Effect of numerical dissipation on the predicted spectra for compressible turbulence

By J. Larsson, S. K. Lele AND P. Moin

1. Motivation and objectives

The present work is aimed towards a detailed study of shock/turbulence interaction, which is of scientific and engineering interest in a wide variety of fields. Among other issues, there are unanswered questions relating to the shock/turbulence interaction physics at large turbulent Mach numbers (large deflections of the main shock), the interaction with passive scalar fields, and the appropriate subgrid scale modeling in large-eddy simulations (LES).

In addition to these more physical issues, the question of how best to numerically handle shock/turbulence interaction is still somewhat open. The core issue is the contradictory requirements of accuracy and stability, coupled with the fact that stability is often ensured by the addition of numerical dissipation. Even in the absense of shockwaves, the discretized compressible Navier-Stokes equations are sensitive to numerical instabilities, often more so than their incompressible counterparts. This is especially true when the convective terms are discretized in conservative form, for example as

$$\frac{\partial \rho u_i u_j}{\partial x_j}\,.$$

It is well-known from numerical experiments that the stability can be improved substantially by use of alternative forms of the convective terms. One example is the splitting by Ducros *et al.* (2000), who proposed discretization of the Euler equations on the form

$$\begin{aligned}
\frac{\partial \rho}{\partial t} &= -\frac{1}{2}\left(\frac{\partial \rho u_j}{\partial x_j} + \rho\frac{\partial u_j}{\partial x_j} + u_j\frac{\partial \rho}{\partial x_j}\right)\,,\\
\frac{\partial \rho u_i}{\partial t} &= -\frac{1}{2}\left(\frac{\partial \rho u_i u_j}{\partial x_j} + \rho u_i\frac{\partial u_j}{\partial x_j} + u_j\frac{\partial \rho u_i}{\partial x_j}\right) - \frac{\partial p}{\partial x_i}\,,\\
\frac{\partial \rho e_0}{\partial t} &= -\frac{1}{2}\left(\frac{\partial \rho h_0 u_j}{\partial x_j} + \rho h_0\frac{\partial u_j}{\partial x_j} + u_j\frac{\partial \rho h_0}{\partial x_j}\right)\,,
\end{aligned} \tag{1.1}$$

which was later shown by Honein & Moin (2005) to lead to a rather robust numerical method in their numerical experiments on under-resolved compressible turbulence. Here ρ, u_i, p, e_0 and h_0 are the density, velocity, pressure, total energy and total enthalpy, respectively.

When shockwaves enter the solution, the use of shock-capturing schemes becomes necessary. Shock-capturing traditionally relies on two essential ingredients. First, a conservative form of the equation is used in order to ensure convergence to the correct weak solution. Secondly, additional numerical dissipation in the vicinity of the shock is needed.

A common and popular numerical method for compressible turbulence (e.g., Rizzetta *et al.* 2001; Cook 2007) is thus to use the conservative form everywhere in the domain, to capture shockwaves by adding some nonlinear (solution-adaptive) dissipation locally,

and to ensure stability away from the shockwaves by adding some linear dissipation everywhere in the domain. The idea of the linear dissipation is to remove energy from the highest (and marginally resolved) wavenumbers only. The linear dissipation can take several different forms, including eddy- and hyper-viscosities acting on $\partial^n u_i/\partial x_j^n$ for $n \geq 2$ or the application of a filter at the end of each time step.

An alternative approach is the class of hybrid methods (e.g., Adams & Shariff 1996; Pirozzoli 2002), where different numerical methods are used around and away from the shockwaves, respectively. This allows for the use of split forms like (1.1) away from the shockwaves, and the increased robustness of the splitting implies a lower (perhaps non-existent) need for stabilizing linear dissipation. The overall dissipation is also lowered since the more dissipative shock-capturing scheme is only used locally.

Hybrid methods thus offer the potential for shock/turbulence calculations with substantially lower dissipation in the system. Since there are additional complications with the hybrid approach, most obviously the question of stability and conservation at the interface, it is worthwhile to ask what the effect of the numerical dissipation is on the computed statistics and spectra for compressible turbulence. In other words, is the benefit of lower dissipation worth the price of additional algorithmic complexity? This can be analyzed analytically for linear cases, but the effect for the full Navier-Stokes equations is less clear. The objective of this paper is to illuminate this effect.

