $\mathbf{j} = (j_x, j_y, j_z)$ are the grid indexes and $\bar{\theta} = (\theta_x, \theta_y, \theta_z)$, $0 \leq |\theta_x|, |\theta_y|, |\theta_z| \leq \pi$ are normalized Fourier frequencies.

High-frequency Fourier modes are the modes satisfying $\max(|\theta_x|, |\theta_y|, |\theta_z|) \geq \frac{\pi}{2}$. For systems of constant-coefficient equations, the measure of h -ellipticity is defined to be the measure of the symbol determinant. For variable-coefficient and nonlinear problems, a measure of h -ellipticity can be computed locally as the measure for a problem with constant (frozen) coefficients locally approximating the target problem. A good measure of h -ellipticity implies that error components with large variations on the scale of the mesh size produce much larger residuals than little varying (smooth) error components of the same amplitude.

Classical multigrid methods, which proved to be efficient for elliptic problems, separate the treatment of oscillatory and smooth error components. The former are efficiently reduced in relaxation; the latter are well approximated on coarse grids and, hence, eliminated through the coarse-grid correction. The difficulties associated with extending TME for solution of the Navier-Stokes equations relate to the fact that these equations are a system of coupled nonlinear equations that is not fully elliptic, even for subsonic Mach numbers, but contains hyperbolic partitions. The efficiency of classical multigrid methods severely degrades for nonelliptic problems because some smooth *characteristic error components* cannot be adequately approximated on coarse grids [2, 9, 11, 14, 28]. These characteristic components are much smoother in the characteristic directions than in other directions.

To be efficient, a multigrid solver for nonelliptic problems has to adequately address three types of errors: (1) high-frequency error components, (2) uniformly smooth error components, (3) characteristic error components. For h -elliptic discretizations, an efficient interior relaxation scheme can usually be found. A few sweeps of such a relaxation can reduce high-frequency error significantly, by an order of magnitude. Coarse-grid correction is usually efficient for uniformly smooth error components. An effective reduction of characteristic error components can be achieved either by designing a relaxation scheme to reduce not only high-frequency but smooth error components as well (which can be done in many cases by downstream ordering of relaxation steps [2, 11]) or by adjusting coarse-grid operators for a better characteristic-component approximation [9, 13, 14].

3 Incompressible Euler Equations

3.1 Primitive Formulation and Solutions

The set of two-dimensional nonconservative incompressible Euler equations is defined as

$$\begin{aligned} uu_x + vu_y + p_x &= f_u, \\ uv_x + vv_y + p_y &= f_v, \\ u_x + v_y &= f_c, \end{aligned} \tag{2}$$

where variables u, v , and p represent the x -directional velocity, the y -directional velocity, and the pressure, respectively. The first and second equations are referred as the u - and v -momentum equations, respectively; the third equation is the continuity equation. A standard set of boundary conditions for this formulation is the velocity components specified at an inflow boundary, the pressure specified at an outflow boundary, and the tangency condition (the normal velocity component is set to zero) defined at a body surface. For several exterior-flow problems, exact analytical solutions of the incompressible Euler equations are known.

Plane-stagnation flow is a flow against a solid plane located at $x = 0$. The exact solution has a form

$$\begin{aligned} u &= -ax, \\ v &= ay, \\ p &= -a^2 \frac{x^2 + y^2}{2}, \end{aligned} \tag{3}$$

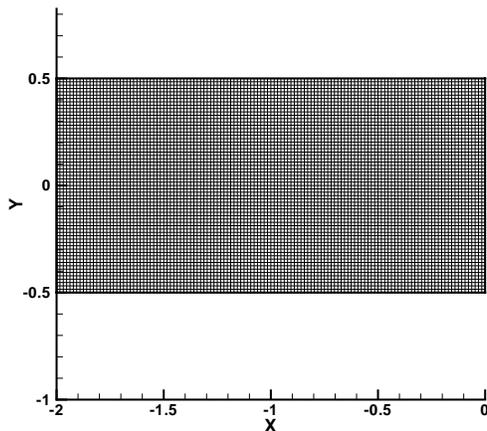


Figure 1. Computational domain for a plane-stagnation flow; 129×65 isotropic grid.

where a is an arbitrary constant. The solution is unbounded at the upstream infinity. Although very simple, the plane-stagnation solution characterizes the behavior of general flows near stagnation. The solution describes *deep stagnation* flows, i.e., flows in (small) neighborhoods of the stagnation point. For *regular stagnation* flows in general geometries, the deep-stagnation regions comprise small parts of computational domains. A typical computational domain for the stagnation plane flow is shown in Figure 1. The inflow boundary is defined at the left vertical edge, the outflow boundary is at the horizontal edges, and the tangency condition is defined at the plane surface $x = 0$.

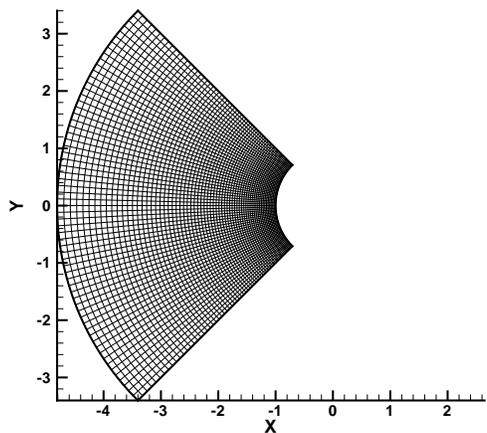


Figure 2. Computational domain for flow around a cylinder; $\theta_{\max} = 45^\circ$, $\theta_{\min} = -45^\circ$; 65×65 isotropic grid.

Cylinder flow is a flow around a cylinder of unit radius centered at the origin $(x, y) = (0, 0)$. The exact solution, with the free stream at infinity characterized by $u_\infty = 1$, $v_\infty = 0$, $p_\infty = \text{const}$, is defined as

$$\begin{aligned}
u &= 1 + \frac{2 \sin^2 \theta - 1}{r^2} + \partial_y \psi, \\
v &= 2 \frac{\sin \theta \cos \theta}{r^2} - \partial_x \psi, \\
p &= p_\infty + 0.5(1 - u^2 - v^2),
\end{aligned} \tag{4}$$

where (r, θ) are the polar coordinates, $r^2 = x^2 + y^2$, $\tan \theta = -\frac{y}{x}$, and $\psi = -\frac{C}{2\pi} \ln(r)$ is the stream function with C being a constant. Throughout the course of this paper $C = 0$.

A typical computational domain with an isotropic orthogonal grid for a regular cylinder flow is shown in Figure 2. The inflow boundary is defined at the external arc corresponding to $r_{\max} \approx 5$; the outflow boundary is defined at the rays corresponding to $\theta = \pm 45^\circ$; the tangency is defined at the inner circle $r_{\min} = 1$. A deep-stagnation cylinder flow is characterized by small positive values of $\theta_{\max}, (-\theta_{\min})$, and correspondingly small extent $(r_{\max} - r_{\min})$.

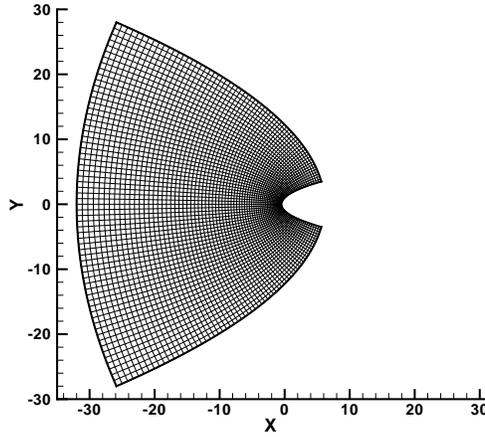


Figure 3. Computational domain for flow around a parabola; $\eta_{\max}/\eta_{\min} = 8$; 65×65 isotropic grid.

Parabola flow is a flow around a parabola with freestream values taken as $u_\infty = 1, v_\infty = 0, p_\infty = \text{const}$. The exact solution is well known and can be found through a conformal transformation $x + iy = z^2$, in which the z -plane is defined as $z = \xi + i\eta$. The computational domain in the physical plane is determined by a rectangular region in the z -plane as

$$(\xi, \eta) \in [\xi_{\min}, \xi_{\max}] \times [\eta_{\min}, \eta_{\max}], \tag{5}$$

where an isotropic Cartesian grid in the z -plane yields an orthogonal and isotropic grid in the physical domain. The surface of the parabola is defined by the line $\eta_{\min} = 1/\sqrt{2}$. The leading edge region is described by a circle of radius $1/2$ centered at the origin $(x, y) = (0, 0)$. The physical point corresponding to a point in the z -plane is given as

$$\begin{aligned}
x &= \xi^2 - \eta^2, \\
y &= 2\xi\eta.
\end{aligned}$$

The exact solution is given in terms of (ξ, η) as

$$\begin{aligned}
u &= 1 - \eta_{\min}\eta/(\xi^2 + \eta^2), \\
v &= \eta_{\min}\xi/(\xi^2 + \eta^2), \\
p &= p_\infty + (1 - u^2 - v^2)/2.
\end{aligned} \tag{6}$$

The upstream extent of the computational domain in the z -plane is determined by the single parameter η_{\max}/η_{\min} . An isotropic, orthogonal grid corresponding to $\eta_{\max}/\eta_{\min} = 8$ is shown in Figure 3, extending upstream approximately sixty leading edge radii. Smaller values of this parameter correspond to domains in closer vicinities of the leading edge. The value $\eta_{\max}/\eta_{\min} = 1.25$ corresponds to the domain extending 25/16 leading edge radii upstream. If extended a similar distance along the parabola curve, the computational domain is in a deep-stagnation region and can be characterized closely by the plane-stagnation solution. For both the cylinder and the parabola flows, the solution asymptotes to the freestream values at upstream infinity.

Throughout discussions in the paper, a rectangular computational domain with Cartesian isotropic grids corresponding to the plane-stagnation flow (see Figure 1) is assumed. For numerical tests reported in Section 6, body-fitted curvilinear coordinate systems (see Figures 2 and 3) have also been used. Practical experience showed that the deep-stagnation regions, which are well described by the plane-stagnation solutions, represent the most difficult part of the flow field. The solution methods found efficient for the plane-stagnation flows in the rectangular geometry proved to be fully extendible to more complicated geometries. Numerical implementations of curvilinear grids were straightforward and did not introduce any significant complications to the Cartesian case.

3.2 Pressure-Equation Formulation

The pressure-equation formulation is obtained from (2) by replacing the continuity equation with an equation for the pressure

$$\begin{aligned} uu_x + vu_y + p_x &= f_u, \\ uv_x + vv_y + p_y &= f_v, \\ (u_x)^2 + 2u_yv_x + (v_y)^2 + \Delta p &= f_p. \end{aligned} \quad (7)$$

The major advantage of this formulation is an easily available h -elliptic collocated-grid discretization. Besides the standard set of boundary conditions (the velocities given at the inflow, the pressure given at the outflow, and the tangency condition given at the body surface), the pressure-equation formulation (7) requires an additional boundary condition at the inflow boundary. To ensure the same solution for the formulations (2) and (7), this additional boundary condition is taken as the continuity equation enforced at the inflow boundary. This boundary condition can be implemented in a number of ways; a convenient choice is through the Neumann condition for the pressure derivative obtained by manipulating the normal (to the boundary) momentum equation and the continuity equation at the inflow boundary:

$$p_x = f_u - uu_x - vu_y = f_u - u(f_c - v_y) - vu_y, \quad (8)$$

where $f_u - u(f_c - v_y) - vu_y$ is a known quantity at the inflow boundary.

The second-order accurate interior discretization, $\mathbf{R}(\mathbf{q})$, of (7) is defined as upwind differencing for the convection operators of the momentum equations and central difference approximations for the pressure gradient; in the pressure equation, the standard h -elliptic five-point discretization, Δ^h , is applied for the Laplacian of the pressure, and central differencing is applied for the velocity derivatives.

$$\mathbf{R}(\mathbf{q}) \equiv \begin{pmatrix} u\partial_x^u u + v\partial_y^u u + \partial_x^c p \\ u\partial_x^u v + v\partial_y^u v + \partial_y^c p \\ (\partial_x^c u)^2 + 2(\partial_y^c u)(\partial_x^c v) + (\partial_y^c v)^2 + \Delta^h p \end{pmatrix} = \begin{pmatrix} f_u \\ f_v \\ f_p \end{pmatrix}, \quad (9)$$

where $\mathbf{q} = (u, v, p)^T$, and superscripts u and c denote second-order accurate upwind and central differences for the first derivatives, respectively.

3.3 Discrete Boundary Conditions

There are two types of discrete boundary conditions. Following [17], discrete approximations to the boundary conditions of the differential equations are called *physical boundary conditions*. For the pressure-equation formulation, the set of physical boundary conditions is defined as the velocity components and the normal pressure derivative given at the inflow boundary, the pressure given at the outflow boundary, and the tangency condition at the body surface.

The boundary conditions required to complete the discrete formulation are called *numerical boundary conditions* or *numerical closure equations*. Numerical boundary conditions are usually implemented as modified discretizations (different from the discretizations in the interior) of the target differential equations. The number of numerical boundary conditions is determined by the stencils involved in the interior discretization. A detailed discussion about relations between the discretization stencils in the interior and the discrete boundary conditions can be found in [16]. For the discretization (9), two numerical closure equations are required at the inflow boundary and one numerical closure equation is required at the outflow boundary and the body surface.

3.3.1 Boundary Conditions at Inflow Boundary

Discretization of the Dirichlet conditions of given velocity components at the inflow boundary is straightforward. The Neumann conditions of the given normal pressure derivative can be implemented explicitly by employing a one-sided difference approximation for the normal derivative. With this implementation, the discrete pressure equation is not defined at the boundary. An alternative implementation uses two discrete equations for pressure – a central approximation for the normal pressure derivative and a modified pressure equation – defined at the same boundary point; the pressure value located in the exterior of the computational domain is eliminated from these two equations. In the modified pressure equation, some velocity derivatives, $u_y, v_y, u_x = f_c - v_y$, are computed analytically from the Dirichlet conditions for the velocities; the derivative v_x is discretized with a one-sided stencil. An advantage of the latter implementation is a more compact stencil obtained for differencing the pressure at the boundary. From our experience, the choice of implementation of the Neumann condition at the inflow boundary is not critical

The numerical boundary conditions at the inflow boundary consist of two modified momentum equations defined at the interior points next to the boundary. The modifications involve replacement of second-order accurate upwind difference approximations of the convection operators with more compact differencing, e.g., central or upwind first-order accurate approximations.

3.3.2 Boundary Conditions at Outflow Boundary

The Dirichlet condition of the given pressure is discretized straightforwardly. The numerical boundary condition is a modified normal-momentum equation at the outflow boundary. In the modified equation, the central approximation to the pressure derivative is replaced with a one-sided approximation.

3.3.3 Boundary Conditions at Body Surface

At the body surface, the physical tangency condition is implemented as explicit assignment of the normal velocity component to zero at the surface. Other quantities at the surface are computed from the tangential-momentum equation and a numerical closure equation.

Away from stagnation, the numerical closure equation can be implemented as a one-sided discretization of the normal-momentum equation, appearing as below for plane stagnation

($u \equiv 0$ along the plane surface $x = 0$)

$$\partial_x^u p = f_u - uu_x - vu_y = f_u. \quad (10)$$

Analogously to the inflow boundary, an alternative numerical closure equation can be derived from the joint formulation of the normal pressure derivative equation and a modified pressure equation at the boundary, leading to a one-sided stencil. For plane stagnation,

$$(\partial_x^u u)^2 + (\partial_y^c v)^2 + \partial_{yy}^h p + \frac{2}{h_x}(f_u - \bar{\partial}_x^u p) = f_p, \quad (11)$$

where ∂_x^u is an upwind differencing that may be first or second order accurate, ∂_y^c and ∂_{yy}^h are central second-order differences approximating the first and second derivatives, respectively, and $\bar{\partial}_x^u$ is a two-point one sided approximation to the first derivative; the term $\frac{2}{h_x}(f_u - \bar{\partial}_x^u p)$ is a second-order accurate approximation to $\partial_{xx} p$ under the condition $\partial_x^c p = f_u$.

