AIAA JOURNAL
Vol. 32, No. 10, October 1994

Comparison of Two Formulations for High-Order Accurate Essentially Nonoscillatory Schemes

Jay Casper*
ViGYAN, Inc., Hampton, Virginia 23666
Chi-Wang Shu†
Brown University, Providence, Rhode Island 02912
and
H. Atkins‡
NASA Langley Research Center, Hampton, Virginia 23681

The finite volume and finite difference implementations of high-order accurate essentially nonoscillatory shock-capturing schemes are discussed and compared. Results obtained with fourth-order accurate algorithms based on both formulations are examined for accuracy, sensitivity to grid irregularities, resolution of waves that are oblique to the mesh, and computational efficiency. Some algorithm modifications that may be required for a given application are suggested. Conclusions that pertain to the relative merits of both formulations are drawn, and some circumstances for which each might be useful are noted.

Introduction

TWO basic formulations, finite volume and finite difference, for the implementation of high-order accurate, essentially nonoscillatory (ENO) shock-capturing schemes have been the subject of considerable interest in recent years. These schemes achieve high-order spatial accuracy in smooth regions by means of a piecewise polynomial approximation operator that is also designed to avoid oscillations associated with interpolation across steep gradients. As such, they are well suited for the study of aeroacoustic and transition-related problems and may serve as an alternative to spectral methods for solving such problems when shocks or complex geometries are involved.

The finite volume implementation, first presented by Harten et al.,[1] is preferred for its strict adherence to the integral form in which conservation laws are defined. The primary motivation for the use of the finite difference approach of Shu and Osher[2] is computational efficiency. These formulations are briefly described, after which results of their numerical implementations are presented for comparison. The intent of this work is to acquaint the reader with the relative merits of both formulations, the circumstances for which each might be useful, and some details of implementation that may be required for a given application. The performances of both algorithms are compared for accuracy, sensitivity to grid irregularities, resolution of waves that are oblique to the mesh, and computational efficiency.

Discrete Formulations

The finite volume and finite difference algorithms compared in this paper differ fundamentally in the way that a system of equations is solved. In both cases, a weak solution of a system of conservation laws is ultimately obtained. The conservation of some quantity U in a spatial domain D can be written

$$\frac{\partial}{\partial t}\int_D U \, dV = -\int_{\partial D} \boldsymbol{F}\cdot\boldsymbol{n} \, dS \tag{1}$$

where $\boldsymbol{F}$ is the flux, ∂D is the boundary of D, dV is a volume element of D, dS is an element of surface area on ∂D, and $\boldsymbol{n}$ denotes the outward unit normal to ∂D.

In the finite volume approach, the conservation law itself in Eq. (1) is approximated. The spatial domain is discretized, $D = \{D_i\}$, which results in

$$\frac{\partial}{\partial t}\overline{U}_i = -\frac{1}{V_i}\int_{\partial D_i} \boldsymbol{F}\cdot\boldsymbol{n} \, dS \tag{2}$$

where V_i is the volume of D_i and

$$\overline{U}_i = \frac{1}{V_i}\int_{D_i} U \, dV \tag{3}$$

is the cell average of U in D_i. Equation (2) is solved for all i, which yields a solution $\{\overline{U}_i\}$ of cell averages.

In the finite difference approach, a pointwise solution is desired. To this end, time differentiation and spatial integration are interchanged in Eq. (2), the divergence theorem is applied on the right-hand side, and in the limit as $V_i \to 0$,

$$\frac{\partial}{\partial t}U_i = -(\nabla\cdot\boldsymbol{F})_i \tag{4}$$

Special care must be taken when this formulation is implemented because flux conservation is not as readily achieved as in Eq. (2).

For integration in time, the method of lines will be employed for Eqs. (2) and (4). High-order accurate Runge-Kutta methods, developed by Shu and Osher,[2] are implemented in the finite volume and finite difference schemes to be compared in this paper. Hence, the brief description of the schemes to follow will concern only the right-hand sides of Eqs. (2) and (4).

Both discrete algorithms involve a reconstruction step followed by an evolution step. What is meant by reconstruction is a high-order accurate polynomial approximation at some point in time. In the finite volume formulation, the solution U is reconstructed from

Received June 29, 1993; presented as Paper 93-3338 at the AIAA 11th Computational Fluid Dynamics Conference, Orlando, FL, July 6–9, 1993; revision received Feb. 10, 1994; accepted for publication Feb. 22, 1994.

*Research Engineer, 30 Research Drive. Member AIAA.
†Associate Professor, Division of Applied Mathematics.
‡Research Scientist, Computational Aerodynamics Branch, Fluid Mechanics Division. Member AIAA.

the cell averages to high order within each cell D_i and evaluated on the boundary ∂D_i. The evolution step involves the solutions of the local Riemann problems that arise from the piecewise continuous reconstruction (see, e.g., Refs. 1 and 3 and the references therein). The spatial integration along the boundary ∂D_i is achieved by a correspondingly high-order quadrature. This method will be referred to throughout the paper as ENO-FV. Shu and Osher[2,4] have proposed the use of the finite difference formulation in Eq. (4), in which the reconstruction operator is applied directly to the pointwise flux values. The evolution step arises from a flux-vector-splitting strategy that is built into the reconstruction. The term ENO-FD will be used to refer to this finite difference scheme. The pointwise nature of this formulation eliminates the need for dealing with cell averages or the integration of a flux over the boundary of a cell. These distinctions between the two formulations become most important with regard to the issue of cost and will be discussed in more detail in a later section.

The most unique aspect of ENO reconstruction operators is their use of adaptive stenciling. At the ith discrete location, the reconstruction set that determines a polynomial $\mathcal{P}_i$ is chosen by a searching algorithm in which decisions are based on local smoothness criteria. For example, in one dimension, where the $\{D_i\}$ are intervals on the real line, Harten et al.[1] have suggested a hierarchical algorithm that is based on local sets of kth divided differences $\{\delta_j^k\}$, where $i-m \le j \le i+m$, $k = 1, 2, \ldots, m$, and m is the desired degree of $\mathcal{P}_i$. If a contiguous one-dimensional stencil for $\mathcal{P}_i$ is defined by its left-most index i_m, then the first-order choice is clearly $i_1 = i$. Then, for $k = 2, \ldots, m$,

$$i_k = \begin{cases} i_{k-1} - 1, & \text{if } \left|\delta^k_{i_{k-1}}\right| < \left|\delta^k_{i_k}\right| \\ i_{k-1}, & \text{otherwise} \end{cases} \tag{5}$$

Because this algorithm allows the reconstruction stencil to shift freely with the detection of any numerical gradient, it will be referred to as "freely adaptive." Reasons for modifying this search will become apparent.

