

A High-Resolution Euler Solver Based on Multigrid, Semi-coarsening, and Defect Correction*

WIM A. MULDER

Department of Mathematics, 405 Hilgard Avenue, University of California, Los Angeles, California 90024-1555

Received August 3, 1989; revised September 14, 1990

In an earlier paper, an $O(N)$ method for the computation of stationary solutions to the Euler equations of inviscid compressible gas dynamics has been described. The method is a variant of the multigrid technique and employs semi-coarsening in all co-ordinate directions simultaneously. It provides good convergence rates for first-order upwind discretisations even in the case of alignment, the flow being aligned with the grid. Here we discuss the application of this scheme to higher-order discretisations. Two-grid analysis for the linear constant-coefficient case shows that it is difficult to obtain uniformly good convergence rates for a higher-order scheme, because of waves perpendicular to stream lines. The defect correction technique suffers from the same problem. However, convergence to a point where the residual of the total error (the sum of the iteration error and the discretisation error) is of the order of the truncation error can be obtained in about seven defect correction cycles, according to estimates for the linear constant-coefficient equations. This result is explored for the nonlinear case by some illustrative numerical experiments. © 1992 Academic Press, Inc.

1. INTRODUCTION

A bottleneck in the application of the multigrid technique to the computation of inviscid stationary flows is alignment [2, 4]. Stream lines following grid lines become decoupled in the direction perpendicular to the flow. A high-frequency iteration error (deviation from the steady state) in that direction cannot be removed by smoothing because there is no coupling, nor by solving the equations on a coarser grid because high frequencies cannot be represented on the coarser grid. The result is slow convergence.

One way to deal with alignment is the use of global relaxation schemes, such as Gauss–Seidel relaxation or line-relaxation. In [10] it was shown that Gauss–Seidel and its symmetric variants can only partly handle alignment. There are still waves for which the multigrid method does not

converge. For the Euler equations in two dimensions, alternating direction damped line Jacobi can provide a uniformly good convergence rate [14].

Another way to tackle alignment is the use of semi-coarsening. This approach is followed in [11]. For arbitrary flows, semi-coarsening must be carried out in all co-ordinate directions simultaneously. To accomplish this, the method described in [11] employs multiple coarser grids on a given level of coarseness, while maintaining its $O(N)$ complexity. A uniformly good convergence rate can be obtained for a first-order upwind discretisation of the two-dimensional Euler equations. The method provides a nonlinear alternative to line-relaxation.

In this paper we will discuss the application of the same method to higher-order upwind discretisations of the Euler equations in two dimensions. The spatial discretisation is based on van Leer's kappa-schemes [20] and provides second- or third-order accuracy. Details are given in Section 2. The multigrid method is reviewed in Section 3. In [12], it has been shown that the convergence rates for higher-order discretisations are poor. This result is explained here in a simpler way. The defect correction technique (cf. [4]) cannot provide good convergence rates either. However, it *can* provide convergence to a point where the residual of the total error (the sum of the iteration error and the discretisation error) is of the order of the truncation error in about seven defect correction cycles, at least in the linear constant-coefficient case. Numerical experiments are carried out to explore this result for the nonlinear equations. Details of the nonlinear implementation are described in Section 4 and results are presented in Section 5.

2. SPATIAL DISCRETISATION

The Euler equations of gas dynamics that describe the flow of a perfect inviscid compressible gas are

$$\frac{\partial w}{\partial t} + \frac{\partial f}{\partial x} + \frac{\partial g}{\partial y} = 0. \quad (2.1a)$$

* A shorter version has been presented at the AIAA 9th Computational Fluid Dynamics Conference, Buffalo, NY, June 1989. This work was supported by NSF Grant DMS88-11863 and ONR Grant N00014-86-K-0691. The author's present address is: Royal Dutch/Shell, Exploration and Production Laboratory, P.O. Box 60, 2280AB Rijswijk, the Netherlands.

Here the vector of states w and the fluxes f and g are given by

$$w = \begin{pmatrix} \rho \\ \rho u \\ \rho v \\ \rho E \end{pmatrix}, \quad f = \begin{pmatrix} \rho u \\ \rho u^2 + p \\ \rho uv \\ \rho uH \end{pmatrix}, \quad g = \begin{pmatrix} \rho v \\ \rho uv \\ \rho v^2 + p \\ \rho vH \end{pmatrix}. \quad (2.1b)$$

The density is denoted by ρ , and u and v are the x - and y -component of the velocity. The energy E , total enthalpy H , pressure p , and sound speed c are related by

$$E = \frac{1}{\gamma - 1} \frac{p}{\rho} + \frac{1}{2} (u^2 + v^2), \quad (2.1c)$$

$$H = E + \frac{p}{\rho}, \quad c^2 = \gamma \frac{p}{\rho}.$$

Our aim is the efficient computation of the steady state $\partial w / \partial t = 0$. The spatial part of the system is discretised by upwind differencing. The vector of state quantities w is represented by cell averages $w_{i,j}$ on a grid consisting of arbitrary quadrilaterals, having four vertices denoted by $(i - \frac{1}{2}, j - \frac{1}{2})$, $(i + \frac{1}{2}, j - \frac{1}{2})$, $(i + \frac{1}{2}, j + \frac{1}{2})$, and $(i - \frac{1}{2}, j + \frac{1}{2})$, and midpoints of faces labelled as $(i, j - \frac{1}{2})$, $(i + \frac{1}{2}, j)$, $(i, j + \frac{1}{2})$, $(i - \frac{1}{2}, j)$. This configuration is sketched in Fig. 1. The integral of $\partial f / \partial x + \partial g / \partial y$ over the cell area is discretised as

$$r_{i,j} = l_{i,j-1/2} T_{i,j-1/2}^{-1} f(T_{i,j-1/2} w(q_{i,j}^{j-}), T_{i,j-1/2} w(q_{i,j-1}^{j+}))$$

$$+ l_{i+1/2,j} T_{i+1/2,j}^{-1} f(T_{i+1/2,j} w(q_{i,j}^{i+}), T_{i+1/2,j} w(q_{i+1,j}^{i-}))$$

$$+ l_{i,j+1/2} T_{i,j+1/2}^{-1} f(T_{i,j+1/2} w(q_{i,j}^{i+}), T_{i,j+1/2} w(q_{i,j+1}^{j-}))$$

$$+ l_{i-1/2,j} T_{i-1/2,j}^{-1} f(T_{i-1/2,j} w(q_{i,j}^{i-}), T_{i-1/2,j} w(q_{i-1,j}^{i+})). \quad (2.2a)$$

This expression will be referred to as the residual. Note that the fluxes f and g can be transformed into each other by rotation. Each rotation matrix T transforms w in such a way that the second component of Tw is the outward momen-




FIG. 1. Quadrilateral cell (i, j) showing the numbering of cell faces and vertices.

tum perpendicular to the cell face, and the third is the momentum parallel to the cell face. In (2.2a), $f(w_l, w_r)$ is a numerical flux corresponding to f in (2.1b) that provides an approximate solution to the Riemann problem in the direction perpendicular to the cell face. In this paper we will use either van Leer's flux-vector splitting [19] or Osher's scheme in the natural ordering [15, 16]. Both are sufficiently smooth (Lipshitz continuous) for our purpose. The first rotation matrix at the midpoint of face $(i, j - \frac{1}{2})$ is given by

$$T_{i,j-1/2} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \phi_{i,j-1/2} & \sin \phi_{i,j-1/2} & 0 \\ 0 & -\sin \phi_{i,j-1/2} & \cos \phi_{i,j-1/2} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (2.2b)$$

where

$$\cos \phi_{i,j-1/2} = (y_{i+1/2,j-1/2} - y_{i-1/2,j-1/2}) / l_{i,j-1/2},$$

$$\sin \phi_{i,j-1/2} = -(x_{i+1/2,j-1/2} - x_{i-1/2,j-1/2}) / l_{i,j-1/2}, \quad (2.2c)$$

$$l_{i,j-1/2} = [(x_{i+1/2,j-1/2} - x_{i-1/2,j-1/2})^2 + (y_{i+1/2,j-1/2} - y_{i-1/2,j-1/2})^2]^{1/2}.$$

The outward normal is $(\cos \phi_{i,j-1/2}, \sin \phi_{i,j-1/2})^T$. The length of the cell face is $l_{i,j-1/2}$. The rotation matrices for the other sides follow in a similar way. The expression $w(q)$ denotes a one-to-one transformation from a set of state quantities q to w . The quantities $q_{i,j}^{\pm}$ and $q_{i,j}^{\pm}$ are values at the cell faces obtained by interpolation from the state $q_{i,j} = q(w_{i,j})$. A first-order-accurate scheme is obtained if the interpolated values simply equal the interior values. Second-order accuracy is obtained by using van Leer's kappa-scheme [1, 20] which lets, in the i -direction,

$$q_{i,j}^{i+} = q_{i,j} + \frac{1}{4} s(\mathcal{A}_{i,j}^-, \mathcal{A}_{i,j}^+) \times [(1 - \kappa) \mathcal{A}_{i,j}^- + (1 + \kappa) \mathcal{A}_{i,j}^+], \quad (2.3a)$$

$$q_{i,j}^{i-} = q_{i,j} - \frac{1}{4} s(\mathcal{A}_{i,j}^-, \mathcal{A}_{i,j}^+) \times [(1 - \kappa) \mathcal{A}_{i,j}^+ + (1 + \kappa) \mathcal{A}_{i,j}^-].$$

Here

$$\mathcal{A}_{i,j}^- = q_{i,j} - q_{i-1,j}, \quad \mathcal{A}_{i,j}^+ = q_{i+1,j} - q_{i,j}. \quad (2.3b)$$

The function $s(\mathcal{A}^-, \mathcal{A}^+)$ is a limiter that prevents numerical oscillations. We use a smooth limiter due to van Albada [18]:

$$s(\mathcal{A}^-, \mathcal{A}^+) = \frac{2(\mathcal{A}^- \mathcal{A}^+ + \varepsilon_a^2)}{(\mathcal{A}^-)^2 + (\mathcal{A}^+)^2 + 2\varepsilon_a^2}. \quad (2.3c)$$

The constant ε_a prevents division by zero. We use $\varepsilon_a^2 = 10^{-16}$ in the numerical experiments of Section 5. Expressions similar to (2.3) are used for the j -direction.