Near stagnation, the choice of numerical closure equation is critical for stability, accuracy and even existence of the solution. As a particular example, for the first-order discrete formulation (9) with the physical tangency condition and a two-point one-sided differencing of (10) applied as the numerical boundary condition at stagnation, the Jacobian of the formulation becomes singular, while implementing either a second-order differencing of (10) or the formulation (11) leads to a well behaved Jacobian. Moreover, as numerical tests show in Section 6.1, in the presence of a curved geometry, even a non-singular Jacobian does not guarantee convergence of the Newton iterations on coarse grids.

We consider two types of numerical closure conditions along the tangency boundary, referred to as non-compact and compact formulations. Both use the formulation (11) everywhere along the boundary except that the compact formulation uses two-point, instead of three-point, one-sided differencing for the $(\partial_x^u u)$ term in (11). Also, at the point closest to stagnation, both numerical closures do not use the nonconservative tangential-momentum equation; instead, a central (finite-volume type) discretization for the quasi-conservative tangential-momentum equation is implemented, given in Cartesian coordinates as

$$(uv)_x + (v^2)_y - v(u_x + v_y) + p_y = f_v. \quad (12)$$

The non-compact formulation uses this quasi-conservative form everywhere on the body surface with one-sided differencing of $(uv)_x$ and u_x , upwinding of $(v^2)_y$, and central differencing of v_y and p_y . The compact formulation uses the nonconservative tangential-momentum equation away from stagnation. At stagnation, the compact formulation applies the quasi-conservative tangential-momentum equation at a location displaced half a mesh size into the interior. Both formulations yield second-order accurate solutions. Numerical tests in Section 6.1 demonstrate that convergence of Newton iterations on coarse grids is much more robust for the compact formulation presented above. A simpler, alternate formulation is to apply the nonconservative form of both the normal and tangential momentum equations everywhere along the body, displaced half a mesh size into the interior. No results are shown for this formulation, but in numerical tests, it has proved as effective as the compact formulation.

4 Relaxation Scheme

4.1 Interior Relaxation

Development of efficient relaxation schemes for nonlinear problems starts from the Newton scheme. A solution update, $\delta \mathbf{q}$, introduced by a Newton iteration of the nonlinear equation

(9) is computed as

$$\frac{\partial \mathbf{R}}{\partial \mathbf{q}} \delta \mathbf{q} = \begin{pmatrix} f_u \\ f_v \\ f_p \end{pmatrix} - \mathbf{R}(\mathbf{q}), \quad \mathbf{q}_{\text{new}} = \mathbf{q} + \delta \mathbf{q}. \quad (13)$$

$\frac{\partial \mathbf{R}}{\partial \mathbf{q}}$ is the full Newton linearization of the operator (9) around a solution $\mathbf{q} = (u, v, p)^T$,

$$\frac{\partial \mathbf{R}}{\partial \mathbf{q}} \equiv \begin{pmatrix} Q^u + (\partial_x^u u) & (\partial_y^u u) & \partial_x^c \\ (\partial_x^u v) & Q^u + (\partial_y^u v) & \partial_y^c \\ 2(\partial_x^c u) \partial_x^c + 2(\partial_x^c v) \partial_y^c & 2(\partial_y^c v) \partial_y^c + 2(\partial_y^c u) \partial_x^c & \Delta^h \end{pmatrix}, \quad (14)$$

where $Q^u = u \partial_x^u + v \partial_y^u$ is an upwind discretization of the linear convection operator.

Solving the full Newton linearization at each relaxation step is too costly for practical applications. To reduce the computational cost of relaxation without compromising efficiency, one can consider a principal linearization. The principal linearization of a scalar equation contains the linearization terms that make major contributions to the residual per a unit change in the unknown variable. A unit change is defined as a perturbation of the solution value at one grid point. The principal terms generally depend on the scale, or mesh size, of interest. For example, the discretized highest derivative terms are principal on grids with small enough mesh size. For a discretized system of differential equations, the principal terms are terms that contribute to the principal part of the system determinant. If the underlying solution, \mathbf{q} , is *smooth* and *nondegenerate* (i.e., the magnitude of the differences between solution values in neighboring grid points is less than the local solution magnitude), one can obtain a local constant-coefficient approximation to (14) by freezing the coefficients at some point in the computational domain

$$\begin{pmatrix} \bar{u} \partial_x^u + \bar{v} \partial_y^u + \overline{(\partial_x^u u)} & \overline{(\partial_y^u u)} & \partial_x^c \\ \overline{(\partial_x^u v)} & \bar{u} \partial_x^u + \bar{v} \partial_y^u + \overline{(\partial_y^u v)} & \partial_y^c \\ \overline{2(\partial_x^c u) \partial_x^c + 2(\partial_x^c v) \partial_y^c} & \overline{2(\partial_y^c v) \partial_y^c + 2(\partial_y^c u) \partial_x^c} & \Delta^h \end{pmatrix}, \quad (15)$$

where bars denote the frozen coefficients. The principal part of the determinant of the matrix (15) is

$$(\bar{Q}^u)^2 \Delta^h, \quad (16)$$

where $\bar{Q}^u = \bar{u} \partial_x^u + \bar{v} \partial_y^u$. For stable discretizations of \bar{Q}^u , e.g., first or second order upwind discretizations, the operator (16) is h -elliptic if the characteristic of \bar{Q}^u is not aligned with the grid, or semi- h -elliptic in the case of alignment. Away from stagnation on fine grids, the solution of (9) is expected to be smooth and nondegenerate; thus, the principal linearization of the nonlinear operator (9) is a diagonal matrix

$$\tilde{\mathbf{L}} \equiv \begin{pmatrix} Q^u & 0 & 0 \\ 0 & Q^u & 0 \\ 0 & 0 & \Delta^h \end{pmatrix}. \quad (17)$$

The property of (semi-) h -ellipticity of (16) indicates that, away from stagnation on fine enough grids, there exists an efficient error smoother for (9) that is based on the principal linearization (17).

A discretization scheme for a system of nonlinear equations with the principal-linearization determinant represented as a product of simple scalar factors is called a *factorizable scheme*. According to this definition, away from stagnation, the discrete scheme (9) is factorizable. Factorizability of a discrete scheme significantly simplifies the task of devising an efficient relaxation procedure for relaxing the system of equations; this task is reduced to a number of much simpler tasks — designing efficient relaxation schemes for scalar factors of the principal-linearization determinant.

One of the main goals of this paper is to demonstrate that, for a factorizable discrete scheme, with proper treatment of special regions, it is possible to develop a multigrid solver that will be as efficient (in terms of convergence per a minimal work unit) as the best available solvers for the discrete scalar factors constituting the principal linearization determinant. With a relaxation scheme employing a downstream order for the convection operators in the momentum equations, the efficiency of such a solver for the system (9) is expected to be similar to the efficiency of the corresponding multigrid solver for the scalar Laplace equation.

Formally the interior relaxation scheme can be described similar to the Newton iterations (13) as

$$\mathbf{L}\delta\mathbf{q} = \begin{pmatrix} f_u \\ f_v \\ f_p \end{pmatrix} - \mathbf{R}(\mathbf{q}), \quad \mathbf{q}_{\text{new}} = \mathbf{q} + \delta\mathbf{q}. \quad (18)$$

The linear operator \mathbf{L} inverted in relaxation is called the driver operator to distinguish it from the target operator, $\mathbf{R}(\mathbf{q})$, used to compute residuals. The driver operator is based on the principal linearization operator (17) with the Laplace operator usually to be relaxed rather than solved. A relaxation scheme with the driver (17) is an efficient smoother on fine enough grids. It may not, however, guarantee stability for smooth error components. For smooth errors, the subprincipal terms of the discrete equations might be as important as the principal terms. In some cases, small smooth-error instability may be tolerated, assuming efficient reduction of smooth errors in coarse-grid correction with only a few relaxation sweeps to be done per a coarse-grid correction step in a multigrid cycle. In general, however, relaxation stability is very important for robustness of multigrid solvers. To improve relaxation stability for smooth components, one can add subprincipal terms to the principal linearization matrix. Subprincipal terms do not affect the smoothing properties of relaxation, but can be very beneficial for improving stability.

One possible driver operator, \mathbf{L}_s , for the interior relaxation of (9), includes all the subprincipal terms of the momentum equations:

$$\mathbf{L}_s = \begin{pmatrix} Q^u + \partial_x^u u & (\partial_y^u u) & \partial_x^c \\ (\partial_x^u v) & Q^u + \partial_y^u v & \partial_y^c \\ 0 & 0 & \Delta^h \end{pmatrix}. \quad (19)$$

This standard driver is called the S-driver to distinguish it from the diagonally dominant driver (D-driver) that proved to be more efficient for FMG applications, preventing initial amplification of algebraic errors in deep-stagnation regions (see numerical tests in Section 6.2). The D-driver is defined as

$$\mathbf{L}_d = \begin{pmatrix} Q^u + |\partial_x^u u| & (\partial_y^u u) & \partial_x^c \\ (\partial_x^u v) & Q^u + |\partial_y^u v| & \partial_y^c \\ 0 & 0 & \Delta^h \end{pmatrix}. \quad (20)$$

With either of these drivers, a relaxation sweep in the interior can be performed in the following order: (1) relaxation of the Laplace equation updating the pressure field; (2) simultaneous downstream marching of the momentum equations updating the velocity field. In special regions, where discretization may switch stencils, (e.g., near the boundaries and discontinuities), or the assumptions of smooth and nondegenerate solutions are violated, (e.g., near stagnation), all terms of the matrix (14) may become principal; then, the driver matrix \mathbf{L} coincides with (14). In these regions all the equations are coupled, and, therefore, the interior relaxation must be complemented with a boundary relaxation procedure.

4.2 Boundary Relaxation

Boundary relaxation represents a general, very important, and often overlooked part of an efficient multigrid solver. Here the term *boundary relaxation* denotes any special procedure performed in some vicinity of the boundaries. There are at least two ways to implement boundary relaxation: (1) direct solution of coupled discrete equations and (2) relaxation passes throughout the boundary relaxation region, possibly with a relaxation scheme different from the scheme in the interior. To avoid uncertainty about sensitivity of multigrid convergence to remnants of residuals in the near-stagnation region, we opted to the direct-solution implementation of the boundary relaxation. In any event, the additional work involved in the procedure is insignificant because the number of grid points (and, therefore, the number of discrete equations) in the vicinity of the boundaries is assumed to be small comparing with the total number of grid points in the entire computational domain. In the multigrid solver for (9) described in this paper, we use a direct solution of all the coupled discrete equations centered within a certain distance (typically, a few mesh sizes) from the boundary. An empirical rule we follow in defining the depth of the boundary relaxation region is to solve simultaneously all the discrete (physical and numerical) boundary conditions as well as all the interior equations that share data with the boundary condition equations.

The first motivation for the use of boundary relaxation is quite obvious. The discretized boundary condition equations are different from the discrete equations in the interior. Combinations of different discrete equations may easily result in a local coupling between the equations that can be resolved only by simultaneous solution. A simple example of such a coupling is a second-order upwind discretization of the convection operator in the interior and a central discretization near the inflow boundary. To start downstream marching in the interior, one has to simultaneously solve the discrete equations defined at the two interior points next to the inflow boundary. For discrete systems of equations, there are instances when the interior relaxation cannot be performed near the boundaries. In particular, with a first-order upwind discretization of the convection operator, the downstream marching of the u -momentum equation linearized around a plane-stagnation solution (3) fails at the grid point $(x, y) = (-h_x, 0)$ next to stagnation. At this point the triangular matrix of the linearized convection operator degenerates having zero at the main diagonal. Some sort of boundary relaxation is often required to provide data, otherwise unavailable, for completing the interior relaxation. This type of boundary relaxation should be considered as a part of the relaxation scheme.

Another, less obvious, but not less important, motivation comes from multigrid applications. Inaccurate residual transfers in the vicinity of the boundaries may significantly hurt the multigrid efficiency. A combination of different discrete equations and a nontrivial geometry may create a situation where the correct weights in a residual restriction operator may be very difficult (or even impossible) to compute. A general approach to avoid such difficulties is to apply boundary relaxation, reducing the size of residuals near the boundary well below the typical size of the interior residuals. With zero (or very small) residuals near the boundary, an accurate residual transfer is no longer important. This type of boundary relaxation may be viewed as a part of the coarse-grid correction scheme. From this standpoint, boundary relaxation reducing the residuals of near-boundary equations may be performed just once, before restricting residuals to a coarser grid, regardless of the total number of relaxation sweeps in a multigrid cycle. Numerical tests performed in Section 5.3 illustrate the significance of boundary relaxation for multigrid efficiency.

A formal definition of the boundary relaxation procedure used in our multigrid solver is solution of the problem (13) over a specified region adjacent to the boundary. A critically important issue is the choice of the data and equations used at the interface between the boundary relaxation region and the interior. Arbitrarily chosen interface conditions might lead to an ill-posed formulation and trigger instability of the entire relaxation process.

Usually, the correct interface conditions can be derived similar to the conditions applied at the actual boundaries of the computational domain.

For the boundary relaxation region near the inflow, the interface is similar to the outflow boundary. The interior solution is usually required to provide the pressure values at the interface; two momentum equations are defined at this interface. In the numerical tests reported in this paper, the inflow boundary relaxation has been implemented as a simultaneous solution of all the discrete equations defined within a block with all necessary downstream data provided by the interior solution. Considering the boundary relaxation near the leading-edge surface and (if needed) near the outflow boundary, the interface is similar to the inflow boundary. The pressure and velocity data are provided by the interior solution and no equations are defined at the interface.

5 Analysis

The analysis of stagnation flows proved to be much more difficult task than anticipated because the simplest model problem representative of stagnation is a two-dimensional system of variable-coefficient equations. In this section, we discuss the limitations of the local-mode Fourier (LMF) analysis in applications to flow problems. Other simplified (one-dimensional or constant-coefficient) analyses as well as more details on the LMF analysis are given in Appendices A – C. The results of these analyses may be misleading, predicting instability of the formulation and divergence of multigrid iterations that contradicted the existing practical experience. Alternative, very general, quantitative analysis methods for multigrid solutions are introduced and tested.

5.1 Limitations of Local-Mode Fourier (LMF) Analysis

The classical LMF analysis applied to a variable-coefficient problem is defined in the following three steps:

1. Form a set of constant-coefficient approximations.
2. Analyze each constant-coefficient problem on the infinite (or periodic) domain.
3. Take the final estimate of the analysis as the worst estimate among all the constant-coefficient problems.

The set of constant-coefficient approximations can be derived by assigning a constant-coefficient problem for each *possible combination* of the variable-coefficient values. It is important to avoid impossible combinations of coefficients. A *possible combination* is a combination that corresponds to freezing the variable coefficients at some point in the computational domain. In applying a constant-coefficient analysis to a variable-coefficient problem, it is a mistake to allow the coefficients to be changed independently. Allowing impossible combinations of coefficients brings in additional computations that are not needed, at least; at worst, it may result in completely misleading estimates.

For many years, the local-mode Fourier (LMF) analysis has been the major practical tool for quantitative evaluation of multigrid convergence properties. The analysis is not expected to predict the convergence rate of each particular iteration (such rates may be strongly dependent on the initial approximation); rather it provides quantitative estimates for extreme cases such as the worst possible amplification factor or the asymptotic convergence rate after many iterations. The most popular multigrid applications of the LMF analysis are estimating the smoothing factor of a relaxation scheme and the amplification factor of a two-grid cycle.

For estimating the smoothing factor, the LMF analysis computes the relaxation symbol, \mathbf{S}_h , which is a square matrix parameterized by the Fourier frequencies. The dimension of this matrix depends on the number of Fourier modes coupled in relaxation. Typically, the smoothing factor is computed as the spectral radius of the product $\mathbf{I}\mathbf{S}_h$, where \mathbf{I} is the indicator of high-frequency modes, i.e., \mathbf{I} is a diagonal matrix that zeroes out all the contributions to the smooth Fourier modes and leaves unchanged all the contributions to the high-frequency modes. High-frequency modes are defined according to the coarsening strategy. The details of the smoothing analysis can be found in the literature [4, 6, 24, 28, 29]

A symbol of a two-grid (ν_1, ν_2) -cycle, \mathbf{M}_h , is a square amplification matrix, which is computed as

$$\mathbf{M}_h = \mathbf{S}_h^{\nu_2} \left(\mathbf{E}_h - \mathbf{P}_H^h \mathbf{L}_H^{-1} \mathbf{R}_h^H \mathbf{L}_h \right) \mathbf{S}_h^{\nu_1}, \quad (21)$$

where \mathbf{S}_h is the relaxation symbol, \mathbf{P}_H^h and \mathbf{R}_h^H are symbols of the prolongation and restriction operators, respectively, \mathbf{E}_h is the symbol of the identity operator, and \mathbf{L}_h and \mathbf{L}_H are symbols of the fine and coarse-grid operators, respectively. Parameters ν_1 and ν_2 are the numbers of relaxation sweeps performed before and after coarse-grid correction. See [6, 28] for detailed discussions on the nature and properties of all these symbols. The asymptotic convergence is estimated as the maximal (over all Fourier frequencies) spectral radius of \mathbf{M}_h . Different norms of \mathbf{M}_h may be computed to estimate convergence in a separate cycle. Numerous successful applications of this LMF analysis have been reported in the literature [4, 28, 30, 31].

