UCLA
COMPUTATIONAL AND APPLIED MATHEMATICS

A High-Resolution Euler Solver

Wim A. Mulder

April 1989

CAM Report 89-08

Department of Mathematics
University of California, Los Angeles
Los Angeles, CA. 90024-1555

A high-resolution Euler solver

Wim A. Mulder
Department of Mathematics
405 Hilgard Avenue
University of California at Los Angeles
Los Angeles, CA 90024-1555
Arpanet: mulder@math.ucla.edu

AIAA paper no. 89-1949, to be presented at the AIAA 9th Computational Fluid Dynamics Conference, Buffalo, New York, June 13-15. 1989.

Work supported by NSF grant DMS88-11863 and ONR grant N00014-86-K-0691.

Printed: April 3, 1989.

A HIGH-RESOLUTION EULER SOLVER

Wim A. Mulder*
Department of Mathematics
University of California at Los Angeles
Los Angeles, CA 90024-1555

Abstract

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 is able to provide 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-level analysis for the linear constant-coefficient case has shown that it is difficult or impossible to obtain uniformly good convergence rates for a higher-order scheme, due to 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 7 defect correction cycles, according to estimates for the linear constant-coefficient equations. Here this result is explored for the nonlinear case by some illustrative numerical experiments.

1. Introduction

A bottleneck in the application of the multigrid technique to the computation of stationary flows is alignment [2,3]. 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 can not be removed by smoothing, because there is no coupling, nor by solving the equations on a coarser grid, because high-frequencies can not 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 [9] 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 convergence. For the Euler equations in two dimensions, Alternating Direction Damped Line Jacobi can provide a uniformly good convergence rate [10].

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 direction 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 [18] and provides second- or third-order accuracy. Details are given in §2. The multigrid method is reviewed in §3. In [12] is 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.[3]) can not 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 7 defect correction cycles, at least in the linear constant-coefficient case. In this paper, numerical experiments are carried out to explore this result for the nonlinear equations. Details of the nonlinear implementation are described in §4 and results are presented in §5.

2. Spatial discretisation

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

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

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

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

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

$$E = \frac{1}{\gamma_1}\frac{p}{\rho} + \tfrac{1}{2}(u^2+v^2), \; H = E + \frac{p}{\rho}, \; c^2 = \gamma\frac{p}{\rho}. \qquad (2.1c)$$

* *Assistant professor, member AIAA*

Here $\gamma_1 = \gamma - 1$. 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 corners 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 corresponding sides by $(i, j-\frac{1}{2})$, $(i+\frac{1}{2}, j)$, $(i, j+\frac{1}{2})$, $(i-\frac{1}{2}, j)$. The discrete residual is

$$\begin{aligned} r_{i,j} =& \ell_{i,j-\frac{1}{2}} T^{-1}_{i,j-\frac{1}{2}} f(T_{i,j-\frac{1}{2}} w(q^{j-}_{i,j}), T_{i,j-\frac{1}{2}} w(q^{j+}_{i,j-1})) \\ &+ \ell_{i+\frac{1}{2},j} T^{-1}_{i+\frac{1}{2},j} f(T_{i+\frac{1}{2},j} w(q^{i+}_{i,j}), T_{i+\frac{1}{2},j} w(q^{i-}_{i+1,j})) \\ &+ \ell_{i,j+\frac{1}{2}} T^{-1}_{i,j+\frac{1}{2}} f(T_{i,j+\frac{1}{2}} w(q^{j+}_{i,j}), T_{i,j+\frac{1}{2}} w(q^{j-}_{i,j+1})) \\ &+ \ell_{i-\frac{1}{2},j} T^{-1}_{i-\frac{1}{2},j} f(T_{i-\frac{1}{2},j} w(q^{i-}_{i,j}), T_{i-\frac{1}{2},j} w(q^{i+}_{i-1,j})). \end{aligned} \tag{2.2a}$$