A few observations that concern the available types of reconstruction operators are in order. Within the Shu-Osher approach, when conservation is desired, the implementation of the high-order reconstruction operator requires a uniform computational mesh. Therefore, the application of the ENO-FD algorithm on a nonuniform physical domain requires a sufficiently smooth transformation to a uniform mesh if third-order or higher accuracy is desired. An analogous form of this transformational reconstruction (TR) can also be implemented within the finite volume formulation. In this case, the reconstruction operator is applied to a set of volume-weighted averages $\{V_i \bar{U}_i\}$, and therefore sufficient smoothness in a spatial transformation is required.[5] However, another option exists for the ENO-FV algorithm. The polynomial approximation can be performed in physical space, which releases such burdensome restrictions on grid smoothness.[6-8] This latter procedure will be referred to as physical reconstruction (PR). Although this PR operator poses no problems in one dimension, its implementation in multidimensional space can be quite complex when the implementation must allow for local stencil adaptation.[7,8] However, the multidimensional, finite volume TR operator can be readily implemented because it is defined as a product of one-dimensional operators.[5] When required, the finite volume algorithms will be distinguished as ENO-FV-TR and ENO-FV-PR.

One-Dimensional Rarefaction Wave

The first test case involves the solution of the Euler equations of gasdynamics in one spatial dimension, as it pertains to the movement of a right-traveling rarefaction wave through a domain of highly varying mesh spacing. The effects of two mesh transformations will be examined. One transformation has the required smoothness for a fourth-order TR operator, and the other does not.

For both transformations, the uniform computational domain is given by $\{-6 \le \xi \le 6\}$. The subset $\{-4 \le \xi \le 4\}$ is divided into five intervals: $[-4, -3]$, $[-3, -1]$, $[-1, 1]$, $[1, 3]$, and $[3, 4]$. In each of these interior intervals, the mapping $x = x(\xi)$ causes the physical mesh width Δx to vary rapidly; the intervals $\{-6 \le \xi < -4\}$ and $\{4 < \xi \le 6\}$ are mapped uniformly.

In the kth interior interval, the first transformation is of the form

$$x(\xi) = \xi_k + \sin\left[\frac{9\pi}{20}(\xi - \xi_k)\right] \tag{6}$$

where ξ_k is an element of the set $\{-4, -2, 0, 2, 4\}$. The mesh spacing Δx in the uniform regions (fore and aft) is determined such that the connections at $\xi = \pm 4$ are smooth.

To generate a smoother grid on $\{-4 \le \xi \le 4\}$, a mapping of the following form is used:

$$x(\xi) = \left(\alpha - \frac{5}{16}\beta\right)\xi + \frac{\beta}{32}\left[\frac{15}{\pi}\sin(\pi\xi) - \frac{3}{\pi}\sin(2\pi\xi) + \frac{1}{3\pi}\sin(3\pi\xi)\right] \tag{7}$$

The parameters α and β are determined so that the ratio of the maximum to minimum values of x_ξ are identical for the two grids and the physical distance between $x(-4)$ and $x(4)$ is the same for both.

Figure 1a illustrates the similar behavior of these two transformations on $0 < \xi < 2$. The seemingly odd formula in Eq. (7) was chosen because its derivative is of the form

$$x_\xi = \alpha - \beta \sin^6\left(\frac{\pi}{2}\xi\right)$$

The transformation derivatives are plotted on the same interval in Fig. 1b. The value $\bar{x}_\xi$ represents the derivative normalized by its maximum value. For the transformation (6), $x_{\xi\xi}$ is discontinuous for $\xi = \pm 1, \pm 3$, which makes for an overall C^1 grid. The mesh produced by Eq. (7) will be C^∞ on the nonuniform region and C^6 overall, because of the connections with the uniform intervals at $\xi = \pm 4$.

The initial solution consists of an isentropic expansion, smoothly distributed on $-6.5 < x < -5.5$, with constant states U_1 and U_2 on the right and left, respectively. The strength of this rarefaction is determined by requiring a mean temperature change of $\pm 5\%$ across the wave and a mean Mach number $\bar{M} = 1.134334$. The problem is nondimensionalized with respect to the mean solution. For $t > 0$, the rarefaction wave moves to the right, and its behavior is monitored within the nonuniform region of the mesh until $t = 4.0$.

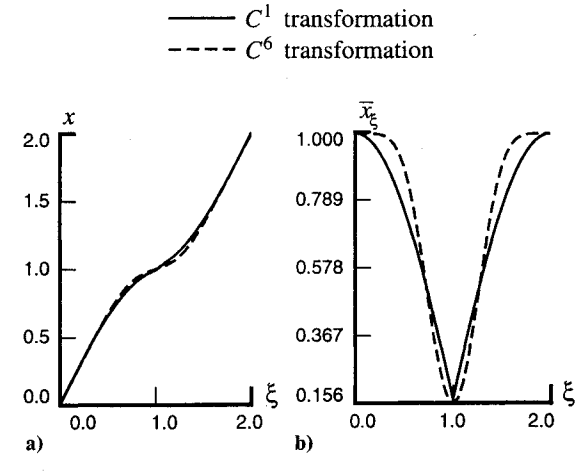

Fig. 1 a) Transformations of Eqs. (6) and (7) and b) normalized derivatives of Eqs. (6) and (7).

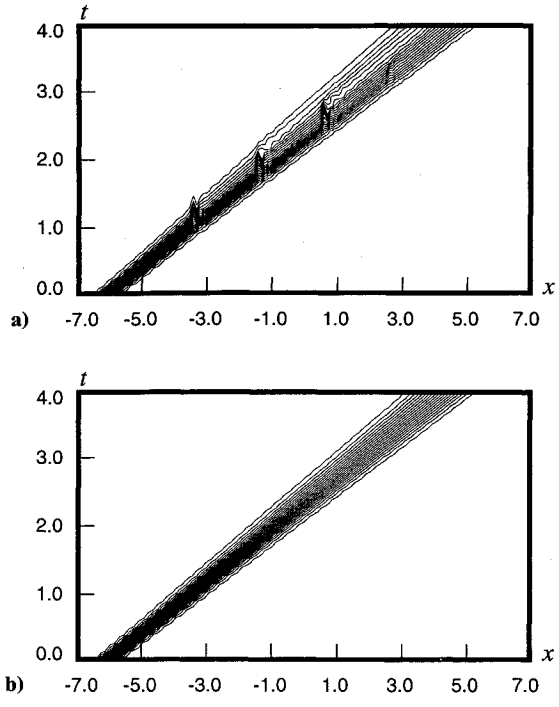

Fig. 2 Density on the C^1 mesh: a) ENO-FD, freely adaptive stencil, and b) ENO-FV-PR, freely adaptive stencil.