Certain limitations of the LMF analysis have also been known for a long time. For a discretized constant-coefficient differential operator, the analysis assumes an infinite (or periodic) domain. It guarantees accuracy only for a Fourier mode satisfying the following *four locality conditions*:

1. The Fourier symbol of the discrete operator does not degenerate for this mode. This condition requires from the discrete operator some measure of ellipticity for the mode. The size of the computational domains at which one can observe the convergence predicted by the LMF analysis depends on the symbol: the smaller the absolute value of the symbol (symbol determinant, for systems) the larger the required computational domain.
2. The typical wave lengths of the mode are much shorter than the corresponding scales of the computational domain. An interpretation of this condition suggests existence of a rectangular subdomain at which the considered mode satisfies the periodic boundary conditions. For characteristic modes, the condition is applied only to the cross-characteristic oscillations. Recall, a Fourier mode is called characteristic for a non-elliptic operator, if it is much smoother in the characteristic direction than in other directions.
3. The mode is not significantly affected by the boundary conditions.
4. The mode is not coupled with nonlocal modes.

Note, that the estimates of the LMF analysis may be accurate for modes that violate some (or all) of the locality conditions; but, then, additional arguments are required to justify such estimates.

The first two and the fourth locality conditions can be directly checked, if a Fourier mode, a discrete operator, and a computational grid are given. The third condition may be enforced algorithmically, through the iterations that use only local operations, e.g., point Jacobi relaxation. In this case, the condition is satisfied for all modes. In iterations with global procedures, such as downstream marching or coarse-grid correction, the third condition is satisfied for the Fourier modes with relatively short typical wave lengths. For characteristic

modes, the third condition implies relatively high-frequency cross-characteristic oscillations. A characteristic mode can be considered as a cross-characteristic wave traveling along the characteristic. One can interpret the third condition as a restriction on the distance that the wave can travel in the discrete solution before its amplitude and/or phase are substantially changed. This distance is called the penetration distance [15, 27] (also termed survival distance in [11]) of the cross-characteristic wave. This penetration distance is determined by the frequency of the wave and by the cross-characteristic interactions (e.g., dissipation and dispersion) in the discrete operator. To satisfy the third locality condition, the penetration distance of the considered Fourier mode in the driver operator should be shorter than (1) the characteristic size of the computational domain and (2) the penetration distance of this mode in the target operator.

The infinite grid assumption makes the LMF analysis intrinsically incapable to predict the effects of boundary conditions. On finite computational domains, the choice of discrete boundary conditions and the way these conditions are treated in iterations can dramatically affect the convergence properties of iterations. A poor choice of boundary conditions may lead to an ill-posed discrete problem, even if the corresponding differential problem is well posed. Significant convergence deterioration may occur in the simplest well-known problems, e.g., the Laplace equation, as a result of improper handling of the boundary conditions. See an example in Section 5.3.1.

For the same reason, the LMF analysis cannot account for the downstream accuracy propagation from the inflow boundary that is an important convergence mechanism for some nonelliptic problems. (See Section 5.3.3.) The computational domain affected by the downstream accuracy propagation may be very large, growing with each iteration. In the case of a downstream marching, it may expand to the entire domain in just one iteration. The estimates of the LMF analysis for nonelliptic flow problems can, however, be accurate in the areas that have yet to be affected by the downstream accuracy propagation from the inflow boundary. (See Section 5.3.2 and [15, 31].) For iterative solvers of hyperbolic problems, e.g., non-marching multigrid methods for the convection equation, LMF analysis estimates are usually accurate for the initial convergence; eventually, the accuracy propagated downstream from the inflow boundary extends over the entire computational domain. Some modifications of the LMF analysis, such as a half-space analysis of the first differential approximation (FDA analysis) [2, 11] and a discrete half-space analysis [15], have been developed to take the boundary conditions into the account.

For uniformly elliptic problems, there is no downstream accuracy propagation; the effect of boundary conditions may be localized and separated from the interior. It has been rigorously proved [5, 6] that the convergence rate predicted by the LMF analysis can be observed in actual computations with a proper treatment of boundary conditions on large enough grids.

Another serious limitation of the LMF analysis is the constant-coefficient assumption itself. A rigorous LMF analysis for uniformly elliptic variable-coefficient problems has been considered in [5]. Accuracy of the LMF analysis extension to general variable-coefficient (or nonlinear) problems is limited by the accuracy of constant-coefficient approximations to the target problem. In this context, approximation means that the values of actual (variable) coefficients do not deviate significantly from a constant state on a certain scale called the *scale of coefficient variation*.

There are many ways to measure how well a constant coefficient formulation approximates a variable-coefficient one in a certain region. Here, we adopt a very strong requirement that limits variation in each variable coefficient separately. Let $C_k(\mathbf{x})$ be variable coefficients of the operator \mathbf{L}_v , and

$$C_k^0 = C_k(\mathbf{x}_0) \tag{22}$$

be values of $C_k(\mathbf{x})$ frozen at the point $\mathbf{x}_0 = (x_0, y_0)$. We regard the constant-coefficient operator \mathbf{L}_c obtained by freezing coefficients of \mathbf{L}_v as approximating \mathbf{L}_v in the vicinity of the point \mathbf{x}_0 with a relative accuracy ϵ , if there is a non-vanishing neighborhood, $\Omega_{\mathbf{x}_0}^{s_1, s_2} = \{\mathbf{x} = (x, y), |x - x_0| < s_1, |y - y_0| < s_2\}$, such that, within this neighborhood,

$$|C_k^0 - C_k(\mathbf{x})| < \epsilon |C_k^0|, \quad (23)$$

for any coefficient $C_k(\mathbf{x})$ of \mathbf{L}_v . The rectangular domain $\Omega_{\mathbf{x}_0}^{s_1, s_2}$ is called the ϵ -neighborhood of the point \mathbf{x}_0 .

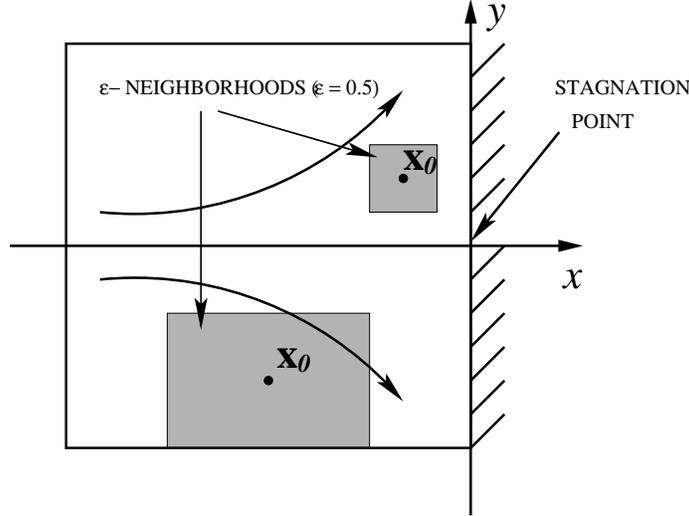


Figure 4. Model variable coefficient operator (24): ϵ -neighborhoods corresponding to $\epsilon = 0.5$

For any chosen ϵ , the values of s_1 and s_2 provide the scales of coefficient variation at which the constant-coefficient problem approximates the variable-coefficient one with the relative ϵ -accuracy. Reasonable values for ϵ are between zero and one, with $\epsilon = 0.5$ considered as a good practical value. See Figure 4 for illustration of typical ϵ -neighborhoods for a variable-coefficient convection equation

$$-xu_x + yu_y = 0. \quad (24)$$

For the variable coefficients frozen at the point $\mathbf{x}_0 = (x_0, y_0)$ one should confine the LMF analysis to a computational subdomain that is located entirely within the ϵ -neighborhood of the point \mathbf{x}_0 . If an ϵ -neighborhood is smaller than several mesh sizes in each dimension, the ϵ -neighborhood is considered vanishing; the LMF analysis should not be applied at this neighborhood. Implementation of the ϵ -neighborhood concept together with an assumption on the minimal number of mesh intervals sets some lower bounds on the coefficient-to-meshsize ratios, e.g., the bounds $\frac{|x_0|}{h} > 2, \frac{|y_0|}{h} > 2$ occur for (24) with $\epsilon = 0.5$ and a minimum of two mesh intervals in each direction per ϵ -neighborhood.

Note that the union of all non-vanishing ϵ -neighborhoods may not cover the entire computational domain. Restricting the scope of the LMF analysis to the non-vanishing ϵ -neighborhoods emphasizes that the analysis avoids the parts of the domain where it might be inaccurate.

A sequence of four steps proposed for obtaining a reliable convergence estimate by the LMF analysis is defined as

1. Go through all the grid nodes and define the ϵ -neighborhood for each node. Retain only non-vanishing ϵ -neighborhoods with the required minimum of mesh intervals in each dimension.
2. Compute the Fourier symbol of the iteration for the operator with frozen coefficients at each ϵ -neighborhood.
3. For each ϵ -neighborhood, define the set of acceptable Fourier modes that satisfy the locality conditions.
4. Compute the convergence estimate at each ϵ -neighborhood as the maximum taken over all the acceptable Fourier modes. The final convergence estimate is taken as the maximal (worst) estimate among all ϵ -neighborhoods.

Examples of the LMF analysis illustrating the limits of its applicability to variable-coefficient problems are shown in Section 5.3. While very useful in many applications, the limitations of the LMF analysis make it inapplicable to many important computational problems, such as problems with large coefficient variation, complex geometries, and unstructured grids. Other, more robust, analysis techniques must be applied. Two such techniques, the IR and ICG analyses, are discussed in the next Section 5.2

5.2 Ideal Relaxation (IR) and Ideal Coarse-Grid (ICG) Iterations

A typical approach to analyzing complex problems is based on the following steps: (1) identification of contributors to different aspects of the problem, (2) formulation of simplified models to isolate each identified contributor, (3) solution of the simplified models, (4) extension of the solutions to actual problems that cannot be directly analyzed. In application to iterative methods, such an approach often leads to dramatic improvements in convergence. However, if the actual difficulty is represented not by a single object or phenomenon, but rather by interactions between several (or many) objects and phenomena, analytically solvable models of such interactions are scarce.

In this section, we propose an alternative approach to analyzing multigrid methods aiming at developing enabling tools for practical CFD solvers. This more generally applicable approach is to apply an available, non-perfect solver to deal with a practical problem in its entirety and, then, to isolate, identify, and improve the parts of the solver responsible for the less-than-optimal performance. The iterative solution method to be analyzed is a two-grid cycle. We focus on analyzing the main complimentary parts of the cycle: relaxation and coarse-grid correction; the restriction and prolongation operators can also be analyzed, but currently are assumed given (as bi-linear prolongation and its adjoint full weighting restriction) and suitable for efficient multigrid solution. In the proposed analysis, idealized iterations probing the two-grid cycle are introduced. In these iterations, one part of the cycle is actual and its complimentary part is replaced with an ideal imitation. For IR iterations, the coarse-grid correction part is actual and an ideal relaxation is employed; for ICG iterations, the relaxation scheme is actual and the coarse-grid correction is imitated. Ideal imitation does not necessarily mean the most efficient operation, it rather reflects our understanding of what role the given part of the cycle plays in the solution process.

The analysis is helpful for identifying what can be done to improve a multigrid solver with unsatisfactory convergence rates. The iterations are very general and can be directly applied in the most complicated situations including highly variable (or nonlinear) coefficients, complex geometries, and unstructured grids. IR iterations may be applied to cycles with many grid-levels as well, not only to a two-grid cycle. The results of this analysis are not single-number estimates, they are rather convergence patterns of the iterations that may either confirm or refute our expectations indicating what part of the actual solver should be improved.

The traditional role for relaxation in a multigrid cycle is to smooth errors. Thus, a possible ideal relaxation (IR) imitation is an explicit averaging procedure, e.g., replacement of the algebraic errors, $e_{i,j}$, with their full-weighted averages, $FWe_{i,j}$; the algebraic errors are zeroed at the boundaries, where the solution values are given, and averaged for other boundaries, where equations are given.

$$e_{i,j}^{\text{new}} = \begin{cases} FWe_{i,j}, & (i,j) \in \text{interior}, \\ 0, & (i,j) \in \text{boundary (given solutions)}, \\ FW_1e_{i,j}, & (i,j) \in \text{boundary (other)}, \end{cases} \quad (25)$$

where FW in the two dimensions is defined as

$$FWe_{i,j} \equiv \frac{1}{4}e_{i,j} + \frac{1}{8}(e_{i+1,j} + e_{i-1,j} + e_{i,j+1} + e_{i,j-1}) + \frac{1}{16}(e_{i+1,j+1} + e_{i+1,j-1} + e_{i-1,j+1} + e_{i-1,j-1}); \quad (26)$$

and FW_1 denotes a one-dimensional averaging along the boundary. This ideal relaxation is operator-independent, i.e., its action does not depend on what discrete operator is relaxed. The smoothing rate of this averaging procedure is 0.5, matching the smoothing rate of the lexicographic Gauss-Seidel relaxation for the five-point Laplacian. At least three averaging sweeps (25) should be performed in an IR iteration (regardless of the type and number of relaxation sweeps in the actual two-grid cycle) to provide error smoothing acceptable for a TME solver.

An ideal coarse-grid (ICG) correction imitation is a coarse-grid representation of the local full-weighted averages of the algebraic errors (25); this ICG correction is prolonged to the fine grid by bi-linear interpolation. The ICG correction is also operator-independent. In the two dimensions, the algebraic errors are updated as

$$\begin{aligned} e_{2i,2j}^{\text{new}} &= e_{2i,2j} - FWe_{2i,2j}, \\ e_{2i-1,2j}^{\text{new}} &= e_{2i-1,2j} - \frac{1}{2}(FWe_{2i,2j} + FWe_{2i-2,2j}), \\ e_{2i,2j-1}^{\text{new}} &= e_{2i,2j-1} - \frac{1}{2}(FWe_{2i,2j} + FWe_{2i,2j-2}), \\ e_{2i-1,2j-1}^{\text{new}} &= e_{2i-1,2j-1} - \frac{1}{4}(FWe_{2i-2,2j} + FWe_{2i-2,2j-2} \\ &\quad + FWe_{2i,2j} + FWe_{2i,2j-2}). \end{aligned} \quad (27)$$

The ICG scheme (27) may serve as a generalized indicator of high-frequency components for smoothing analysis (cf. operator \mathbf{I} in the LMF smoothing analysis, Section 5.1). It also closely relates to the compatible relaxation scheme where the coarse-grid variables are chosen as full-weighted averages.

There are only two requirements for application of the IR and ICG iterations: The first condition is that the exact discrete solution should be available on demand for computing algebraic errors. This requirement is not very restrictive because for analysis purposes one can create a manufactured solution of the target operator with a modified source function. In particular, for linear problems, one can conveniently analyze the case where the exact solution is identical zero. The second condition is that an averaging procedure, such as FW in (26), should be defined at each grid node. A similar procedure must be derived for the residual restriction anyway. In many cases, some weighted combination of the neighboring values would be acceptable. With given restriction and prolongation operators, R_h^H and P_H^h , the ICG correction can be written as

$$\mathbf{q}_h^{\text{new}} = \mathbf{q}_h^{\text{old}} + P_H^h R_h^H (\mathbf{q}_h^{\text{exact}} - \mathbf{q}_h^{\text{old}}). \quad (28)$$

In the IR iterations, (part of) the usual relaxation in a two-grid cycle is replaced by the explicit averaging of algebraic errors (25) applied three times. The coarse-grid correction is

the actual coarse-grid correction employed by the cycle. In the ICG iterations, the actual cycle is simulated with the actual relaxation sweeps; the coarse-grid correction is replaced with the explicit redistribution (27) of local averages of algebraic errors.