The standard second-order upwind scheme is obtained for $\kappa = 0$. Central differencing is obtained with $\kappa = 1$, if the limiter is not used ($s(\Delta^-, \Delta^+) = 1$). The choice $\kappa = -1$ provides a one-sided upwind scheme. For $\kappa = \frac{1}{3}$, we obtain third-order accuracy in a point-wise sense, but not in a volume-averaged sense, because there is a second-order difference between point-values and volume-averages. The limiter may cause the accuracy to reduce to first-order at isolated points. Also, the flux of the average state at the mid-point of the cell face is not equal to the average of the flux over the cell face. Furthermore, steady discontinuities are smeared out over at least two cells, which results in a local $O(1)$ error. Thus, we have at most second-order accuracy in smooth regions of the flow and first-order or even zero-order accuracy at isolated points or lines. For this reason, the discretisation is referred to as a *high-resolution* scheme. An additional problem occurs if the grid is not locally Cartesian or if cell sizes vary strongly from one cell to another. Then the one-dimensional interpolation (2.3) should be corrected for stretching and curvature. Here we will assume that the grid is locally Cartesian without significant stretching from cell to cell. Note that the aspect ratio of the cells is not involved in this discussion; it may be far away from 1.

So far, we have assumed that the independent variables are the conserved variables w of (2.1). It has been pointed out by the authors of [6] that another set can be used as well. Here we use

$$W = (\rho, u, v, c)^T, \quad (2.4)$$

as the set of independent variables. Other choices, not considered here, are the entropy S instead of ρ , or the pressure p instead of c . The variables for the interpolation (2.3) in the kappa-scheme are chosen to be $q = (S, u, v, c)^T$. This choice makes it easier to perform the correct characteristic switching near boundaries. For outgoing characteristics in the direction perpendicular to the boundary, we extrapolate differences of the characteristic variables, whereas for incoming characteristics, the differences are computed with respect to the given exterior variables. After this has been done, we transform back to differences of q .

The residual is measured by defining a quantity

$$R_{i,j} = \max_{k=1,\dots,4} \left(\frac{|r_{i,j,k}|}{|W_{i,j,k}| + h_{i,j,k}} \right), \quad (2.5)$$

where $h_{i,j,2} = h_{i,j,3} = c_{i,j}$ and $h_{i,j,1} = h_{i,j,4} = 0$, computing its l_1 -norm by adding $R_{i,j}$ over all cells, and dividing by the total area of the computational domain.

3. MULTIGRID AND DEFECT CORRECTION

The multigrid variant used in this paper has been presented in [11]. The method employs semi-coarsening in two directions simultaneously (for a two-dimensional problem). Figure 2 shows the various grids employed if the finest grid has 8×8 points and the coarsest has 1×1 . For a problem in d dimensions, the total number of points on all grids is $2^d N$, where N is the number of points on the finest grid. The cost of a V-cycle is proportional to this number, whereas the cost of an F-cycle is proportional to $(d+1) 2^d N$. For a W-cycle, the $O(N)$ complexity is lost. A definition of V-, F-, and W-cycles can be found in [2].

The usual restriction and prolongation operators have to be modified to handle input from more than one grid. If one grid needs data from two finer grids, the two sets of data obtained by the restriction from each finer grid are averaged with equal weights. For prolongation, the correction is computed with respect to the latest fine-grid solution available, which now may be different from the one used during restriction.

This method is useful for any problem with strong anisotropy. It cannot, however, handle alignment at 45° . (A method developed by Hackbusch [5] treats this case correctly at an increased complexity of $O(N \log N)$. It is not clear, however, how to construct stable coarse-grid operators for his scheme in the nonlinear case.) In [11], it is shown that alignment at 45° is not a problem for the Euler equations of gas dynamics if first-order upwind differencing is used. Two-grid analysis for the linearised Euler equations with constant coefficients leads to a worst-case *two-grid* convergence rate of 0.5 per cycle, if damped Point-Jacobi is used as smoother. Numerical experiments on the nonlinear Euler equations show *multigrid* convergence rates better than 0.5.

A fundamental problem encountered in extending the method to higher-order discretisations is the fact that the




FIG. 2. Arrangement of finest (8×8) and coarser grids, that leads to an $O(N)$ multigrid method for problems with alignment. The arrows indicate how the grids are linked by restriction (downward) and prolongation (upward).

exact operator vanishes for waves perpendicular to stream lines. Alignment, as mentioned earlier, is one of the results. Another one has been described in [12]: for a scheme of order p , the worst-case two-grid convergence rate cannot be better than $1 - 2^{-p}$. Note that the first-order ($p = 1$) solver described above actually obtains this value as its worst-case two-grid convergence rate.

This lower limit can be explained as follows: Consider the linear scalar operator

$$L = u \frac{\partial}{\partial x} + v \frac{\partial}{\partial y}, \quad \text{with } u > 0, v \geq 0. \quad (3.1a)$$

The symbol of this operator is

$$\hat{L} = i(u\omega_x + v\omega_y), \quad (3.1b)$$

where ω_x and ω_y are frequencies. It vanishes if $v\omega_y = -u\omega_x$, i.e., for waves perpendicular to stream lines (characteristics). Suppose this operator is discretised on a grid with cell size h . This yields a discrete operator L^h , which has a truncation error τ^h of $O(h^p)$. Let the solution be computed with a two-grid method, involving a fine and a coarse grid. The coarse-grid correction operator is

$$K = I - I_{2h}^h (L^{2h})^{-1} I_h^{2h} L^h. \quad (3.2)$$

First the fine-grid residual L^h is restricted to a coarser grid, using a restriction operator I_h^{2h} . Then the coarse-grid problem, which involves the residual operator L^{2h} , is solved exactly. Finally, the fine-grid solution is corrected by the coarse-grid result, using a prolongation operator I_{2h}^h . The coarse-grid correction operator should remove low-frequency iteration errors, and is used in combination with a smoother that removes high-frequency errors. Because the smoother is usually inefficient for low frequencies, the convergence of low-frequency iteration errors depends almost entirely on K . For the lowest frequencies (ω_x and ω_y close to zero), the restriction and prolongation operators have practically no effect, so

$$K \simeq I - (L^{2h})^{-1} L^h. \quad (3.3)$$

If the exact operator vanishes, the discrete operator L^h equals the truncation error τ^h . For those waves, $K \simeq 1 - \tau^h/\tau^{2h}$. If the scheme is of order p , then $\tau^h/\tau^{2h} \simeq h^p/(2h)^p$, so $K \simeq 1 - 2^{-p}$. This implies that the worst-case convergence rate is at best $1 - 2^{-p}$. For a first-order scheme, the resulting 0.5 is acceptable as the convergence rate. Values larger than that are not attractive, because one must carry out many cycles on coarser levels to obtain a sufficiently accurate solution of the coarse-grid equations. This may increase the complexity beyond $O(N)$.

For higher-order schemes, the situation is actually worse than suggested by the above estimate. If first-order restriction (volume-averaging) and prolongation (piecewise constant interpolation) operators are chosen, the coarse-grid correction operator for the second- and third-order scheme is unstable. The third-order scheme ($\kappa = \frac{1}{3}$) can be stabilized by using a third-order restriction operator. The second-order scheme can be stabilized by using residuals on coarser grids that are first-order in the direction of semi-coarsening. This, however, increases the lower limit for the worst-case convergence rate to 1 (no convergence at all). The reader is referred to [12] for details.

The above implies that one cannot design a multigrid scheme with a uniformly good convergence rate for a spatial discretisation based on second-order upwind differencing. For a third-order scheme, one might be able to obtain a convergence rate of at best $\frac{7}{8}$, which is not very impressive. However, these conclusions are too pessimistic, as the values are dominated by those waves for which the exact operator becomes of the order of the truncation error. Because it does not make much sense to require convergence below the truncation error, the numbers found are not representative for the performance of the multigrid method.

There does not appear to be a simple way to measure convergence relative to the size of the truncation error in the context of local mode analysis. It can, however, be easily accomplished in the framework of the defect correction technique (cf. [4]).

It should be noted that defect correction cannot provide a uniformly good convergence rate for the present problem. Consider a linear operator L_p^h with an order of accuracy $p > 1$ and an exact first-order solver. The iteration operator for defect correction is

$$I - (L_1^h)^{-1} L_p^h. \quad (3.4)$$

If the exact operator vanishes, we obtain a convergence rate $1 - \tau_p^h/\tau_1^h = 1 - O(h^{p-1})$. However, this results is again dominated by the truncation error.