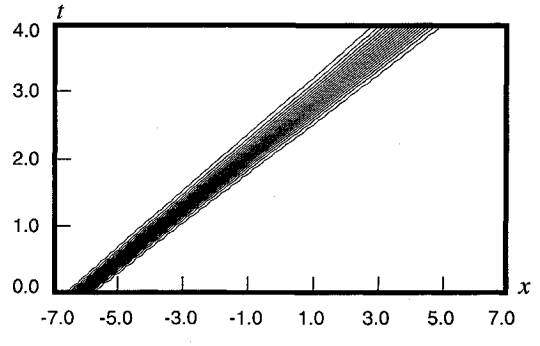

Fig. 3 Density on the C^6 mesh, ENO-FD, freely adaptive stencil.

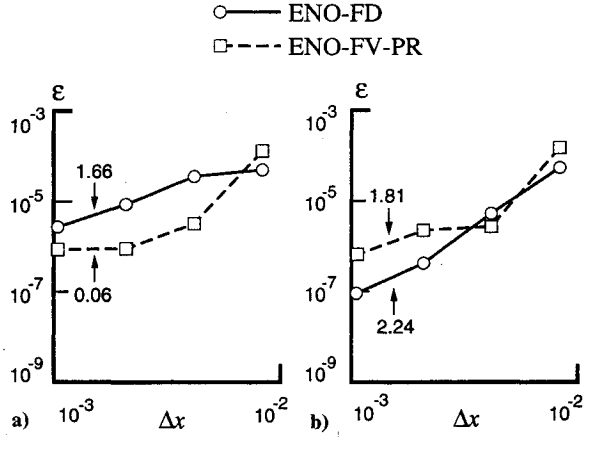

Fig. 4 L_1 density error, $t = 4.0$, C^6 grid: a) freely adaptive stencil and b) biased stencil ($\bar{\sigma} = 2.0$).

Both algorithms used here are fourth-order accurate in the L_1 sense. Figure 2 compares solutions for the two formulations by x-t contour plots of the density, for $0.0 < t < 4.0$, on a mesh of 120 intervals with the C^1 transformation given by Eq. (6). The freely adaptive stencil algorithm in Eq. (5) is employed. All of the contour plots in this section represent the entire solution range, $\rho_1 < \rho < \rho_2$, with 20 equally spaced levels. The ENO-FD solution exhibits significant qualitative error in neighborhoods about the second-derivative discontinuities in the mesh. In Fig. 2b, such problems are not apparent when the solution is computed with the freely adaptive ENO-FV-PR algorithm.

Figure 3 depicts the use of the same fourth-order ENO-FD algorithm on the C^6 grid generated by Eq. (7). With the required mesh smoothness to support the reconstruction, no visible distortions appear as on the C^1 grid. However, the fourth-order design accuracy is not achieved for this solution when density error is measured with respect to the L_1 norm in a region to the left of the rarefaction at $t = 4.0$. The same accuracy problem was found with the freely adaptive ENO-FV-PR scheme. The results of both mesh-refinement studies are shown on a log-log plot in Fig. 4a. The spatial discretizations employed in this study are 120, 240, 480, and 960 intervals. The numbers on the loci represent the computational order of accuracy as measured between the two finest meshes.

A similar loss-of-accuracy phenomenon was reported by Rogerson and Meiberg,[9] which prompted a response from Shu[10] that the problem arises by allowing the stencil to adapt too freely. Shu has suggested that the stencil adaptation algorithm be modified to bias the stencil toward one that is stable, in the sense of linear stability analysis. In the present application, the resulting stencil is one that is upwind biased. This biasing can be done by implementing a factor σ in the stencil search in Eq. (5), viz.,

$$i_k = \begin{cases} i_{k-1} - 1, & \text{if } \sigma_L \left|\delta^k_{i_{k-1}}\right| < \sigma_R \left|\delta^k_{i_k}\right| \\ i_{k-1}, & \text{otherwise} \end{cases} \tag{8}$$

where $(\sigma_L, \sigma_R) = (1, \bar{\sigma})$ or $(\bar{\sigma}, 1)$, for biasing to the left or right, respectively, with $\bar{\sigma} > 1$.

Figure 4b shows grid-refinement results using the biased stencil with ($\bar{\sigma} = 2.0$). Although the convergence is improved, the fourth-order design accuracy is still not achieved. The convergence problem persists because, in regions on either side of the wave, the solution is essentially a constant and numerical gradients are essentially zero. In such regions, the magnitudes of adjacent differences can easily vary by an order of magnitude, and thus a factor of 2.0 will have little, if any, effect. In Ref. 11, Atkins has suggested a very different adaption algorithm having the property that the stencil tends toward a preferred, linearly stable stencil, whenever the solution is smooth. This philosophy is implemented in the present case as

$$i_k = \begin{cases} i_{k-1} - 1, & \text{if } \sigma_L \left(\left|\delta^k_{i_{k-1}}\right| + \bar{\varepsilon} \right) < \sigma_R \left(\left|\delta^k_{i_k}\right| + \bar{\varepsilon} \right) \\ i_{k-1}, & \text{otherwise} \end{cases} \tag{9}$$

where $\bar{\varepsilon} = \varepsilon / k! \bar{h}^k$, ε is a small positive parameter, and $\bar{h}$ is a local average mesh width.

Another grid-refinement study was performed with Eqs. (8) and (9) wherein $\bar{\sigma} = 2.0$ and $\varepsilon = 0.01$. The results are shown in Fig. 5a and are much more consistent with the scheme's design accuracy. In fact, both algorithms, coupled with the stenciling modifications in Eqs. (8) and (9), yield fourth-order computational accuracy on the C^1 grid, which is shown in Fig. 5b. The high-order convergence of the ENO-FD algorithm is a fortuitous result inherent to this particular test case. In general, a nonsmooth grid will adversely affect the results of the ENO-FD method. This will be demonstrated in the next section in which a two-dimensional test case is considered. Figure 6 shows that the qualitative error in the ENO-FD solution on the C^1 grid, which was previously shown in Fig. 2a, has been significantly reduced with these stencil biasing modifications.