Considering the situation in which convergence of the actual two-grid cycle is not satisfactory, the following conclusions can be made from the two possible outcomes of the IR analysis:

- *The convergence of IR iterations is unsatisfactory.* The coarse-grid correction does not sufficiently reduce at least some of the smooth error components and needs improvements. As shown in examples from Sections 5.3, boundary relaxation may help to overcome this difficulty.
- *The convergence of IR iterations is satisfactory.* The coarse-grid correction scheme works well. The relaxation smoothing rates of the actual cycle need to be improved. Alternative relaxation scheme may be considered. Another option is to perform more relaxation sweeps.

The possible outcomes of the ICG analysis are

- *The convergence of ICG iterations is unsatisfactory.* The relaxation scheme does not sufficiently smooth the algebraic errors and needs to be improved.
- *The convergence of ICG iterations is satisfactory.* The relaxation scheme works well. The coarse-grid correction scheme of the actual cycle needs to be improved.

As shown by examples in Section 5.3, the analysis methods employing ideal imitations are sensitive to very delicate details of the tested two-grid cycles, clearly indicating what part of the algorithm could be improved.

5.3 Examples of Analyses

5.3.1 Example 1: Laplace Equation

In this section, the IR and ICG analyses are applied to the Laplace equation defined on a square domain $(x, y) \in [0, 1] \times [0, 1]$ with one Neumann boundary condition at $x = 1$ and Dirichlet conditions at other boundaries. The Laplacian is discretized on a uniform 65×65 grid with the standard five-point stencil. The Neumann boundary condition is discretized with a three-point one-sided differencing. The solution method is a two-grid (1, 2)-cycle with lexicographic point Gauss-Seidel relaxation, full-weighting residual restriction, and bi-linear prolongation. At the end of the relaxation sweep, the residuals of the Neumann equations at the boundary are zero. Three sweeps of the averaging procedure (25) are applied in an IR iteration. The ideal coarse-grid correction is described in (27).

Figure 5 shows the residual convergence rates for the actual two-grid (1, 2)-cycle and for the corresponding IR and ICG iterations. For the actual two-grid (1, 2)-cycle, the LMF analysis predicts the convergence rate of 0.12 per cycle. The observed convergence rate of the actual two-grid cycle is 0.20, significantly worse. Better rates observed in *both* the ICG and IR iterations indicate that there are two different ways to accelerate convergence of the two-grid cycle. One possible approach is to design a local relaxation procedure at the Neumann boundary that smoothes the error at the boundary (not necessarily enforcing zero residuals) without producing large residual spikes at the neighboring interior points. This approach would correspond to an improved relaxation procedure. An alternative approach we follow is improving the coarse-grid correction by performing boundary relaxation near the Neumann boundary.

In the version of boundary relaxation, BR_K , used in this algorithm, all the equations defined at K grid lines adjacent to the Neumann boundary are relaxed simultaneously; depth

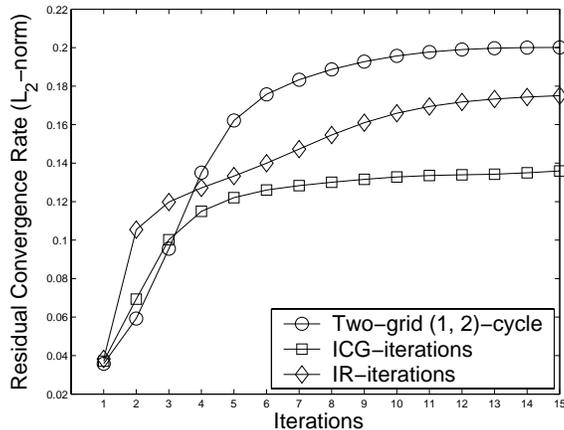


Figure 5. Convergence rates for the Laplace equation with Neumann boundary conditions. No boundary relaxation.

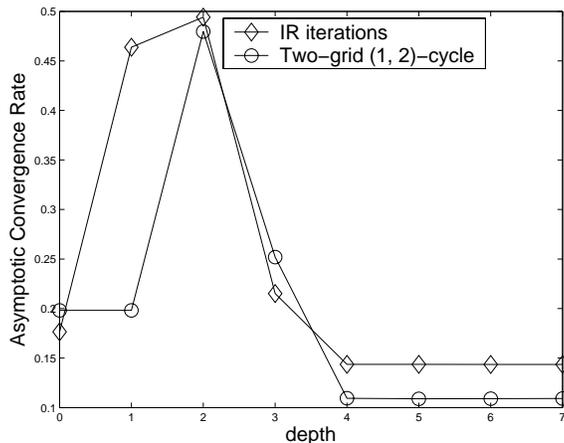


Figure 6. Asymptotic convergence rates for the Laplace equation with Neumann boundary conditions. Boundary relaxation is applied once before residual restriction

$K = 1$, corresponds to simultaneous relaxation of just the boundary equations. Without boundary relaxation, the full-weighting residual restriction near the boundary is inaccurate and degrades the quality of the coarse-grid correction. With zero residuals in a boundary vicinity, an accurate residual restriction is no longer important. The boundary relaxation procedure is performed once per cycle before restricting residuals to the coarse grid. Figure 6 illustrates the sensitivity of the asymptotic convergence rate to the depth of the boundary relaxation region for the actual two-grid (1, 2)-cycle and the IR iterations; the ICG iterations are not sensitive at all to the presence of boundary relaxation. The test results can be summarized as follows: For small depths of the boundary relaxation region, $0 \leq K \leq 3$, the convergence rates of both iterations are not satisfactory. For $K \geq 4$ the desired convergence rates of 0.12 per cycle are achieved. The performance of iterations is not sensitive to further expansions of the boundary relaxation region beyond $K = 4$ and, thus, the value $K = 4$ is considered optimal. Note that convergence rates of the two-grid (1, 2)-cycle are better than the rates of the idealized IR iterations. This is not actually surprising: For the intermediate error components, that are neither high-frequency nor very smooth, the coarse-grid

correction has a limited efficiency. The overall convergence rates for these components are significantly dependent on relaxation. The Gauss-Seidel relaxation outperforms the ideal averaging procedure (25) in reducing the intermediate components.

5.3.2 Example 2: Constant Coefficient Convection Equation

Convergence deterioration in multigrid solutions for nonelliptic equations related to poor coarse-grid correction for the characteristic error components have been observed and explained long ago [2,9,11,14]. Recall that the characteristic components are smooth functions that are much smoother in the characteristic directions than in other directions. The reason for the convergence slowdown is that the main residual contribution from a characteristic error component is made through the numerical cross-characteristic interactions in the discretization. These interactions are typically much larger on coarser grids. Thus, the amplitude of the correction to the characteristic error component computed on a coarse grid is much smaller than the amplitude of the error itself. This explains why the coarse-grid correction is limited for characteristic error components. The theoretical estimates of multigrid performance for nonelliptic operators can be computed, for example, by the LMF analysis [31] or by FDA analysis [2,9,11,14].

However, many researchers observed convergence rates of multigrid solvers for nonelliptic operators that are much better than predicted by the theory. Even textbooks report results that seemingly outperform the theoretical predictions. (See, for example, [28] Section 7.2.1 for convergence of a multigrid solver for a first-order discretization of convection equation. The convergence of about seven orders of magnitude over 14 multigrid cycles is significantly better than the convergence of 0.5 per cycle predicted by the theory.)

The discrepancies between theoretical estimates and convergence rates observed in numerical tests for nonelliptic equations usually occur when a test problem fails to isolate characteristic errors. In multigrid solutions of hyperbolic equations, there exist two parallel processes affecting convergence of iterations: (1) a local error reduction through smoothing and coarse-grid correction and (2) a downstream accuracy propagation from the inflow boundary. The theoretical bounds for the multigrid performance based on estimates for convergence of characteristic error components relate exclusively to the first process of local error reduction. However, the impact of the accuracy propagation process may be very significant, overshadowing the local error reduction. In the extreme case of downstream marching, the accuracy may propagate throughout the entire computational domain in just one iteration. Even relaxation schemes that do not rely on a downstream order usually provide a downstream accuracy propagation of a few (or one) mesh sizes in a relaxation sweep. An analysis estimating the number of defect-correction iterations required to propagate accuracy across the domain has been reported in [15]. After enough relaxation sweeps, the accuracy may propagate from the inflow boundary to the outflow boundary; the (asymptotic) convergence of multigrid iterations may become extremely fast, up to several orders of magnitude in error reduction per iteration. This fast convergence, however, is offset in practice because the number of iterations required to arrive at this asymptotic convergence regime grows with the size of the problem. Away from the inflow boundary, in the regions that are not reached by the downstream accuracy propagation, the multigrid convergence is bounded by the convergence of the characteristic error components and closely follows the theoretical prediction. In typical numerical tests, the distance from the inflow to the outflow boundaries along the characteristics may be different in different parts of the domain. Thus, the overall multigrid convergence is a mixed result of both the processes; the downstream error propagation determines convergence at the (growing) region adjacent to the inflow boundary, and the local error reduction provides convergence in the regions remote from the inflow boundary.

In this section, we analyze a multigrid solver for the standard first-order accurate upwind

discretization of the constant-coefficient convection equation

$$\partial_x^h u + \partial_y^h u = 0, \quad (29)$$

where ∂_x^h and ∂_y^h are two-point upwind differencing for the x and y derivatives, respectively. The rectangular computational domain is covered by a uniform Cartesian grid with the characteristics aligning with the diagonal direction. The solver is a two-grid (1, 2)-cycle including full coarsening, a red-black relaxation scheme with an under-relaxation parameter $w = 0.8$, full-weighting residual restriction, and bi-linear prolongation. In the first experiments, the coarse-grid discretization is identical to the discretization on the fine-grid. For this two-grid cycle, the LMF analysis predicts convergence rates limited by coarse-grid correction for the characteristic error components to be 0.5 per cycle, while the relaxation smoothing factor is 0.52.

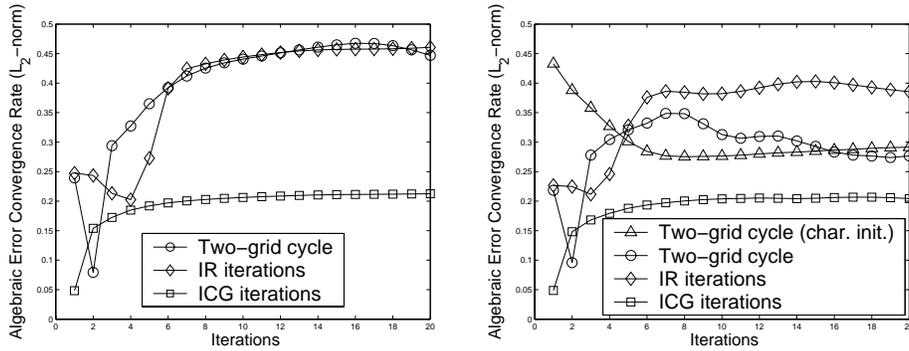


Figure 7. Convergence rates for a first-order discretization of the convection equation with periodic boundary conditions in the y -direction. Standard coarse-grid discretization. Grids 1025×65 (left) and 65×65 (right)

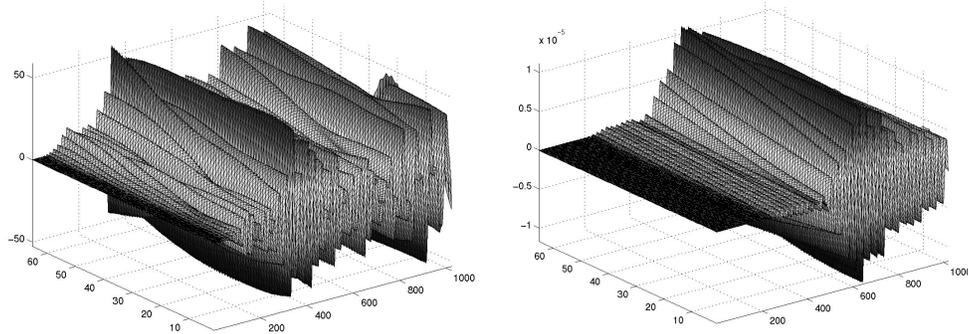


Figure 8. Error surface after 20 (left) and 40 (right) two-grid (1, 2)-cycles for a first-order discretization of the convection equation with periodic boundary conditions in the y -direction. Grid 1025×65

The numerical tests including the two-grid (1, 2)-cycle, IR iterations, and ICG iterations have been performed on two uniform grids, 1025×65 (domain 16×1) and 65×65 (domain 1×1), with the zero inflow boundary condition at $x = 0$ and periodic boundary conditions in the y -direction. The initial approximation is random except for one two-grid-cycle test on the grid 65×65 . In this latter test, the initial approximation is a collection of characteristic

errors. The ideal relaxation and ideal coarse-grid correction have been implemented as (25) and (27), respectively. Convergence rates of the two-grid (1, 2)-cycle on the larger grid (left plot in Figure 7) are close to the 0.5 per cycle predicted in the LMF analysis. Convergence on the smaller grid (right plot in Figure 7) is significantly better than the LMF analysis estimate.

The surface of the algebraic errors after many iterations (see Figure 8 for error profiles after 20 and 40 iterations for the larger grid) reveals the reason for this discrepancy: the LMF analysis ignores the accuracy propagation from the inflow boundary. The characteristic error components dominate the region remote from the inflow boundary that is not affected by the downstream accuracy propagation. Note the highly compressed scale in the x -directions; the error surfaces are composed of smooth waves slowly oscillating in the cross-characteristic direction. The accuracy propagation affects a significantly larger region as more iterations are performed; it becomes a dominating convergence mechanism on small grids. Note that the LMF analysis may still be used for estimating a “worst-case” scenario. The right plot in Figure 7 shows one test on a 65×65 grid with the initial error imposed as a collection of characteristic components; the initial convergence is close to the 0.5 predicted by the LMF analysis.

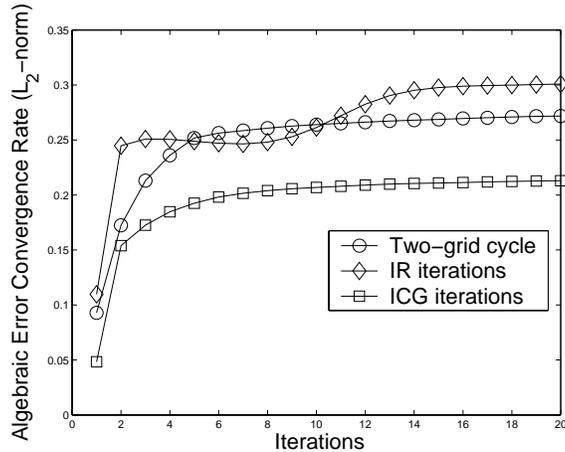


Figure 9. Convergence rates for a first-order discretization of the convection equation with periodic boundary conditions in the y -direction. Grid 1025×65 . Modified coarse-grid discretization

IR and ICG iterations provide very useful information about components of the cycle. Convergence of IR iterations for the analyzed two-grid (1, 2)-cycle displayed in Figure 7 is not satisfactory. Therefore, significant acceleration may be achieved with a better coarse-grid correction. As shown in [9, 13, 14, 18], one can improve the coarse-grid correction scheme by modifying the coarse-grid operator to match the fine-grid (characteristic-component) first differential approximation. In tests demonstrated in Figure 9, the coarse-grid discretization has been modified as

$$\bar{\partial}_x^h = \partial_x^h + \frac{h}{4} \partial_{xx}^u, \quad (30)$$

where $\bar{\partial}_x^h$ is a modified discretization of the x -derivative, ∂_x^h is the first-order upwind discretization, and ∂_{xx}^u is the first-order accurate three-point upwind discretization of the second derivative. The discretization of the y -derivative has been modified accordingly. Convergence rates of the two-grid (1, 2)-cycle are improved significantly to 0.27 per cycle; the rates of IR iterations are also improved (see Figure 9).