Better estimates can be obtained if convergence down to machine-zero is abandoned. For the linear constant-coefficient case, it can be shown that the residual of the total error (the difference between the current guess of the numerical steady state and the projection of the exact solution to the same grid) is bounded by an $O(1)$ constant times the truncation error τ_p^h (see the Appendix). This bound can be reached within 10% in about seven defect correction steps, implying that seven steps are sufficient to obtain a higher-order solution. To derive this result, it is assumed that the problem is linear with constant coefficient and periodic boundary conditions. Successive grid refinement is used, starting with a fully converged solution on the first coarsest grid and interpolating the solution after a fixed

number of defect correction steps to subsequently finer grids. The first-order solver used for defect correction is assumed to have a convergence rate 0.5.

An important assumption for the above estimate is that new *global* structures should not suddenly appear after interpolation to a finer grid (unless they are of the order of the discretisation error). This assumption will be violated, if for instance, there are details on the boundary of the computational domain that are not resolved on coarser grids and suddenly appear after grid refinement. One could think of a totally different inflow condition in one cell on a boundary with a constant inflow rate in all other cells. This will give rise to a large-scale structure along a stream line that requires many more defect correction cycles than the seven mentioned above. An alternative is the use of a global relaxation scheme on the higher-order residual that converges rapidly for this structure. In [3], Gauss–Seidel relaxation in the flow direction is suggested. Unfortunately, this scheme is unstable when used with higher-order upwind differencing, unless full upwind weighting is used ($\kappa = -1$). For the numerical experiments in Section 5, we will use constant inflow boundaries, so this issue will not be important.

If *locally* large errors show up after grid refinement, the estimate of seven defect correction cycles is not valid either. However, local errors can be easily removed by some extra smoothing steps on the higher-order residual. Details will be given in Section 5.

It should be noted that defect correction has been applied to the Euler equations in earlier work [8, 9], using van Leer’s flux–vector splitting. This scheme smears slip lines, and therefore does not suffer from alignment if damped symmetric Gauss–Seidel is used as smoother [10]. A fundamentally different approach is chosen in [7, 17], using a first-order solver described in [6]. These authors start out with a first-order solution on the finest grid. Because their first-order solver is not exact, not even for the long waves, at least $O(\log N)$ iterations with the defect-correction technique are required, thus leading to an $O(N \log N)$ complexity. The situation is actually worse, because their first-order solver suffers from alignment.

4. NONLINEAR IMPLEMENTATION

Here we discuss some details of the nonlinear implementation of the defect correction method with the first-order solver based on semi-coarsening [11]. The method starts on a not too coarse initial grid. Defect correction steps with a first-order solver are carried out until the higher-order numerical solution has converged. The result is interpolated to a grid with twice the number of cells in each co-ordinate direction. Then, a *fixed* number of defect correction steps is performed. The result is interpolated to the next finer grid, etc. We choose to perform eight defect correction cycles on each subsequently finer level—one more than suggested by

the linear analysis. The aim is to show that eight cycles with the first-order solver are sufficient to obtain a higher-order solution. Between defect correction cycles, one smoothing step with a two-stage scheme is carried out on the higher-order residual to remove locally large errors near shocks or other singularities.

The upwind discretisation has been outlined in Section 2. We use van Leer’s flux–vector splitting [19] or Osher’s scheme in the natural ordering [15, 16] as an approximate Riemann solver. Only $\kappa = \frac{1}{3}$ (“third order”) is considered for the high-resolution discretisation.

In a defect correction step, the first-order solver acts on the sum of the most recent first-order residual $r_1(W')$ and a correction term Q , which is the difference between the higher-order residual $r_p(W)$ and first-order residual $r_1(W)$ at the beginning of the cycle. That is, the first-order solver finds the approximate solution W' to

$$r_1(W') + Q = 0, \quad Q = r_p(W) - r_1(W). \quad (4.1)$$

After this has been done, the higher-order solution is updated by $W := W'$. Because the high-resolution residual, in general, ceases to converge if the iteration error becomes of the order of the truncation error, examining its convergence rate does not provide much information. It is useful to monitor the convergence rate of the first-order solver (at the cost of computing an extra fine-grid residual). If a cycle provides a first-order convergence rate worse than 0.6, some extra smoothing with damped Point–Jacobi is applied. If this does not help, the cycle is repeated. In the examples presented in Section 5, this never occurred except in the initial stages of successive grid refinement at very coarse grids.

The first-order solver has been described briefly in Section 3 and in detail in [11]. As already noted in Section 3, a standard multigrid approach can be used for van Leer’s flux–vector splitting, because it smears stream lines and does not suffer from alignment. Here we choose to use the more complex first-order solver based on semi-coarsening, because it can handle Osher’s scheme as well.

The restriction and prolongation operators for the first-order solver are volume-averaging and piecewise constant interpolation, respectively. Both are first-order operators. F-cycles are used, with one post-smoothing step on each grid.

The smoother for the first-order solver is damped Point–Jacobi. In the nonlinear case, this scheme determines $\tilde{W}_{i,j}$ from the equation

$$r_{1i,j}(\tilde{W}_{i,j}, W_{i-1,j}, W_{i+1,j}, W_{i,j-1}, W_{i,j+1}) + Q_{i,j} = 0 \quad (4.2a)$$

for all i and j and then updates the solution according to

$$W := \frac{1}{2}(W + \tilde{W}). \quad (4.2b)$$

A less costly (but potentially less robust) approach is obtained by performing just one Newton step,

$$W_{i,j} := W_{i,j} - \frac{1}{2} N_{i,j}^{-1} (r_{1i,j} + Q_{i,j}), \quad (4.3)$$

where $N_{i,j} = \partial r_{1i,j} / \partial W_{i,j}$. This scheme is adopted here.

For van Leer's flux-vector splitting, the matrix $N_{i,j}$ is non-negative. To make it positive, we replace it by

$$\tilde{N}_{i,j} = N_{i,j} + \sigma I, \quad (4.4)$$

where σ is a positive scalar. Following [21], the parameter σ in (4.4) is chosen to be $\sigma = \varepsilon_\sigma^{-1} \max_{i,j} R_{i,j}$, with $R_{i,j}$ given in (2.5). This decreases the change in the solution if the residual is large and helps to avoid negative values of ρ and c , although this is not guaranteed. The parameter ε_σ controls the relative change in the solution and is chosen to be 1. Smaller values can be used for complex flows with strong shocks, but then the number of relaxation sweeps must be increased.

For Osher's scheme, this approach fails because the matrix $N_{i,j}$ can be extremely ill-conditioned. It can be made non-negative by using an approximate linearisation. This is done as follows. For the first-order scheme, the first term on the right-hand side of (2.2a) contributes a matrix

$$l_{i,j-1/2} T_{i,j-1/2}^{-1} A_{i,j-1/2} T_{i,j-1/2}, \quad (4.5a)$$

where

$$A_{i,j-1/2} = \frac{\partial f_{i,j-1/2}}{\partial W'_{i,j}}, \quad W'_{i,j} = T_{i,j-1/2} W_{i,j}. \quad (4.5b)$$

For Osher's scheme, we replace $A_{i,j-1/2}$ by $A_{i,j-1/2}^+(W'_{i,j})$, where $A^+(W)$ is obtained from $A(W) = df(W)/dW$ by transforming to its diagonal form, setting the negative eigenvalues to zero, and transforming back. This approximate linearisation performed satisfactorily in [11]. The resulting non-negative matrix can be made positive by (4.4).

Near boundaries we use a linearisation that is the same as for interior cells. Any special dependencies of boundary values on solution and exterior values are ignored.

For the grid refinement of the high-resolution solution we use third-order interpolation. Within each cell on the coarser grid, a linear distribution of states is computed in a manner similar to (2.3), but with $\kappa = 0$. The initial guess on the finer grid is obtained by evaluating the values of the linear distribution at the four new cell centers.

After grid refinement and between defect correction cycles, extra smoothing can be applied on the higher-order residual. This is necessary if the solution is not smooth. We use a two-stage scheme:

$$\begin{aligned} W^* &:= W - \beta_1 \tilde{N}(W)^{-1} r_p(W), \\ W &:= W - \beta_2 \tilde{N}(W^*)^{-1} r_p(W^*). \end{aligned} \quad (4.6)$$

A less costly version uses $\tilde{N}(W)$ instead of $\tilde{N}(W^*)$ for the second step. In the computation of \tilde{N} , we ignore the interpolation from cell centers to cell faces. The parameters β_1 and β_2 can be chosen as to provide an optimal smoothing rate. A (conjectured) optimal choice for the scalar convection equation $uw_x + vw_y = 0$ with third-order upwind differencing ($\kappa = \frac{1}{3}$) under the requirement that the smoother is total variation non-increasing and linearly stable, is $2\beta_1 = \beta_2 = 6(7 - 2\sqrt{5})/29 = 0.523$. The resulting smoothing rate is $\bar{\mu} = 1 - \frac{1}{3}\beta_2(1 - \frac{1}{3}\beta_1) = 0.841$. Details are given in the Appendix. In the examples of the next section, scheme (4.6) is applied once between multigrid defect correction cycles.

