

Stencil Adaptation Properties of a WENO Scheme in Direct Numerical Simulations of Compressible Turbulence

Ellen M. Taylor¹ and M. Pino Martín¹

Received December 21, 2005; accepted (in revised form) May 18, 2006; Published online February 16, 2007

Weighted essentially non-oscillatory (WENO) methods can simultaneously provide the high order of accuracy, high bandwidth-resolving efficiency, and shock-capturing capability required for the detailed simulation of compressible turbulence. However, rigorous analysis of the actual versus theoretical error properties of these non-linear numerical methods is difficult. We use a bandwidth-optimized WENO scheme to conduct direct numerical simulations of two- and three-dimensional decaying isotropic turbulence, and we evaluate the performance of quantitative indicators of local WENO adaptation behavior within the resulting flow fields. One aspect of this assessment is the demarcation of shock-containing and smooth regions where the WENO method should, respectively, engage its adaptation mechanism and revert to its linear optimal stencil. Our results show that these indicators, when synthesized properly, can provide valuable quantitative information suitable for statistical characterization.

KEY WORDS: Shock capturing; adaptive numerical method; WENO; DNS; compressible turbulence; isotropic turbulence.

1. INTRODUCTION

The detailed simulation of turbulent fluid flow requires numerical methods that avoid excessive damping of spatial features over a large range of length scales, including those as small as the grid spacing. Such methods demand high order of accuracy and high bandwidth-resolving efficiency. Compressibility, however, imposes a competing constraint due to the possible presence of shocks and shocklets (small transient shocks). The

¹ Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA. E-mail: pmartin@princeton.edu

application of uniformly high-order numerical methods across shock-containing regions causes spurious oscillations, which are often indistinguishable from turbulent fluctuations and may lead to numerical instabilities and inaccuracies. Weighted essentially non-oscillatory (WENO) schemes offer an alternative approach that combines high order of accuracy and bandwidth-resolving efficiency with shock-capturing for the unsteady shockwaves in compressible turbulence.

WENO schemes [6] compute numerical fluxes at a point in space via several different candidate stencils and form a final flux approximation by summing weighted contributions from each stencil. Smoothness measurements cause stencils that span large flow field gradients to be assigned small relative weights so that a nearly discontinuous shock would provide a weight of almost zero to any stencil containing it. In smooth regions, the relative values of the weights are designed to be optimal by some gauge such as maximum order of accuracy or maximum bandwidth.

Jiang and Shu [3] cast the WENO methodology into finite-difference form and provide an efficient implementation of a robust and high-order-accurate WENO scheme. Unfortunately, this scheme often generates excessive numerical dissipation for detailed simulations of turbulence, especially for large-eddy simulations (LES) [2]. WENO dissipation arises from two distinct sources: (i) the optimal stencil, which on its own describes a linear scheme, and (ii) the adaptation mechanism, which drives the final numerical stencil away from the optimal one. Bandwidth optimization can reduce the dissipation of the optimal stencil [9, 8]; and Martin *et al.* [8] demonstrate that such a bandwidth-optimized WENO method, WENO-SYMBO, indeed reduces numerical dissipation and provides accurate results for direct numerical simulations (DNS) of isotropic turbulence and turbulent boundary layers.

Although the WENO-SYMBO scheme successfully reduces numerical dissipation due to the optimal stencil, the adaptation mechanism still leads to significant dissipation. Decreasing the grid spacing can compensate for this; however, in certain cases, increasing the number of grid points is not feasible. According to Martin [7], LES of turbulent boundary layers with WENO-SYMBO fail because of insufficient distinction between shock-containing and smooth regions on typical LES grids. Additionally, Wu *et al.* [10] encounter local disparities between DNS and experiments of shock/turbulent boundary layer interactions and find that the numerical dissipation of the WENO-SYMBO method, even at the highest possible resolutions, is responsible for the disagreement. Decreasing the dissipation inherent in the WENO adaptation mechanism is challenging because potential deficiencies relevant primarily to the damping of turbulent fea-

tures may become apparent only in realistic simulations and must then be examined locally.

The purpose of this paper is to provide and assess general techniques for probing the local behavior of the WENO stencil adaptation mechanism in turbulent flow fields. Specifically, we investigate and quantitatively summarize characteristics of the final combined stencils produced by the WENO method in realistic simulations. This ability can aid future endeavors to improve the WENO methodology. We conduct direct numerical simulations of isotropic turbulence and focus on comparing shock-containing and smooth flow regions, where WENO adaptation should, respectively, engage and remain disengaged. In the first half of the paper, we present the governing equations, background on the WENO-SYMO method, and details of the numerical simulations. Next, we explain our approaches for isolating shocklets and distilling relevant WENO adaptation information. Finally, we discuss the results and formulate the conclusions from this work.