Here $f(w_l, w_r)$ is a numerical flux corresponding to f in (2.1b) that provides an approximate solution to the Riemann problem. In this paper we will use either van Leer's Flux-Vector Splitting [17] or Osher's scheme in the natural ordering [13,14]. Both are sufficiently smooth (Lipshitz continuous) for our purpose. The first rotation matrix is given by

$$T_{i,j-\frac{1}{2}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\phi_{i,j-\frac{1}{2}} & \sin\phi_{i,j-\frac{1}{2}} & 0 \\ 0 & -\sin\phi_{i,j-\frac{1}{2}} & \cos\phi_{i,j-\frac{1}{2}} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{2.2b}$$

where

$$\begin{aligned} \cos\phi_{i,j-\frac{1}{2}} &= (y_{i+\frac{1}{2},j-\frac{1}{2}} - y_{i-\frac{1}{2},j-\frac{1}{2}})/\ell_{i,j-\frac{1}{2}}, \\ \sin\phi_{i,j-\frac{1}{2}} &= -(x_{i+\frac{1}{2},j-\frac{1}{2}} - x_{i-\frac{1}{2},j-\frac{1}{2}})/\ell_{i,j-\frac{1}{2}}, \\ \ell_{i,j-\frac{1}{2}} &= [(x_{i+\frac{1}{2},j-\frac{1}{2}} - x_{i-\frac{1}{2},j-\frac{1}{2}})^2 + \\ &\qquad (y_{i+\frac{1}{2},j-\frac{1}{2}} - y_{i-\frac{1}{2},j-\frac{1}{2}})^2]^{\frac{1}{2}}. \end{aligned} \tag{2.2c}$$

Here the outward normal is $(\cos\phi_{i,j-\frac{1}{2}}, \sin\phi_{i,j-\frac{1}{2}})^T$. 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 $q^{i\pm}_{i,j}$ and $q^{j\pm}_{i,j}$ are values at the cell-boundaries 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 [1,18], which lets, in the i-direction,

$$\begin{aligned} q^{i+}_{i,j} &= q_{i,j} + \tfrac{1}{4} s(\Delta^{i-}_{i,j}, \Delta^{i+}_{i,j}) \left[(1-\kappa)\Delta^{i-}_{i,j} + (1+\kappa)\Delta^{i+}_{i,j}\right], \\ q^{i-}_{i,j} &= q_{i,j} - \tfrac{1}{4} s(\Delta^{i-}_{i,j}, \Delta^{i+}_{i,j}) \left[(1-\kappa)\Delta^{i+}_{i,j} + (1+\kappa)\Delta^{i-}_{i,j}\right]. \end{aligned} \tag{2.3a}$$

Here

$$\Delta^{i-}_{i,j} = q_{i,j} - q_{i-1,j}, \quad \Delta^{i+}_{i,j} = q_{i+1,j} - q_{i,j}. \tag{2.3b}$$

The function $s(\Delta^-, \Delta^+)$ is a limiter that prevents numerical oscillations. Here we use a smooth limiter due to van Albada [16]:

$$s(\Delta^-, \Delta^+) = \frac{2(\Delta^-\Delta^+ + \varepsilon_a^2)}{(\Delta^-)^2 + (\Delta^+)^2 + 2\varepsilon_a^2}. \tag{2.3c}$$

The constant ε_a prevents division by zero. We use $\varepsilon_a = 10^{-16}$ in the numerical experiments of §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 fully one-sided upwind scheme. The limiter may cause the accuracy to reduce to first-order at isolated points. For $\kappa = 1/3$, 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. Also, the flux of the average state is not equal to the average of the flux over one part of the cell-boundary. 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.

3. Multigrid

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 1 shows the various grids employed if the finest grid has 8×8 points and the coarsest 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.

The usual restriction and prolongation operators have to be modified to handle input from more than one grid. This is done as follows. 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 can not, however, handle alignment at 45°. It is shown in [11] that this is not a problem for the Euler equations of gas dynamics if first-order upwind differencing is used. Two-level analysis for the linearised Euler equations with constant coefficients leads to a worst-case 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.