A problem in any nonlinear multigrid method is the occurrence of solution values outside the admissible range. Negative densities and sound speeds may occur after prolongation and relaxation. Experiments with monotone prolongation did not lead to satisfactory convergence rates for the first-order solver. In our code, we reject the entire coarse-grid correction if the result is inadmissible, even if this happens only in one cell. Given the redundancy apparent in Fig. 2, eliminating parts of the data structure does not necessarily lead to loss of convergence. If relaxation leads to inadmissible values, we simply do not update the corresponding cell. A different nonlinear multigrid method by Hackbusch [4] may provide a more robust code, but this has not been explored.

Another problem is the occurrence of strong discontinuities. These can cause large $O(1)$ iteration errors after prolongation and grid refinement. Since these errors are local, they can be removed by additional smoothing. In a




FIG. 3. (a) Mach-lines for Mach 0.5 inflow through a channel with a \sin^2 bump on a 128×64 grid. Contours are 0.025 apart. (b) Restriction to 64×32 of the 128×64 grid used for Fig. 3a.

one-dimensional study reported in [13], the additional smoothing is carried out only locally if the residuals are large, at small extra cost. Local relaxation has not been used in the following section, as no solutions with strong discontinuities are considered, but it is expected to be necessary in general. An alternative approach to avoid large errors after prolongation and grid refinement is interpolation based on the operator (cf. [4], Section 10.3). In the present context, this would lead to interpolation of fluxes. The computation of states from fluxes involves a decision between a supersonic and subsonic solution. This remains to be investigated.

5. NUMERICAL EXPERIMENTS

As a first example, we consider a smooth flow through a channel with a $\sin^2(\pi x)$ bump, having a thickness 0.1 over a length 1. The length of the channel is 5, its height 2. The free-stream Mach-number is 0.5. At the left boundary,

inflow conditions are given by the free-stream values for total enthalpy, entropy, and inflow angle (0°). At the outlet, the free-stream value of the static pressure is imposed. It should be noted that characteristic in- and outflow conditions result in a weak boundary layer near the outlet and therefore are not used. Characteristic boundary conditions are used at the walls. Figure 3a shows Mach-lines for a 128×64 grid, using Osher's scheme. The restriction of the grid to 64×32 is presented in Fig. 3b, the finer grid being

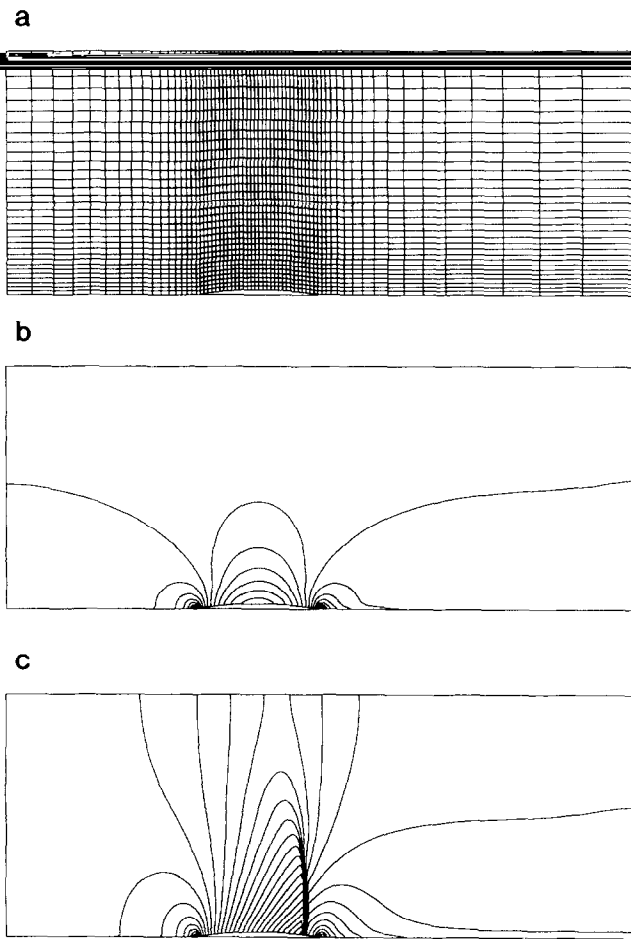


FIG. 4. (a) Restriction to 64×32 of the 128×64 grid used for Figs. 4b and 4c. (b) Mach-lines for Mach 0.5 inflow through a channel with a non-smooth bump on a 128×64 grid. Contours are 0.01 apart. (c) Mach-lines for Mach 0.85 inflow through a channel with a non-smooth bump on a 128×64 grid. Contours are 0.025 apart.




FIG. 5. (a) Mach-lines for Mach 1.4 inflow through a channel with a non-smooth bump on a 128×64 grid. Contours are 0.025 apart. (b) As Fig. 5a, but after 100 defect correction cycles instead of 8. (c) As Fig. 5a, but for a first-order discretisation. (d) Restriction to 64×32 of the 128×64 grid used for Figs. 5a-c.

too dense to be properly displayed. The coarsest grid used has size 4×2 . The relative order of accuracy of the solution can be estimated from

$$p = \log_2 \left(\frac{\|W^{4h} - I_{2h}^{4h} W^{2h}\|}{\|W^{2h} - I_h^{2h} W^h\|} \right). \quad (5.1)$$

Here I_h^{2h} and I_{2h}^{4h} are restriction operators based on cell averaging. For the example in Fig. 3 we find 2.5 for ρ , 2.5 for u , 2.2 for v , and 2.5 for c in the l_1 -norm. In the l_∞ -norm, values between 1.5 and 2 are obtained. The multigrid convergence rate of the first-order solver is well below 0.5. The defect correction convergence rate for the high-resolution residual on the finest grid, averaged over eight cycles, is 0.68. This experiment suggests that the linear estimates provide a reasonable description of the nonlinear scheme for smooth solutions.

Next we consider flow through a channel with a non-smooth bump. The bump is a circular arc with a thickness of 4.2% of the chord. The length of the channel is 5, its height 2. The 128×64 grid is clustered near the bump. Its restriction to 64×32 is shown in Fig. 4a. Figure 4b shows Mach-lines for Mach 0.5 inflow, using Osher's scheme. The relative orders of accuracy in the l_1 -norm are 1.3, 1.1, 1.4, and 1.3 in ρ , u , v , and c . These low values are not due to the iteration error; the same results are found after 100 (instead of 8) defect correction cycles when the residual has decreased with a mean convergence rate of about 0.93. The singularities at the endpoints of the bump lead to sharp peaks in the solution, which are responsible for the decrease of accuracy in l_1 . These singularities also pollute the downstream solution near the lower wall somewhat, as can be seen from Fig. 4b.

Mach-line for Mach 0.85 on the same grid are shown in

TABLE I

Lift and Drag for a NACA0012 Airfoil on Various Grids for Osher's Scheme and van Leer's Flux-Vector Splitting (FVS)

M_∞, α	Grid	Lift(8)	Drag(8)	$\lambda(8)$	Lift(100)	Drag(100)	$\lambda(100)$
0.80	32×16	0.3158	0.0488	0.77	0.3097	0.0481	0.88
1.25°	64×32	0.3465	0.0286	0.77	0.3497	0.0284	0.93
Osher	128×64	0.3517	0.0233	0.76	0.3525	0.0235	0.92
0.80	32×16	0.3356	0.0561	0.75	0.3313	0.0564	0.95
1.25°	64×32	0.3630	0.0311	0.74	0.3675	0.0312	0.94
FVS	128×64	0.3679	0.0247	0.74	0.3675	0.0247	0.80
1.20	32×16	0.506	0.169	0.71	0.508	0.169	0.80
7.00°	64×32	0.530	0.160	0.71	0.522	0.159	0.82
Osher	128×64	0.525	0.155	0.75	0.524	0.155	0.90
1.20	32×16	0.510	0.175	0.72	0.510	0.173	0.92
7.00°	64×32	0.527	0.160	0.72	0.523	0.160	0.85
FVS	128×64	0.523	0.155	0.80	0.524	0.155	0.97

Note. M_∞ denotes the free-stream mach-number and α the angle of attack. Results are given for 8 and 100 defect correction cycles per grid refinement level. Also included is the mean convergence rate λ .

Fig. 4c. Figure 5a shows Mach 1.4 inflow on a different 128×64 grid (see Fig. 5d). Both results are obtained after eight defect correction cycles per level, using Osher's scheme. The relative orders of accuracy do not reflect the dramatic improvement with respect to the first-order solution (cf. [11] and Fig. 5c). For Mach 0.85 inflow, we obtain estimates of p around 1.2 in l_1 ; for Mach 1.4 inflow values around 0.6 are obtained. Because the estimates are dominated by the discontinuities, these results are not very illuminating.