2. GOVERNING EQUATIONS

The motion of a compressible fluid is governed by the Navier–Stokes equations describing conservation of mass, momentum, and energy,

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_k} (\rho u_k) = 0, \tag{2.1}$$

$$\frac{\partial}{\partial t} (\rho u_i) + \frac{\partial}{\partial x_k} (\rho u_i u_k + p \delta_{ik} - \sigma_{ik}) = 0, \tag{2.2}$$

$$\frac{\partial}{\partial t} (\rho e) + \frac{\partial}{\partial x_k} [(\rho e + p) u_k - \sigma_{kl} u_l + q_k] = 0 \tag{2.3}$$

in which ρ is density, u_i is velocity in the i direction, and p is pressure. The shear stress tensor σ_{ij} is assumed to obey a linear stress-strain relationship,

$$\sigma_{ij} = \mu \left[\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) - \frac{2}{3} \frac{\partial u_k}{\partial x_k} \delta_{ij} \right] \tag{2.4}$$

in which viscosity μ depends only on temperature T through a power law,

$$\mu = \mu_0 \left(\frac{T}{T_0} \right)^n. \tag{2.5}$$

Our working fluid is air, which we treat as a perfect gas with gas constant R and constant specific heat capacity (at constant volume) c_v . Total energy per unit mass e is given by

$$e = c_v T + \frac{1}{2} u_k u_k. \quad (2.6)$$

The heat flux in the i direction q_i arises from conduction through Fourier's law,

$$q_i = -\kappa \frac{\partial T}{\partial x_i} \quad (2.7)$$

in which thermal conductivity κ is directly related to μ as

$$\kappa = \left(\frac{5}{2} c_v - \frac{3}{2} R \right) \mu. \quad (2.8)$$

3. BANDWIDTH-OPTIMIZED WENO METHOD

We summarize the bandwidth-optimized WENO method [9,8] in the context of the scalar, one-dimensional advection equation,

$$\frac{\partial u}{\partial t} + \frac{\partial}{\partial x} f(u) = 0. \quad (3.1)$$

This model equation represents the decoupled forms of Eqs. (2.1)–(2.3) of Sec. 2 after a transformation from physical into characteristic space. If the spatial domain is discretized such that $x_i = i\Delta$, in which Δ is the grid spacing, and $u_i = u(x_i)$, Eq. (3.1) may be cast into the semidiscretized form

$$\frac{du_i}{dt} = -\frac{1}{\Delta} \left(\hat{f}_{i+\frac{1}{2}} - \hat{f}_{i-\frac{1}{2}} \right) \quad (3.2)$$

in which $\hat{f}_{i+1/2}$ is a numerical approximation of $f(u(x_{i+1/2}))$. Once the right-hand side of this expression has been evaluated, numerical techniques for solving ordinary differential equations, such as Runge–Kutta methods, may be employed to advance the solution in time. In order to ensure stability, procedures that approximate $f(u)$ generally split it into $f^+(u)$, which has a strictly non-negative derivative, and $f^-(u)$, which has a strictly non-positive one.

The WENO schemes compute $\hat{f}_{i+1/2}^+$ through interpolating polynomials on a number of candidate stencils each containing r grid points. In the modified WENO method, there are in total $(r+1)$ stencils. The one fully upwinded stencil ranges from $(i-r+1)$ to i , the one fully downwinded stencil ranges from $(i+1)$ to $(i+r)$, and the other stencils fall in between

these two extremes. Figure 1 provides a schematic of this arrangement. Note that, unlike the stencils of Jiang and Shu [3], the present collection is symmetric about the point $(i + \frac{1}{2})$. If the flux approximation on stencil k is designated q_k and the weight assigned to that stencil is ω_k , the final numerical approximation becomes

$$\hat{f}_{i+\frac{1}{2}}^+ = \sum_{k=0}^r \omega_k q_k \tag{3.3}$$

ω_k are normalized forms of weights α_k defined as

$$\alpha_k = \frac{C_k}{(\varepsilon + IS_k)^p} \tag{3.4}$$

in which ε prevents division by zero, IS_k is a smoothness measurement that becomes large when discontinuities are present within stencil k , and p may be varied to increase or decrease WENO adaptation sensitivity. In completely smooth regions, each stencil is equally desirable, and ω_k revert to the optimal weights C_k .

The corresponding stencil diagram for $\hat{f}_{i+1/2}^-$ is simply a mirror image of Fig. 1. Because the total number of data points available to the modified WENO algorithm is $2r$, its maximum order of accuracy is also $2r$; but the bandwidth-optimization procedure [9,8] reduces its guaranteed order of accuracy to r . Thus n th-order accuracy implies $r = n$ and vice versa. The bandwidth-optimization process also introduces a small amount of artificial dissipation to an otherwise neutrally stable optimal stencil to enhance its stability. In practice, the weight of the fully downwinded stencil ω_r is artificially constrained to be no greater than the least of the others so that other independent adverse stability effects are avoided.