2. Numerical method

A summary of the methods used is listed in Table 1, along with their abbreviations. The baseline method (S2, S8) uses central difference schemes of second- or eighth-order accuracy to discretize the convective terms on the split form (1.1). The viscous terms are discretized using eighth-order accurate central schemes on non-conservative form (cf. Honein & Moin 2005) to increase the viscous dissipation at the highest wavenumbers. The fluid is assumed to be a calorically perfect gas with

$$\begin{aligned} p &= \rho R T \,, \\ R &= c_p - c_v = \text{constant} \,, \\ \gamma &= \frac{c_p}{c_v} = \text{constant} \,, \end{aligned}$$

and the viscosity is assumed to follow a power-law

$$\mu = \mu_{\text{ref}} \left(\frac{T}{T_{\text{ref}}} \right)^{0.75} .$$

Here T is the temperature, μ is the viscosity, and γ and R are gas constants.

To capture shockwaves, a seventh-order accurate weighted essentially non-oscillatory (WENO, cf. Jiang & Shu 1996) scheme is used locally around the shockwaves, together with Roe flux splitting for minimal numerical dissipation (methods H2 and H8 in the table). To determine in which regions the WENO scheme is to be used, the sensor

$$s = \frac{-\theta}{|\theta| + \sqrt{\omega_i \omega_i}}$$

is computed at every time step, where $\theta = \partial u_j/\partial x_j$ is the rate of dilatation and ω_i is the vorticity. Regions where $s > 0.65$ are solved with the WENO scheme, whereas the central scheme is used elsewhere.

Code	Style	Convective flux scheme
S2	$\cdots$	2nd order central on split form
S8	——	8th order central on split form
S8D8	$-\cdot-$	S8 with 8th order linear dissipation
W7	$--$	7th order WENO with Roe flux splitting
C2	$+$	2nd order central on conservative form
C8	$\times$	8th order central on conservative form
H2	$\cdots$	S2 with local W7 around shocklets
H8	——	S8 with local W7 around shocklets
H8D8	$-\cdot-$	S8D8 with local W7 around shocklets

TABLE 1. Summary of numerical methods. Note that the local WENO is never activated for the Taylor-Green vortex, making H2, H8, H8D8 equivalent to S2, S8, S8D8 for this problem.

In addition to the baseline method outlined above, some variations are also considered. The method W7 uses the seventh-order WENO scheme everywhere in the domain, and is representative of approaches relying solely on a shock-capturing scheme. Indeed, WENO methods have been used for shock/turbulence interaction problems in the literature (e.g., Martin *et al.* 2006).

To assess the effect of the linear dissipation that is commonly added for stability purposes (e.g., as a filter), an eighth-order dissipation term of the form

$$\psi\lambda h^7\left(\frac{\partial^8 q}{\partial x^8}+\frac{\partial^8 q}{\partial y^8}+\frac{\partial^8 q}{\partial z^8}\right)$$

is added in methods S8D8 and H8D8, where $\lambda=\sqrt{u_iu_i}+c$ is the largest eigenvalue magnitude, h is the grid spacing, and $q=\rho,\rho u_i,\rho e_0$ are each of the conserved variables. The coefficient ψ is set to 0.002, which is about one-third of the stability limit. This dissipation is highly biased towards the largest wavenumbers, with very little direct effect (in the linear sense) on the largest scales.

Finally, the convective terms are discretized in the fully conservative form (C2 and C8) simply to illustrate the issue of numerical instability associated with this form.

For all cases, the semi-discrete system is integrated in time using classical four-stage Runge-Kutta. While the temporal discretization of course contributes to the numerical dissipation, this effect is not considered in this work.

3. Results: the inviscid Taylor-Green vortex

To assess the effect of dissipation in the absense of shockwaves, the inviscid Taylor-Green vortex is considered. From a well-resolved initial condition, the vortex begins stretching and producing ever smaller scales. It thus constitutes a non-regularized problem with no lower bound on the length scale, and is solved here without any regularization

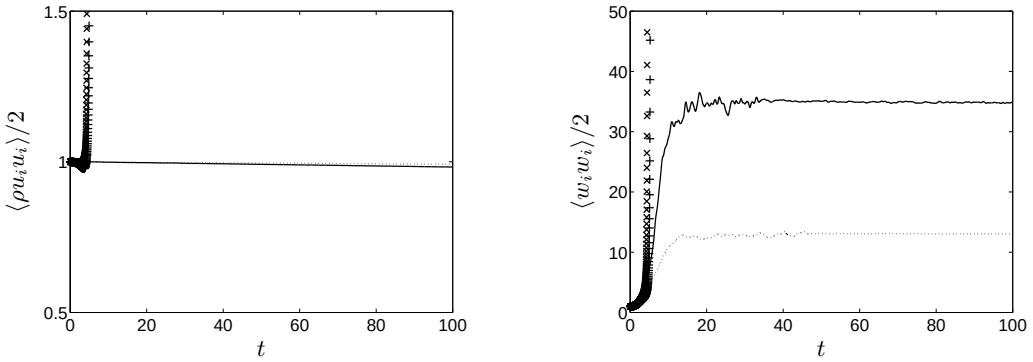

(a) Kinetic energy. (b) Enstrophy.