ICG iterations converge with a better rate of 0.21 per cycle. One can consider this rate as an ideal rate that can be achieved by improving the coarse-grid approximation but maintaining the original relaxation scheme, coarsening strategy (full coarsening), and prolongation operator. To achieve the ICG rates, the coarse-grid operator should be further modified to provide an even more accurate approximation to the fine-grid operator. One should carefully consider the trade-off between the improved convergence rates and additional complexity of the operators.

5.3.3 Example 3: System of Constant-Coefficient Partial Differential Equations

In this section we analyze a two-grid solver for a system of constant-coefficient equations (15) related to the the full Newton linearization around the plane-flow solution (3) with $a = 1$. The system of constant-coefficient differential equations is formulated as

$$\mathbf{L}(\mathbf{q}) \equiv \begin{pmatrix} \bar{Q}^u - 1 & 0 & \partial_x^c \\ 0 & \bar{Q}^u + 1 & \partial_y^c \\ -2\partial_x^c & 2\partial_y^c & \Delta^h \end{pmatrix} \mathbf{q} = 0, \quad (31)$$

where $\mathbf{q} = (u, v, p)^T$, $\bar{Q}^u = \bar{u}\partial_x^u + \bar{v}\partial_y^u$, and \bar{u} and \bar{v} are positive constants. The computational domain is a rectangle $(x, y) \in [-\bar{u} - X, -\bar{u} + X] \times [\bar{v} - Y, \bar{v} + Y]$ with

$$0 < X \leq \frac{\bar{u}}{2}, \quad 0 < Y \leq \frac{\bar{v}}{2}. \quad (32)$$

The latter conditions to be referred as *compatible-coefficient restrictions* are introduced to insure that the constant-coefficient problem provides a good relative ϵ -accuracy ($\epsilon = 0.5$) in approximating the full-Newton linearization of the plane flow in the vicinity of $\mathbf{x}_0 = (-\bar{u}, \bar{v})$. A uniform Cartesian grid with a minimum of four mesh intervals in each spatial direction is applied. The physical boundary conditions are the following: u, v , and $\partial_x p$ are given at the inflow $x = -\bar{u} - X$; p is given at the outflow $x = -\bar{u} + X$; periodicity is assumed in the y -direction ($\mathbf{q}(x, \bar{v} - Y) = \mathbf{q}(x, \bar{v} + Y)$). The normal derivative of p at the inflow boundary is discretized with the three-point one-sided stencil. The numerical closure equations include central differencing for the x -derivative in the operator \bar{Q}^u at the interior grid points adjacent to the inflow boundary; the x -derivative of the pressure in the first momentum equation is discretized with the three-point upwind stencil at the outflow. The exact solution of the problem is $\mathbf{q} \equiv 0$.

The relaxation scheme consists of three steps:

- (i) Line-implicit Gauss-Seidel relaxation is performed for the Laplace operator in the pressure equation with frozen u and v . In this relaxation, all the pressure equations defined at a vertical grid line with the same x -coordinate are relaxed simultaneously; this y -implicit relaxation scheme proceeds in the *upstream* direction, from the outflow to the inflow.
- (ii) The boundary relaxation, BR_5 , is performed near the inflow boundary. All the equations defined at the five grid lines adjacent to the inflow boundary (the boundary included) are solved simultaneously. Values of u, v , and p in the interior beyond the boundary relaxation region remain unchanged.
- (iii) The momentum equations are *downstream* marched with the S-driver (19) for u and v with p being fixed.

At the end of the relaxation sweep, the residuals of the momentum equations are zero.

The Fourier symbol of this relaxation scheme is computed as

$$\hat{\mathbf{S}}_h = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} - \begin{pmatrix} \hat{Q} - 1 & 0 & \hat{\partial}_x^c \\ 0 & \hat{Q} + 1 & \hat{\partial}_y^c \\ 0 & 0 & G\hat{S}^h \end{pmatrix}^{-1} \begin{pmatrix} \hat{Q} - 1 & 0 & \hat{\partial}_x^c \\ 0 & \hat{Q} + 1 & \hat{\partial}_y^c \\ -2\hat{\partial}_x^c & 2\hat{\partial}_y^c & \Delta^h \end{pmatrix}, \quad (33)$$

where

$$\begin{aligned} \hat{Q} &= \frac{\bar{u}}{h} \left(\frac{3}{2} - 2e^{-i\theta_x} + \frac{1}{2}e^{-i2\theta_x} \right) + \frac{\bar{v}}{h} \left(\frac{3}{2} - 2e^{-i\theta_y} + \frac{1}{2}e^{-i2\theta_y} \right), \\ \hat{\partial}_x^c &= \frac{i}{h} \sin(\theta_x), \\ \hat{\partial}_y^c &= \frac{i}{h} \sin(\theta_y), \\ \Delta^h &= \frac{1}{h^2} (2 \cos(\theta_x) + 2 \cos(\theta_y) - 4), \\ G\hat{S}^h &= \frac{1}{h^2} (e^{i\theta_x} + 2 \cos(\theta_y) - 4), \end{aligned} \quad (34)$$

θ_x and θ_y are normalized Fourier frequencies, $|\theta_x| \leq \pi, |\theta_y| \leq \pi$. The smoothing factor, μ , is

$$\mu = \max \rho(\hat{\mathbf{S}}_h), \quad (35)$$

where $\rho(\hat{\mathbf{S}}_h)$ is the spectral radius of the matrix $\hat{\mathbf{S}}_h$, and the maximum is taken over the set of high-frequency Fourier modes $\{(\theta_x, \theta_y), \max(|\theta_x|, |\theta_y|) \geq \pi/2\}$.

The unrestricted LMF analysis considers the infinite domain with no compatible-coefficient restrictions (32) applied. This analysis predicts divergence of the relaxation scheme for high-frequency error components. For $\theta_x = \pi/2, \theta_y = -\pi/2$, and $\frac{\bar{u}}{h} = \frac{\bar{v}}{h} = \frac{1}{2}$, the convection symbol $\hat{Q} = 1$ and the smoothing factor $\rho(\hat{\mathbf{S}}_h) = \infty$. Imposing the compatible-coefficient restrictions (32) together with the assumption that there are at least four mesh intervals in each dimension effectively implies $\frac{\bar{u}}{h} \geq 4, \frac{\bar{v}}{h} \geq 4$. With these restrictions, the LMF analysis predicts the smoothing factor of 0.53 on the 4×4 grid and factors better than 0.5 at finer grids converging to $\frac{1}{\sqrt{5}}$; the latter is the smoothing factor of the line relaxation for the scalar Laplace equation.

For $\bar{u} = \bar{v} = 1$, the LMF analysis restricted by (32) predicts the amplification of the two-grid cycle to be greater than one indicating divergence. This amplification factor is predicted for specific, isolated smooth characteristic components $\theta_x = -\theta_y, 0 < |\theta_x| < \frac{\pi}{4}$ that may vary as the grid is refined. The LMF analysis does not account for the accuracy propagation from the inflow boundary, that is a major convergence mechanism for this two-grid cycle. Employment of a downstream marching in the relaxation scheme on a domain restricted as (32) ensures that the characteristic modes, for which the LMF analysis predicts divergence, are strongly affected by the inflow boundary.

Figure 10 shows the convergence rates of the L_2 -norm of the pressure-equation residuals for the four iterative methods solving the problem (31) with $\bar{u} = \bar{v} = 1$ on the computational domain $(x, y) \in [-\frac{3}{2}, -\frac{1}{2}] \times [\frac{15}{16}, \frac{17}{16}]$ covered by the uniform isotropic 257×33 Cartesian grid. The plot marked ‘‘Two-grid (1, 2)-cycle’’ shows the convergence rates of a cycle solving (31). The two-grid cycle includes the three-step relaxation scheme described above, full coarsening, full weighting residual restriction, bi-linear prolongation, and the coarse-grid discretization identical to (31). ICG iterations keep the actual relaxation scheme, but employ an ideal coarse-grid correction for all variables computed as in (27). The IR iterations employ the ideal relaxation procedure (25) instead of the actual relaxation sweeps. The IR iterations of the first type, IR-p iterations, idealize smoothing for the pressure only. This ideal relaxation replaces the step (i) in the relaxation scheme with averaging algebraic errors in the pressure, employing three iterations of (25). To illustrate the critical role played by the downstream marching in the cycle, we experimented with another type of IR iterations, IR-all iterations, where all three steps (i)-(iii) are replaced by averaging (25) for each of u, v , and p ; thus, no downstream marching is performed.

ICG iterations and IR-p iterations provide a very accurate estimate for the convergence of the actual cycle. The convergence rates of the ICG iterations and the two-grid cycle

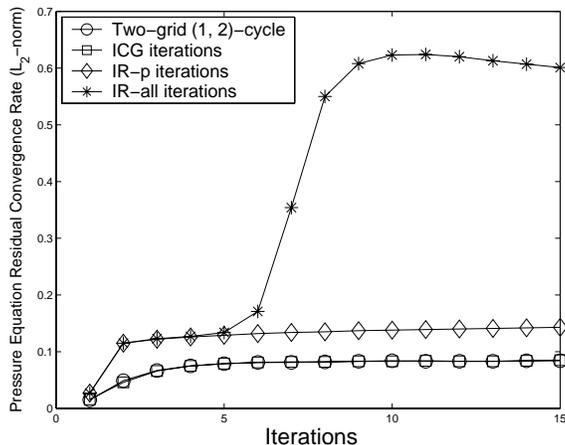


Figure 10. Convergence rates of the pressure equation residuals in solutions of the constant-coefficient problem (31). Grid 257×33 .

are virtually indistinguishable; the IR-p convergence rates are predictably somewhat worse because, as mentioned previously, line relaxation damps intermediate error components significantly better than the ideal relaxation. With (not shown) progressive grid refinement, the convergence rates of the IR-all iterations tend to the value 0.75; this is in agreement with the LMF analysis predictions for convergence rates of full-coarsening multigrid methods applied to a second-order discretization of the convection equation.

Although not shown in the figure, on larger domains (e.g., $[-5, 3] \times [\frac{1}{2}, \frac{3}{2}]$), accuracy propagation from the inflow boundary does not dominate over the entire computational domain; divergence of the two-grid cycle is indeed observed in numerical tests as predicted by the LMF analysis. However, on such domains, the compatible-coefficient restrictions (32) are violated, and, therefore, the formulation (31) does not provide a relevant constant-coefficient approximation to the variable-coefficient problem (14).

6 Numerical Tests

6.1 Convergence of Newton Iterations

Certain difficulties with solution of the coarse-grid discretizations have been observed in numerical tests, even if the Newton method is used as the iteration scheme. The larger discretization errors on coarse grids adversely affect the convergence of Newton iterations, if solutions can be found. For example, a transient solution may violate intended steady state boundary conditions, such as outflow conditions at a downstream boundary. Sometimes, solutions cannot be found with the usual initial approximations, such as freestream values, exact continuous solutions, or coarser-grid solutions.

Typical convergence results are shown in Table 1 for Newton iterations for flow over a parabola with the computational domain extending 32 nose radii upstream and a comparable distance from the leading edge to the downstream boundary. The maximum discretization error in u and the number of iterations to converge residuals to the machine zero are presented versus the grid size. Results are shown for non-compact and compact formulations of numerical closures at the tangency boundary described in Section 3.3.3. The initial approximation for the Newton iterations is either injected from the exact solution of the differential problem, or interpolated from the solution on the coarser grid.

Non-Compact Numerical Closure			
Grid	Discr. Error	Newton Iterations, Initial Approximation:	
		Differential Solution	Coarse-Grid Solution
9×9	No solution	No convergence	Not applied
17×17	0.4605	10	Not applied
33×33	0.3627	10	No convergence
65×65	0.02458	6	No convergence
129×129	0.00575	5	7

Compact Numerical Closure			
Grid	Discr. Error	Newton Iterations, Initial Approximation:	
		Differential Solution	Coarse-Grid Solution
9×9	0.29	6	Not applied
17×17	0.15	6	6
33×33	0.055	6	7
65×65	0.01796	5	6
129×129	0.00536	5	5

Table 1. Parabola flow: L_∞ -norm of the discretization error in u and the number of Newton iterations to converge residuals to machine zero.

For the non-compact numerical closure, a solution on the coarsest grid (9×9) could not be found. The Newton iterations on the finer grids were able to converge starting from the exact differential solution. On the grids 33×33 and 65×65 , however, the initial approximations interpolated from the coarser grid solutions led to transient solutions that violated the outflow conditions. For the compact numerical closure, the iterations converged to machine-zero residuals on all of the grids with either initial approximation. Note that on fine enough grids, Newton iterations converge fast for both the formulations. The particular examples cited here are rather extreme; the angle subtended by the grid at the leading edge is on the order of seventy degrees on the 17×17 grid. However, the examples are representative of problems encountered during the course of this study.

6.2 Amplification of Algebraic Errors in Stagnation Region

Fast asymptotic convergence of a multigrid solver does not guarantee TME; the algebraic errors should be reduced below the level of the discretization errors in the very first multigrid cycle on a given grid. The algebraic error reduction in the first multigrid cycle may depend on the initial solution approximation. For instance, if the initial solution approximation contains large pressure gradient errors, the velocity updates computed from relaxing the momentum equations with a fixed pressure field may significantly amplify the algebraic errors in velocities. Numerical tests for deep-stagnation flows in curved geometries with the S-driver (19) relaxation scheme and pointwise Gauss-Seidel relaxation for the Laplacian have indicated insufficient accuracy of FMG-1 solutions; FMG-2 solvers have been required to achieve the solutions with algebraic errors less than the discretization errors. Below in this section, we discuss the reasons for the error amplification in relaxation and propose adjustments in the downstream marching scheme to contain the amplification and recover the accuracy of FMG-1 solutions.

As discussed in Section 4, the relaxation scheme for the nonlinear system (9) computes solution updates by iterating a system of linear equations (13) with the right-hand side computed from the current nonlinear solution approximation, $\hat{\mathbf{q}}$, and fixed throughout the relaxation procedure. The system (13) is not solved to zero residuals; an approximate

solution, $\tilde{\delta}\mathbf{q} = (\tilde{\delta}u, \tilde{\delta}v, \tilde{\delta}p)$, with certain algebraic errors $\epsilon = (\epsilon^u, \epsilon^v, \epsilon^p)$, is updated in the relaxation procedure. A typical initial assignment is $\tilde{\delta}\mathbf{q} \equiv 0$. Recall, that the algebraic errors are differences between the exact and an approximate solutions to a discrete problem. The algebraic errors ϵ closely approximate the algebraic errors of the approximate nonlinear solution $\tilde{\mathbf{q}} + \tilde{\delta}\mathbf{q}$, within the accuracy of Newton linearizations approximating the nonlinear equations (9).

At some stages of relaxation, the residuals of particular linearized discrete equations may become zero, imposing some conditions on the transient algebraic errors. Relaxation of the linearized pressure equation performed at step (1) of the interior relaxation scheme (see Section 4.1), updates the algebraic error in pressure, ϵ^p ; the velocity errors, ϵ^u and ϵ^v , remain unchanged. If the Laplace operator, Δ^h , were to be “solved” to zero residuals, rather than relaxed, the obtained algebraic errors would satisfy

$$2\left((\partial_x^c u)\partial_x^c + 2(\partial_x^c v)\partial_y^c\right)(\epsilon_{\text{old}}^u) + 2\left((\partial_y^c v)\partial_y^c + 2(\partial_y^c u)\partial_x^c\right)(\epsilon_{\text{old}}^v) + \Delta^h(\epsilon_{\text{new}}^p) = 0, \quad (36)$$

where discrete operators Δ^h , ∂_x^c , and ∂_y^c are defined in Section 3.1. The equation (36) implies that the algebraic errors in the pressure of the transient solution, $\tilde{\delta}\mathbf{q}_{\text{new}}$, would not depend on the initial (“old”) pressure errors; they would be completely defined by the initial distribution of the velocity algebraic errors. Significant error amplification can, thus, occur depending on the relative size of the initial errors in the pressure and velocities. This amplification, if it occurs, is a one-time event, i.e., successive “solutions” of the Laplacian would not update the algebraic error distribution until the velocity errors are changed. Usually the Laplace operator of the pressure equation is not “solved” but relaxed with typically point or line relaxation methods. These relaxations mostly affect the high-frequency part of the pressure error; the smooth error components are hardly changed.