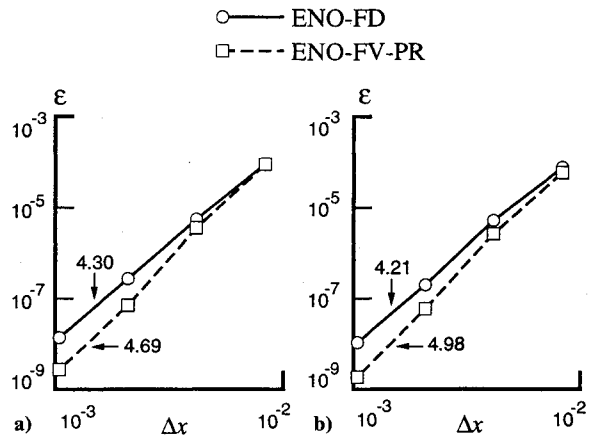

Fig. 5 L_1 density error, $t = 4.0$: a) C^6 grid, biased stencil ($\bar{\sigma} = 2.0$, $\varepsilon = 0.01$), and b) C^1 grid, biased stencil ($\bar{\sigma} = 2.0$, $\varepsilon = 0.01$).

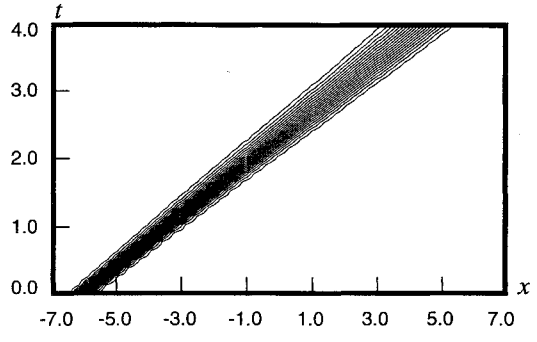

Fig. 6 Density on the C^1 mesh, ENO-FD, biased stencil ($\bar{\sigma} = 2.0$, $\varepsilon = 0.01$).

At this point, the slightly improved performance of the ENO-FV algorithm can be attributed to the fact that this formulation allows for a physical reconstruction that is not available for the finite difference algorithm. However, as previously noted, multidimensional extensions for this more generalized operator can be complicated and costly. Therefore, some results produced with the ENO-FV-TR algorithm will be discussed. In Fig. 7a, when the freely adaptive version of this algorithm is implemented on the C^6 grid, significant error exists that was not evident in the analogous case for the ENO-FD solution in Fig. 3. This error is due to the reconstruction of volume-weighted averages, in which a rapidly varying mesh will have an inordinate effect on the stencil choice in a region of small solution gradients. However, biasing the reconstruction stencils with Eqs. (8) and (9) is not enough to entirely rid the solution of these grid distortion errors, as shown in Fig. 7b. The problem is still related to the reconstruction of volume-weighted averages, in particular to the fact that a uniform flow is not preserved by this procedure if the analytic transformation is used for the necessary discrete values $\{x_{\xi,i}\}$. Therefore, instead of using Eq. (7b) to compute the derivatives, these derivatives are numerically approximated in a manner that enables the reconstruction operator to preserve the freestream exactly.[5] The resulting solution, with this final modification, is shown in Fig. 7c. In this final form, the numerical accuracy of the ENO-FV-TR algorithm is found to perform to design as shown in Fig. 8a. The ENO-FD and ENO-FV-PR results from Fig. 5a are repeated in Fig. 8b for comparison.

Two-Dimensional Channel Flow

The ENO-FD and ENO-FV algorithms are now compared in two spatial dimensions. The test case involves a steady, subsonic flow in a channel of varying area. Although high-order ENO schemes are clearly designed with unsteady solutions in mind,

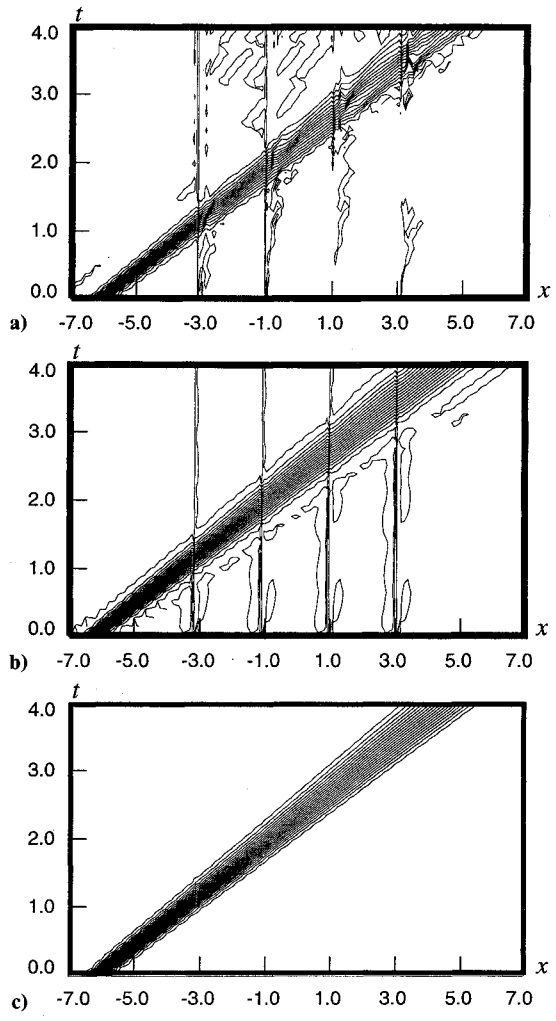

Fig. 7 Density on the C^6 mesh, ENO-FV-TR: a) freely adaptive stencil, b) biased stencil ($\bar{\sigma} = 2.0$, $\varepsilon = 0.01$), and c) biased stencil, freestream preserving.

many such solutions of interest can be considered as the imposition of perturbations upon a steady flow. It is therefore essential that a steady flow be accurately predicted to obtain meaningful results from unsteady problems of an aeroacoustic or transitional nature. Unless otherwise stated, the remainder of the applications of the ENO-FV algorithm will involve the TR type of reconstruction.