FIGURE 1. Inviscid Taylor-Green vortex on 32^3 grid, quantities normalized by their initial values. Split forms of eighth- (S8, solid) and second- (S2, dotted) order accuracy, conservative forms of eighth- (C8, cross) and second- (C2, plus) order accuracy.

other than the numerical method. The initial condition is

$$\begin{aligned} \rho &= 1\,, \\ u &= \sin x \cos y \cos z\,, \\ v &= -\cos x \sin y \cos z\,, \\ w &= 0\,, \\ p &= 100 + \frac{(\cos(2z)+2)(\cos(2x)+\cos(2y))-2}{16}\,, \end{aligned}$$

with $\gamma = 5/3$ on a triply periodic domain $[0, 2\pi)^3$. The mean pressure is chosen sufficiently high to make the problem essentially incompressible, thus allowing for comparison to the semi-analytical results for the enstrophy growth by Brachet *et al.* (1983).

Consider first the influence on stability by the form of the convective terms. Figure 1 shows the temporal evolution for some different schemes. The conservative forms quickly lead to unbounded growth and divergence, whereas the split form maintains stability even for long times. The velocity spectrum at $t = 100$ is equipartioned with $E \sim k^2$, which is the correct spectrum for this unregularized inviscid problem.

Now consider the influence of the dissipation on this problem. The same quantities are shown in Fig. 2 for a different combination of schemes for one grid. Note that the results with different grids are very similar.

The kinetic energy decays slowly with time for the split forms – for this nearly incompressible case, one should not expect perfect conservation of this quantity which is only conserved for infinite Mach number. The seventh-order WENO scheme is certainly stable, but at the expense of a large dissipation that causes under-prediction of both the kinetic energy and the enstrophy for $t \gtrsim 3$. The result with the eighth-order dissipation is somewhat better than, but essentially similar to, the WENO result. Despite the rather high order of the dissipation term, which has a k^8 wavenumber scaling, the term quickly dissipates most of the energy.

Brachet *et al.* (1983) used a semi-analytical method to find the enstrophy growth for $t \leq 4$, where $t \leq 3.5$ was considered well converged. The eighth-order central scheme produces enstrophy growth very similar to the results of Brachet *et al.* (1983), and continues to generate very large enstrophy until it levels off when the solution reaches an equipartioned

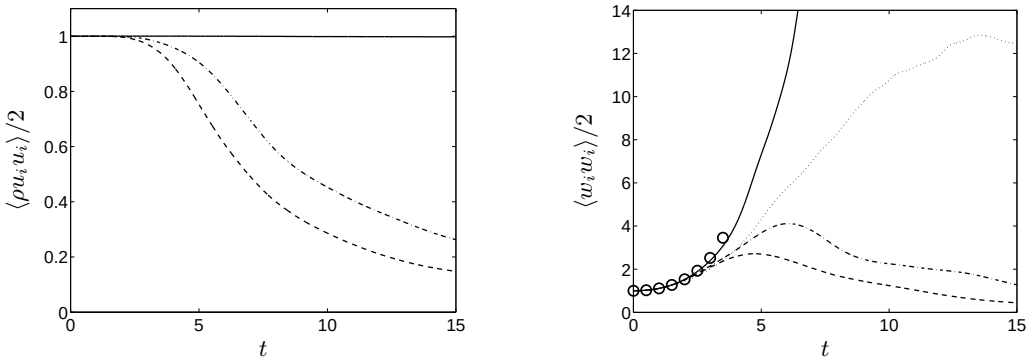

(a) Kinetic energy. (b) Enstrophy.

FIGURE 2. Inviscid Taylor-Green vortex on 32^3 grid, quantities normalized by their initial values. Split forms of eighth- (S8, solid) and second- (S2, dotted) order accuracy, with eighth-order dissipation (S8D8, dash-dotted), and seventh-order WENO (W7, dashed). Semi-analytical results of Brachet *et al.* (1983) (circles).