The downstream marching of the momentum equations updates all (high-frequency and smooth) velocity components; it does not affect the errors in the pressure. If the marching has been performed with the S-driver, the algebraic errors of the updated solution satisfy

$$\begin{aligned} \left(Q^u + (\partial_x^u u)\right)(\epsilon_{\text{new}}^u) + (\partial_y^u u)(\epsilon_{\text{new}}^v) + \partial_x^c(\epsilon_{\text{old}}^p) &= 0, \\ \left(Q^u + (\partial_y^u v)\right)(\epsilon_{\text{new}}^v) + (\partial_x^u v)(\epsilon_{\text{new}}^u) + \partial_y^c(\epsilon_{\text{old}}^p) &= 0, \end{aligned} \quad (37)$$

where discretized derivatives are defined in Section 3.1 and Q^u is the linear convection operator defined in Section 4. The “new” algebraic errors in the velocities are determined by the “old” algebraic errors in the pressure and are independent of the initial velocity errors.

The convection operator linearized around a plane-stagnation solution has an unbounded fundamental solution. The velocity updates computed from the S-driver marching of the momentum equations toward stagnation are, thus, very sensitive to the errors in the pressure and prone to error amplification. To demonstrate this sensitivity, we consider a differential approximation to the u -momentum equation linearized around the plane-stagnation solution (3) with $a = 1$ and the inflow boundary defined at $x = -L$. For this example, $\check{\delta}\mathbf{q} = (\check{\delta}u, \check{\delta}v, \check{\delta}p)$ represents an approximate continuous solution with errors $\check{\epsilon} = (\check{\epsilon}^u, \check{\epsilon}^v, \check{\epsilon}^p)$ representing its deviation from the exact solution. At the symmetry line ($y = 0$), the pre-marching residuals of the linearized u -momentum equation are defined as

$$r^u = \check{f}_u - \left(-x\partial_x(\check{\delta}u_{\text{old}}) - \check{\delta}u_{\text{old}} + \partial_x(\check{\delta}p_{\text{old}})\right) = -x\partial_x(\check{\epsilon}_{\text{old}}^u) - \check{\epsilon}_{\text{old}}^u + \partial_x(\check{\epsilon}_{\text{old}}^p), \quad (38)$$

where \check{f}_u is a continuous approximation to the right-hand side of the relaxation u -momentum equation (the first equation of (13)). The update, $\delta(\check{\delta}u)$, obtained in the downstream march-

ing with the S-driver is computed as

$$\begin{aligned} -x\partial_x\delta(\check{\delta}u) - \delta(\check{\delta}u) &= r^u, \\ \check{\delta}u_{\text{new}} &= \check{\delta}u_{\text{old}} + \delta(\check{\delta}u). \end{aligned} \quad (39)$$

The errors of the updated solution satisfy the equation

$$-x\partial_x(\check{\epsilon}_{\text{new}}^u) - \check{\epsilon}_{\text{new}}^u + \partial_x(\check{\epsilon}_{\text{old}}^p) = 0. \quad (40)$$

For illustration purposes, let us consider an initial approximation with a linearly changing algebraic error in the pressure,

$$\partial_x(\check{\epsilon}_{\text{old}}^p) = D. \quad (41)$$

The error, $\check{\epsilon}_{\text{new}}^u$, of the updated solution is, then, given as

$$\check{\epsilon}_{\text{new}}^u = D \left(1 + \frac{L}{x} \right); \quad (42)$$

it is independent of $\check{\epsilon}_{\text{old}}^u$ and unbounded as x tends to zero.

One possible approach to control this velocity-error amplification in the deep-stagnation zone is to replace downstream marching with a more local relaxation scheme (e.g., red-black) that would efficiently reduce high-frequency error components, but would not affect the smooth errors. Derivation of good local relaxation smoothers for convection-type operators is possible [14, 18, 19], but may be cumbersome. Moreover, rejecting downstream marching might require changes in the coarse-grid discrete formulation to better approximate characteristic error components.

On the other hand, the interior relaxation scheme that includes relaxation of the Laplace operator of the pressure equation followed by downstream marching of the convection operators of the momentum equations proved to be very efficient away from stagnation [20–23]. It would be preferable to develop techniques allowing the downstream marching in stagnation proximity, but containing the error amplification.

A possible technique is to change the driver operator used in the downstream marching. We propose the D-driver operator (20). Away from stagnation, the D-driver has the same principal linearization (17) as the S-driver operator (19); at stagnation, the D-driver does not possess an unbounded fundamental solution. Marching with the D-driver can be very beneficial, but cannot assuredly prevent error amplification. In particular, the D-driver counterpart of (39),

$$\begin{aligned} -x\partial_x\delta(\check{\delta}u) + \delta(\check{\delta}u) &= r^u, \\ \check{\delta}u_{\text{new}} &= \check{\delta}u_{\text{old}} + \delta(\check{\delta}u), \end{aligned} \quad (43)$$

results in the approximate solution with $\check{\epsilon}_{\text{new}}^u$ determined by a certain combination of initial pressure and velocity errors. After the downstream marching with the D-driver, the errors satisfy the equation

$$-x\partial_x(\check{\epsilon}_{\text{new}}^u) + \check{\epsilon}_{\text{new}}^u - 2\check{\epsilon}_{\text{old}}^u + \partial_x(\check{\epsilon}_{\text{old}}^p) = 0. \quad (44)$$

If the initial errors satisfy

$$-2\check{\epsilon}_{\text{old}}^u + \partial_x(\check{\epsilon}^p) = D, \quad (45)$$

where D is constant, the $\check{\epsilon}_{\text{new}}^u$ is given as

$$\check{\epsilon}_{\text{new}}^u = -D \left(1 + \frac{x}{L} \right). \quad (46)$$

The error $\check{\epsilon}_{\text{new}}^u$ is bounded, but still may be larger than $\check{\epsilon}_{\text{old}}^u$.

Figure 11 illustrates the amplification of algebraic errors in relaxation for the u -component of velocity along the symmetry line of a deep-stagnation cylinder flow spanning two degrees

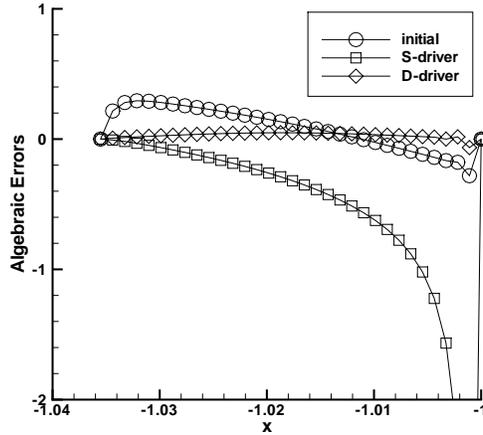


Figure 11. The algebraic error in u along the symmetry line for the deep-stagnation cylinder flow. Stagnation point is located at $x = -1$.

of arc sector. The uniform body-fitted 33×33 grid is applied. The initial approximation is the exact continuous solution (6). On Figure 11, the initial u -velocity algebraic error is compared with algebraic errors obtained after the downstream marching of the u -momentum equation with the S-driver and with the D-driver. Marching with the S-driver leads to very significant error amplification; marching with the D-driver results in significant reduction of the algebraic errors. For specific initial approximations obtained by solution interpolation from a coarser grid within an FMG algorithm, the D-driver relaxation proved to be very efficient, always leading to textbook efficient FMG-1 solutions.

6.3 Multigrid Description

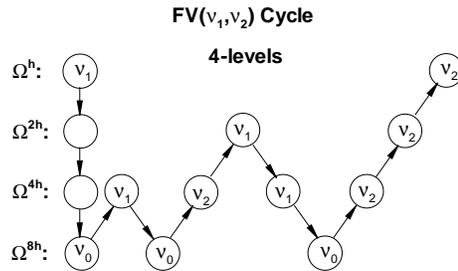


Figure 12. Schematic of the FV-cycle for 4-level multigrid where ν_0 denotes the number of relaxations on the coarsest mesh (Ω^{8h}).

The basic cycle employed in the multigrid algorithms reported in this section is an FV(ν_1, ν_2) cycle sketched in Figure 12. Unless otherwise specified, the multigrid solutions have been obtained as follows: At each grid, an FAS-FV(2,1) multigrid cycle has been used. The nonlinear equations at the fixed coarsest grid have been solved exactly; thus finer grids

used progressively more levels of multigrid. Within the FAS cycle, a D-driver relaxation scheme, full-weighting restriction of residuals, and bilinear prolongation of corrections have been used with direct injection of the nonlinear solution to the coarser level. The FMG- k algorithm has started from an exact solution on the coarsest grid; cubic interpolation from coarser grid solutions has been used to form initial approximations on finer grids.

Compact formulations for discrete boundary conditions at the inflow and tangency boundaries have been implemented (see Section 3.3). The D-driver relaxation scheme on isotropic grids has included three steps: (1) point Gauss-Seidel relaxation of the pressure equation performed in a downstream order, (2) 2×2 -block line-implicit downstream relaxation of momentum equations, (3) boundary relaxation BR_3 at the inflow and tangency boundaries.

For anisotropic grids with stretching, the relaxation of the pressure equation at step (1) has been replaced with line-implicit alternate direction relaxation. In general situations, it is necessary to account explicitly for boundary coupling effects in these line relaxations. For example, the equations governing line relaxation for the pressure can be augmented with other equations to solve for other variables in some block near the boundary. This was not necessary for the test cases considered in this paper. Note, solving the momentum equations coupled at the boundary is a simple illustration of this general consideration since the boundary conditions pertain to the normal and tangential velocities, which differ from the Cartesian velocities near a curved body.

6.4 Isotropic Grids

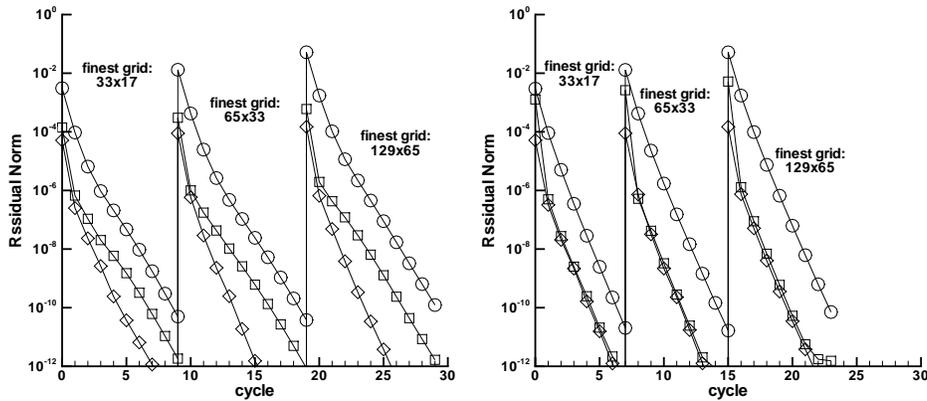


Figure 13. Residual convergence history. Plane stagnation flow with offset 0 (left plot) and offset 10 (right plot).

6.4.1 Convergence Rates

The convergence of the multigrid scheme is shown in Figures 13-15 for plane, cylindrical, and parabolic stagnation. The body-fitted grids are orthogonal and isotropic. Tests are performed for deep and regular stagnation flows. The L_2 norms of residuals are shown. Circles correspond to the residuals of the pressure equation; squares and diamonds denote residuals of the u - and v -momentum equations, respectively. The largest residuals at each cycle are those of the pressure equation with the momentum residuals generally 1-2 orders of magnitude smaller.

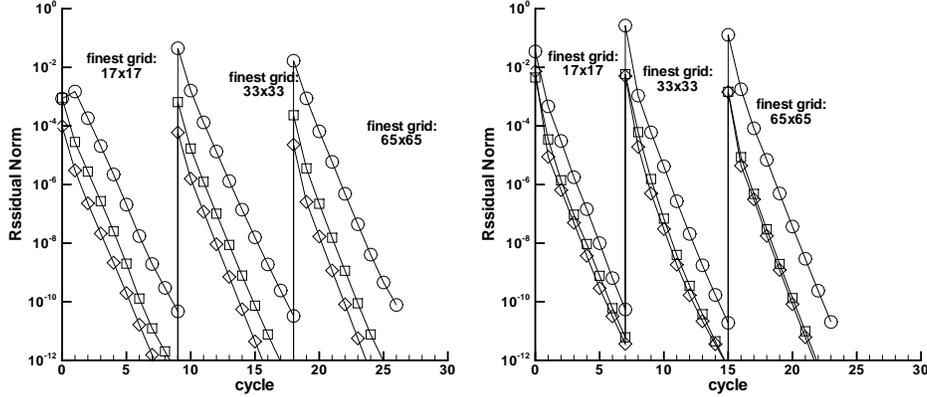


Figure 14. Residual convergence history. Cylindrical stagnation flow. Left plot — deep within stagnation ($\theta_{\max} = -\theta_{\min} = 5^\circ$). Right plot — regular stagnation ($\theta_{\max} = -\theta_{\min} = 45^\circ$).

The plane stagnation results (Figure 13) were obtained on the 2×1 domain shown in Figure 1. The computational domain is either adjacent to the plane with the tangency conditions defined at the right vertical boundary or shifted upstream with outflow conditions given at the right vertical boundary with offset $x = -10$. The finest isotropic grid was 129×65 and the coarsest grid level was 17×9 . Since there is no discretization error for the plane-flow equations, the initial approximations on all grids were taken as random perturbations of the exact solutions for all variables. The results with the computational domain shifted away from the origin show the expected behavior — fast convergence rates of more than an order of magnitude per cycle. The asymptotic convergence rate of 0.09 per cycle is identical to the rate of a similar multigrid applied to the scalar five-point Laplacian operator. When the stagnation point is a part of the computational boundary, the convergence rates are grid-independent but show less regularity. Depending on initial approximations, the observed convergence rates may be different, but are always between 0.13 and 0.2. Note also, that for computational domains including stagnation, the results were somewhat sensitive to the length-to-width ratio of the computational domain; the 2×1 domain chosen is the worst of the ranges investigated.

The cylinder results (Figure 14) were obtained for a grid spanning 90 degrees of arc sector (regular stagnation) and for one spanning 10 degrees of arc sector (deep stagnation). The finest isotropic grid was 65×65 , and the coarsest grid level was 9×9 . The initial approximations on each grid, except for the coarsest one, were obtained using cubic solution prolongation from the coarser grid. The regular-stagnation tests show grid-independent convergence rates (better than 0.1 per cycle) comparable to the rates for the scalar five-point Laplacian operator. The convergence for deep-stagnation flows is grid-independent, but slightly slower; asymptotic convergence rates are about 0.17 per cycle.

The parabola results (Figure 15) were obtained with η_{\max}/η_{\min} values of 8 and 1.25, corresponding to regular and deep stagnation flows, respectively. The finest isotropic grid was 129×129 , and the coarsest grid level was 9×9 . The initial approximations on fine grids were obtained by cubic interpolations of coarser-grid solutions. The results are quite similar to those for the previous case — grid-independent asymptotic convergence rates better than 0.1 per cycle for regular stagnation and about 0.17 per cycle for deep-stagnation flows.

Although not shown, the asymptotic rates for deep-stagnation flows were better for 2-

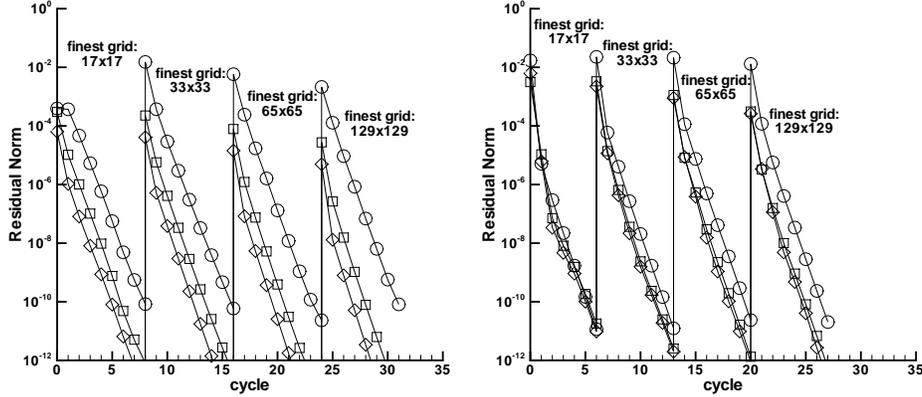


Figure 15. Residual convergence history. Parabolic stagnation flow. Left plot — deep within stagnation ($\eta_{\max}/\eta_{\min} = 1.25$). Right plot — regular stagnation ($\eta_{\max}/\eta_{\min} = 8$)

level cycles but did not recover the convergence rates found for regular stagnation. Also, the convergence rates can be significantly improved with other, more efficient, relaxation schemes for the pressure equation, e.g., line-implicit or red-black relaxation schemes; the improved convergence rates are still somewhat slower than the (also improved) convergence rates of the corresponding multigrid cycle for the five-point Laplace operator. In the extreme case of *solution* (with a separate multigrid cycle) of the pressure equation for a fixed velocity field, convergence rates of 0.3 per relaxation (0.027 per cycle) were observed.