This channel flow solution is assumed to be governed by the two-dimensional Euler equations and is computed with both the ENO-FD and ENO-FV algorithms on two different geometries. For the two channels under consideration, the length-to-height ratio is $L/H = 1.5$, with constant-area sections fore and aft of a section of varying area. Each constant-area section has length $L/5$; at the throat, the constriction is 10% of H. The most significant difference between the two geometries is the order of smoothness to which the constant-area sections are connected to the center section. Both channels can be described as follows. Let the rectangle $\{0 \le \xi \le L\} \times \{0 \le \eta \le H\}$ denote the computational domain for each channel. The identity maps (ξ, η) to the physical point (x, y), for ξ in $[0, L/5] \cup [4L/5, L]$ and all η. On the interval $\{L/5 < \xi < 4L/5\}$, the geometry for the varying-area section is given by a transformation of the form

$$x = \xi$$
$$y = \left(1 - \frac{\eta}{H}\right) y_1(\xi) + \frac{\eta}{H} y_2(\xi) \tag{10}$$

where $y_1(\xi)$ and $y_2(\xi)$ are, respectively, the equations for the lower and upper walls.

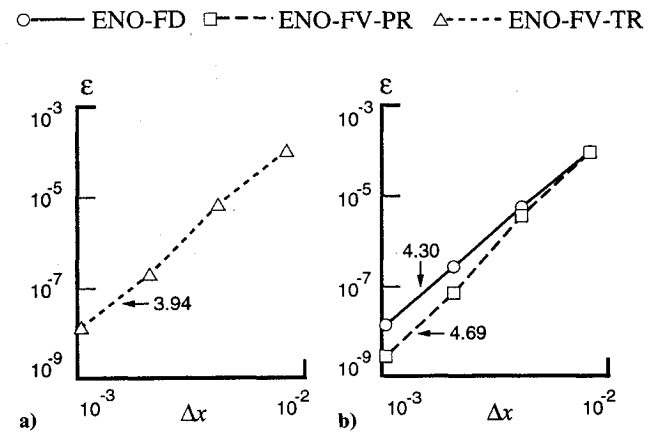

Fig. 8 L_1 density error, $t = 4.0$, C^6 grid: a) biased stencil, freestream preserving, and b) biased stencil ($\bar{\sigma} = 2.0$, $\varepsilon = 0.01$).

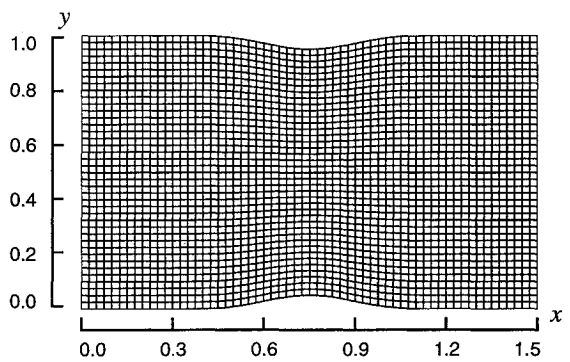

Fig. 9 C^3 geometry of Eqs. (10) and (11).

The first middle section considered is determined by walls of the form

$$y_i(\xi) = a_i \sin^4(b_i\xi + c_i) + d_i, \qquad i = 1, 2 \tag{11}$$

where the values of the coefficients are set for the desired throat constriction and a smooth connection to the constant-area sections. Three continuous ξ derivatives of y exist at $\xi = L/5, 4L/5$, for all η. Therefore, Eqs. (10) and (11) generate a C^3 geometry. Figure 9 illustrates this geometry on a 60×40 mesh. The second channel geometry differs from the first only in that the walls in the middle section are given by polynomials of the form

$$y_i(\xi) = a_i + b_i\xi^2 + c_i\xi^4, \qquad i = 1, 2 \tag{12}$$

In this case, the connections of the middle to the outer sections are continuous in ξ to only one derivative, for $0 < \eta < H$.

The desired solution is that of a steady-state flow that is caused by a uniform, parallel freestream entering the channel at $x = 0$. This solution is achieved by solving the time-dependent Euler equations with the fourth-order algorithms. At $t = 0$, the solution in the entire channel is set to freestream with a Mach number of 0.3. Tangency is imposed on the walls, and a nonreflecting boundary condition[12] is applied at the inflow and outflow. The stencil-biasing modifications in Eqs. (8) and (9) are essential for convergence of the solution to a steady state and are implemented with $\bar{\sigma} = 2.0$ and $\varepsilon = 0.01$. Flux residuals were readily driven to machine zero in both test cases.

These solutions were computed by both algorithms on a sequence of successively refined grids, and the solutions' errors were determined by deviation from isentropy, as measured by the quantity $S = P/\rho^\gamma$. The refinement sequence employed here is 15×10, 30×20, 45×30, 60×40, and 90×60. The results of the grid-refinement study on the C^3 geometry are shown in Figs. 10a–c. The "global" error is computed over the entire computational domain, the "wall" error is computed only at the points along one wall, and the "interior" error is computed on $\{L/4 < \xi < 3L/4\} \cup \{H/4 < \eta < 3H/4\}$. Both algorithms evidently perform at or near design accuracy, as expected for fourth-order schemes on a C^3 mesh.