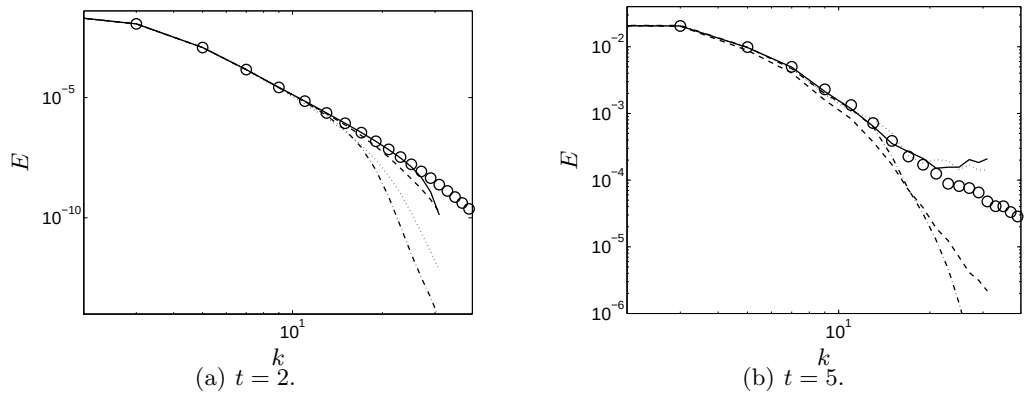

(a) $t = 2$. (b) $t = 5$.

FIGURE 3. Velocity spectra of inviscid Taylor-Green vortex on 64^3 grid. Split forms of eighth- (S8, solid) and second- (S2, dotted) order accuracy, with eighth-order dissipation (S8D8, dash–dotted), and seventh-order WENO (W7, dashed). Converged spectrum on 256^3 (circles).

equilibrium (Fig. 1). The second-order scheme is qualitatively similar, but of course with lower enstrophy for the under-resolved fields. Note that the second-order scheme was used to compute the enstrophy for the plots (thus the difference between S2 and S8 is due both to different instantaneous fields and different post-processing numerics).

Of greater interest are the velocity spectra, which are shown in Fig. 3 at two different times. For reference, the spectrum is first computed on 128^3 and 256^3 grids with and without the eighth-order dissipation, and it is found that the spectrum converges for the wavenumbers and times of interest here ($k \leq 32$, $t \leq 5$). This converged spectrum is used as reference.

At $t = 2$ the field is still quite well resolved, and the kinetic energy has not yet begun to decrease for the dissipative methods. At this time the spectra from the seventh-order WENO and the eighth-order central schemes are similar, with both schemes yielding well-resolved spectra up to high wavenumbers. The second-order central scheme predicts less energy at the highest wavenumbers, indicating a somewhat lower nonlinear transport

of energy in the cascade. The eight-order dissipation causes lower spectral density for $k \gtrsim 16 = N/4 = k_{\max}/2$. Thus, despite the high order of the term and the fact that the solution is still quite well resolved, the dissipation causes a measurable effect on the upper half of the spectrum. The fact that the WENO method is similar to the eighth-order central scheme at this time is due to the well-resolved solution – the WENO adaptation mechanism results in the optimal scheme being used primarily.

At $t = 5$ the fields are no longer well resolved, which corresponds more closely to the case of under-resolved turbulence in LES. Both purely central schemes now have a pile-up of energy for $k \gtrsim 2k_{\max}/3$, which is the range directly affected by aliasing errors. The linear dissipation yields a similar result as at $t = 2$, with under-prediction for $k \gtrsim k_{\max}/2$. Most interesting is the WENO result, with an under-predicted spectrum for a large range of scales, say $k \gtrsim k_{\max}/3$. This shows that the adaptation of the stencils that WENO performs in order to capture discontinuities is in full force, leading to dissipation over a large range of scales. Since seventh-order WENO is equivalent to an eighth-order central scheme with an eighth-order dissipation in its optimal stencil, it is clear that the optimal stencil says little about the performance of WENO on under-resolved realistic flowfields.

Overall, the non-dissipative central schemes that rely on the split convective form rather than dissipation for stability are accurate for about twice as large wavenumbers as WENO, and 50% higher wavenumbers than with the linear dissipation. In addition to being of interest for the direct purpose of accuracy, this also has implications for subgrid scale modeling in LES. Even the gentle eighth-order dissipation is enough to remove the need for any explicit subgrid scale model, and indeed the results would likely get worse should one be added.