As an alternative, we can make a comparison between the size of the residual r^h after eight cycles and the truncation




FIG. 6. (a) C_p -curves for a NACA0012 airfoil ($M_\infty = 0.8$, $\alpha = 1.25^\circ$) for Osher's scheme and van Leer's FVS (dashed line). (b) As Fig. 6a, but for $M_\infty = 1.2$ and $\alpha = 7^\circ$.

error τ^h . The latter can be estimated from the relative truncation error

$$\tau_h^{2h} = I_h^{2h} r^h(W^h) - r^{2h}(I_h^{2h} W^h), \quad (5.2)$$

using $\tau^h \simeq (2^p - 1) \tau_h^{2h}$. Define the ratios

$$\psi_1 = \frac{1}{4} \sum_{k=1}^4 \frac{\|r_k^h\|_1}{\|\tau_{hk}^{2h}\|_1}, \quad (5.3)$$

$$\psi_\infty = \max_{k=1,\dots,4} \frac{\|r_k^h\|_\infty}{\|\tau_{hk}^{2h}\|_\infty},$$

where k denotes each of the four components, and either the l_1 - or the l_∞ -norm is used. For the Mach 0.85 example, we find $\psi_1 = 0.19$ and $\psi_\infty = 0.45$, whereas for the Mach 1.4 problem, 0.22 and 0.23 are found. Thus, the residual has converged to the level of the truncation error, which is consistent with the linear analysis.

To illustrate the effect of the remaining iteration error (which is of the order of the discretisation error) on the solution, Fig. 5b shows the solution for Mach 1.4 inflow after 100 defect correction cycles. The residual is an order of

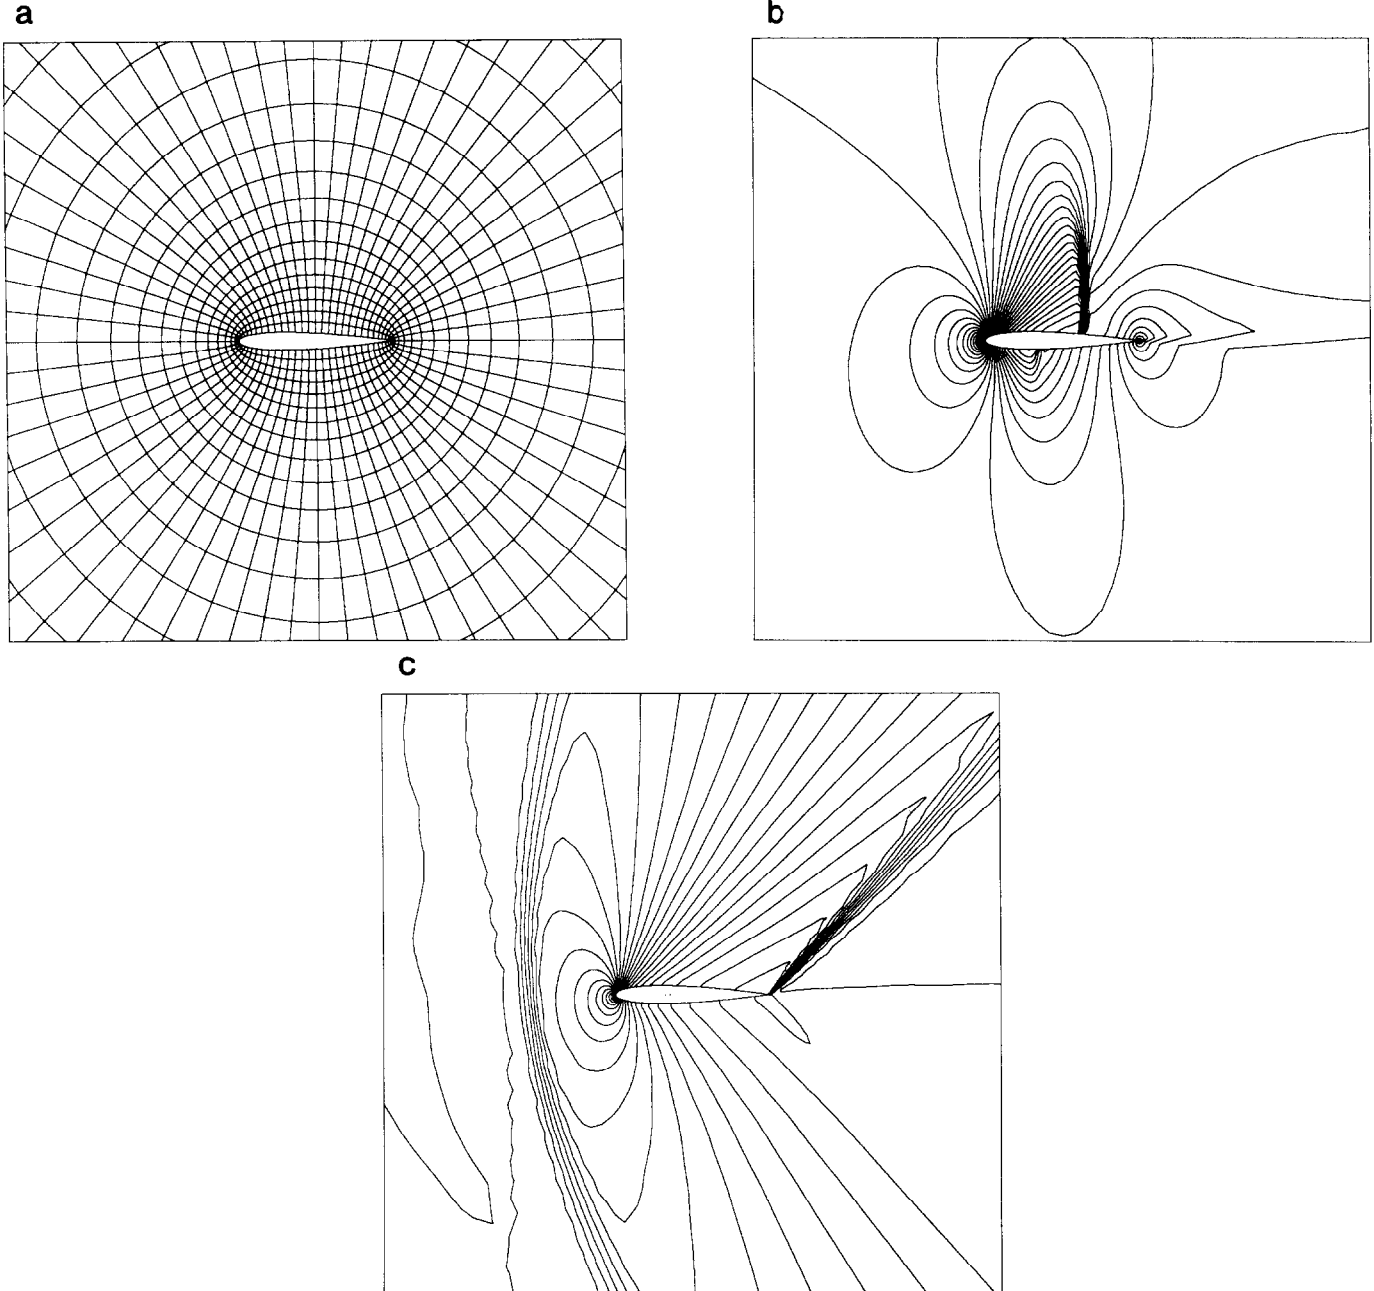


FIG. 7. (a) Part of the restriction to 64×32 of the 128×64 grid used for the airfoil computations. (b) Mach-lines for Mach 0.8 inflow around a NACA0012 airfoil at 1.25° angle of attack. Contours are 0.025 apart. (c) Mach-lines for Mach 1.2 inflow around a NACA0012 airfoil at 7° angle of attack. Contours are 0.05 apart.

magnitude smaller than after eight cycles, the relative truncation error is practically the same. There are differences on a small scale, but the large-scale properties, including shock positions, are identical. Figure 5c shows a first-order solution for comparison.

The third set of examples involves flow over a NACA0012 airfoil. A fairly orthogonal grid has been used, with a circle as outer boundary at 50 chord lengths (see Fig. 7a). Characteristic boundary conditions are applied at the inner and

outer boundary. Table I lists lift and drag for two types of flow, using Osher's scheme and van Leer's flux-vector splitting (FVS). The results, computed for eight defect correction cycles, agree very well with those in [1, 7]. The values of ψ_1 and ψ_∞ are 0.12 and 0.13 for the first example in Table I, 0.041 and 0.16 for the second, 0.73 and 0.17 for the third, and 0.97 and 0.17 for the fourth. The effect of the iteration error can be seen by comparing to the result after 100 defect correction cycles. The size of the differences due to the iteration error is of the same order as that of the differences between values for subsequent grids. The latter reflect the effect of the discretisation error. We conclude that the iteration error after eight cycles is of the same order of magnitude as the discretisation error.

Pressure curves for the first case of Table I are shown in Fig. 6a. The solutions for Osher's scheme and FVS agree reasonably well. The pressure curves for the second case (Fig. 6b) are practically indistinguishable. Mach-lines for Osher's scheme are shown in Figs. 7b and 7c.

6. CONCLUSIONS

The performance of a multigrid method for higher-order discretisations of the steady Euler equations is limited by the hyperbolicity of the equations. Waves perpendicular to a stream line are difficult to remove by the coarse-grid correction operator and this leads to poor convergence rates. Because these waves are related to the truncation error (they do not appear in the exact differential equations), the convergence rates obtained by two-grid estimates are too pessimistic. It does not make sense to require iteration errors to be much smaller than the discretisation error.

The defect correction technique allows us to obtain estimates of the iteration error in terms of the discretisation error. Linear analysis shows that about seven cycles can provide a residual of the order of the truncation error, if higher-order upwind differencing is used and if the first-order solver has a convergence rate of at worst $\frac{1}{2}$. The convergence rate of the first-order solver can be monitored in a computer code. Measuring the convergence rate of the higher-order residual does not necessarily provide any useful information, because convergence will be slow or absent if the residual becomes of the order of the truncation error.