6.4.2 FMG Performance

Simulation	Grid	Discr. Err.	FMG-1 Err.	AE/DE
Deep stagnation, Cylinder flow	33×33	0.7303×10^{-5}	0.8107×10^{-5}	0.11
	65×65	0.1907×10^{-5}	0.2058×10^{-5}	0.08
Regular stagnation, Cylinder flow	33×33	0.1681×10^{-2}	0.1660×10^{-2}	0.01
	65×65	0.4274×10^{-3}	0.4302×10^{-2}	0.01
Deep stagnation, Parabola flow	65×65	0.9699×10^{-6}	0.1043×10^{-5}	0.07
	129×129	0.2488×10^{-6}	0.2627×10^{-6}	0.06
Regular stagnation, Parabola flow	65×65	0.3739×10^{-2}	0.3739×10^{-2}	0.02
	129×129	0.8978×10^{-3}	0.9305×10^{-3}	0.04

Table 2. L_{∞} norm of the discretization error in p and FMG-1 convergence for stagnation flows (isotropic grids).

The convergence of FMG-1 algorithm (one FV(2,1) cycle per grid level) for the cylinder and parabola, are shown in Table 2. The L_{∞} norms of discretization errors are compared with the norms of the errors after FMG-1 algorithm for four different simulations and two different grids; the ratios of the norms of algebraic errors to the norms of discretization errors are exhibited in the last column. On all grid levels, the algebraic errors after the FMG-1 algorithm are at least an order of magnitude smaller than the discretization errors. The total computational work of the algorithm is less than 8 minimal work units. For all computations, every variable shows second order convergence of discretization error in all

norms. The residuals of the discretized continuity equation evaluated from the discrete solution exhibit second-order convergence in the L_2 -norm and first-order convergence in the L_∞ -norm. Only the maximum discretization error in pressure is tabulated in Table 2, showing a factor of 4 reduction between the coarser and finer grid levels.

6.5 Stretched Grids

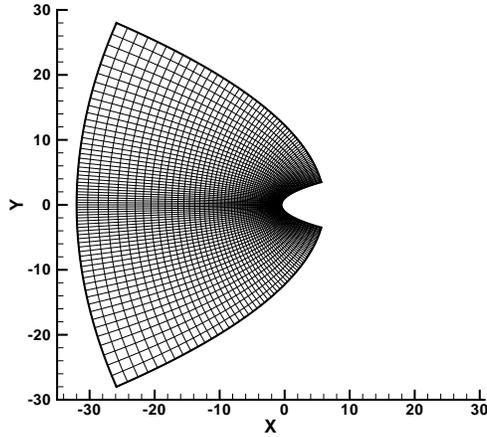


Figure 16. Computational domain for parabolic stagnation flow, stretched 65×65 grid

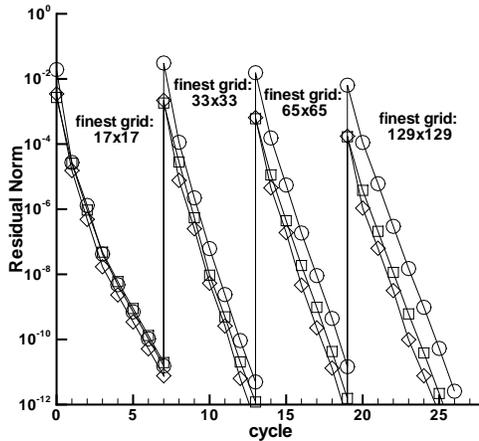


Figure 17. Convergence history on stretched grids

Grid stretching is often used in applications to cluster grid points in certain regions, such as the leading or trailing edges of an airfoil or in viscous boundary layers. Stretching may cause unintended anisotropies in the discrete equations; line-implicit relaxation can be used in combination with full-coarsening multigrid to account for this grid-induced anisotropy in two dimensions. A stretched grid for the parabola is shown in Figure 16, where grid-induced

anisotropies are evident far upstream of the leading edge and also downstream along the body surface. The stretching ratio on the parabolic grid relates to the stretching factor $\beta = 1.0125$ in the transformed plane on the 65×65 grid, resulting in a cell aspect ratio of approximately 4. This anisotropy is significant enough to cause a deterioration in smoothing properties of point Gauss-Seidel relaxation; asymptotic convergence rates around 0.73 per relaxation were observed in FV(2, 1) cycles solving the Laplace equation with the Dirichlet boundary conditions on the 129×129 grid with stretching factor $\beta = \sqrt{1.0125}$.

The multigrid convergence on the stretched parabolic grids is shown in Figure 17. The finest grid is 129×129 with the stretching ratio $\beta = 1.00625$, and the coarsest grid level is 9×9 . Relaxation for the pressure equation includes three line-implicit sweeps: first, lines parallel to the body surface are solved in downstream order, then, lines normal to the body surface are marched in alternating directions. The multigrid FV(2, 1) cycle exhibits grid-independent convergence rates of about 0.064 per cycle on the three finest grids and somewhat slower convergence on the coarsest grid (17×17). The convergence rates for the Laplace equation alone with multigrid cycle employing the triple-sweep line-implicit relaxation scheme described above are about 0.008 per cycle, somewhat faster than for the full system. The slower convergence rate on the coarsest grid in Figure 17 is symptomatic of a lack of leading edge resolution that was also observed with larger grid stretching rates for the same numbers of grid points. This deterioration is attributed to the ensuing lack of discretization accuracy on coarse grids with high stretching rates. The adverse effect of poor accuracy on convergence of Newton iterations has been discussed in Section 6.1. In particular, with a stretching ratio $\beta = 1.018$ on the 129×129 grid, the FMG solution becomes inaccurate. The convergence of Newton's method on the corresponding 33×33 grid ($\beta \approx 1.08$) is quite erratic even with the initial approximation obtained from the exact differential solution.

Grid	Variable	Discr. Err.	FMG Err.	AE/DE
65×65	u	0.1096×10^{-2}	0.1044×10^{-2}	0.047
	v	0.1026×10^{-2}	0.9824×10^{-3}	0.043
	p	0.5279×10^{-3}	0.4987×10^{-3}	0.055
129×129	u	0.2893×10^{-3}	0.2730×10^{-3}	0.056
	v	0.2851×10^{-3}	0.2665×10^{-3}	0.065
	p	0.1344×10^{-3}	0.1268×10^{-3}	0.057

Table 3. L_2 norm of the discretization errors and FMG-1 convergence for parabola-stagnation flows (stretched grids).

The L_2 -norm of the discretization errors in each of u , v , and p and the total (discretization plus algebraic) errors of the FMG-1 algorithm are shown in Tables 3 for the two finest grids. Second-order convergence of the discretization errors is obtained and the algebraic errors of the FMG-1 solutions are far less than the discretization errors. The total computational cost of this FMG-1 algorithm is about 12 minimal work units.

7 Conclusions

Textbook multigrid efficient solvers have been developed for a pressure-equation formulation of the incompressible inviscid equations, although the conclusions are expected to apply to more general situation involving viscosity and compressibility. The relaxation strategy, based largely on analysis and computation for plane stagnation, enables convergence of algebraic errors below discretization errors in one multigrid cycle and asymptotic convergence rates approaching those for the scalar elliptic factor.

During the course of this work the following four principal difficulties associated with achieving TME for stagnation-flow problems have been identified and overcome:

1. The accuracy of the discrete solution is very sensitive to boundary condition formulations in the vicinity of stagnation point. Compact numerical closure conditions leading to sufficient solution accuracy on all uniform, including anisotropic, grids have been found and applied.

2. Strong coupling between the equations near the boundaries and associated difficulties with restricting fine-grid residuals to coarser grids led to a significant slowdown of multigrid convergence rates. This difficulty has been solved by implementing the concept of boundary relaxation.

3. Initial amplification of the algebraic errors in relaxation was an obstacle to achieving the TME goal. A special relaxation scheme has been developed to avoid amplification of the specific errors arising in the FMG solver.

4. Lack of reliable quantitative analytical tools hampered our abilities to optimize multigrid solvers. New, very general, analysis methods for multigrid solutions of stagnation flows have been developed.

References

1. A. BRANDT, *Multi-level adaptive solutions to boundary-value problems.*, Math. Comp., 31 (1977), pp. 333–390.
2. A. BRANDT, *Multigrid solvers for non-elliptic and singular-perturbation steady-state problems.* Manuscript (unpublished). The Weizmann Institute of Science, Rehovot, Israel, December 1981.
3. A. BRANDT, *Guide to multigrid development*, in Multigrid Methods, W. Hackbusch and U. Trottenberg, eds., Lecture Notes in Math. 960, Springer-Verlag, Berlin, 1982, pp. 220–312.
4. A. BRANDT, *Multigrid techniques: 1984 guide with applications to fluid dynamics*, in Lecture Notes for the Computational Fluid Dynamics, Lecture Series at the Von-Karman Institute for Fluid Dynamics, The Weizmann Institute of Science, Rehovot, Israel, 1984. ISBN-3-88457-081-1, GMD-Studien Nr. 85, Available from GMD-AIW, Postfach 1316, D-53731, St. Augustin 1, Germany. Also available from Secretary, Department of Mathematics, University of Colorado at Denver, CO 80204-5300.
5. A. BRANDT, *Rigorous local mode analysis of multigrid*, in Proc. Fourth Copper Mountain Conf. on Multigrid Methods, Copper Mountain, Colorado, April 1989.
6. A. BRANDT, *Rigorous quantitative analysis of multigrid, I: Constant coefficients two-level cycle with L_2 -norm*, SIAM J. Num. Anal., 31 (1994), pp. 1695–1730.
7. A. BRANDT, *Appendix C: Recent developments in multigrid efficiency in computational fluid dynamics*, in Multigrid, Academic Press, London, 2000, pp. 573–589. Ulrich Trottenberg, C. W. Oosterlee, and A. Schüler.
8. A. BRANDT AND N. DINAR, *Multi-grid solutions to elliptic flow problems*, in Numerical Methods for Partial Differential Equations, S. Parter, ed., Academic Press, New York, 1979, pp. 53–147.
9. A. BRANDT AND B. DISKIN, *Multigrid solvers for nonaligned sonic flows*, SIAM J. Sci. Comp., 21 (2000), pp. 473–501.

10. A. BRANDT, B. DISKIN, AND J. L. THOMAS, *Textbook multigrid efficiency for computational fluid dynamics simulations*, AIAA Paper 2001-2570, 15th AIAA CFD Conference, Anaheim, CA, June 2001.
11. A. BRANDT AND I. YAVNEH, *On multigrid solution of high-Reynolds incompressible entering flow*, J. Comput. Phys., 101 (1992), pp. 151–164.
12. W. L. BRIGGS, S. F. MCCORMICK, AND V. E. HENSON, *Multigrid Tutorial, 2nd edition*, SIAM, USA, 2000.
13. B. DISKIN, *Multigrid algorithm with conditional coarsening for the non-aligned sonic flow*, Electronic Trans. Num. An., 6 (1997), pp. 106–119.
14. B. DISKIN, *Efficient multigrid methods for solving upwind-biased discretizations of the convection equation*, Applied Mathematics and Computation, 123 (2001), pp. 343–379. (Also ICASE Report 99-25, NASA CR-1999-209355).
15. B. DISKIN AND J. L. THOMAS, *Half-space analysis of the defect-correction method for Fromm discretization of convection*, SIAM J. Sci. Comp., 22 (2000), pp. 633–655.
16. B. DISKIN AND J. L. THOMAS, *Analysis of boundary conditions for factorizable discretizations of the Euler equations*, ICASE Report 2002-13, NASA CR-2002-211648, May 2002.
17. C. HIRSCH, *Numerical computation of internal and external flows. Vol.1, Fundamentals of numerical discretization*, A Wiley-Interscience publication, John Wiley & Sons, Inc., 605 Third Avenue, New York, NY 10158-0012, USA, 1988.
18. I. M. LLORENTE, M. PRIETO-MATIAS, AND B. DISKIN, *An efficient parallel multigrid solver for 3-D convection and convection-diffusion problems*, Parallel Computing, 27 (2001), pp. 1715–1741.
19. C. W. OOSTERLEE, F. J. GASPAR, T. WASHIO, AND R. WIENANDS, *Multigrid line smoothers for higher order upwind discretizations of convection-dominated problems*, J. Comput. Phys., 139 (1998), pp. 274–307.
20. T. W. ROBERTS, D. SIDILKOVER, AND R. C. SWANSON, *Textbook multigrid efficiency for the steady Euler equations*, AIAA Paper 97-1949, 13th AIAA CFD meeting, Snowmass Village, CO, June-July 1997.
21. T. W. ROBERTS AND R. C. SWANSON, *Extending ideally converging multigrid methods to airfoil flows*, AIAA Paper 99-3337, 14th AIAA CFD meeting, Norfolk, VA, June-July 1999.
22. T. W. ROBERTS, R. C. SWANSON, AND D. SIDILKOVER, *An algorithm for ideal multigrid convergence for the steady Euler equations*, Computers and Fluids, 28 (1999), pp. 427–442.
23. D. SIDILKOVER AND U. ASHER, *A multigrid solver for the steady-state Navier-Stokes equations using the pressure-Poisson formulation*, Matematica Aplicada e Computacional, 14 (1995), pp. 21–35.
24. K. STÜBEN AND U. TROTTEBERG, *Multigrid methods: Fundamental algorithms, model problem analysis and application*, in Multigrid Methods, W. Hackbusch and U. Trottenberg, eds., Lecture Notes in Math. 960, Springer-Verlag, Berlin, 1982, pp. 1–176.

25. R. C. SWANSON, *Towards optimal multigrid efficiency for the Navier-Stokes equations*, AIAA Paper 2001-2574, 15th AIAA CFD Conference, Anaheim, CA, June 2001.
26. S. TA'ASAN, *Essentially optimal multigrid method for steady state Euler equations*, AIAA Paper 95-0209, 33rd Aerospace Sciences Meeting and Exhibit, January 1995.
27. J. L. THOMAS, B. DISKIN, AND A. BRANDT, *Textbook multigrid efficiency for fluid simulations*, Ann. Rev. Fluid Mech., 35 (2003), pp. 317–340. An extended version has been published as ICASE Report 2002-16.
28. U. TROTTEBERG, C. W. OOSTERLEE, AND A. SCHÜLER, *Multigrid*, Academic Press, London, 2000.
29. P. WESSELING, *A robust and efficient multigrid method*, in Multigrid Methods, W. Hackbusch and U. Trottenberg, eds., Lecture Notes in Math. 960, Springer-Verlag, Berlin, 1982, pp. 614–630.
30. P. WESSELING, *An introduction to multigrid methods*, Pure and Applied Mathematics, John Wiley & Sons, Chichester, 1992.
31. R. WIENANDS AND C. W. OOSTERLEE, *On three-grid Fourier analysis for multigrid*, SIAM J. Scient. Comp., 23 (2001), pp. 651–671.

Appendix A

One-dimensional analysis

In this appendix, we apply a simplified one-dimensional analysis to nonlinear problems associated with stagnation flows. First, we consider a one-dimensional nonlinear convection equation

$$uu_x = f, \quad u \geq 0, \quad (\text{A1})$$

defined on the interval $x \in [-a, -b]$, $a > b \geq 0$, where stagnation can be enforced through the force term. The solution u of (A1) is assumed nonnegative, therefore, the standard boundary condition for this equation is defined at the left (inflow) end of the interval $[-a, -b]$.

$$u(-a) = U_a. \quad (\text{A2})$$

Upstream boundary conditions for convection equations are usually justified by considering a constant-coefficient problem corresponding to the free-stream linearization of convection-diffusion equations, in the limit of vanishing diffusion. Nonlinear stagnation flow problems usually have highly variable, (at least in the vicinity of stagnation) solutions that are very different from the the (nearly constant) free-stream solutions. The convection equation linearized around a spatially varying solution may favor another (more stable) boundary condition. Such a situation occurs when the velocity and the velocity derivative along the characteristic direction have the opposite signs (decelerating flow). Choosing the standard, but less stable, upstream boundary condition for a deceleration flow regime may lead to a reduced convergence basin (or even unconditional divergence) for the Newton iterations.