4. Results: isotropic turbulence

To assess the effect of numerical dissipation in a more realistic setting, isotropic decaying turbulence is considered next. The relevant parameters are the turbulent Mach number M_t and the Reynolds number Re_λ based on the Taylor length scale λ defined as

$$M_t = \frac{\sqrt{3}\, u_{\rm rms}}{\langle c \rangle}\,,$$
$$Re_\lambda = \frac{\langle \rho \rangle u_{\rm rms} \lambda}{\langle \mu \rangle}\,,$$

where

$$u_{\rm rms} = \sqrt{\frac{\langle u_i u_i \rangle}{3}}\,,$$
$$\lambda = \frac{1}{3}\left(\frac{u^2}{u_x^2} + \frac{v^2}{v_y^2} + \frac{w^2}{w_z^2}\right)\,,$$
$$c = \sqrt{\gamma R T}\,.$$

The initial density and pressure fields are taken as constants ρ_0 and p_0, leading to constant initial temperature T_0 and speed-of-sound c_0. The initial velocity field is randomized under the constraints

$$\frac{\langle u_i u_i \rangle}{2} = \frac{3u_0^2}{2} = \int_0^\infty E(k) dk\,,$$
$$E(k) \sim k^4 \exp\left(-2(k/k_0)^2\right)\,,$$

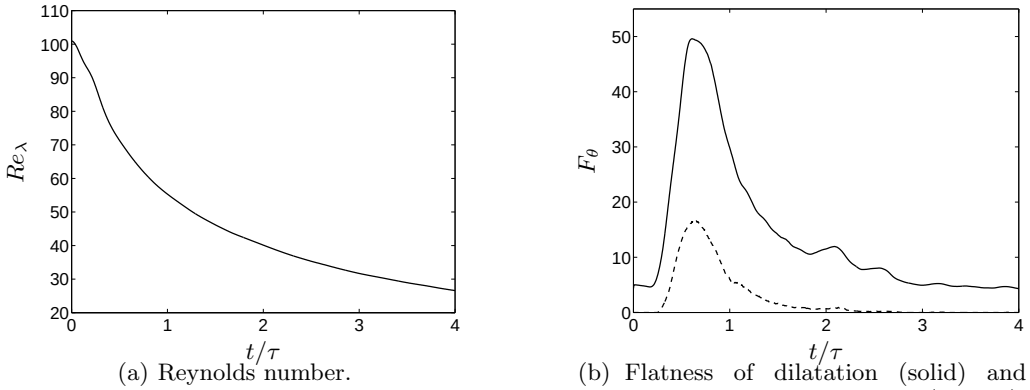

(a) Reynolds number.

(b) Flatness of dilatation (solid) and 1000× fraction of WENO points (dashed).

FIGURE 4. Isotropic turbulence on 256^3 grid using the hybrid central/WENO method (H8).

with $k_0 = 4$. The initial parameters are taken as $Re_{\lambda,0} = 100$ and $M_{t,0} = 0.6$, and $\gamma = 1.4$ and $Pr = 0.7$. The eddy-turnover time is $\tau = \lambda_0/u_0$, and the initial Taylor scale is $\lambda_0 = 2/k_0$. This combination of parameters is such that the turbulence spontaneously generates shocklets, and thus a shock-capturing scheme is necessary on realistic grids. Thus the combination of methods considered here is H8 as the baseline; W7 to represent the use of shock-capturing everywhere; H8D8 to represent methods that add linear dissipation for stability; and H2 to assess the effect of the order of accuracy of the convective terms. The purpose of including the H2 method here is to gain some insight into the importance of the order of accuracy for this problem, with a view towards future implementations in lower-order unstructured solvers.

The initial field is generated on a 256^3 grid and run up to $t/\tau = 4$. The development of two important quantities are shown in Fig. 4. The flatness factor F_θ (cf. Pope 2000) of the dilatation θ is a good measure of the presence of shocklets (Lee *et al.* 1991), and this begins rising rapidly at $t/\tau \approx 0.2$ and peaks at $t/\tau \approx 0.6$. The fraction of WENO points in the domain of the hybrid method is shown in the same plot, and this curve closely resembles F_θ. Note that the hybrid method has fewer than 1% of the grid points treated by WENO during most of the run. To avoid the most severe initial transient on the coarser grids, the field is spectrally filtered to coarser resolutions at $t/\tau = 0.8$, and this is used as the initial condition for all subsequent runs. Thus the Reynolds number during the time of interest $0.8 \leq t/\tau \leq 4$ is roughly $60 \gtrsim Re_\lambda \gtrsim 25$.