Fig. 1. Candidate stencils for the numerical flux $\hat{f}_{i+1/2}^+$ when $r = 3$ for the bandwidth-optimized WENO method [9,8].

4. INDICES OF STENCIL ADAPTATION

The continuity of the WENO weighting process allows the performance characteristics of the final numerical stencil to fall anywhere between those of the least favorable candidate stencil and those of the optimal stencil, which is most favorable. In order to gauge this variation quantitatively but efficiently in a full flow field, Weirs [9] proposed two measures known as the centrality and non-linearity indices. Both compare the weights α_k to the optimal weights C_k and can be incorporated into an existing WENO code with little effort or performance penalty.

The centrality index CI is a measure of stencil symmetry.

$$CI = \left[\sum_{k=0}^r \frac{k (\alpha_k/C_k)}{\sum_{k=0}^r (\alpha_k/C_k)} \right] - \frac{r}{2}. \quad (4.1)$$

A value of zero indicates that the final numerical stencil is perfectly symmetric although it may still deviate substantially from the optimal stencil. If CI is negative, upwinding occurs along with associated dissipation, so dissipation increases with negative departure from CI=0. Similarly, positive values suggest downwinding to some extent. As we previously noted, preferential selection of the stencil lying entirely downwind is prohibited, but no other constraints exist to force the final stencil center to be upwind. For the bandwidth-optimized WENO scheme with $r=3$, the possible range of CI is $[-1.5:1.5]$.

The non-linearity index NI is a measure of the degree of departure from the optimal stencil, which, without adaptation, would lead to a linear numerical method.

$$NI = \left(\sum_{k=0}^r \left[1 - \frac{(r+1) (\alpha_k/C_k)}{\sum_{k=0}^r (\alpha_k/C_k)} \right]^2 \right)^{1/2}. \quad (4.2)$$

This definition forces NI to always be non-negative, and only the optimal stencil can provide a value of zero. NI reaches its theoretical maximum when any one candidate stencil is chosen exclusively to form the final numerical stencil. Ideally, in smooth regions where WENO adaptation is unnecessary, NI should remain much less than this maximum so that the favorable performance capabilities of the optimal stencil are realized. Quantifying the upper limit of these acceptable values is a subject of ongoing research beyond the scope of this paper. The possible range of NI for $r=3$ is approximately $[0:3.5]$.

5. DIRECT NUMERICAL SIMULATIONS

Decaying isotropic turbulence is a canonical flow field that realistically represents the small scales of many turbulent flows. Because the WENO adaptation mechanism relies on local flow information only, studies of adaptation behavior in isotropic turbulence should therefore apply to other turbulent configurations as well. We first concentrate on two-dimensional rather than three-dimensional DNS due to more straightforward visualization and faster simulation turn-around. The density contours of Fig. 2 are typical of these simulations and illustrate the presence of shocklets where gradients are steep. Then, we consider three-dimensional turbulence, which is arguably the only truly physically meaningful form.

The physical domain for isotropic turbulence is either a two-dimensional square or three-dimensional cube with periodic boundary conditions and an edge length such that the large-scale turbulence statistics are sufficiently uncorrelated between the center and edges. An evenly-spaced Cartesian grid discretizes this domain into, respectively, N^2 or N^3 points. The following four parameters govern the generation of an initial field: the average density $\langle \rho \rangle$ and temperature $\langle T \rangle$; the turbulent Mach number




Fig. 2. Density contours produced by a typical direct numerical simulation of two-dimensional isotropic turbulence.

$$M_t = \frac{q}{\langle a \rangle} \quad (5.1)$$

in which $\langle a \rangle$ is the average speed of sound and q is the root-mean-squared velocity summed over all directions; and the Reynolds number based on the Taylor microscale

$$Re_\lambda = \frac{\langle \rho \rangle u'_{\text{rms}} \lambda}{\langle \mu \rangle} \quad (5.2)$$

in which u'_{rms} is the root-mean-squared velocity and λ is the Taylor microscale. Note that these parameters do not remain constant throughout a simulation. For both two and three dimensions, the global strength of turbulent fluctuations decays over time without external forcing, which we have not included.

Initialization of isotropic turbulence begins with uniform density and temperature across the domain and completely random velocity values in the range $[-0.5:0.5]$. The physical velocity field is then transformed into Fourier space where its compressible part is removed and the remainder is manipulated to provide the kinetic energy distribution

$$E(k) = A k^4 e^{-2(k/k_0)^2} \quad (5.3)$$

in which k is non-dimensional wavenumber magnitude, k_0 is the wavenumber at which the energy spectrum peaks, and A is a constant of proportionality that depends on q and k_0 . This is an approximation of the observed distribution in isotropic turbulence, and a numerical simulation starting from Eq. (5.3) must be allowed sufficient time to develop a fuller spectrum. Blaisdell and Ristorcelli [1] have presented methods by which to add consistent dilatation and density and pressure fluctuations to the flow field, but these rely on perturbation techniques that are strictly valid only for small M_t . Because this is not true of the conditions we consider, our initial fields are solenoidal.