The boundary condition (A2) is in agreement with the linearization ($\bar{u}\partial_x(\delta u)$) around a constant (free-stream) solution ($u = \bar{u}, f = 0$). The fundamental solution of the homogeneous equation

$$\bar{u}\partial_x(\delta u) = 0 \quad (\text{A3})$$

is $\delta u \equiv 1$; therefore, a stable formulation can be obtained with boundary condition assigned at either side of the interval $x \in [-a, -b]$. The situation is very different if the linearization is performed around a stagnation-like decelerating solution $u = -x, f = x$. The linearized operator is defined now as

$$-x\partial_x(\delta u) - \delta u, \quad (\text{A4})$$

and the fundamental homogeneous solution, $\delta u = \frac{1}{x}$, grows to infinity as x tends to zero. For equation (A1), with a decelerating flow solution, a more stable boundary condition would be defined downstream, at the right (outflow) end of the interval,

$$u(-b) = U_b. \quad (\text{A5})$$

With the boundary condition (A5), the nonlinear convection problem (A1) is well posed, and the Newton iterations are converged unconditionally. With the boundary condition (A2), the problem (A1) is also well posed as long as $b > 0$. Convergence rate of the quasi-Newton iterations with operator (A4) as a driver is $\mu = \frac{a^2\epsilon^2}{b^2}$, where ϵ is a measure of the difference between the initial approximation and the exact solution. Obviously, if $b = 0$, the quasi-Newton iterations always diverge.

For one-dimensional system of equations

$$\begin{aligned} uu_x + p_x &= 0, \\ u_x^2 + p_{xx} &= 0, \end{aligned} \quad (\text{A6})$$

with solution

$$u = -x, \quad p = -\frac{x^2}{2}, \quad (\text{A7})$$

the linearization is given by the matrix

$$\begin{pmatrix} -x\partial_x - 1 & \partial_x \\ -2\partial_x & \partial_{xx} \end{pmatrix}, \quad (\text{A8})$$

and the fundamental homogeneous solutions, $\mathbf{q} = \begin{pmatrix} \delta u \\ \delta p \end{pmatrix}$, are

$$\mathbf{q}_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad \mathbf{q}_2 = \begin{pmatrix} 1 \\ x \end{pmatrix}, \quad \mathbf{q}_3 = \begin{pmatrix} x \\ x^2 \end{pmatrix}.$$

The quasi-Newton iterations with driver (A8) unconditionally converge quadratically for the standard set of boundary conditions:

$$u(-a) = a, \quad p_x(-a) = a, \quad p(-b) = -\frac{b^2}{2}; \quad (\text{A9})$$

With the S-driver defined as

$$\begin{pmatrix} -x\partial_x - 1 & \partial_x \\ 0 & \partial_{xx} \end{pmatrix}, \quad (\text{A10})$$

the fundamental homogeneous solutions are

$$\mathbf{q}_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad \mathbf{q}_2 = \begin{pmatrix} 1 \\ x \end{pmatrix}, \quad \mathbf{q}_3 = \begin{pmatrix} \frac{1}{x} \\ 0 \end{pmatrix}.$$

The S-driver iterations are only conditionally stable for $b > 0$ and diverge for $b = 0$. For two dimensional stagnation flows, no asymptotic divergence has been found in numerical iterations with the S-driver; however, an initial amplification of algebraic errors has been observed. This discrepancy between the analysis and the practical experience suggests that one-dimensional analysis is not capable to accurately predicting the two-dimensional behavior.

With the D-driver defined as

$$\begin{pmatrix} -x\partial_x + 1 & \partial_x \\ 0 & \partial_{xx} \end{pmatrix}, \quad (\text{A11})$$

the fundamental homogeneous solutions are

$$\mathbf{q}_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad \mathbf{q}_2 = \begin{pmatrix} -1 \\ x \end{pmatrix}, \quad \mathbf{q}_3 = \begin{pmatrix} x \\ 0 \end{pmatrix}.$$

The iterations with the D-driver are stable for (A9). This observation was the first motivation for considering D-driver schemes versus S-driver schemes.

Appendix B

Constant-Coefficient Analysis

As we have already mentioned in Section 5.1, the constant-coefficient approximations to the stagnation flow equations are valid only in restricted regions that do not include the major flow features such as the characteristic curve going into the stagnation point. Nevertheless, for completeness, we present here a stability analysis of the differential formulation for the constant-coefficient approximation to the plane-stagnation flow equations. This analysis does not take into account the compatible-coefficient restrictions (32).

The constant-coefficient differential equations are defined as

$$\mathbf{L}(\mathbf{q}) \equiv \begin{pmatrix} \bar{Q} - 1 & 0 & \partial_x \\ 0 & \bar{Q} + 1 & \partial_y \\ -2\partial_x & 2\partial_y & \Delta \end{pmatrix} \mathbf{q} = 0, \quad (\text{B1})$$

where $\mathbf{q} = (u, v, p)^T$, $\bar{Q} = \bar{u}\partial_x + \bar{v}\partial_y$, and \bar{u} and \bar{v} are positive constants; Δ is the Laplace operator. The computational domain is a strip $(x, y), x_0 \leq x \leq x_1$. The boundary $x = x_0$ is called the inflow boundary, and $x = x_1$ is the outflow boundary.

The fundamental solutions to the homogeneous equations (B1) are sought in the form

$$\mathbf{q}_k = \begin{pmatrix} u_k \\ v_k \\ p_k \end{pmatrix} e^{\lambda_k x + i\omega y}, \quad (\text{B2})$$

where u_k, v_k, p_k, λ_k are constants possibly dependent on the parameter ω . The coefficients λ_k are found as roots of the characteristic polynomial

$$((\bar{u}\lambda + i\bar{v}\omega)^2 + 1)(\lambda^2 - \omega^2) + 2(\bar{u}\lambda + i\bar{v}\omega)(\lambda^2 + \omega^2) = 0. \quad (\text{B3})$$

obtained after substituting $e^{\lambda_k x + i\omega y}$ into the equation

$$\det \mathbf{L}(\mathbf{q}) = 0. \quad (\text{B4})$$

For any ω , for which $Re(\lambda_k) \neq 0$, a formal necessary condition for stability of boundary conditions is that the number of roots $\lambda_k (Re(\lambda_k) < 0)$ should be equal to the number of the boundary conditions defined at the inflow boundary; the number of roots $\lambda_k (Re(\lambda_k) > 0)$ should be equal to the number of the boundary conditions defined at the outflow boundary. Unfortunately, the number of roots with a given sign of the real parts depends on ω . For $\bar{u} = \bar{v} = 1$, there exist an $\omega^*, \frac{\pi}{4} < \omega^* < \frac{\pi}{3}$, such that, for $0 < \omega < \omega^*$, there are three negative real parts and one positive real part; for $\omega > \omega^*$, there are one negative real part and three positive real parts. Thus, with standard set of boundary conditions defined as u, v , and $\partial_x p$ given at the inflow and p given at the outflow, the modes $\omega > \omega^*$ are unstable. Instability means that there are homogeneous solutions that grow exponentially from the boundary toward the interior; thus, small changes in boundary data may cause large changes in the solution.

Appendix C

Local Mode Analysis: Details

In this appendix, we present details pertaining to the LMF analyses of the relaxation scheme and two-level cycle for the formulation introduced in Section 5.3.3. The system of equations is defined in (31). The results were determined from parametric investigations using various mesh intervals and values of $v/u = \tan \alpha$; in each case, an ϵ -neighborhood was considered with $\epsilon = 0.5$. Uniform isotropic grids were used, so in general the considered domains were rectangular with different numbers of mesh intervals, N_x and N_y , in the x - and y -directions, respectively. The normalized Fourier frequencies, θ_x and θ_y , span the interval $[-\pi, \pi]$ with the increments $\delta\theta_x = 2\pi/N_x$ and $\delta\theta_y = 2\pi/N_y$, respectively.

The symbol of the S-driver relaxation is defined in (33). For lexicographic pointwise Gauss-Seidel relaxation of the Laplace operator in the pressure equation, the symbol, $\hat{G}S^h$, is replaced as

$$\hat{G}S^h = \frac{1}{h^2} (e^{-i\theta_x} + e^{-i\theta_y} - 4); \quad (\text{C1})$$

several other relaxation schemes of the Laplace operator have also been considered, including anti-lexicographic pointwise Gauss-Seidel relaxation, y -line-implicit Gauss-Seidel relaxation (both lexicographic and anti-lexicographic order) and a direct solve. All other symbols are defined in (34). The symbol of the D-driver relaxation is very similar to the S-driver relaxation symbol:

$$\hat{S}_h = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} - \begin{pmatrix} \hat{Q} + 1 & 0 & \hat{\partial}_x^c \\ 0 & \hat{Q} + 1 & \hat{\partial}_y^c \\ 0 & 0 & \hat{G}S^h \end{pmatrix}^{-1} \begin{pmatrix} \hat{Q} - 1 & 0 & \hat{\partial}_x^c \\ 0 & \hat{Q} + 1 & \hat{\partial}_y^c \\ -2\hat{\partial}_x^c & 2\hat{\partial}_y^c & \hat{\Delta}^h \end{pmatrix}, \quad (\text{C2})$$

Scheme	S-driver	D-driver	$h \rightarrow 0$
Point Lexicographic	0.514	0.512	0.500
Point Anti-lexicographic	0.514	0.513	0.500
y -Line Lexicographic	0.547	0.519	0.444
y -Line Anti-lexicographic	0.444	0.444	0.444
Solve	0.117	0.222	0.0

Table C1. Smoothing factors for various relaxations of the pressure equation

Table C1 summarizes the worst smoothing factors for the S-driver and D-driver from a set of parametric LMF analyses; the minimum number of mesh intervals was 4 and increased as powers of 2 and the values of $\tan \alpha$ were 0.5, 1, and 2. The additional entries (labeled $h \rightarrow 0$) correspond to the limiting case of very fine grids for either driver. The worst smoothing factors are somewhat higher than those on the very fine grids and depend on the order of relaxation because of the coupling between the equations that is much stronger on coarser grids. This latter dependence is most noticeable for y -line relaxation, which has better smoothing when the pressure is relaxed in the upstream order. However, all the smoothing factors are well bounded from unity.

The relaxation amplification factors from the parametric analyses are bounded by 1 everywhere except for the characteristic errors. As an example, the Table C2 collects the amplification and smoothing factors for the S- and D-drivers for the flow at 45 degrees with lexicographic pointwise Gauss-Seidel relaxation. As noted previously, the smoothing factors

S-Driver: N = 128, Maximal Amplification = 1.8030 D-Driver: N = 128, Maximal Amplification = 2.0000

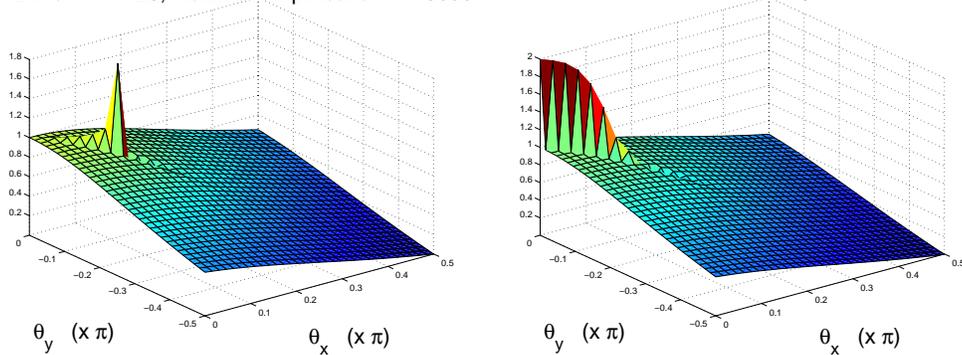


Figure C1. Relaxation amplification factors

N	S-driver		D-driver	
	Smoothing	Amplification	Smoothing	Amplification
8	0.51	1.00	0.51	2.00
16	0.51	1.14	0.51	2.00
32	0.50	1.17	0.50	2.00
58	0.50	18.66	0.50	2.00
64	0.50	1.46	0.50	2.00
128	0.50	1.80	0.50	2.00

Table C2. Smoothing and amplification factors for the ϵ -neighborhood corresponding to $\bar{u} = \bar{v} = 1, \epsilon = 0.5$.

are good. The amplification factor for the S-driver is erratic on the grids with a relatively big mesh size. On finer grids, the amplification factor tends to infinity for certain smooth (but not the smoothest) characteristic frequencies of the convection operator ($\theta_x = -\theta_y$), for which the symbol \hat{Q} approaches 1. On the coarse grids, the smallest possible value of $\hat{Q} - 1$ is determined by the meshsize h and the frequency increment $\delta\theta$; that explains the irregular behavior of the amplification factor. For illustration purposes, the results for a non-standard 59×59 grid are added to the table. At this grid, $h = 1/58$, and, for $\theta_x = -\theta_y = \frac{8}{58}\pi$, $\hat{Q} = 0.9909$. The amplification factor of the D-driver is 2 on all the grids. Figure C1 illustrate the amplification factor over a part of the frequency domain. Note that the large amplification factor for the S-driver occurs for very specific isolated smooth characteristic components of the convection operator. The amplification of the D-driver is also much larger for the smooth characteristic components than for other components, reaching its maximum at the smoothest component.

Table C3 compares the amplification factors of the two-grid (1,2) cycle with either S-driver or D-driver. The factors for the grid 65×65 are illustrated in Figure C2. Again, huge amplification factors are observed for isolated characteristic frequencies in the cycle with the S-driver. More regular but still diverging amplification factors are predicted for the cycle with the D-driver. The predicted instabilities of the two-grid cycle correspond to characteristic errors with moderate oscillations in the cross-characteristic direction. The characteristic errors with smoothest cross-characteristic oscillations are well reduced by the coarse grid correction but the correction for the moderate oscillations can be shown to be 0.75 and is not enough to suppress the instability in relaxation. In multigrid solvers for hyperbolic problems, the effects of boundary conditions on characteristic errors would be

N	S-driver	D-driver
8	0.14	0.36
16	1.71	1.46
32	2.60	1.22
58	31703.52	1.31
64	8.15	1.42
128	24.23	1.43

Table C3. Amplification factors for 2-level (1,2) cycle with different relaxation schemes

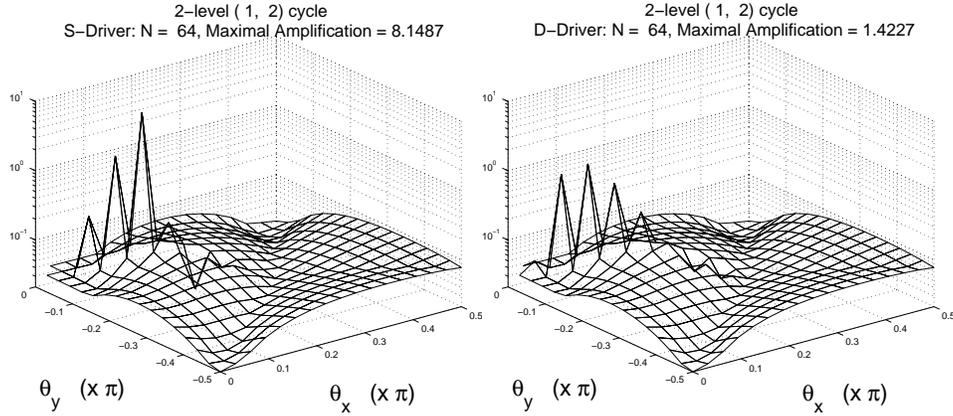


Figure C2. Amplification factors of a two-grid cycle

quite pronounced and, thus, violate the locality conditions discussed in Section 5.1. In this appendix, we do not rigorously enforce the locality conditions nor do we account for boundary conditions, that would be possible with a half-space analysis. Rather, we regard the results of the LMF analyses only as an indicator of possible amplifications and rely on other analyses (ICG and IR) with fewer restrictions developed herein for quantitative predictions.