Numerical experiments on a variety of flows show that

acceptable results are obtained in a full multigrid code with eight defect correction cycles at each level of nesting. The convergence rate of the underlying first-order solver was better than 0.5 in all examples. The high-resolution residual converged to the level of the estimated truncation error.

APPENDIX

With defect correction, a solution with a total error of the order of the discretisation error can be obtained in a fixed number of steps, at least in the linear constant coefficient case. The number of steps required for the linearised Euler equations has been estimated in [12]. This estimate will be repeated below for completeness.

The following linearised version of (2.1) is adopted:

$$\frac{\partial w'}{\partial t} + A \frac{\partial w'}{\partial x} + B \frac{\partial w'}{\partial y} = 0, \quad (\text{A1a})$$

where

$$A = \begin{pmatrix} u & 0 & c & 0 \\ 0 & u & 0 & 0 \\ c & 0 & u & 0 \\ 0 & 0 & 0 & u \end{pmatrix}, \quad B = \begin{pmatrix} v & 0 & 0 & 0 \\ 0 & v & c & 0 \\ 0 & c & v & 0 \\ 0 & 0 & 0 & v \end{pmatrix}, \quad \delta w' = \begin{pmatrix} \delta u \\ \delta v \\ \delta p/(\rho c) \\ \delta S \end{pmatrix}. \quad (\text{A1b})$$

Here $S = \log(p/\rho^\gamma)$ is the specific entropy. The fourth equation describes the convection of entropy along stream lines. The remaining 3×3 system represents the combination of convection and sound waves. In the isentropic case, the fourth equation can be dropped and the third component of w' becomes $2c/\gamma_1$. For the steady-state problem, we consider the linear operator

$$L = A \frac{\partial}{\partial x} + B \frac{\partial}{\partial y}, \quad (\text{A2})$$

with constant coefficients and periodic boundary conditions. This operator is elliptic in the subsonic and hyperbolic in the supersonic case. The fourth equation, considered by itself, is hyperbolic.

To accomplish the upwind differencing, the matrix A is diagonalised by Q_1 , according to

$$A = Q_1 A_1 Q_1^{-1}, \quad A_1 = \text{diag}(u - c, u, u, u + c),$$

$$Q_1 = \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & -1 & 0 & 0 \\ -1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}. \quad (\text{A3a})$$

For B we have

$$B = Q_2 A_2 Q_2^{-1}, \quad A_2 = \text{diag}(v - c, v, v, v + c),$$

$$Q_2 = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 \\ -1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}. \quad (\text{A3b})$$

The matrix A_k ($k = 1, 2$) is split into A_k^+ and A_k^- , which contain the positive and negative elements of A_k , respectively. This implies

$$A_k^+ + A_k^- = A_k, \quad A_k^+ - A_k^- = |A_k|. \quad (\text{A4})$$

Now define

$$A^\pm \equiv Q_1 A_1^\pm Q_1^{-1}, \quad B^\pm \equiv Q_2 A_2^\pm Q_2^{-1}. \quad (\text{A5})$$

It follows that

$$A = A^+ + A^-, \quad |A| \equiv Q_1 |A_1| Q_1^{-1} = A^+ - A^-; \quad (\text{A6})$$

$$B = B^+ + B^-, \quad |B| \equiv Q_2 |A_2| Q_2^{-1} = B^+ - B^-.$$

The discrete linear operator becomes

$$L^{h_x, h_y}(T_x, T_y) = L_x^{h_x}(T_x) + L_y^{h_y}(T_y), \quad (\text{A7a})$$

where

$$L_x^{h_x}(T_x) = \frac{1}{h_x} [A^+ D(T_x, s_x, \kappa) - A^- D(T_x^{-1}, s_x, \kappa)], \quad (\text{A7b})$$

$$L_y^{h_y}(T_y) = \frac{1}{h_y} [B^+ D(T_y, s_y, \kappa) - B^- D(T_y^{-1}, s_y, \kappa)],$$

and

$$D(T, s, \kappa) = (1 - T^{-1})[1 + \frac{1}{4}s((1 - \kappa)(1 - T^{-1}) + (1 + \kappa)(T - 1))] \quad (\text{A7c})$$

Here T_x and T_y are shift operators: $T_x^k v_{i,j} = v_{i+k,j}$, $T_y^k v_{i,j} = v_{i,j+k}$. We assume that the grid is Cartesian with a constant aspect ratio h_y/h_x . First-order accuracy is obtained for $s_x = s_y = 0$. For $s_x = s_y = 1$ and $|\kappa| \leq 1$, we have a high resolution scheme. The scheme is third-order for $\kappa = \frac{1}{3}$ in a point-wise sense. Other values of κ lead to second-order accuracy.

The analysis will be carried out in Fourier space. We consider Fourier modes of the form

$$\exp[-i(\theta_x x + \theta_y y)], \quad i = 1, \dots, N_1, \quad j = 1, \dots, N_2, \quad (\text{A8a})$$

where the frequencies on a $N_1 \times N_2$ grid are

$$\theta_x = 2\pi \frac{l_1}{N_1}, \quad \theta_y = 2\pi \frac{l_2}{N_2}, \quad (\text{A8b})$$

$$l_1 = 0, \dots, N_1 - 1, \quad l_2 = 0, \dots, N_2 - 1.$$

The symbols of the shift operators are $\hat{T}_x = \exp(i\theta_x)$ and $\hat{T}_y = \exp(i\theta_y)$. The symbol of the operator is obtained from (A7) by replacing T_x and T_y with \hat{T}_x and \hat{T}_y . In that case, we have

$$\text{Re } \hat{D}(\hat{T}, s, \kappa) = 2\zeta[1 - s + s(1 - \kappa) \sin^2(\frac{1}{2}\theta)], \quad (\text{A9})$$

$$\text{Im } \hat{D}(\hat{T}, s, \kappa) = \sin(\theta)[1 + s(1 - \kappa) \sin^2(\frac{1}{2}\theta)].$$

The singularities of the discrete linear operator \hat{L}^{h_x, h_y} are listed next.

LEMMA. *The linearised operator \hat{L}^{h_x, h_y} , with $0 \leq s_x \leq 1$, $0 \leq s_y \leq 1$, and $\kappa < 1$, is singular only in each of the following cases:*

- (i) $\hat{T}_x = 1, \hat{T}_y = 1$
- (ii) $\hat{T}_x \neq 1, \hat{T}_y = 1: u = -c$ or $u = 0$ or $u = c$;
- (iii) $\hat{T}_x = 1, \hat{T}_y \neq 1: v = -c$ or $v = 0$ or $v = c$;
- (iv) $\hat{T}_x \neq 1, \hat{T}_y \neq 1: u = v = 0$.

The corresponding null-spaces do not depend on s_x, s_y , and κ .

The proof is straightforward and can be found in [12]. In the following, the pseudoinverse of the residual operator will be denoted by $(\hat{L}^{h_x, h_y})^\dagger$.

Now consider defect correction. We will use a version based on successive grid refinement. First, the higher-order problem is solved exactly on a coarsest grid at level $l = 0$. The solution is interpolated to a finer grid at level $l = 1$, and a fixed number of defect correction steps with a first-order solver are carried out. The resulting solution is interpolated to the next finer grid, where the same number of defect correction steps is performed. This is repeated until the (approximate) solution at the desired level is found.

Let the linear problem be $Lu = f$. The discrete operator on a grid at level l is L_p^l with order of accuracy p . The coarsest grid corresponds to level $l = 0$, and finer grids have grid spacings $h_x^l = \frac{1}{2}h_x^{l-1}$, $h_y^l = \frac{1}{2}h_y^{l-1}$. The discrete representation of the exact solution at level l is $\tilde{T}^l u$, and $f_p^l = I_p^l f$ is the representation of f with order of accuracy p . The discrete solution of the higher-order problem $L_p^l \tilde{u}^l = f_p^l$ has a discretisation error $e_p^l = \tilde{T}^l u - \tilde{u}^l$, and the corresponding truncation error of the residual is defined by $\tau_p^l = L_p^l e_p^l = (L_p^l \tilde{T}^l - I_p^l L)u$. The current guess of the solution is u_i^l , where i is the iteration count. The iteration error $v_i^l = \tilde{u}^l - u_i^l$, and the total error is $z_i^l = \tilde{T}^l u - u_i^l$. The current residual $r_i^l = f_p^l - L_p^l u_i^l = L_p^l v_i^l$.

In our case, defect correction is used for the higher-order problem with an iterative scheme M^l that solves the first-order problem. Let

$$\tilde{M}^l \equiv L_1^l M^l (L_1^l)^\dagger, \quad \|\tilde{M}^l\| \leq \bar{\lambda}, \quad (\text{A10})$$

where the convergence rate $\bar{\lambda}$ should be well below 1. The similarity transform based on L_1^l is introduced to suppress waves in the null-space of L_1^l , which are not damped. The iteration error for the higher-order solution obeys

$$v_{i+1}^l = [I - (I - M^l)(L_1^l)^\dagger L_p^l] v_i, \quad i \geq 0. \quad (\text{A11})$$

Note that a necessary condition for boundedness of $(L_1^l)^\dagger L_p^l$ is that the null-space of L_1^l is the same as, or a subset (for $\kappa = 1$) of, the null-space of L_p^l . This condition is met, according to the earlier lemma.