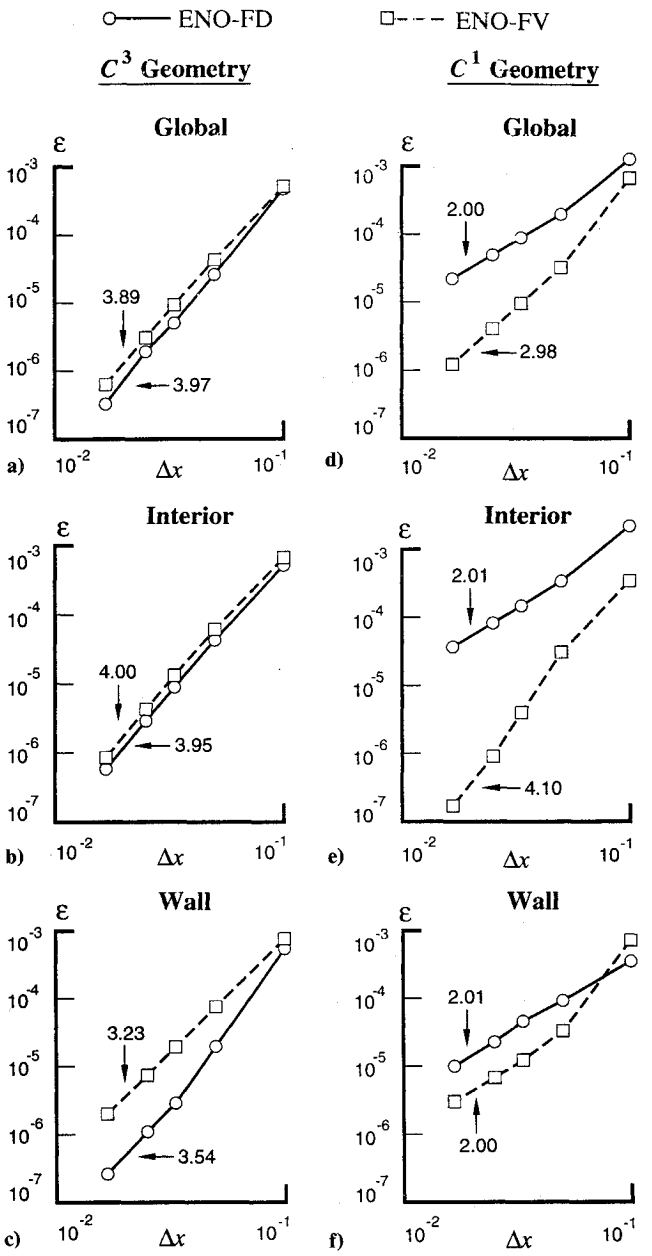

Fig. 10 L_1 entropy errors, steady two-dimensional channel flow.

Figures 10d–10f show the grid-refinement results for the C^1 mesh. As expected, second-order results are obtained on the wall for both algorithms. However, the finite volume algorithm performs at third-order accuracy with respect to the global error and at design accuracy on the interior, whereas the finite difference algorithm shows second-order accuracy in all three measures. These results suggest that, for the ENO-FD algorithm, the second-order entropy error that arises from the nonsmoothness at the section connections is propagating into the interior. This assumption is supported by Fig. 11, in which the quantity $\log S$ is plotted along the centerline of a 60×40 mesh. Clearly, in the finite difference solution, there is a lower-order entropy error within the middle section than exists in the constant-area sections.

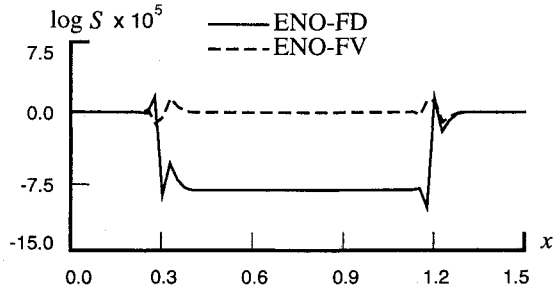

Fig. 11 Centerline entropy, C^1 grid.

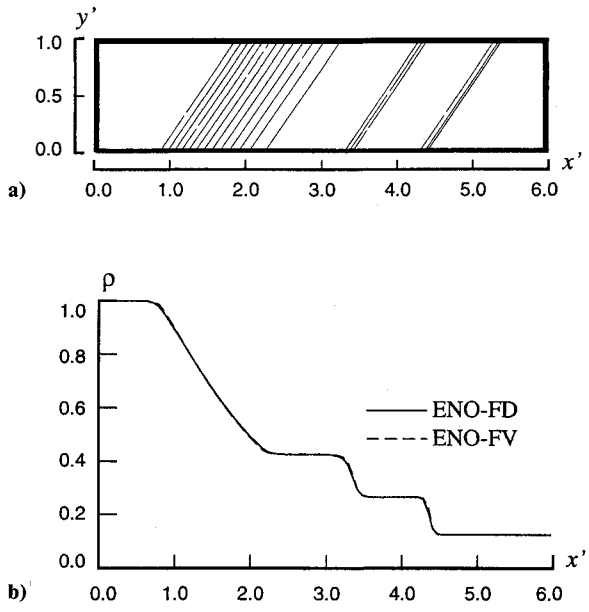

Fig. 12 a) Density contours, ENO-FV, θ = arctan 1 and b) density, $y' = 0.0$, θ = arctan 1.

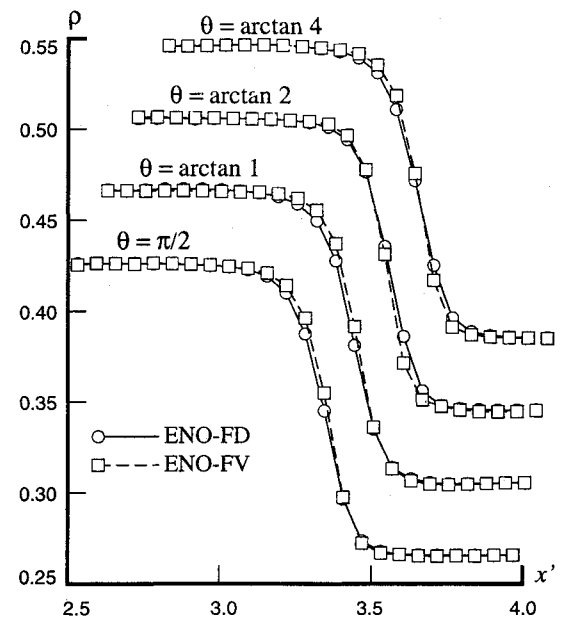

Fig. 13 Density, contact discontinuity, $y' = 0.0$.

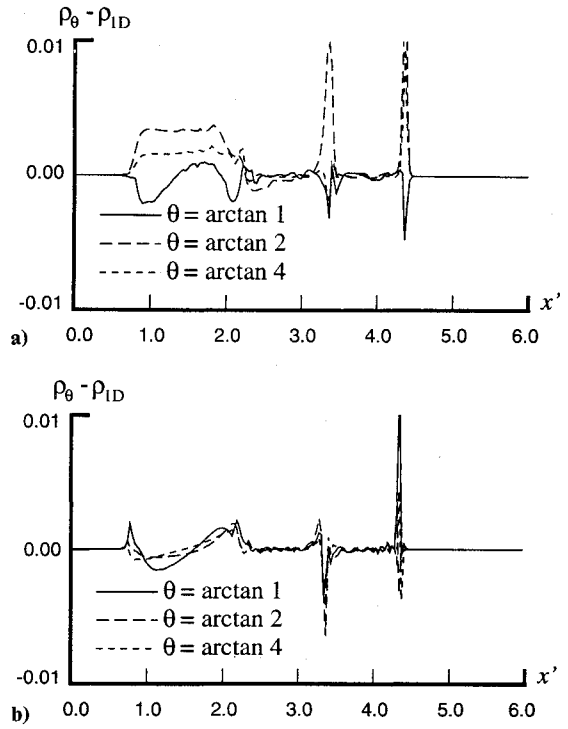

Fig. 14 a) Density deviation from one dimension, ENO-FD, and b) density deviation from one dimension, ENO-FV.