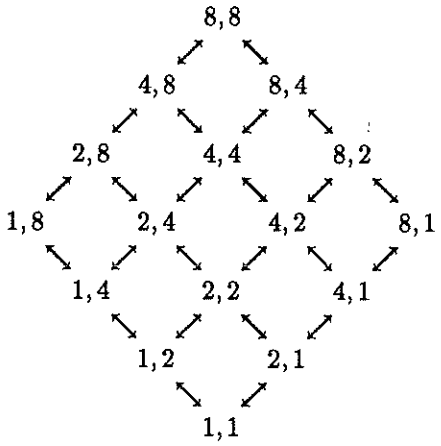

Fig. 1. *Arrangement of finest* (8×8) *and coarser grids, that leads to an* $O(N)$ *multigrid method for problems with alignment. The arrows indicate how the grids are linked by restriction (downward) and prolongation (upward).*

A fundamental problem encountered in extending this method to higher-order discretisations is the fact that the exact operator vanishes for waves perpendicular to stream lines. Alignment, as mentioned above, is one of the results. Another one has been described in [12]: for a scheme of order p, the worst-case two-level convergence rate can not better than $1 - 2^{-p}$. Note that the first-order ($p = 1$) solver described above actually obtains this value as its worst-case two-level convergence rate.

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

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

The symbol of this operator is

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

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 operator L^h, which has a truncation error τ^h of $O(h^p)$. Let the solution be computed with a two-level method, involving a fine and a coarse grid. The coarse-grid correction operator is

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

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-frequencies 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

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

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 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 one chooses a first-order restriction operator (volume-averaging) and prolongation operator (piecewise constant interpolation), the coarse-grid correction operator for the second- and third-order scheme becomes unstable. The third-order scheme ($\kappa = 1/3$) 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) [12].

The above implies that one can not 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 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.[3]). Defect correction can not 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. Then the iteration operator is

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

If the exact operator vanishes, we obtain a convergence rate $1 - \tau_p^h/\tau_1^h = 1 - O(h^{p-1})$. However, this result is again dominated by the truncation error. Good results can be obtained if convergence down to machine zero is abandoned. An estimate for the linear constant-coefficient case is presented in [12]. The difference between the residual after a number of defect correction steps and the residual of the restriction of the exact solution of the differential equations, can be bounded by some factor times the truncation error τ_p^h. The asymptotic value (after many steps) of this factor is about 4. A value within 10% can be reached in about 7 defect correction steps, if $\kappa = 0$ or $\kappa = 1/3$. This result is obtained for the linear constant-coefficient case and assumes that the solution is sufficiently smooth. If not, extra smoothing steps can be carried out on the high-resolution residual between defect correction steps. Other assumptions are: the first-order solver has a worst-case convergence rate of 0.5, and successive grid refinement is used on the higher-resolution solution, starting with an exact higher-resolution solution on the first coarsest grid.

It should be noted that defect correction has been applied to the Euler equations in earlier work [6,7], using van Leer's Flux-Vector Splitting. This scheme smears slip-lines, and therefore does not suffer from alignment. A fundamentally different approach is chosen in [15,5], which is based on a first-order solver described in [4]. 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 upwind discretisation has been outlined in §2. We use van Leer's Flux Vector Splitting [17] or Osher's scheme in the natural ordering [13,14] as approximate Riemann solver. Only $\kappa = 1/3$ is considered for the high-resolution discretisation.

It must be noted that a standard multigrid approach can be used for van Leer's Flux Vector Splitting, because it smears streamlines and does not suffer from alignment. Damped Symmetric Gauss-Seidel is a good smoother [9]. 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 smoother for the first-order scheme is damped Point-Jacobi. In the nonlinear case, one has to determine $\tilde{w}_{i,j}$ from the equation

$$r(\tilde{w}_{i,j}, w_{i-1,j}, w_{i+1,j}, w_{i,j-1}, w_{i,j+1}) = 0 \tag{4.1a}$$

for all i and j, and then update the solution according to

$$w := \tfrac{1}{2}(w + \tilde{w}). \tag{4.1b}$$

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

$$w_{i,j} := w_{i,j} - \tfrac{1}{2} N_{i,j}^{-1} r_{i,j}, \tag{4.2}$$

where $N_{i,j} = \partial r_{i,j}/\partial w_{i,j}$. For van Leer's Flux Vector Splitting, the matrix $N_{i,j}$ is non-negative. To make it positive, we replace it by

$$\tilde{N}_{i,j} = N_{i,j} + \sigma I, \tag{4.3}$$

where σ is a positive scalar, to be specified later.