The equivalent of Eq. (A11) for the residual is

$$r_{i+1}^l = L_p^l [I - (I - M^l)(L_1^l)^\dagger L_p^l] (L_p^l)^\dagger r_i^l, \quad i \geq 0. \quad (\text{A12a})$$

Using $r^l = L_p^l z^l - \tau_p^l$, this can be rewritten as

$$\begin{aligned} L_p^l z_{i+1}^l &= L_p^l (L_1^l)^\dagger [\tilde{M}^l L_p^l z_i^l + (I - \tilde{M}^l) \tau_p^l] \\ &\quad + L_p^l [I - (L_1^l)^\dagger L_p^l] z_i^l. \end{aligned} \quad (\text{A12b})$$

The first term on the right-hand side in square brackets clearly allows for convergence down to the truncation error. The last term may cause problems, though. Because the null-space of L_1^l is a subset of the null-space of L_p^l , the last term can be rearranged as

$$(I - L_p^l (L_1^l)^\dagger) L_p^l z_i^l = L_p^l (L_1^l)^\dagger (L_1^l - L_p^l) z_i^l. \quad (\text{A12c})$$

With successive grid refinement, it may be assumed that the initial guess of the solution has an error of the order of the discretisation error. Therefore, we let $\|L_p^l z_i^l\| = \alpha_i^l \|\tau_p^l\|$, where α_i^l is assumed to be $O(1)$. Also, we assume that $\|L_p^l (L_1^l)^\dagger\| \leq C_{2,1}$, where $C_{2,1}$ is $O(1)$. It follows that

$$\begin{aligned} \alpha_{i+1}^l \|\tau_p^l\| &\leq \bar{\lambda} C_{2,1} \alpha_i^l \|\tau_p^l\| + (1 + \bar{\lambda}) C_{2,1} \|\tau_p^l\| \\ &\quad + C_{2,1} \|(L_1^l - L_p^l) z_i^l\|. \end{aligned} \quad (\text{A13})$$

The first term on the right-hand side decreases after each iteration if $\bar{\lambda} C_{2,1} < 1$. The second term is of the order of the truncation error. The last term can be neglected if

$$\|(L_1^l - L_p^l) z_i^l\| \leq \|L_1^l - L_p^l\| \|z_i^l\| = O(h^{1+p}). \quad (\text{A14})$$

This estimate is valid only if z_i^l is $O(h^p)$ and sufficiently smooth. It will be violated, for instance, if completely new large-scale structures of a size larger than the discretisation

error appear in the solution after grid refinement, because of, e.g., non-smooth boundary conditions or singularities in the flow field. If large errors occur only *locally*, some additional smoothing after grid refinement and between defect correction cycles may help to remove them.

Assuming that (A14) holds and that $\|\tau_p^l\| = O(h^p) \neq 0$, which excludes trivial solutions, and neglecting higher-order terms in h , we find

$$\begin{aligned} \alpha_i^l &\leq (\bar{\lambda} C_{2,1})^i \alpha_0^l + (1 + \bar{\lambda}) C_{2,1} \beta(\bar{\lambda} C_{2,1}, i), \\ \beta(x, i) &= \frac{1 - x^i}{1 - x}. \end{aligned} \quad (\text{A15})$$

Therefore, the defect-correction technique can provide a solution with a total error which is a factor $(1 + \bar{\lambda}) C_{2,1} \beta(\bar{\lambda} C_{2,1}, i)$ times the truncation error.

Next consider successive grid refinement. We would like to obtain an estimate of α_i^l for a fixed iteration count. Given a solution u_{i-1}^l , we obtain an initial guess on level l through interpolation: $u_0^l = \Pi_{i-1}^l u_{i-1}^{l-1}$. Here we assume that the order of interpolation is at least $(p+1)$. If $\|\tau_p^l\| = 2^p \|\tau_p^{l-1}\|$, then $\alpha_0^l = 2^p \alpha_{i-1}^{l-1}$. If we assume that the problem on the first coarsest level is fully converged, implying that $\alpha_\infty^0 = 1$, we find for a fixed number of defect correction steps i on finer levels $l > 0$ that

$$\alpha_i^l \leq \zeta^l + \phi(\bar{\lambda}, p, C_{2,1}, i, l), \quad (\text{A16a})$$

where

$$\begin{aligned} \zeta &= 2^p (\bar{\lambda} C_{2,1})^i, \\ \phi(\bar{\lambda}, p, C_{2,1}, i, l) &= (1 + \bar{\lambda}) C_{2,1} \beta(\bar{\lambda} C_{2,1}, i) \beta(\zeta, l). \end{aligned} \quad (\text{A16b})$$

Here we have ignored terms of $O(h)$ or smaller.

To obtain quantitative estimates, we consider a second-order discretisation ($p=2$) and a first-order solver with a multigrid convergence rate $\bar{\lambda} = \frac{1}{2}$. Note that the multigrid convergence rate is usually larger than the two-grid convergence rate. The number 0.5 obtained for the method in Section 3 corresponds to the two-grid rate. However, the numerical experiments in [11] show a multigrid convergence rate well below 0.5, and this motivates the present choice for $\bar{\lambda}$.

According to the lemma, L_1 and L_p have the same null-space for $\kappa < 1$. If $\kappa = 1$, the null-space of L_2 contains the one of L_1 , but is larger. Equation (A12c) still holds, but the estimate (A14) may break down because of the checkerboard mode. This mode may result in an initial guess after grid refinement which is not smooth enough. Therefore, the use of a scheme with $\kappa = 1$ (central differencing) is not recommended. We will nevertheless include it in our estimates

TABLE II

The Constant $C_{2,1} \geq \|L_1^\dagger L_2\|$ Determines a Bound on the Accuracy Achievable with the Defect Correction Technique

κ	$C_{2,1}$	$\phi(\frac{1}{2}, *, C_{2,1}, \infty, \infty)$	$i (< 10\%)$
1	1.00	(3.00)	(6)
$\frac{1}{3}$	1.06	3.39	6, 7
0	1.15	4.10	7
-1	2.00	—	—

Note. This accuracy is measured by $\phi(\tilde{\lambda}, p, C_{2,1}, i, l)$ times the norm of the truncation error. It is assumed that $\tilde{\lambda} = \|\tilde{M}\| = 0.5$. The third column gives the asymptotic values, which are independent of p , the order of accuracy of the discretisation. The last column gives the number of cycles needed to obtain this accuracy within 10%. The convergence estimate for $\kappa = 1$ can not be justified and is only included for reference. The two values for $\kappa = \frac{1}{3}$ refer to $p = 2$ and 3, respectively. The other values of κ result in second-order accuracy ($p = 2$).

Results are listed in Table II. For $C_{2,1}$, we have used the semi-norm

$$\|\cdot\| \equiv \max\{\rho(\tilde{L}_p(\cdot)) : u/c, v/c, h_y/h_x, \theta_x, \theta_y\}, \quad (\text{A17})$$

where $\rho(\cdot)$ denotes the spectral radius. The similarity transform based on \tilde{L}_p masks out the effect of singularities. The values of $C_{2,1}$ in Table II are obtained numerically by using the norm (A17) on $\tilde{L}_1^\dagger \tilde{L}_p$. For the fourth equation of the system (A2), we obtain $C_{2,1} = 1$ for $\kappa = 1$, $C_{2,1} = \sqrt{9/8}$ for $\kappa = \frac{1}{3}$, $C_{2,1} = \sqrt{4/3}$ for $\kappa = 0$, and $C_{2,1} = 2$ for $\kappa = -1$, which happen to agree with the numerical results for the full system. The bound for convergence of the residual in terms of the truncation error is asymptotically determined by $\phi(\frac{1}{2}, *, C_{2,1}, \infty, \infty)$. The asterisk indicates that this asymptotic value is independent of p . It should be noted that the result for $\kappa = 1$ cannot be justified, as already mentioned. For $\kappa = -1$, the product $\tilde{\lambda}C_{2,1}$ becomes 1 and convergence is lost. Applying two defect correction cycles without updating the second-order residual does provide convergence, but this is not considered here. The last column of Table II lists the number of iterations required to reach the asymptotic result within 10%. This, or a slightly larger value, can be used as the number of cycles to be carried out in practical computations. The result $i (< 10\%) = 6$ for $\kappa = \frac{1}{3}$ is obtained with $p = 2$. For $p = 3$, we obtain $i (< 10\%) = 7$.