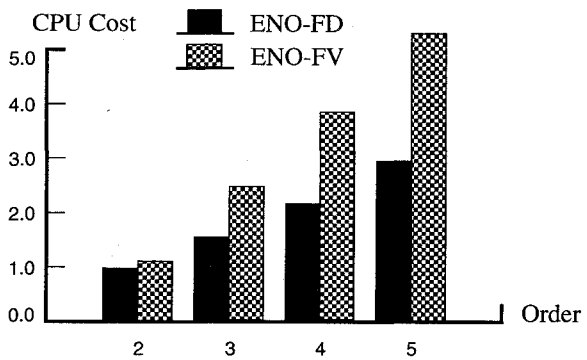

Fig. 15 Relative CPU cost of the two algorithms in two spatial dimensions.

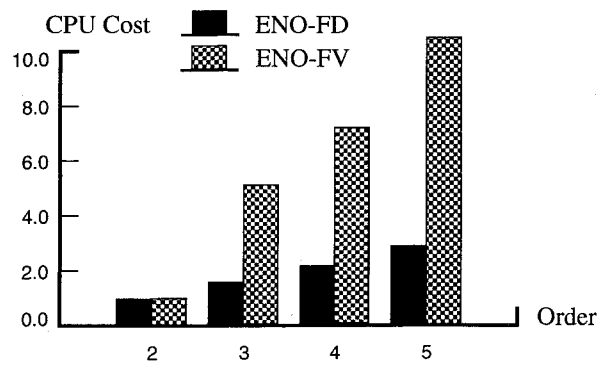

Fig. 16 Relative CPU cost of the two algorithms in three spatial dimensions (estimated).

Oblique Sod's Problem

The final test under consideration is selected to compare the capabilities of the two algorithms to resolve waves that are oblique to the computational mesh. Sod's problem[13] will be solved in two-dimensional space so that the planar waves produced will propagate at various angles of incidence with respect to a rectangular grid. The intent is not only to inspect the qualitative resolution of the oblique waves but also to quantify the manner in which each algorithm detects an oblique wave with respect to its detection of a wave that is normal to the mesh.

Sod's problem is a Riemann problem that is subject to the solution of the Euler equations and determined by initial conditions that consist of a thermodynamic discontinuity imposed upon a fluid at rest. The magnitudes of the density and pressure jumps are, from left to right, $\rho_L/\rho_R = 8$ and $P_L/P_R = 10$, respectively. Such an initial solution is installed on a two-dimensional Cartesian grid so that the discontinuity is a straight line that makes an angle θ with the x axis. The physical mesh varies with θ in the following manner. A computational mesh is defined on the rectangle $[0, L] \times [0, H]$, with a length-to-height ratio $L/H = 6$. If $0 < \theta \leq \pi/2$, then the physical domain is $\{0 < x < L_\theta\} \times \{0 < y < H_\theta\}$ and is determined by $H_\theta = H/\sin\theta$ and $L_\theta = L/\sin\theta$. This scaling achieves the same grid resolution normal to the wave propagation on a given mesh at some fixed time for all choices of θ. The mesh discretization in each direction is chosen such that $\Delta x = \Delta y$.

At $t = 0$, the initial discontinuity is positioned at $(x, y) = (3L/8, 0)$ and inclined at the angle θ. The angles of inclination are chosen so that $\tan\theta$ is an integer and that the boundary conditions at $y = 0$ and $y = H_\theta$ can be determined in a "shifted-periodic" manner. In particular, the test angles are $\theta = \arctan 1$, $\arctan 2$, and $\arctan 4$, in addition to the one-dimensional problem $\theta = \pi/2$. The solution is computed to $t = 1.2$. The stencil biasing in Eq. (8) is used in the reconstruction procedures with $\bar{\sigma} = 2.0$.

Figure 12a represents the solution at $t = 1.2$, on a 96×16 grid, with $\theta = \arctan 1$, with the use of the ENO-FV algorithm. The three wave structures are, from left to right, an expansion wave, a contact discontinuity, and a shock. The axis variables x' and y' represent the physical coordinates scaled by $\sin\theta$. Figure 12b depicts the same solution as in Fig. 12a, for both algorithms, along the grid line $y' = 0$. On this level, the qualitative difference between the solutions is barely detectable. Therefore, in Fig. 13, the solutions are examined in the vicinity of the contact discontinuity for all test values of θ. (The successive loci in this plot have been shifted upward and to the right.) Although the ENO-FV scheme produces a slightly steeper numerical gradient at the location of the discontinuity, such a small difference can be judged insignificant. Figure 13, then, suggests that both algorithms perform equally well with respect to oblique waves irrespective of θ. On a more quantitative level, the differences between the oblique cases and the one-dimensional case, for each algorithm, have been measured. These results are plotted in Figs. 14a and 14b. By this measure, the finite volume scheme performs marginally better, particularly with respect to the rarefaction wave.

Cost Comparison

As with any numerical algorithm, the cost of implementing either the ENO-FD or ENO-FV scheme is a major concern. Conceptually, the two formulations can be made equally cost effective for high-order solutions of one-dimensional problems or for first- or second-order-accurate solutions in multiple dimensions. However, for third- or higher-order accuracy in two or more dimensions, the algorithms are radically different. This difference translates to a significant disparity in cost.