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

$$\ell_{i,j-\frac{1}{2}} T_{i,j-\frac{1}{2}}^{-1} A_{i,j-\frac{1}{2}} T_{i,j-\frac{1}{2}}, \tag{4.4a}$$

where

$$A_{i,j-\frac{1}{2}} = \frac{\partial f_{i,j-\frac{1}{2}}}{\partial w'_{i,j}}, \quad w'_{i,j} = T_{i,j-\frac{1}{2}} w_{i,j}. \tag{4.4b}$$

For Osher's scheme, we replace $A_{i,j-\frac{1}{2}}$ by $A^+_{i,j-\frac{1}{2}}(w'_{i,j})$, 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 (4.3).

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.

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 [3] that another set can be used as well. Here we use

$$W = (\rho, u, v, c)^T, \tag{4.5}$$

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 choice (4.5) requires an additional multiplication of $N_{i,j}$ with the Jacobian dw/dW, which can be done directly when evaluating the contribution of each cell-face. 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

$$R_{i,j} = \max_{k=1,\ldots,4} \left(\frac{|r_{i,j,k}|}{|W_{i,j,k}| + h_{i,j,k}} \right), \tag{4.6}$$

where $h_{i,j,2} = h_{i,j,3} = c_{i,j}$ and $h_{i,j,1} = h_{i,j,4} = 0$, and computing its ℓ_1-norm. Following [19], the parameter σ in (4.3) is chosen to be $\sigma = \varepsilon_\sigma^{-1} \max_{i,j} R_{i,j}$. 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 as 1. Smaller values can be used for complex flows with strong shocks, but then the number of relaxation sweeps must be increased.

High-resolution solutions are computed by a Full Multigrid method. On each level, 8 F-cycles are performed with the first-order solver, using one post-smoothing step of damped Point-Jacobi on each grid. The first-order solver acts on the sum of most recent first-order residual $r_1(W')$ and a source term, which is the difference between the higher-order residual $r_p(W)$ and and first-order residual $r_1(W)$ at the begin of the cycle. That is, the first-order solver finds the (approximate) solution W' to

$$r_1(W') + [r_p(W) - r_1(W)] = 0. \tag{4.7}$$

After this has been done, the new solution becomes $W := W'$. Because the high-resolution residual, in general, should stop converging 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 §5, this never occurred except in the initial stages of successive grid refinement at very coarse grids.

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:

$$\begin{aligned} W^* &:= W - \beta_1 \tilde{N}(W)^{-1} r_p(W), \\ W &:= W - \beta_2 \tilde{N}(W^*)^{-1} r_p(W^*). \end{aligned} \tag{4.8}$$

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-boundaries.

The parameters β_1 and β_2 can be chosen as to provide an optimal smoothing rate. Straightforward analysis of scheme (4.5) applied to a high-resolution discretisation of the one-dimensional scalar equation $au_x = 0$ using limiter (2.3c) shows that a necessary condition for (4.7) to be Total Variation Diminishing (TVD) is

$$\beta_{1,2} \leq \left(1 + \tfrac{1}{2}\kappa + \tfrac{1}{2}\sqrt{2(1+\kappa^2)}\right)^{-1}, \tag{4.9}$$

for $2 - \sqrt{6} \leq \kappa \leq 1$. For $\kappa = 1/3$, we find $\beta_{1,2} \leq 0.523$. Linear analysis of the two-dimensional scalar equation $au_x + bu_y = 0$ under the TVD constraint provides an optimal smoothing rate $\overline{\mu} = 0.840$ for $\beta_1 = 0.251$, $\beta_2 = 0.523$. This choice provides both linear and TVD stability. In the examples of the next section, scheme (4.5) is applied once between multigrid defect correction cycles.