Table I lists the governing parameters for the simulations that we perform. In each case, $\langle \rho \rangle = 1 \text{ kg/m}^3$, $\langle T \rangle = 300 \text{ K}$, and $k_0 = 4$. We evolve the initial fields via DNS with a third-order Runge–Kutta time integration technique. The convective numerical fluxes are computed using a third-order, bandwidth-optimized WENO scheme [9,8], and the viscous fluxes, which are diffusive in nature, are computed using a fourth-order central standard (order-optimized) scheme. We deem a field to have reached a proper state of isotropic turbulence when the skewness of the velocity derivatives becomes relatively steady. If we define a reference time $\tau_t = \lambda/u'_{\text{rms}}$, we find that our two-dimensional simulations can be meaningfully

Table I. Initial Turbulent Mach Number M_t and Reynolds Number Based on the Taylor Microscale Re_λ for each Simulation Run. Two-dimensional Simulations are on the Left and Three-dimensional on the Right

Case	M_t	Re_λ	Case	M_t	Re_λ
2D1	0.4	50	3D0	0.1	20
2D2	0.6	50	3D1	0.4	35
2D3	0.8	50	3D2	0.5	35
2D4	1.0	50	3D3	0.6	35
2D5	1.2	50	3D4	0.7	35
2D6	1.4	50	3D5	0.8	35
2D7	1.6	50	3D6	0.9	35
2D8	1.8	50	3D7	1.0	35
2D9	2.0	50	3D8	1.1	35

interrogated for $t/\tau_t \geq 2.0$ and our three-dimensional simulations for $t/\tau_t \geq 1.5$.

We include the validation case 3D0, for which M_t is low, in order to compare the performance of the WENO method to that of a well-established sixth-order Padé scheme [5,4]. Figure 3(a) demonstrates that for $N = 128$ the two techniques generate an identical temporal evolution of turbulent kinetic energy. However, although the Padé scheme produces grid-converged results for $N = 64$, the WENO method requires $N = 128$. In Fig. 3(b), the instantaneous energy spectra of the two methods also agree well when $N = 128$, but for $N = 64$ the WENO scheme leads to pronounced aliasing errors at high wavenumbers. We find that for all three-dimensional flow conditions in Table I, $N = 128$ is sufficient to obtain grid-converged turbulence statistics, and for all two-dimensional conditions, $N = 256$ is sufficient.

6. SHOCK DETECTION

The purpose of WENO methods is to provide an adequate reduction of order of accuracy across the near-discontinuities caused by shocks and shocklets while retaining favorable order-of-accuracy and bandwidth properties elsewhere. We have already confirmed such behavior in canonical test cases [8]. In isotropic turbulence, however, we cannot so easily distinguish shocklets from smooth regions. In order to examine the local WENO adaptation mechanism in realistic flow fields, we require an automated means of determining which grid points lie within shocklets and which do not.




Fig. 3. Comparison of a WENO and a Padé scheme for initial $M_t = 0.1$. (a) Temporal evolution of the average turbulent kinetic energy q^2 (normalized by its initial value). (b) Instantaneous kinetic energy spectra (normalized by the total kinetic energy) at time $t/\tau_t = 1.5$.

Our procedure focuses on the construction of instantaneous streamlines. Although the path of a fluid particle eventually diverges from a streamline with the same starting location, the two are sufficiently close during the typical amount of time that the particle is influenced by a shocklet. We first narrow the field of shock candidates by considering only those grid points at which the magnitude of the pressure gradient is greater than an adjustable threshold value. In theory, this gradient should

be nearly infinite for length scales greater than a few mean free paths, but in practice numerical simulations cause a shock to become smeared due to the necessary reduction in order of accuracy. The gradient test is less computationally expensive than the full streamline construction and reduces the number of potential false positives.

A grid point that has passed the gradient test becomes the anchor for both a forward and a backward streamline. To advance forward, our algorithm moves a small distance, typically a fraction of the grid spacing, in the direction of the local velocity vector. Each flow variable is then computed at the new endpoint via bi- or tri-linear interpolation. In a regular two-dimensional grid, every possible location lies within a grid rectangle, the four corners of which contain known variable values. Since only three points are necessary for linear interpolation, the fourth may be employed to reduce error. Similarly, the eight corners of a cube are more than the four points required for linear interpolation in three dimensions, so, again, error reduction is feasible. Thus the error due to this portion of our shock-detection algorithm is better than linear. In this way a forward streamline emerges incrementally, and receding from the anchor point in the opposite direction of the local velocity vector produces a backward streamline.