With regard to cost, the ENO-FD algorithm has a clear advantage when applied to multidimensional problems. This advantage can be entirely attributed to the fact that the finite difference operator solves a system of equations in a pointwise manner. In this case, a high-order multidimensional reconstruction can be accomplished in a dimension-by-dimension fashion. If the algorithm is rth order accurate in the L_1 sense, then the reconstruction stencil in k dimensions contains $k(r-1) + 1$ points. The dimension-by-dimension approach to reconstruction is not possible within the finite volume formulation, if third-order or better accuracy is required. Because the solution at any time is in the form of cell averages, the multidimensional ENO-FV-TR reconstruction is implemented as a product of one-dimensional operators and thereby requires r^k cells in a reconstruction stencil. Moreover, the Gaussian quadrature used to integrate the flux on each cell boundary requires multiple solutions of the Riemann problem on each cell interface. The number of points (n_q) required in this quadrature on an interface is $n_q = \{\text{int}[(r + 1)/2]\}^{k-1}$, where int$(\cdot)$ denotes integer truncation. This high-order quadrature is entirely avoided within the finite difference approach. It should be noted that both formulations require the same amount of logic in the stencil adaptation, which is the most expensive part of either algorithm.

Figure 15 depicts the relative CPU cost of the two algorithms, for solutions of the Euler equations in two dimensions, for second-, third-, fourth-, and fifth-order-accurate applications. CPU times were measured on a Cray Y-MP and normalized with respect to the second-order ENO-FD cost. An estimated cost comparison for three-dimensional problems can be obtained by the following reasoning. For a specified order of accuracy, let C_1 denote the cost of solving one equation at a point in one dimension; this is the same for either algorithm. Then, because of its dimension-by-dimension implementation, the estimated cost incurred by the ENO-FD algorithm in solving the Euler equations in k dimensions on a grid of N^k points is

$$C_{FD} = C_1(k + 2)kN^k \tag{13}$$

This formula predicts a cost-per-point increase factor of 8/3 when the one-dimensional ENO-FD algorithm is extended to two dimensions. This value has been supported by computer measurements. The corresponding factor for the extension from two to three dimensions is 15/8. The estimation of a three-dimensional cost for the ENO-FV algorithm cannot be done by such a simple linear extrapolation. There are two significant costs in addition to the base value given by Eq. (13). One is the extension of the finite volume reconstruction to a three-dimensional product. The other arises from the additional $3(n_q-1)$ flux computations that are required in the integration over the surface of each cell. These additional costs were extrapolated from two-dimensional CPU measurements on a Cray Y-MP. Figure 16 represents the estimated three-dimensional cost comparison where, again, the data are normalized by the cost for the second-order ENO-FD algorithm.

Concluding Remarks

For accuracy on a sufficiently smooth mesh and resolution of oblique waves, both algorithms perform equally well. The finite volume implementation is less sensitive to derivative discontinuities, whether in the computational mesh or in the solution. In particular, the ENO-FV-PR algorithm has the capacity to perform at design accuracy, independent of the mesh. Although the generalized multidimensional adaptive-stenciling implementation of this algorithm has not sufficiently matured, some promising results can be found in Ref. 8. Either of the multidimensional finite volume algorithms is significantly more costly than the finite difference algorithm. Therefore, if, for a given application, the computational domain is known to be sufficiently smooth and can be suitably structured, then the ENO-FD algorithm is the method of choice. However, for problems with complex geometries, it might pay to use the more expensive algorithm if the grid is significantly less costly to generate in a less restrictive fashion.

Acknowledgments

The first author acknowledges support under NASA Contract NAS1-19672. The second author acknowledges support under NASA Grant NAG1-1145 and NASA Contract NAS1-19480 while in residence at the Institute for Computer Applications in Science and Engineering, NASA Langley Research Center, Hampton, VA, as well as support from the U.S. Army Research Office, Grant

DAAL03-91-G-0123, the National Science Foundation, Grant DMS 9211820, and the Air Force Office of Scientific Research, Grant 93-0090.

References

[1]Harten, A., Engquist, B., Osher, S., and Chakravarthy, S., "Uniformly High Order Accurate Essentially Non-Oscillatory Schemes III," *Journal of Computational Physics*, Vol. 71, No. 2, 1987, pp. 231–323.

[2]Shu, C., and Osher, S., "Efficient Implementation of Essentially Non-Oscillatory Shock-Capturing Schemes," *Journal of Computational Physics*, Vol. 77, No. 2, 1988, pp. 439–471.

[3]Harten, A., "Preliminary Results on the Extension of ENO Schemes to Two-Dimensional Problems," *Proceedings, International Conference on Nonlinear Hyperbolic Problems* (Saint-Etienne, France), Lecture Notes in Mathematics, edited by C. Carasso, P. A. Raviart, and D. Serre, Berlin, 1987, pp. 23–40.

[4]Shu, C., and Osher, S., "Efficient Implementation of Essentially Non-Oscillatory Shock-Capturing Schemes II," *Journal of Computational Physics*, Vol. 83, No. 1, 1989, pp. 32–78.

[5]Casper, J., and Atkins, H., "A Finite Volume High-Order ENO Scheme for Two-Dimensional Hyperbolic Systems," *Journal of Computational Physics*, Vol. 106, No. 1, 1993, pp. 62–76.

[6]Barth, T. J., and Frederickson, P. O., "High-Order Solution of the Euler Equations on Unstructured Grids Using Quadratic Reconstruction," AIAA Paper 90-0013, Jan. 1990.

[7]Harten, A., and Chakravarthy, S., "Multi-Dimensional ENO Schemes for General Geometries," NASA CR 187637, Inst. for Computer Applications in Science and Engineering, ICASE Rept. 91-76, Sept. 1991.

[8]Abgrall, R., "Design of an Essentially Non-Oscillatory Reconstruction Procedure on Finite Element Type Meshes," NASA CR 189574, Inst. for Computer Applications in Science and Engineering, ICASE Rept. 91-84, Dec. 1991.

[9]Rogerson, A. M., and Meiberg, E., "A Numerical Study of the Convergence Properties of ENO Schemes," *Journal of Scientific Computing*, Vol. 5, No. 2, 1990, pp. 151–167.

[10]Shu, C., "Numerical Experiments on the Accuracy of ENO and Modified ENO Schemes," *Journal of Scientific Computing*, Vol. 5, No. 2, 1990, pp. 127–150.

[11]Atkins, H., "High-Order ENO Methods for the Unsteady Compressible Navier-Stokes Equations," AIAA Paper 91-1557, June 1991.

[12]Atkins, H., and Casper, J., "Non-Reflective Boundary Conditions for High-Order Methods," *AIAA Journal*, Vol. 32, No. 3, 1994, pp. 512–518.

[13]Sod, G. A., "A Survey of Several Finite-Difference Methods for Systems of Nonlinear Hyperbolic Conservation Laws," *Journal of Computational Physics*, Vol. 27, No. 1, 1978, pp. 1–31.