The restriction operator used in the first-order solver is volume-averaging, prolongation is piecewise constant interpolation. Both operators are first-order. 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 4 new cell-centers.

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 do 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. 1, 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 [3] 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 one-dimensional study reported in [8], 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 solution with strong discontinuities are considered, but is expected to be necessary in general.

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. (Characteristic in- and outflow conditions result in a weak boundary layer near the outlet). Characteristic boundary conditions are used at the walls. Figure 2 shows mach-lines for a 128×64 grid, using Osher's scheme. The coarsest grid used has size 4×2. The relative order of accuracy of the solution can be estimated from

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

and equals 2.5 for ρ, 2.4 for u, 2.3 for v and 2.5 for c. 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 8 cycles, is 0.67. This experiment suggests that the linear estimates of §4 provide a reasonable description of the nonlinear scheme for smooth solutions.

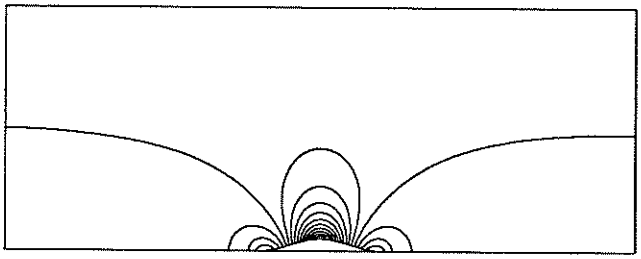

Fig. 2. *Mach-lines for mach 0.5 inflow through a channel with a* $\sin^2$ *bump on a* 128×64 *grid. Contours are* 0.025 *apart.*

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. Figure 3a 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 with 100 instead of 8 defect correction cycles. The singularities at the begin and end 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 Figure 3a.

Figure 3b and 3c show mach-lines for mach 0.85 and mach 1.4 inflow, using Osher's scheme and 8 defect correction cycles. The grid for Fig. 3b is clustered around the bump, the grid for Fig. 3c is fairly uniform. The relative orders of accuracy do not reflect the dramatic improvement with respect to the first-order solutions (cf.[11]). For mach 0.85 inflow, we obtain estimates of p around 1, in ℓ_1, for mach 1.4 inflow values around 0.6 are obtained. Because the norms 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 8 cycles and the truncation error τ^h. The latter can be estimated from the relative truncation error

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

using $\tau^h \simeq (2^p - 1)\tau_h^{2h}$. The linear analysis in [12] predicts that the norm of the residual should be at most 3.7 times the norm of the truncation error after 7 defect correction cycles for $\kappa = 1/3$. Assuming a second-order solution, this implies that $\|r^h\| \leq 1.23\|\tau_h^{2h}\|$ after 7 cycles. For the mach 0.85 example, we find $\|r^h\|/\|\tau_h^{2h}\| = 0.15$. for the mach 1.4 problem this ratio is 0.21, using the ℓ_1-norm. These ratios are even smaller in ℓ_∞. Thus, the residual has converged below the truncation error, and much better than predicted by 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. 3d shows the solution for mach 1.4 inflow after 100 defect correction cycles. The residual is about 10 times smaller than after 8 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. Note that the accuracy is mainly determined by the long waves, which determine the large-scale properties.