Figure 5 shows some quantities for a 64^3 grid with the different schemes. For comparison the results of the 256^3 grid are first filtered to 64^3 resolution, and then the statistics are computed using the eighth-order central scheme. Only the WENO method underpredicts the kinetic energy, whereas the remaining methods agree well with the filtered reference data. The enstrophy is more dependent on the smaller scales and is thus a more sensitive quantity. Here the hybrid method agrees well with the filtered data, whereas the eighth-order dissipation causes a measurable under-prediction. The pure WENO method is even more dissipative. The skewness and flatness of dilatation are less sensitive to the numerical dissipation, with all methods generating reasonable results. Note that the flatness factors are everywhere larger than 3, the value for a Gaussian distribution. This shows that the velocity fields do carry some dilatational structure, and it can be inferred that there are some weak shocklets, at least for the earliest times. The skewness of the

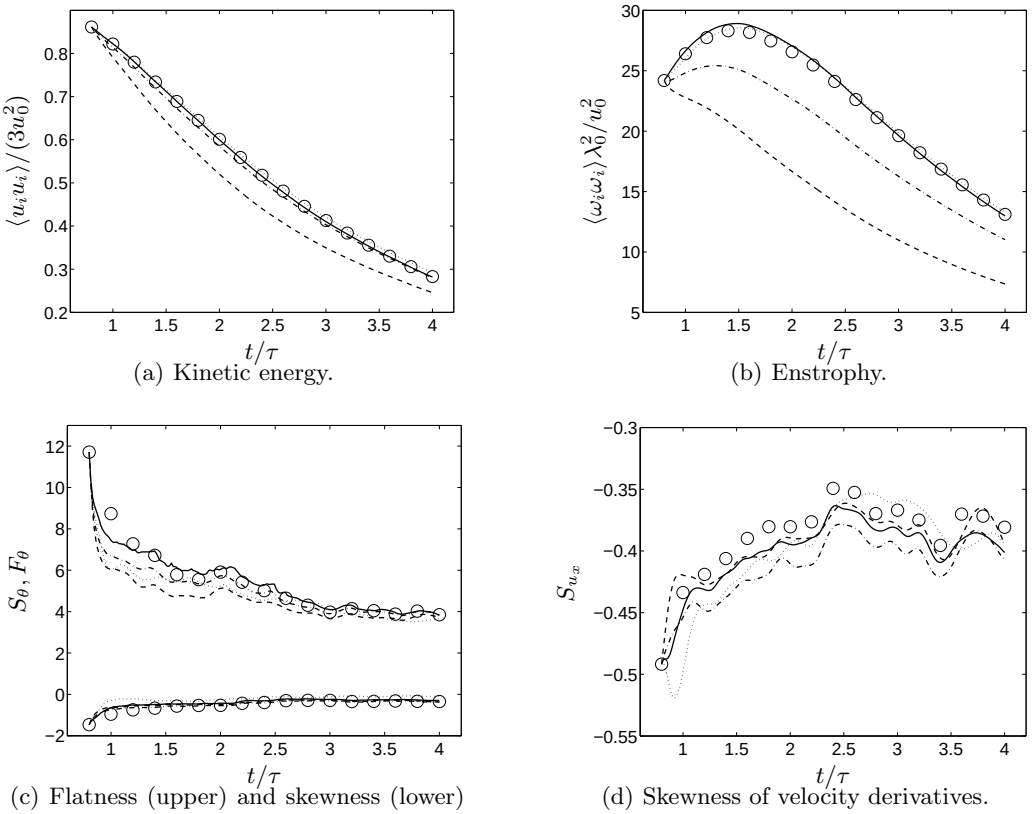

(a) Kinetic energy.

(b) Enstrophy.

(c) Flatness (upper) and skewness (lower) of dilatation.

(d) Skewness of velocity derivatives.

FIGURE 5. Isotropic turbulence on 64^3 grid, starting from filtered field at $t/\tau = 0.8$. Hybrid methods with eighth- (H8, solid) and second- (H2, dotted) order central schemes, with eighth-order dissipation (H8D8, dash-dotted), and seventh-order WENO (W7, dashed). Filtered 256^3 results (circles).

(normal) velocity derivatives tells a similar story, with all methods agreeing with the reference data. Thus it appears that the structure of the turbulent cascade is reasonably predicted with all schemes.

Interestingly, replacing the eighth-order central scheme in the hybrid method with its second-order counterpart hardly affects the results. This suggests that numerical dissipation is a more critical issue than the resolving power of the scheme, at least for this particular problem.

Figure 6 shows the spectra of velocity and vorticity on different grids comparing the four methods. As reference the spectra from a 256^3 grid is used, at which resolution the spectra have converged for the wavenumbers of interest. Note that there is no real inertial range, likely partly due to the low Reynolds number (between 25 and 60, see Fig. 4) and partly due to the short time of the simulation (giving insufficient time for the spectral cascade to develop).