Once the streamline is complete, the algorithm searches for the first local maximum of pressure downstream and the first local minimum upstream. The density and pressure ratios between these two points are then tested against the Rankine–Hugoniot jump conditions that hold across a shock. If a subscript R denotes a ratio of like quantities, this condition is

$$\rho_R - \frac{(\gamma + 1) p_R + (\gamma - 1)}{(\gamma - 1) p_R + (\gamma + 1)} = 0 \quad (6.1)$$

in which γ is the specific heat capacity ratio. A deviation smaller than an adjustable threshold value (in our present study 0.5–1.0%) indicates that the anchor point lies within a shocklet. If this is true, the anchor point is flagged and the downstream and upstream points are treated as the numerical boundaries of the shocklet. Note that Eq. (6.1) is also satisfied when both the density and pressure ratios are unity (the trivial case). We have already guarded against this possibility by applying the gradient test described above, and, as added insurance, the algorithm rejects any “shocks” with a numerical width of less than two grid spacings, which is a theoretical minimum.

Figure 4 shows density and pressure profiles along an instantaneous streamline through one of the grid points that our search algorithm flagged as a shocklet in Case 2D4. The distance along the streamline L , relative to the point of interest, is normalized by the grid spacing Δ , which




Fig. 4. Density and pressure profiles along an instantaneous streamline through one of the grid points of Case 2D4 ($N=512$) flagged as a shocklet. Flow is from left to right, and symbols mark the computed numerical boundaries of the shock.

is the same in all Cartesian directions. The computed forward and backward edges of the shock are marked by symbols. We provide a similar plot in Fig. 5 for Case 3D4. If we define shock strength as the incoming normal Mach number in a frame of reference where the shock is stationary, the shocklets in these two examples have strengths of 1.5 and 1.4, respectively. Strengths in our two-dimensional simulations are typically less than 2.0 and in three dimensions less than 1.5.

Because our procedure for detecting shocklets errs on the side of caution, some discarded points may, in fact, constitute especially weak or numerically smeared shocks. However, since the number of points near or within shocklets is quite small compared to the total number of grid points, global statistics concerning the smooth regions are nearly equivalent to global statistics involving the entire domain. Thus, we will later refer to adaptation behavior regarding “shocklets” versus the “aggregate” flow, which substitutes for smooth regions alone.

7. DISTILLING ADAPTATION INFORMATION

The centrality and non-linearity indices of Sec. 4 are not quite yet able to encapsulate local WENO stencil adaptation behavior within two simple numbers. This is because Eq. (3.1), a model equation, arises in varying forms multiple times during the simulation of the full Navier–




Fig. 5. Density and pressure profiles along an instantaneous streamline through one of the grid points of Case 3D4 ($N=128$) flagged as a shocklet. Flow is from left to right, and symbols mark the computed numerical boundaries of the shock.

Stokes Eqs. (2.1)–(2.3). Each conserved variable, such as mass, has an associated governing equation in characteristic space, and each of these in turn contains expressions applying to each cardinal direction separately (in the context of convection). Since our three-dimensional simulations, for example, cover three directions and five characteristic variables, we would face up to fifteen different values apiece for CI and NI. We must distill this information further.

To eliminate directional distinctions, we simulate for each variable an equivalent single CI and NI as if the local flow velocity was exactly aligned with one of the cardinal directions. That is, the indices associated with each direction are weighted by the components of the velocity vector and combined, and the combination is then re-normalized. The resulting values quantify the nature of stencil adaptation that a virtual fluid particle momentarily “feels” rather than just “sees” as it travels through a point in space.

We have computed these modified non-linearity indices, now stripped of their directional information, for each of the five characteristic variables at every grid point of Cases 2D5 ($N=256$) and 3D6 ($N=64$). Table II lists the linear least-squares correlation coefficients for the two NI values associated with all possible pairs of variables, and none of the pairs is strongly correlated or even substantially more correlated than any other. In other words, the degree of stencil adaptation that the WENO method

Table II. Linear Least-Squares Correlation Coefficients R_{mn} for the Non-linearity Indices Associated with all Possible Pairs of Characteristic Variables

	Case	
	2D5	3D6
R_{12}	0.57	0.36
R_{13}	0.57	0.35
R_{14}	N/A	0.36
R_{15}	0.58	0.32
R_{23}	0.56	0.32
R_{24}	N/A	0.32
R_{25}	0.55	0.28
R_{34}	N/A	0.32
R_{35}	0.56	0.28
R_{45}	N/A	0.27

has deemed suitable for one variable is only loosely related over the entire flow field to the degree deemed suitable for another. Because we can discern no basis by which to favor certain variables over others, we choose to form our final CI and NI by simply averaging all five of their intermediate values.