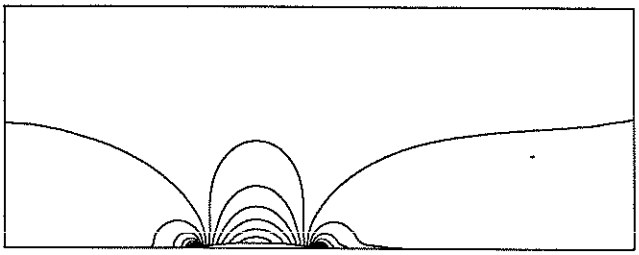

Fig. 3a. *Mach-lines for mach 0.5 inflow through a channel with a non-smooth bump on a 128 × 64 grid. Contours are 0.01 apart.*

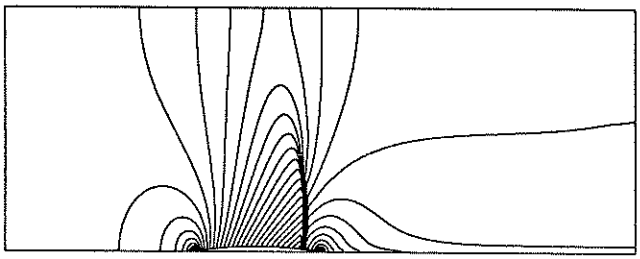

Fig. 3b. *Mach-lines for mach 0.85 inflow through a channel with a non-smooth bump on a 128 × 64 grid. Contours are 0.025 apart.*

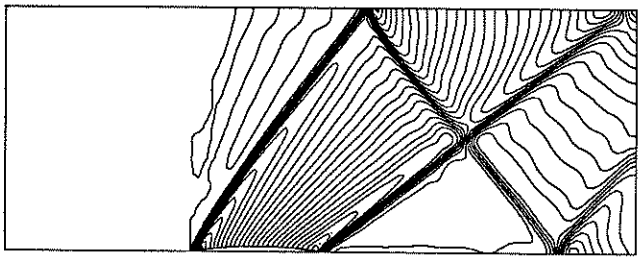

Fig. 3c. *Mach-lines for mach 1.4 inflow through a channel with a non-smooth bump on a 128 × 64 grid. Contours are 0.025 apart.*

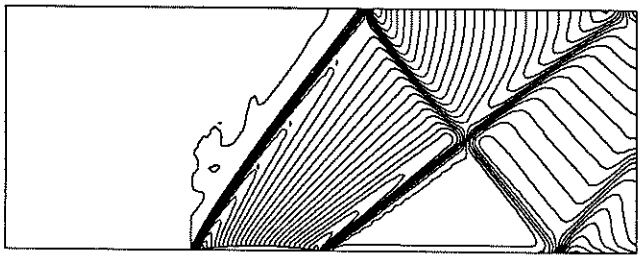

Fig. 3d. *As Fig. 3c, but after 100 defect correction cycles instead of 8.*

M_∞, α	grid	lift	drag
0.80	32 × 16	0.3122	0.0493
1.25°	64 × 32	0.3472	0.0283
Osher	128 × 64	0.3508	0.0234
0.80	32 × 16	0.3353	0.0558
1.25°	64 × 32	0.3642	0.0310
FVS	128 × 64	0.3676	0.0247
1.20	32 × 16	0.510	0.169
7.00°	64 × 32	0.525	0.159
Osher	128 × 64	0.524	0.155
1.20	32 × 16	0.511	0.173
7.00°	64 × 32	0.522	0.160
FVS	128 × 64	0.524	0.155

Table 1. *Lift and drag for NACA0012 airfoil on various grids for Osher's scheme and van Leer's Flux-Vector Splitting (FVS).*

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. Characteristic boundary conditions are applied at the wall and outer boundary. Table 1 lists lift and drag for 2 types of flow, using Osher's scheme and van Leer's Flux Vector Splitting. The results, computed for 8 defect correction cycles, agree very well with those in [1] and [5]. The improvement of lift and drag with grid refinement suggests that higher-order accuracy is indeed obtained.

Pressure curves for the first case of Table 1 are shown in Fig. 4. The solutions for Osher's scheme and FVS agree reasonably well, except for the positions of the shock. Mach-lines for Osher's scheme are shown in Fig. 5a and Fig. 5b.