The 64^3 and 128^3 spectra tell a similar story, in that the hybrid method yields spectra that begin deviating downwards for $k \gtrsim 2k_{\max}/3$ whereas the pure WENO spectra begin deviating downwards for $k \gtrsim k_{\max}/3$ or so – a doubling of the cut-off length

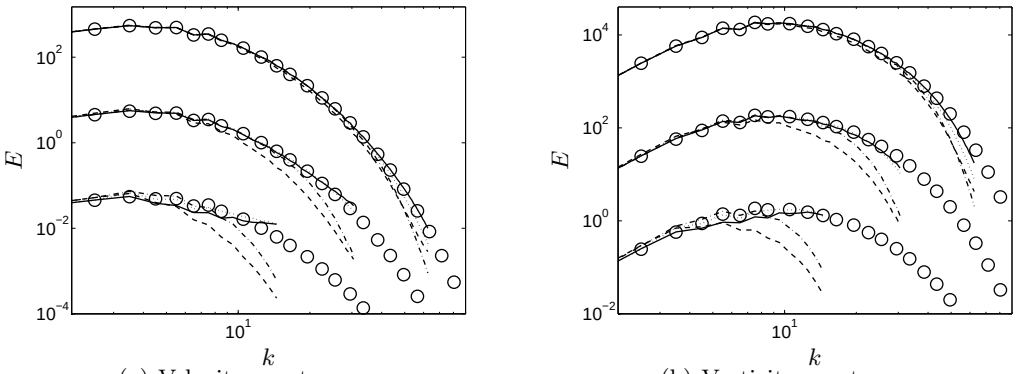

(a) Velocity spectrum. (b) Vorticity spectrum.

FIGURE 6. Spectra of isotropic turbulence at $t/\tau = 4$. Hybrid methods with eighth- (H8, solid) and second- (H2, dotted) order central schemes, with eighth-order dissipation (H8D8, dash-dotted), and seventh-order WENO (W7, dashed). Converged spectrum on 256^3 grid (circles). Grids from bottom to top (with offset): 32^3, 64^3 and 128^3.

scale. The eighth-order dissipation is somewhere in between, with under-predictions for $k \gtrsim k_{\max}/2$. There is little pile-up at the highest wavenumbers even without the added dissipation, showing that the local use of WENO around the shocklets provides sufficient dissipation to avoid the pile-up (at least for this particular problem). The fact that the WENO dissipation affects a much larger range of scales compared to the eighth-order dissipation, despite it being equivalent in its optimal stencil, provides further evidence that linear analysis provides at best incomplete information about the behavior on under-resolved turbulence. Also, the eighth-order dissipation affects the upper half of the range of wavenumbers, despite it being of formally high order.

Neither method yields very good results on the coarsest grid, although WENO is still clearly too dissipative. Interestingly, both the added dissipation and the lower-order convective scheme somewhat improve the spectra on this grid. One hypothesis that could explain this is that the higher-order scheme over-predicts the spectral transfer of energy towards the smallest scales, since it yields too low energy at the intermediate scales. The hypothesis is then that the dissipation and the lower-order convection decreases this transfer, thereby increasing the intermediate scale energy.

On the finer grids the second-order scheme is similar to the eighth-order one, except for a slightly smaller energy near $k_{\max}$. Again, the numerical dissipation has a greater effect on the solution than the accuracy of the convective scheme for this problem.

5. Summary

The effect of numerical dissipation on the accuracy of predicted compressible turbulence has been assessed using two canonical problems. The base method uses eighth-order accurate central difference schemes to treat the convective and viscous terms, and relies on a split form of the convective terms for stability. For flows with shockwaves, this base method uses a seventh-order accurate WENO scheme locally around the shockwaves. Since the shockwaves typically occupy less than 1% of the grid points, this implies that the resulting hybrid method is minimally dissipative.

To assess the effect of minimizing the dissipation in this manner, the results are compared to two different methods. First, the seventh-order WENO method is used for the

whole flowfield. This is roughly representative of approaches where a shock-capturing scheme is applied in the whole domain, as was indeed the original motivation of ENO and WENO. The second method is the base hybrid method with an eighth-order dissipation term added. This dissipation is highly skewed towards the largest wavenumbers, and is in some sense representative of commonly used methods to maintain stability away from shockwaves, such as filtering of the solution after each time step.