These two simplifications allow each point in a flow field to have exactly one associated CI and one NI. Shocks, however, span multiple grid points due to numerical smearing, so one further distillation step is necessary when considering shocklets. Figures 6 and 7 again show the streamlines of Figs. 4 and 5 but this time with profiles of NI included. NI achieves a distinct maximum immediately upwind of the center of each identified shocklet and thereby indicates that the WENO adaptation mechanism is engaging strongly as it senses the presence of a near-discontinuity ahead. Since this behavior is consistent with the intent of WENO methods in general, we designate both CI and NI at the spatial location of that maximum as the representative values for the shocklet.

8. TWO-DIMENSIONAL RESULTS

Our two-dimensional simulations provide statistically reliable shock data for Cases 2D3 and onward, which correspond to initial $M_t=0.8$ (final $M_t \simeq 0.64$) and greater. As we explained in Sec. 6, we may freely designate two threshold values during the search for shocklets. We choose to restrict our attention to shocks that obey the jump condition of Eq. (6.1) to within 1.0% and then adjust the gradient threshold to such a value that




Fig. 6. Instantaneous streamline of Fig. 4 with the addition of the non-linearity index profile. The symbol marks the value of NI associated with this shocklet.




Fig. 7. Instantaneous streamline of Fig. 5 with the addition of the non-linearity index profile. The symbol marks the value of NI associated with this shocklet.

the total number of flagged shocklets is roughly equal for each simulation run. This helps to mitigate selection bias when comparing shock-related statistics between cases.

Figure 8 displays the probability densities of the non-linearity index for Case 2D4 over the aggregate flow field and also restricted to the sets




Fig. 8. Probability density of the non-linearity index for Case 2D4.

of identified shocklets. We observe that the aggregate distributions peak at significantly larger values of NI with decreasing resolution while the shocklet distributions remain nearly identical. In fact, when $N = 128$, the aggregate distribution begins to substantially overlap the shocklet distribution.

We find similar qualitative results for the remainder of the simulation cases, so we do not include the corresponding figures. Instead, Table III summarizes the quantitative variations by listing the mean values of instantaneous turbulent Mach number M_t , shock strength M_\perp , non-linearity index associated with shocklets NI_S , and non-linearity index concerning the aggregate flow NI_A . Note that M_t is approximately equal for all resolutions of a given case so that overall fluctuation strength has a negligible impact on differences due only to resolution. Unsurprisingly, average shock strength (as captured by the WENO method) increases slightly as M_t rises, but it also increases as resolution becomes more fine. Despite the escalation of shock strength, NI_S decreases under the same transformations. However, the drop in NI_A as resolution increases is far more dramatic, and, unlike NI_S , NI_A rises with increasing M_t .

In Fig. 9, we present a plot analogous to that of Fig. 8 but regarding the centrality index. Note the x -axis values; they are all positive, and this is no mistake. Apparently, the adaptation mechanism of the third-order WENO scheme that we employ strongly favors downwind-centered stencils when it engages. In contrast, Fig. 10 shows that a fourth-order bandwidth-optimized WENO scheme [9,8] leads to overwhelmingly negative values of CI and hence upwind-centered stencils. With neither of the two schemes do the shocklet distributions of CI at various resolutions

Table III. Mean Values of Instantaneous Turbulent Mach Number M_t , Shock Strength M_\perp , Non-linearity Index Concerning the Aggregate Flow NI_A , and Non-linearity Index Associated with Shocklets NI_S for Varying Resolutions N of Select Two-dimensional Cases

N	Case	M_t	M_\perp	NI_S	NI_A
128	2D3	0.32	1.22	1.66	1.11
	2D4	0.73	1.24	1.59	1.15
	2D5	0.83	1.23	1.57	1.18
	2D6	0.93	1.27	1.56	1.20
	2D7	1.00	1.28	1.52	1.20
256	2D3	0.64	1.28	1.72	0.74
	2D4	0.72	1.29	1.59	0.77
	2D5	0.84	1.35	1.64	0.79
	2D6	0.94	1.29	1.44	0.83
	2D7	1.06	1.35	1.45	0.83
512	2D3	0.66	1.35	1.49	0.44
	2D4	0.75	1.40	1.48	0.50
	2D5	0.87	1.53	1.54	0.49
	2D6	0.95	1.37	1.18	0.50
	2D7	1.02	1.46	1.33	0.54

match one another like those of NI in Fig. 8. The peaks of *both* the shocklet and aggregate distributions shift leftward to lesser values of CI as resolution decreases. However, the overlap among the various distributions is significant, especially for the third-order scheme; so these resolution trends for CI are not robust.