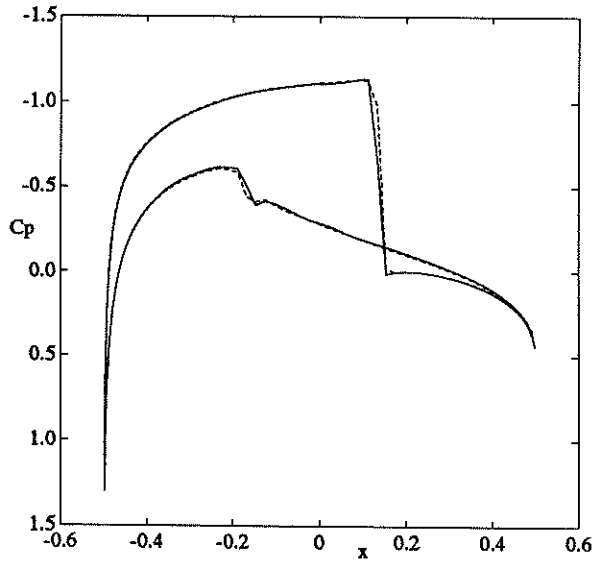

Fig. 4. *C_p-curves for a NACA0012 airfoil ($M_\infty = 0.8$, $\alpha = 1.25°$) for Osher's scheme and van Leer's FVS (dashed line).*

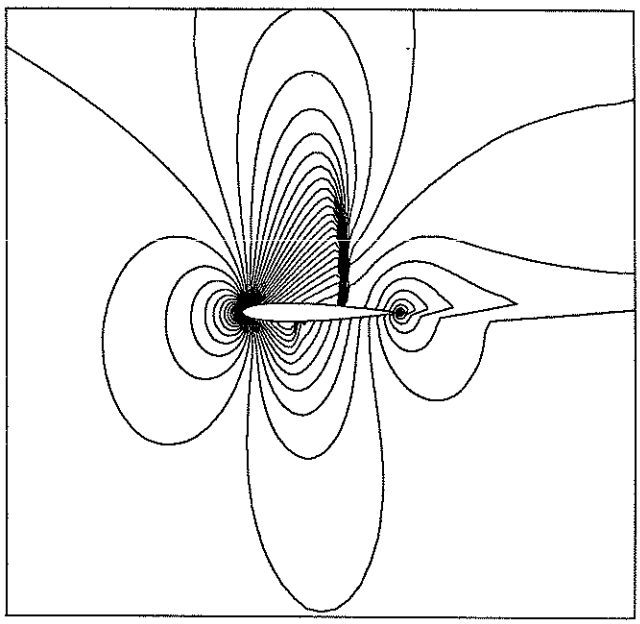

Fig. 5a. *Mach-lines for mach 0.8 inflow at 1.25° around a NACA0012 airfoil. Contours are 0.025 apart.*

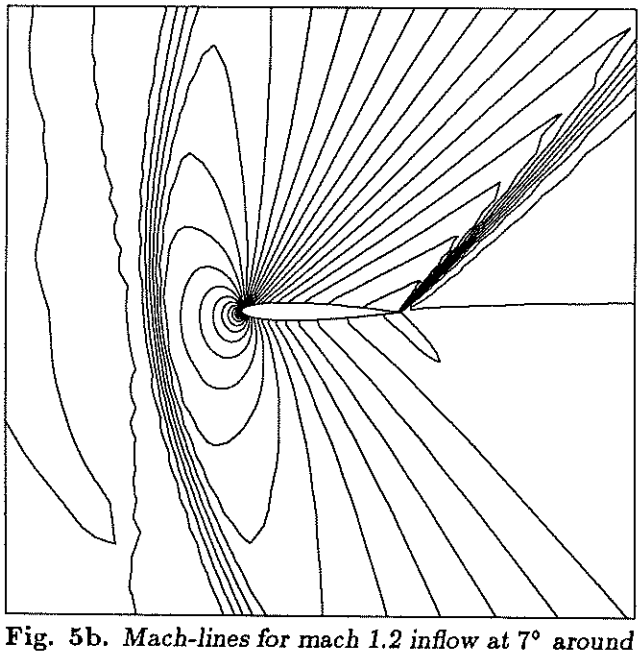

Fig. 5b. *Mach-lines for mach 1.2 inflow at 7° around a NACA0012 airfoil. Contours are 0.05 apart.*

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 streamline 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-level 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 7 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 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 lost of 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 8 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 well below the estimated truncation error.