The tests show two things. First, that the added numerical dissipation corrupts a larger range of wavenumbers than are corrupted by aliasing errors in the base method. This has an effect especially on the predicted spectra and on statistical quantities that involve derivatives, like the enstrophy. The second point is that the behavior of WENO is different from what would be inferred from linear analysis. A linearization of both the underlying equations and the WENO method itself leads to the WENO method adopting its optimal stencil, which in the present case is an eighth-order central scheme with an eighth-order dissipation term. While the magnitude of the dissipation term implied by WENO would be different from what was used here, the fact that the present results show a large difference between WENO and the "central plus dissipation" method indicate that WENO behaves in ways that are more complex than linear analysis can predict. This was also concluded by Pirozzoli (2006), who analyzed shock-capturing schemes in an expanded quasi-linear framework.

A secondary comparison is made by replacing the eighth-order convective scheme by a second-order one, keeping everything else the same (including the viscous terms). This is done as a crude test of what could be expected in future implementations in unstructured codes, and the results show that the lower order yields very similar results for the problems considered here.

Overall it can be concluded that the property of minimal dissipation is valuable and worth pursuing in a general method for shock/turbulence interactions. This is especially true for LES, where the flowfields are necessarily under-resolved. Assessments of subgrid scale models using dissipative numerical methods are questionable, since in all cases studied here the numerical dissipation is large enough to remove the small scale motions – thus the "best" model for dissipative numerics is likely none at all. Only with minimally dissipative numerics can subgrid scale models be studied in an adequate framework.

In conclusion, the advantages of minimally dissipative numerics are well worth the "cost" that arises in the form of additional complications in a hybrid method, like the need to interface different numerical methods and the need to accurately find regions of shockwaves. While these issues are non-trivial, progress is being made. Pirozzoli (2002) showed how to interface the methods in a conservative manner, while Larsson & Gustafsson elsewhere in this volume present an analysis of stability of the coupled method.

There is much remaining work to be done. The presently used shock sensor works well for traditional shock/turbulence interaction, but does not find contact discontinuities and fails entirely for vorticity-free problems. The fact that the computational cost of WENO is several times greater than the base cost of the method has the advantage of making the hybrid method very fast (since WENO is used in small regions only), but also the drawback of making load balancing non-trivial in a parallel setting. Finally, the hybrid approach should be implemented in an unstructured framework, both to allow for more general geometries to be considered and to make use of the ability to perform local grid refinement around the shocks.

Acknowledgments

Financial support has been provided by the Department of Energy Scientific Discovery through Advanced Computing (SciDAC) program and the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

Adams, N. A. & Shariff, K. 1996 A high-resolution hybrid compact-ENO scheme for shock-turbulence interaction problems. *J. Comput. Phys.* **127**, 27–51.

Brachet, M. E., Meiron, D. I., Orszag, S. A., Nickel, B. G., Morf, R. H. & Frisch, U. 1983 Small-scale structure of the Taylor-Green vortex. *J. Fluid Mech.* **130**, 411–452.

Cook, A. W. 2007 Artificial fluid properties for large-eddy simulation of compressible turbulent mixing. *Phys. Fluids* **19**, 055103.

Ducros, F., Laporte, F., Souleres, T., Guinot, V., Moinat, P. & Caruelle, B. 2000 High-order fluxes for conservative skew-symmetric-like schemes in structured meshes: application to compressible flows. *J. Comput. Phys.* **161**, 114–139.

Honein, A. E. & Moin, P. 2005 Numerical aspects of compressible turbulence simulations. Report TF-92. Department of Mechanical Engineering, Stanford University.

Jiang, G.-S. & Shu, C.-W. 1996 Efficient implementation of weighted ENO schemes. *J. Comput. Phys.* **126**, 202–228.

Lee, S., Lele, S. K. & Moin, P. 1991 Eddy shocklets in decaying compressible turbulence. *Phys. Fluids* **3** (4), 657–664.

Martin, M. P., Taylor, E. M., Wu, M. & Weirs, V. G. 2006 A bandwidth-optimized WENO scheme for the effective direct numerical simulation of compressible turbulence. *J. Comput. Phys.* **220**, 270–289.

Pirozzoli, S. 2002 Conservative hybrid compact-WENO schemes for shock-turbulence interaction. *J. Comput. Phys.* **178**, 81–117.

Pirozzoli, S. 2006 On the spectral properties of shock-capturing schemes. *J. Comput. Phys.* **219**, 489–497.

Pope, S. B. 2000 *Turbulent Flows*. Cambridge University Press.

Rizzetta, D. P., Visbal, M. R. & Gaitonde, D. V. 2001 Large-eddy simulation of supersonic compression-ramp flow by high-order method. *AIAA J.* **39** (12), 2283–2292.