9. THREE-DIMENSIONAL RESULTS

Our three-dimensional simulations produce results that are qualitatively identical to those of the previous section. Thus we emphasize the few differences and reiterate only some of the more general agreements. In three dimensions, statistically reliable shock data emerge for Cases 3D3 and onward, which correspond to initial $M_t = 0.6$ (final $M_t \simeq 0.44$) and greater. The jump condition threshold for Eq. (6.1) is now 0.5%, and, as before, we adjust the gradient threshold so that roughly equal numbers of shocklets are identified for each simulation case.

Figure 11 is the three-dimensional analog to Fig. 8 and shows probability densities of the non-linearity index for Case 3D4. Again, the aggregate distributions peak at significantly larger values of NI as resolution decreases while the shocklet distributions remain nearly identical.




Fig. 9. Probability density of the centrality index for Case 2D4 with a third-order-accurate bandwidth-optimized WENO method.

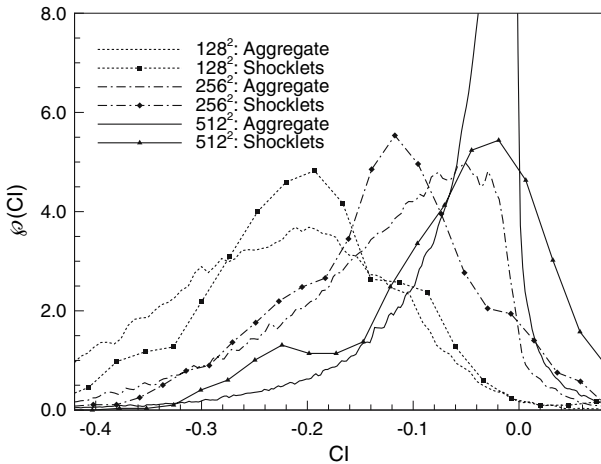


Fig. 10. Probability density of the centrality index for Case 2D4 with a fourth-order-accurate bandwidth-optimized WENO method.

For other cases, Table IV presents, like Table III, the mean values of instantaneous turbulent Mach number M_t , shock strength M_\perp , non-linearity index associated with shocklets NI_S , and non-linearity index concerning the aggregate flow NI_A . The only notable difference between the trends of the two tables is that NI_S rises with increasing M_t for the range




Fig. 11. Probability density of the non-linearity index for Case 3D4.

Table IV. Mean Values of Instantaneous Turbulent Mach Number M_t , Shock Strength M_\perp , Non-linearity Index Concerning the Aggregate Flow NI_A , and Non-linearity Index Associated with Shocklets NI_S for Varying Resolutions N of Select Three-dimensional Cases

N	Case	M_t	M_\perp	NI_S	NI_A
64	3D3	0.43	1.13	1.56	1.39
	3D4	0.49	1.17	1.60	1.40
	3D5	0.55	1.22	1.64	1.40
	3D6	0.61	1.23	1.68	1.40
	3D7	0.65	1.29	1.70	1.41
128	3D3	0.45	1.21	1.50	1.00
	3D4	0.51	1.29	1.57	1.02
	3D5	0.58	1.36	1.65	1.02
	3D6	0.62	1.39	1.65	1.04
	3D7	0.68	1.40	1.61	1.05

of M_t available in the three-dimensional simulations. All other prior observations are unchanged.

The probability densities of the centrality index in three dimensions are displayed in Fig. 12 for the third-order WENO scheme and Fig. 13 for the fourth-order method. As before, the third-order scheme leads entirely to downwind-centered final numerical stencils while the fourth-order scheme leads to upwind-centered stencils. Also, the peaks of both

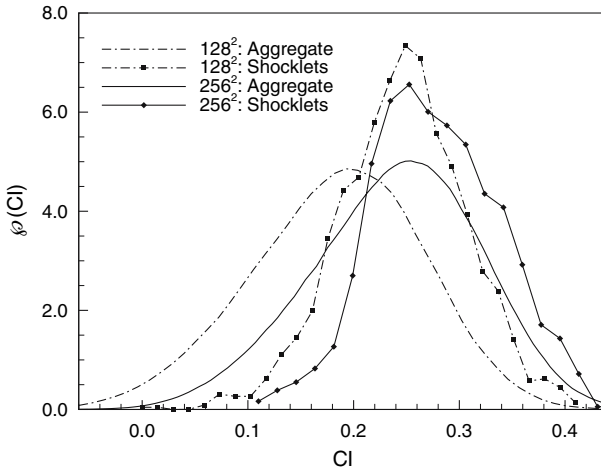


Fig. 12. Probability density of the centrality index for Case 3D4 with a third-order-accurate bandwidth-optimized WENO method.




Fig. 13. Probability density of the centrality index for Case 3D4 with a fourth-order-accurate bandwidth-optimized WENO method.

the aggregate and shocklet distributions shift to lesser values of CI with decreasing resolution. Again, however, significant overlap among the distributions, especially for the third-order method, calls the reliability of these resolution effects into question.