References

[1] W. K. Anderson, J. L. Thomas, B. van Leer, *A comparison of finite volume flux vector splittings for the Euler equations*, AIAA Paper No. 85-0122, presented at the AIAA 23rd Aerospace Sciences Meeting, January 1985, Reno, Nevada.

[2] A. Brandt, *Guide to Multigrid Development*, Lecture Notes in Mathematics 960 (1981), 220-312.

[3] W. Hackbusch, *Multi-Grid Methods and Applications*, Springer Series in Computational Mathematics 4 (1985), Springer Verlag, Berlin/Heidelberg.

[4] P. W. Hemker and S. P. Spekreijse, *Multiple grid and Osher's scheme for the efficient solution of the steady Euler equations*, Appl. Num. Math. 2 (1986), 475-493.

[5] B. Koren, *Defect Correction and Multigrid for an Efficient and Accurate Computation of Airfoil Flows*, J. Comp. Phys. 77 (1988), 183-206.

[6] W. A. Mulder, *Multigrid Relaxation for the Euler Equations*, J. Comp. Phys. 60 (1985), 235-252.

[7] W. A. Mulder, *Computation of the Quasi-Steady Gas Flow in a Spiral Galaxy by Means of a Multigrid Method*, Astron. Astrophys. 156 (1986), 354-380.

[8] W. A. Mulder, *Multigrid for the one-dimensional inviscid Burgers' equation*, SIAM J. Sci. Stat. Comput., in press.

[9] W. A. Mulder, *Analysis of a multigrid method for the Euler equations of gas dynamics in two dimensions*, in *Multigrid Methods: Theory, Applications, and Supercomputing*, ed. S. McCormick, Marcel Dekker, New York (1988).

[10] W. A. Mulder, *A note on the use of Symmetric Line Gauss-Seidel for the upwind differenced Euler equations*, SIAM J. Sci. Stat. Comput., in press.

[11] W. A. Mulder, *A new multigrid approach to convection problems*, J. Comp. Phys., in press.

[12] W. A. Mulder, *Multigrid, alignment, and Euler's equations*, presented at the Fourth Copper Mountain Conference on Multigrid Methods, Colorado, April 1989. To be published by SIAM.

[13] S. Osher and F. Solomon, *Upwind difference schemes for hyperbolic systems of conservation laws*, Math. Comp. 38 (1982), 339-374.

[14] M. M. Rai and S. R. Chakravarty, *An Implicit Form for the Osher Upwind Scheme*, AIAA J. 24 (1986), 735-743.

[15] S. P. Spekreijse, *Multigrid Solution of the Steady Euler Equations*, Ph. D. Thesis, Centrum voor Wiskunde en Informatica, Amsterdam (1987).

[16] G. D. van Albada, B. van Leer, and W. W. Roberts, *A Comparative Study of Computational Methods in Cosmic Gas Dynamics*, Astron. Astrophys. 108 (1982), 76-84.

[17] B. van Leer, *Flux-vector splitting for the Euler equations*, Lecture Notes in Physics 170 (1982), 507-512.

[18] B. van Leer, *Upwind-difference methods for aerodynamic problems governed by the Euler equations*, in *Large-Scale Computations in Fluid Mechanics*, Lectures in Applied Mathematics 22, eds. B. E. Engquist, S. Osher, R. C. J. Somerville, American Math. Society, Providence, Rhode Island, 1985, 327-336 (part 2).

[19] B. van Leer and W. A. Mulder, *Relaxation Methods for Hyperbolic Conservation Laws*, in *Numerical Methods for the Euler Equations of Fluid Dynamics*, eds. F. Angrand, A. Dervieux, J. A. Desideri, and R. Glowinski, SIAM, Philadelphia (1985), 312-333.