Some additional smoothing can be performed on the higher-order residual between defect correction cycles if the assumption (A14) is violated because of *locally* large errors. Point-Jacobi relaxation is linearly unstable when used with higher-order upwind differencing, so we use a two-stage scheme:

$$\begin{aligned} w'^* &:= w' - \beta_1(|A| + |B|)^{-1} L_p w', \\ w' &:= w' - \beta_2(|A| + |B|)^{-1} L_p w'^*. \end{aligned} \quad (\text{A18})$$

The damping parameters β_1 and β_2 are chosen such that the scheme (A18) is (i) total variation non-increasing (TVNI) in the one-dimensional case, (ii) linearly stable, and (iii) has an optimal smoothing rate for the high-frequencies in the two-dimensional linear *scalar* case. The analysis is carried out for the equation $uw_x + vw_y = 0$, with $u > 0$ and $v \geq 0$. It should be noted that for (ii) and (iii) we consider the linear scheme without the limiter ($s = 1$), whereas (i) is based on the limiter and is therefore a nonlinear scheme. We start with the TVNI property and consider the one-dimensional scalar equation $uw_x = 0$. Define $\Delta_{i+1/2} = w_{i+1} - w_i$, where w_i is the discrete solution in cell i on a grid with cell size h . Let $\rho_i = \Delta_{i+1/2}/\Delta_{i-1/2}$. The limiter (2.3c) at cell i lets $s_i = 2\rho_i/(1 + \rho_i^2)$. The first step of the two-stage scheme is

$$w_i^* = w^i - \beta_1 \left(\frac{u}{h}\right)^{-1} \frac{f_{i+1/2} - f_{i-1/2}}{h}, \quad (\text{A19a})$$

where the flux

$$\begin{aligned} f_{i+1/2} &= u(w_i + F(\rho_i, \kappa) \Delta_{i-1/2}), \\ F(\rho, \kappa) &= \frac{1}{4} \frac{2\rho}{(1 + \rho^2)} [(1 - \kappa) + (1 + \kappa)\rho]. \end{aligned} \quad (\text{A19b})$$

Define

$$\begin{aligned} H_i &= \frac{f_{i+1/2} - f_{i-1/2}}{u \Delta_{i-1/2}} \\ &= 1 + F(\rho_i, \kappa) - F(\rho_{i-1}, \kappa)/\rho_{i-1}. \end{aligned} \quad (\text{A20})$$

Then

$$w_i^* = (1 - \beta_1 H_i) w_i + (\beta_1 H_i) w_{i-1}, \quad (\text{A21})$$

which is TVNI for $0 \leq \beta_1 H_i \leq 1$. Using the fact that $|\kappa| \leq 1$,

$$\begin{aligned} (1 + \kappa) - \sqrt{2(1 + \kappa^2)} \\ \leq 4F(\rho, \kappa) \leq (1 + \kappa) + \sqrt{2(1 + \kappa^2)}, \end{aligned} \quad (\text{A22a})$$

and

$$F(\rho, \kappa) = \rho F(1/\rho, -\kappa), \quad (\text{A22b})$$

we find that

$$\begin{aligned} 1 + \frac{1}{2}\kappa - \frac{1}{2}\sqrt{2(1 + \kappa^2)} \\ \leq H_i \leq 1 + \frac{1}{2}\kappa + \frac{1}{2}\sqrt{2(1 + \kappa^2)}. \end{aligned} \quad (\text{A23})$$

For $2 - \sqrt{6} \leq \kappa \leq 1$, we must have

$$\begin{aligned} 0 \leq \beta_1 \leq \beta_{\max}, \\ \beta_{\max} = [1 + \frac{1}{2}\kappa + \frac{1}{2}\sqrt{2(1 + \kappa^2)}]^{-1}. \end{aligned} \quad (\text{A24})$$

Next, consider the two steps combined:

$$\begin{aligned} w_i := & [1 - \beta_2 H_i^* (1 - \beta_1 H_i)] w_i \\ & + [\beta_2 H_i^* (1 - \beta_1 (H_i + H_{i-1}))] w_{i-1} \\ & + [\beta_2 H_i^* \beta_1 H_{i-1}] w_{i-2}. \end{aligned} \quad (\text{A25})$$

A sufficient condition for non-negative coefficients, which implies that the scheme is TVNI, is

$$\begin{aligned} 0 \leq \beta_1 \leq \frac{1}{2} \beta_{\max}, \quad 0 \leq \beta_2 \leq \beta_{\max}, \\ \text{for } \kappa \in [2 - \sqrt{6}, 1]. \end{aligned} \quad (\text{A26})$$

We continue with linear stability, condition (ii). In the scalar case, the amplification factor of the two-stage scheme obeys

$$|\hat{G}|^2 = (1 - \beta_1 \beta_2 f_2)^2 + \beta_2 (1 - 2\beta_1 f_0) (\beta_2 f_2 - 2f_0), \quad (\text{A27a})$$

where

$$\begin{aligned} f_0 = \frac{h_1}{1 + \alpha} (z_1^2 + \alpha z_2^2), \quad f_1 = \frac{v_1 \pm \alpha v_2}{1 + \alpha}, \\ f_2 = f_0^2 + f_1^2, \end{aligned} \quad (\text{A27b})$$

with

$$h_1 = 2(1 - \kappa), \quad \alpha = \frac{v}{h_y} \Big/ \frac{u}{h_x}, \quad (\text{A27c})$$

and

$$\begin{aligned} z_k = \sin^2(\theta_k/2), \\ v_k = (2 + h_1 z_k) \sqrt{z_k(1 - z_k)}, \quad k = 1, 2. \end{aligned} \quad (\text{A27d})$$

A necessary (but possibly not sufficient) condition for linear stability is

$$0 \leq \beta_2 \leq 2\beta_1 \leq \frac{1}{1 - \kappa}, \quad \text{for } \kappa < 1. \quad (\text{A28})$$

To obtain an optimum smoothing rate, as required by (iii), we have to determine

$$\bar{\mu} = \min_{\beta_1, \beta_2} \max_{z_1, z_2, \alpha} |\hat{G}|, \quad (\text{A29})$$

subject to conditions (A26) and (A28). The high-frequency range is defined by $z_k \in [\frac{1}{2}, 1]$, for $k = 1, 2$, and because of symmetry it is sufficient to consider $\alpha \in [0, 1]$. Numerical optimisation of (A29) suggests an optimal smoothing rate for $2\beta_1 = \beta_2 = \beta_{\max}$ if $\kappa = \frac{1}{3}$, namely, $\bar{\mu} = 1 - \frac{1}{3} \beta_{\max} (1 - \frac{1}{6} \beta_{\max})$. Using the same program with this choice of β_1 and β_2 , we find the $\max_{z_1, z_2, \alpha} |\hat{G}| \leq 1$ for $z_k \in [0, 1]$, $k = 1, 2$, implying linear stability.

REFERENCES

1. W. K. Anderson, J. L. Thomas, and B. van Leer, AIAA Paper No. 85-0122, AIAA 23rd Aerospace Sciences Meeting, Reno, Nevada, Jan. 1985 (unpublished).
2. A. Brandt, in *Lecture Notes in Mathematics*, Vol. 960 (Springer-Verlag, Berlin/Heidelberg/New York, 1982), p. 220.
3. A. Brandt, in *Proceedings of the Fourth Copper Mountain Conference on Multigrid Methods*, edited by J. Mandel *et al.* (SIAM, Philadelphia, 1989), p. 13.
4. W. Hackbusch, *Multi-Grid Methods and Applications*, Computational Mathematics, Vol. 4 (Springer-Verlag, Berlin/Heidelberg/New York, 1985).
5. W. Hackbusch, in *Proceedings, First International Conference on Industrial and Applied Mathematics*, edited by J. McKenna and R. Temam (SIAM, Philadelphia, 1988).
6. P. W. Hemker and S. P. Spekreijse, *Appl. Numer. Math.* **2**, 475 (1986).
7. B. Koren, *J. Comput. Phys.* **77**, 183 (1988).
8. W. A. Mulder, *J. Comput. Phys.* **60**, 235 (1985).
9. W. A. Mulder, *Astron. Astrophys.* **156**, 354 (1986).
10. W. A. Mulder, in *Multigrid Methods: Theory, Applications, and Supercomputing*, edited by S. F. McCormick (Dekker, New York, 1988), p. 467.
11. W. A. Mulder, *J. Comput. Phys.* **83**, 303 (1989).
12. W. A. Mulder, in *Proceedings of the Fourth Copper Mountain Conference on Multigrid Methods*, edited by J. Mandel *et al.* (SIAM, Philadelphia, 1989), p. 348.
13. W. A. Mulder, *SIAM J. Sci. Stat. Comput.* **11**, 33 (1990).
14. W. A. Mulder, *SIAM J. Sci. Stat. Comput.* **11**, 389 (1990).
15. S. Osher and F. Solomon, *Math. Comput.* **38**, 339 (1982).
16. M. M. Rai and S. R. Chakravarty, *AIAA J.* **24**, 735 (1986).
17. S. P. Spekreijse, *Multigrid Solution of the Steady Euler Equations*, Ph.D. thesis (Centrum voor Wiskunde en Informatica, Amsterdam, 1987).
18. G. D. van Albada, B. van Leer, and W. W. Roberts, *Astron. Astrophys.* **108**, 76 (1982).
19. B. van Leer, in *Lecture Notes in Physics*, Vol. 170 (Springer-Verlag, Berlin/Heidelberg/New York, 1982), p. 507.
20. B. van Leer, in *Large-Scale Computations in Fluid Mechanics*, Lectures in Applied Mathematics 22 (part 2), edited by B. E. Engquist, S. Osher, and R. C. J. Somerville (Am. Math. Soc., Providence, RI, 1985), p. 327.
21. B. van Leer and W. A. Mulder, in *Numerical Methods for the Euler Equations of Fluid Dynamics*, edited by F. Angrand, A. Dervieux, J. A. Desideri, and R. Glowinski (SIAM, Philadelphia, 1985), p. 312.