10. CONCLUSIONS

The two- and three-dimensional results produce a number of consistent trends. The mean non-linearity index of the aggregate field NI_A increases with increasing M_t . This leads to the sensible interpretation that stronger turbulent fluctuations induce greater degrees of WENO adaptation. As resolution becomes coarser, calculated shock strengths decrease, NI_A increases sharply, and the mean non-linearity index of identified shocklets NI_S rises slightly but much less significantly than NI_A . The first effect reflects the enhanced ability of a finer grid to faithfully capture near-discontinuities. Additionally, the aggregate probability distribution of NI at the lowest simulated resolution substantially overlaps the shocklet distributions at all resolutions. These resolution trends collectively indicate that on coarser grids the WENO adaptation mechanism misidentifies normal variations as discontinuities and engages the vigorous adaptation intended for shocks. Although over-adaptation deficiencies of WENO methods on coarse grids have previously been inferred in simulations of turbulent boundary layers [7], the present data demonstrate such behavior quantitatively.

Reduction of resolution also causes both the aggregate and shocklet probability distributions of the centrality index to peak at lesser values. This holds for the third-order WENO scheme, which produces overwhelmingly positive CI values, as well as for the fourth-order scheme, which leads to predominantly negative values. Recall that positive CI indicates that the final non-linear numerical stencil is centered downwind of the point of interest and negative CI indicates that the final stencil is centered upwind. Because the focus of this study is to evaluate CI and NI against known or hypothesized numerical behaviors and not to discover new ones, the possible relevance of downwind- versus upwind-centered stencils to the error characteristics of the WENO adaptation mechanism is a subject for future investigations. Nonetheless, the apparent dominance of downwinding in the third-order WENO scheme is troubling. It suggests that this scheme cannot be fully stable, and perhaps the strength of the inherent WENO dissipation is the only reason that instabilities do not appear among computed results.

We therefore conclude that the non-linearity and centrality indices are functional tools not only in theory but also in practice through statistical characterization. Although NI often provides the most readily interpretable information about WENO stencil adaptation, we cannot afford to discount the contributions of CI. For example, an analysis of NI alone would suggest that the adaptation mechanism treats shocklets nearly equally regardless of resolution. However, the corresponding analysis of CI

demonstrates that, despite similar overall degrees of preferential selection, the particular candidate stencils that the mechanism favors (e.g., upwind- or downwind-centered) do vary significantly.

In this paper, we have presented two theoretical indicators of WENO adaptation behavior, a means to procure equivalent values from a realistic three-dimensional fluid simulation, and a general approach to separating information specifically related to shocklets from the aggregate flow. We have conducted this analysis in decaying isotropic turbulence, and the canonical nature of such turbulence implies that more complex turbulent flows should produce qualitatively similar results.

ACKNOWLEDGMENTS

This work was sponsored by the National Science Foundation under Grant CTS-0238390. Computational resources were provided by the CRoCCo Laboratory at Princeton University.

REFERENCES

1. Blaisdell, G. A., and Ristorcelli, J. R. (1996). Consistent initial conditions for the DNS of compressible turbulence. In Proc. 49th Annual Meeting of the Division of Fluid Dynamics of the American Physical Society, November 24–26.
2. Garnier, E., Mossi, M., Sagaut, P., Comte, P., and Deville, M. (1999). On the use of shock-capturing schemes for large-eddy simulations. *J. Comput. Phys.* **153**, 273–311.
3. Jiang, G.-S., and Shu, C.-W. (1996). Efficient implementation of weighted ENO schemes. *J. Comput. Phys.* **126**, 202–228.
4. Lee, S., Lele, S. K., and Moin, P. (1991). Eddy shocklets in decaying compressible turbulence. *Phys. Fluids A* **3**, 657–664.
5. Lele, S. K. (1992). Compact finite-difference schemes with spectral-like resolution. *J. Comput. Phys.* **103**, 16–42.
6. Liu, X.-D., Osher, S., and Chan, T. (1994). Weighted essentially non-oscillatory schemes. *J. Comput. Phys.* **115**, 200–212.
7. Martín, M. P. (2000). Shock-capturing and the LES of high-speed flows. In *Annual Research Briefs*, Center for Turbulence Research, pp. 193–198.
8. Martín, M. P., Taylor, E. M., Wu, M., and Weirs, V. G. (2006). A bandwidth-optimized WENO scheme for the direct numerical simulation of compressible turbulence. *J. Comput. Phys.* in press.
9. Weirs, V. G. (2005). A numerical method for the direct simulation of compressible turbulence. Ph.D thesis, University of Minnesota.
10. Wu, M., Taylor, E. M., and Martín, M. P. (2005). Assessment of STBLI DNS data and comparison against experiments. Paper 2005-4895, American Institute of Aeronautics and Astronautics.